

Mathematical modelling and image processing for some challenges in the In Vitro Fertilisation clinic



Timothy Robert Webb Ostler

School of Mathematics

Cardiff University

A thesis submitted for the degree of

Doctor of Philosophy

3rd November 2023

Dedicated to Rebecca Wallace.

Acknowledgements

This project could not have taken place without financial support from KESS2, and from my industrial partner, the London Women's Clinic (LWC). KESS2 is a pan-Wales higher level skills initiative led by Bangor University on behalf of the HE sector in Wales. It is part funded by the Welsh Government's European Social Fund (ESF) convergence programme for West Wales and the Valleys.

First and foremost, I wish to thank my supervisors, Dr Thomas Woolley and Dr Katerina Kaouri for their infectious enthusiasm, endless ideas and freely shared advice throughout the last four years. Beyond the work you've invested to help me develop as a mathematician, the kindness and humanity with which you approached this project have made it a privilege to be your student.

I would also like to thank Prof. Karl Swann, for the insightful discussions and guidance which have driven this project. I also want to thank my supervisors at LWC over the years, Dr Helen Priddle, Andrew Thomson and Giles Palmer. You were all a joy to work with, and always provided invaluable insight into matters far beyond my understanding. Additionally, I want to thank Prof. Wolfgang Langbein, Prof. Paola Borri, Dr Yisu Wang and Emily Lewis for gathering data which has been a core part of this work.

To the PhD students in the School of Mathematics, I extend my heartfelt gratitude. The generously shared wisdom of those who came before me, and the excitement and energy of those who followed, have guided and motivated me along the way. Though research can be lonely, I could always count on the Procrastination Squad to cheer me up and cheer me on. It would take a separate thesis to mention you all by name, but my special thanks go to Josh, Lizzie, Michela¹, Matthew and Layla, for listening to me moan about how tired I am, enabling my addiction to tea, and always being up for Greggs.

My journey into mathematics could not have been possible without the unconditional love and support of my parents, who only asked that I do what makes me happy. I also want to thank my witty and wonderful brother, James, who only ever asked that I do what makes me lots of money, and is happy I am leaving academia to work in finance.

¹Told you that you'd get a mention.

Last but certainly not least, I want to thank Rebecca, who has been my foundation. Sharing the highs and lows of PhD life together has made me truly happy, and I look forward to seeing what life has in store for us both.

Abstract

In Vitro Fertilisation (IVF) is a treatment involving the fertilisation of human egg cells in the laboratory to create embryos, which are transferred to the uterus of the patient in the hope that they become pregnant. Only 19% of treatments worldwide result in a live birth. This thesis involves interdisciplinary collaboration with experimentalists and the London Women's Clinic, to identify and tackle a series of mathematical challenges and barriers to improving success rates in IVF clinics.

We first explore the use of Differential Dynamic Microscopy (DDM) as a non-invasive oocyte health assessment tool. Identifying challenges that limit the clinical usefulness of DDM, we develop a methodology, using synthetic data, to enhance and validate parameter fitting in DDM. We optimise the selection of synthetic data parameters, and present a new pipeline for generating parameter fitting. After showing existing non-linear curve fitting algorithms are inaccurate in DDM applications, we establish a new machine learning parameter fitting pipeline, trained exclusively on synthetic data and applied in real datasets.

We, subsequently, explore the application of DDM to phase-contrast microscopy. Phase-contrast images exhibit shadowing, leading to anisotropy in the DDM matrix and invalidating a key assumption of DDM. We derive an analytic expression describing this anisotropy, and conclude for isotropic motions that shadowing does not affect parameter fitting. For anisotropic motion, we also outline conditions on the microscope setup and imaged behaviour that affect fitting error.

The second part of this thesis considers challenges related to cryopreservation of oocytes and embryos through a rapid-cooling technique, vitrification. We numerically simulate the process of vitrification (rapid freezing) and show that cooling rates are unaffected by the number or arrangement of embryos or oocytes on the device, which validates current protocol. Additionally, the challenge of predicting embryo viability from time-lapse images of post-thaw embryo re-expansion is tackled. We implement machine-learning image segmentation to measure the cross-sectional area of the embryo, and identify non-linear re-expansion as a new

metric indicating lower implantation rates by as much as 3% in a sample of clinical data.

List of Acronyms

ART = Assisted Reproductive Technologies

IVF = In Vitro Fertilisation

UK = United Kingdom

LWC = London Women's Clinic

COS = Controlled ovarian stimulation

ICSI = Intra-cytoplasmic sperm injection

ESHRE = The European Society of Human Reproduction

FET = Frozen embryo transfers

HFEA = The Human Fertilisation & Embryology Authority

NICE = The National Institute for Health and Care Excellence

OHSS = Ovarian hyper-stimulation Syndrome

SPT = Single particle tracking

PIV = Particle image velocimetry

DDM = Differential dynamic microscopy

eSET = Elective single embryo transfer

ICM = Inner cell mass

CPA = Cryoprotective agent

ISF = Intermediate Scattering Function

LSI = Linear space-invariant

PSF = Point spread function

Contents

1	Introduction	1
1.1	Overview of In Vitro Fertilisation (IVF)	2
1.1.1	Cryopreservation (oocyte or embryo freezing)	3
1.1.2	Time-lapse imaging	4
1.1.3	The need for improving IVF	5
1.2	Mathematical challenges in IVF	9
1.2.1	Personalising controlled ovarian stimulation (COS)	9
1.2.2	Health assessment and selection in IVF	11
1.2.2.1	Oocyte assessment	11
1.2.3	Sperm selection	18
1.2.4	Embryo selection	19
1.2.5	Optimal cryopreservation protocols	23
1.2.6	Post-thaw embryo assessment	24
1.2.7	IVF logistics, planning and optimisation	27
1.3	Thesis outline	29
2	Differential Dynamic Microscopy (DDM)	30
2.1	Introduction	30
2.1.1	Time-lapse images of mouse oocytes	31
2.1.2	Image analysis algorithms for quantifying movement parameters	32
2.1.2.1	SPT	32
2.1.2.2	PIV	33
2.1.2.3	DDM	34
2.2	DDM Theory	35
2.2.1	Inferring particle movement from image data	35
2.2.2	The DDM tensor, $\mathcal{D}(\mathbf{q}, \Delta t)$	42
2.2.3	The radial average of the DDM tensor	44
2.2.4	Example: Brownian motion	47
2.2.5	Example: Bacterial advection-diffusion	48
2.2.6	Constructing ISFs for multiple independent processes	50

2.2.7	Parameter fitting	53
2.2.8	Algorithm optimisation	55
2.2.8.1	Time lag sampling	55
2.2.8.2	Choosing a spatial frequency interval for fitting .	57
2.2.8.3	Image windowing	58
2.2.8.4	Assuming the absence of boundaries	59
2.2.9	Summary of assumptions	60
2.3	Preliminary exploration of DDM in oocyte data	62
2.3.1	Brownian motion describes cytoplasmic movement in FCCP treated oocytes, but active diffusion is insufficient to describe healthy oocytes	64
2.3.2	The bacteria model performs poorly in both control and FCCP oocytes.	66
2.3.3	Cytoplasmic movement is anisotropic	69
2.3.4	Diffusion coefficients fall following poisoning, but not significantly	69
2.3.5	Identifying areas of improvement for DDM in oocyte data	71
3	Optimal design of synthetic data in DDM	75
3.1	Introduction	75
3.1.1	Simulation construction	76
3.2	Controlling a and σ to optimise the appearance of simulated images	80
3.3	Balancing the particle number, N_p , and number of frames, N_t . .	84
3.4	Summary and conclusion	90
4	Simulation-guided approaches to Differential Dynamic Microscopy in IVF	92
4.1	Introduction	92
4.1.1	Example: Brownian motion in a square domain with reflective boundary conditions	95
4.2	Simulation-determined fitting functions in DDM analysis.	96
4.3	Simulations highlight poor performance of fitting algorithms when some movement sources dominate others	106
4.4	Introducing a simulation-driven machine learning pipeline for parameter fitting in DDM	114
4.5	Summary and Discussion	126

5	Differential Dynamic Microscopy can be applied to Differential Interference Contrast images despite shadowing effects	131
5.1	Introduction	131
5.2	An analytic expression for the DDM matrix of a DIC image stack for isotropic motion	136
5.2.1	Verifying the use of DDM in DIC through simulations and colloid images	142
5.2.1.1	Validation with simulations	142
5.2.1.2	Validation with experimental data	144
5.2.2	DDM in isotropic motion can extract the direction, but not the magnitude, of the shear	146
5.3	Anisotropy	150
5.3.1	Fitting the DDM matrix to a simulated advection-diffusion process	156
5.3.2	Approximating advection and DIC shears as separable effects	159
5.4	Exploring the use of the ‘bright-field’ ISF in DDM analysis of DIC images	164
5.4.1	Finding parameter regions where error in the fitted velocity is predictable	170
5.4.1.1	Case 1: $\Phi = 0$	172
5.4.1.2	Case 2: $\Phi = \pi/2$	173
5.4.2	Estimating the error factor γ	174
5.5	Discussion	176
6	Modelling the effect of vitrification protocols on oocyte or embryo cooling rates	180
6.1	Cryopreservation and vitrification protocols	180
6.2	Mathematical model of the vitrification process	184
6.2.1	Assumptions	189
6.2.1.1	Liquid nitrogen is a stationary, isothermal liquid	189
6.2.1.2	Samples are spheres	190
6.2.1.3	Thermophysical properties and parameters	191
6.2.1.4	Cryoprotectant droplets are hemispheres	195
6.2.1.5	Samples have the same chemical properties as the cryoprotectant solution	196
6.2.1.6	The transition from liquid to glass-like state occurs instantaneously, no moving boundary effects	196
6.2.2	Computational modelling	198
6.2.3	Spatial arrangements of oocytes/embryos	201

6.3	Results	204
6.4	Discussion	209
7	Statistical analysis of time-lapse images of post-thaw embryos	214
7.1	Introduction	214
7.2	Machine-learning image segmentation	215
7.3	Non-linear expansion rates predict negative clinical pregnancy . .	220
7.4	Summary and discussion	227
8	Summary, discussion and future work	232
A	Experimental details regarding mouse oocyte collection	244
B	Single Particle Tracking & Particle Image Velocimetry	245
B.1	Single Particle Tracking	245
B.2	Particle Image Velocimetry	247
C	DDM applied to bright-field oocyte images	248
C.1	The theory of the formation of a Differential Interference Contrast (DIC) image	253
C.2	Approximation of the ISF for DIC images as a weighted average of parallel and orthogonal components	255
	Bibliography	258

List of Figures

1.1	Diagrammatic representation of a general IVF process, as described in [4]. Freezing and thawing of oocytes and embryos are optional stages, with cycles using cryopreserved embryos are known as frozen embryo transfers (FET). Created with BioRender.com. .	3
1.2	A collection of frames from a time-lapse movie obtained by the London Women’s Clinic, Wales and Bristol. A single human embryo is depicted, which re-expands after being thawed from cryopreservation. More detail on this image stack, and analysis of similar movies, can be found in Chapter 7	5
1.3	Flowchart of the reportable outcomes of IVF treatment. Whilst the overall objective of IVF is live birth, success can be reported as reaching any of the positive milestone stages that are pre-requisite for live birth, and statistics on each stage can be complicated by the exclusion of patients who did not reach the previous milestone.	6
1.4	The live birth per embryo transferred rate, reported by the Human Fertility and Embryology Authority (UK) between 1991 and 2019. Reproduced from [7].	7
1.5	The human oocyte: (a) schematic diagram. (b) microscope images.	12
1.6	Schematics of the cytoplasmic streaming behaviour in various cells, alongside annotated microscope images of each showing the direction of the local velocity field.(a) Chara corallina algae, schematic adapted from [154] and microscope image taken from [154], (b) mouse oocyte, schematic adapted from [151] and microscope image taken from [155] (c) Drosophila (fruit fly) oocyte, schematic adapted from [142, 154] and microscope image taken from [150]. .	16

1.7	The morphology of a typical embryo at the blastocyst stage. The cells of the embryo are contained inside the zona pellucida, with the trophectoderm surrounding the inner cell mass (ICM) and fluid-filled blastocoel. (Left) Diagram produced using BioRender. (Right) Image of an embryo graded as 3AA using the Veeck and Zaninovic system, a slightly modified form of the Gardner system (using which, this embryo would be classified as 4AA [36]) . Taken from [190].	20
1.8	An illustration by senior embryologist, Diana Tain, depicting some of the observable patterns which groups of four oocytes have formed when vitrified together on the same Cryotop [®] device. Taken with permission from [239].	24
2.1	The first frame of a bright-field time-lapse movie depicting eleven mouse oocytes. Analysis of this data is presented in Section 2.3. Images provided by Emily Lewis and Karl Swann in the School of Biosciences, Cardiff University.	31
2.2	Example of the construction of a Linear Space Invariant (LSI) image. By definition, these images would exhibit no shadowing or obstruction when objects appear above or below one another in the z axis. Reproduced from [267].	37
2.3	An overview of the algorithm which converts the image stack $I(\mathbf{r}, t)$ into the DDM matrix $\mathcal{D}_R(q, \Delta t)$, from which movement parameters in the van Hove function $G(\Delta\mathbf{r}, \Delta t)$ may be determined. Step (I) constructs image differences over all pairs of images, which can be grouped by time lag. The time-averaged absolute square of the spatial Fourier transform (Equation (2.2.12)) constructs the DDM tensor $\mathcal{D}(\mathbf{q}, \Delta t)$, which is reduced to a matrix by radial averaging (Equation (2.2.16)).	43
2.4	(Left) The contour plot of $\mathcal{D}(\mathbf{q}, \Delta t)$ for a fixed Δt , for simulated Brownian motion data, as described in Chapter 3. Each dashed circle represents an annulus over which the radial average is applied. (Right) The curve generated by the radial average, where the average over the smaller purple annulus produces the circular point in the right hand figure, and the average over the larger red annulus produces the diamond point.	45

2.5	Example of oversampling of pixels to improve the accuracy of the radial average. (Left) The example image is made up of pixels of brightness 0, 1 or 2. In the original image, the 2 brightness pixels intersect the red circular contour over a shorter distance than the 1 brightness pixels, but have equal weight in the calculation of the average. (Right) Oversampling splits pixels into smaller grids of equal brightness, allowing us to weight contributions towards the radial average based on the length of the contour intersecting the original pixel. The 1 brightness pixels now have a greater contribution to the average than the 2 brightness pixels, resulting in a drop in the average brightness.	46
2.6	A plot of the Schulz distribution for $\bar{v} = 5$, when $Z = 0.2$ (blue solid line) and 4 (red dashed line).	49
2.7	A typical parameter fitting diagram, taken from [164], which shows the fitting results obtained when attempting to classify Brownian motion in a colloidal dispersion. (a) The parameter $\tau_D(q)$ is plotted as a function of spatial frequency. Fitting is performed over region (2), where the log-log plot of τ_D is linear with gradient -2. (b) The fitted values of $A(q)$ (red dashed line) and $B(q)$ (blue dotted line) are plotted, alongside estimates of these parameters obtained by taking short and long time behaviours of $\mathcal{D}_R(q, \Delta t)$	56
2.8	In movies with equally spaced frames, such that the time lag can be mapped to a natural number m denoting frame difference, we plot the sampling rate of time lags when using the logarithmic spacing given in Equation (2.2.29). As b increases, both the maximum sampled time and the density of time sampling increase.	57
2.9	Visualising the benefit of image windowing, reproduced from [267]. When particles touch the boundary of the image (top left), the discontinuity in brightness causes smearing in the Fourier transform (top right). If we premultiply by a window function such Equation (2.2.30) (Bottom left), smearing is reduced (bottom right).	59
2.10	Contour plot of the Blackman-Harris window given in Equation (2.2.30), for a square image of length $L = 100$. The brightness is 1 at the centre and 0 at the edges, with a continuous slope between.	60
2.11	Subdomains of the cytoplasm of a <i>Drosophila</i> oocyte, over which DDM was applied in [142], in order to avoid imaging the nucleus (labelled n) and the cell walls. Taken from [142].	61

2.12	Fitting results for a given oocyte (7 in Appendix C), using the Brownian model in both the control condition 2.12a and the FCCP condition 2.12b. The FCCP condition yields better fitting quality under the Brownian model than in control conditions.	65
2.13	Fitting results for the an oocyte (7 in Appendix C), using the bacterial model in both the control experiment (2.13a) and the FCCP condition 2.13b. The solid and dashed red lines represent the curves fit to τ_d and τ_v respectively, generated by linear regression as presented in Equations (2.2.27) and (2.2.28). The performance of the bacterial model is poor in both cases.	66
2.14	The motility fraction α fitted at each frequency q for oocyte 7 under both the control (2.14a) and FCCP (2.14b) conditions. Vertical dashed bars correspond to the lower and upper bounds of the fitting region, as in Figures 2.12 and 2.13.	68
2.15	Contour plots of the DDM tensor $\mathcal{D}(\mathbf{q}, \Delta t)$ for $t = 1$ s, for the same oocyte (2) in both the control (2.15a) and FCCP (2.15a) datasets. A slight anisotropy is observed in the control dataset, with the positive diagonal axis achieving lower values than the negative diagonal. This effect is reduced in the FCCP dataset, but still present.	70
2.16	Box plots of the fitted diffusion coefficients extracted from the control and FCCP mouse oocyte datasets, using a Brownian model to describe cytoplasmic movement. The diffusion coefficients all drop following poisoning with FCCP.	71
3.1	Simulations are performed over a larger simulation domain (solid black) with periodic boundaries to maintain a constant image brightness. The images themselves are taken as a central region of the simulation domain (red dashed square), to minimise the correlation of particles which exit the frame and on side and rejoin on the other. The simulation and image domain can be mapped onto a torus.	78

3.2	Diagram showing loss of conservation of mass in simulated particles when the beam spread parameter, σ , is too small. The left and right images use the same particle brightness, $a = 50$, and the same particle locations. The top left particle is located at the centre of a pixel, whilst the bottom right particle is located at a corner between four pixels. When $\sigma = 3.0$, the particle appearance is invariant of the particle position, but when $\sigma = 0.5$, the particle appearance changes depending on the location within the frame.	82
3.3	The error observed when fitting the diffusion coefficient, D , for different combinations of particle brightness, a , and beam spread, σ , in the same simulated Brownian motion. Error is constant everywhere, except for $\sigma < 0.1$, where the error is large and has high variance.	83
3.4	Visualisation of the image differences ΔI for a single particle undergoing a displacement of 10 pixels. In the left hand side, $\sigma = 3$, and in the right hand side, $\sigma = 50$. As the beam spread increases, the brightness gradient in the image difference decreases over small distances, resulting in insensitivity of DDM at higher frequency.	84
3.5	Plots of the fitting parameter τ_D for the same simulation, performed using $a = 50$ at different values of $\sigma \in \{0.1, 0.5, 1.0, 3.0, 5.0\}$ px. When $\sigma \geq 0.5$ px, a suitable fitting region can be drawn, but the upper limit of the fitting interval (right-hand grey dashed line) begins to decrease with increasing σ	85
3.6	Plots of the mean diffusion coefficient, D , fitted to simulated Brownian motion images containing varying numbers of particles, with error bars showing \pm one standard deviation. The mean diffusion coefficient stays constant at around $0.48 \text{ px}^2/\text{frame}$, but with increasing particle number, the standard deviation decreases.	89
4.1	The process flow for the simulation-driven pipeline to fit the diffusion coefficient from Brownian motion in a circular domain with reflective boundaries. The objective is to create a continuous numerical polynomial approximation of the ISF, by first generating it numerically from simulated particle trajectories. These polynomials can be used to extract the diffusion coefficient from the DDM matrix.	98
4.2	Re-ordering allows us to express the ISF at a single frequency q , with varying time and diffusion coefficient, Δt and D respectively.	98

4.3	Visualisation of the scheme used to refine the image domain length iteratively, defining subdomains (squares) that inscribe circles which, in turn, inscribe the squares defining the previous subdomain. . .	100
4.4	Fitted diffusion coefficients from simulated Brownian motion inside of a reflective circular domain of diameter 500 pixels, contained within an image window with length 512 pixels. Fitting is performed both using a classical windowed DDM approach (black-pink gradient denoting varying window size), and using the spline interpolation method proposed (x curve). The ‘true value’ curve is also given (red dash).	101
4.5	Fitted diffusion coefficients from simulated Brownian motion inside of a reflective circular domain of diameter 250 pixels, contained within an image window with length 512 pixels. Comparing to Figure 4.4, almost the exact same distribution of diffusion coefficients is observed despite the diameter of the bounded domain being halved. Additionally, the spline fitting under-estimates the diffusion coefficients more significantly, particularly when $D > 2.25$, compared to Figure 4.4.	102
4.6	For a particle located inside a circular domain which undergoes a horizontal displacement of Δx , collision occurs if the particle was initially located inside of the mono-colour ‘crescent moon’ area. .	102
4.7	Visualisations of the probability of boundary interactions for particles undergoing Brownian motion inside of a circular domain. These plots demonstrate that, for the small time lags investigated during DDM, collision is unlikely.	104
4.8	Fitting error for simulations of advection-diffusion behaviour, expressed as a percentage of the true value of each parameter. Input diffusion coefficients and velocity magnitude range between $0.1 \leq D \leq 2$ px ² /frame and $0.1 \leq v \leq 2$ px/frame, respectively. .	111
4.9	Scatter plot of the percentage error for fits of the velocity v and diffusion coefficient D in simulations, as a function of the Péclet number, Pe , defined in Equation (4.3.4). As Pe increases, the fit of the velocity improves, whilst D is fit optimally for $Pe \approx 10$. . .	112
4.10	A very simple neural network architecture, featuring four layers of neurons. The first layer is determined by the input into the network, and subsequent layers are defined defined sequentially by summing weighted contributions from previous layers, with an additional bias.	115

4.11	Plot of the fitted diffusion coefficients, D , in the test set of simulated Brownian motions, fitted by either both classical fitting approaches outlined in Section 2.2.7 (purple filled square) or machine learning (open circle).	123
4.12	Plot of the diffusion coefficients fitted to the three colloidal image datasets with varying kinematic viscosity ($\nu \in \{1.46, 2.10, 3.20\}$ mPa s), using the classical fitting approach (purple filled square) and machine learning (open circle). The Stokes-Einstein relation, showing the expected diffusion coefficient for each dataset, is given by the blue line.	124
4.13	A proposed future workflow for designing new DDM fitting methods for use in the IVF clinic. A small amount of real data is studied using invasive imaging techniques with injected tracer particles, to motivate fluid mechanical models from which synthetic particle trajectories may be generated. The synthetic data may be used to train a neural network, which may then be applied to the real data for non-invasive extraction movement statistics.	125
5.1	An example of the relationship between OPL and its spatial gradient $\nabla(\text{OPL})$ in a homogeneous semi-spherical particle. The OPL and $\nabla(\text{OPL})$ are plotted along a single line oriented parallel to the DIC shear direction, with the square and diamond on the plots corresponding to the location on the image of the particle. Where $\nabla(\text{OPL})$ is positive on the left of the particle, the particle is light, and gets darker as it approaches the right hand edge. Based on a similar diagram from [268].	133
5.2	Image of a single mouse oocyte taken with a DIC microscope setup.	134
5.3	Isocontours of $\mathcal{D}(\mathbf{q}, \Delta t)$ for (5.3a) bright field [164] and (5.3b) DIC images of colloidal dynamics. Both sets of images are displayed for the smallest available time separation: $\Delta t = 2.5$ ms for bright field and 0.12s for DIC. The bright field images yield concentric circular isocontours, but in DIC images, there is a clear asymmetry aligned with the direction of the DIC microscope shear, denoted by the white arrow \mathbf{s}	134

5.4	A schematic for a series of rotating Janus particles, each half coated with two chemicals ‘A’ and ‘B’. It is possible to tether these particles in such a way that rotation occurs only in the axis parallel to the imaging plane, ensuring that in the resulting images, each particle has exactly half its visible surface coated with each type of chemical [340].	136
5.5	The first frame of the simulated phase image, X , and the corresponding DIC image, I , which is generated by $X(\mathbf{r} + \mathbf{s}, t) - X(\mathbf{r}, t)$.	142
5.6	Snapshots of $\mathcal{D}_{I,R}(q, \Delta t)$ and the rescaled form $2(1 - J_0(qs))\mathcal{D}_{X,R}(q, \Delta t)$, plotted at fixed times over all spatial frequencies. There is excellent agreement for all frequencies $q < 2\mu\text{m}^{-1}$ hence the relationship (5.2.21) is valid over this interval.	143
5.7	The upper plot depicts the fitted values of τ_{D_X} and τ_{D_I} for the image stacks X and I , respectively, in a single simulation, on a logarithmic scale. The linear regression (2.2.27) is also plotted, and agreement of the intercepts demonstrates equivalent diffusion coefficients are determined. The lower plot visualises the fitting parameters $A(q)$ and $B(q)$ for each stack.	144
5.8	For 150 simulations of Brownian motion, we generate phase stacks $X(\mathbf{r}, t)$ and corresponding DIC images $I(\mathbf{r}, t)$. DDM analysis on each stack yields a diffusion coefficient, D_X and D_I , respectively. (5.8a) Although the fitted diffusion coefficients have some spread around the true value $D = 0.5\mu\text{m}^2/\text{s}$, (5.8b) the difference $D_X - D_I$ is small, so the DIC shear has minimal effect on the fitting accuracy.	145
5.9	DDM analysis of the diffusion coefficients for fitted DIC images, from colloidal suspensions, yields a fitted diffusion coefficient of $0.676\mu\text{m}^2/\text{s}$. The fitting region $1.1\mu\text{m}^{-1} < q < 3.3\mu\text{m}^{-1}$, marked between the vertical blue dashed lines, is chosen so that the signal ($A_I(q)$) to noise ($B_I(q)$) ratio is sufficiently high.	146
5.10	Scatter plots of the left hand side of (5.2.24), with fitted curves as described in the right hand side of Equation (5.2.24), applied to the DDM tensor generated by DIC colloidal data. Curves are plotted for $q \in \{0.40, 0.78, 1.16, 1.54\}$	148
5.11	Plots of the fitted parameters $\underline{b}(q)$ in (5.2.24), compared to the known parameter values given by the magnitude and directions of the DIC shear, x_0 and ϕ respectively.	149

5.12	The image difference ΔI of one DIC imaged particle will contain four spots, corresponding to the position of the particle at times t_1 and t_2 , both before and after shearing. The DDM tensor is formed by the spatial correlations between spots, which can be broken down into terms T_1 - T_4 in Equations (5.3.3)-(5.3.6). The vectors between spot pairs at different times and shearing are described by R_{\pm}	153
5.13	Fits of the three key parameters in Equation (5.3.14); (5.13a) τ_q (5.3.16), (5.13b) τ_v (5.3.17) and (5.13c) Φ (5.3.18), each at different true values of the angle difference Φ . The fit for τ_D and τ_v is universally good, but Φ is always poorly fit.	157
5.14	The expected ISF $f_I(q, \Delta t)$ (5.3.15) in our simulated data for DIC shear values of $s \in \{0.238, 1.000, 2.000, 5.000\}$. For small shears, the functions are tightly constrained and indiscernible from one another, but both the function range and the significance of Φ increase as the shear increases.	158
5.15	Fits of the movement parameters τ_D and τ_v as in Figure 5.13 using the simplified Equation (5.3.25) to fit the DDM matrix as if no DIC shear existed. Diffusion is well fit in all cases, but although the velocity curves are fitted more smoothly than using other fitting functions, the gradients are less accurately fit when $\Phi = 0$ or $\pi/2$	162
5.16	For 10,000 simulated instances of the DIC ISF, $f_I(q, \Delta t)$, with randomly generated input parameters (Equations (5.4.7)-(5.4.9)), we fit the output parameters κ_{out} and λ_{out} by minimising the residual function E_I in Equation (5.4.3).	167
5.17	Scatter plot values of the fitted scaling term \hat{A} , alongside the curve $1 - J_0(\xi)$	167
5.18	Paired distribution histograms of estimation errors $\kappa_{\text{out}} - \kappa_{\text{in}}$ and $\lambda_{\text{out}} - \lambda_{\text{in}}$, against the input parameters κ_{in} , λ_{in} and Φ . On the right-most plots where ξ is on the horizontal axis, vertical bars denote the values of ξ for which $J_2(\xi) = 0$, coinciding with the minimisation of both estimation errors.	168
5.19	For the limited parameter region $\lambda < j_{0,1}$ and $\xi < j_{1,1}$, the predicted value of the velocity scaling term γ (Figure 5.19a) aligns with the numerically determined scaling term (Figure 5.19a).	176

6.1	A labelled image of the Open Cryotop [®] device used to vitrify oocytes and embryos, modified from the Kitazato website [350]. The blue plastic part on the right hand is the handle of the device. The embryo is mounted onto the thin transparent plastic (polypropylene) plate on the left, just below the black diamond-shaped tip of the device.	181
6.2	Diagram depicting the process of transferring a sample from a dish onto the plate of the Open Cryotop [®] , and removing excess cryoprotectant through aspiration to leave a ‘minimal droplet’. Adapted from the Open Cryotop [®] user manual [248].	182
6.3	A cross section of the model geometry.	184
6.4	Experimental temperature reading of the Open Cryotop [®] device after being plunged into liquid nitrogen, taken from [356]. Initially at ambient room temperature, the device is first submerged in liquid nitrogen at around 4 seconds, cooling completely within less than half a second.	189
6.5	The diffusion coefficient for each the different thermophysical parameter choices in Section 6.2.1.3: ethylene glycol, vitreous water and water (with variable thermophysical parameters), under the assumption of constant pressure. The vertical axis is on a log scale. For water, the sharp transition occurs at the freezing point, when water will solidify into ice during cooling. Since we model the cooling problem, temperature goes from right to left.	192
6.6	A schematic diagram of the Cryotop [®] system, including dimensions of the three key domains: the sample, the plate and the cryoprotectant droplet. The entire system is submerged in liquid nitrogen.	196
6.7	The moving boundary problem encountered during modelling of a cooling problem where phase change occurs in the domain. As the temperature falls to the freezing point u_m of the liquid, parts of the domain solidify, resulting in a solid-liquid interface, with position $X(t)$ at time t and fixed temperature $u = u_m$. The interface moves through the domain as more of the liquid solidifies.	197

6.8	A snapshot of the model rendered in COMSOL for the three sample ‘Line’ arrangement described in Section 6.2.3. The snapshot is taken at 0.3 seconds into the simulation, with temperature in Kelvin, corresponding to colour, determined by solving Equations (6.2.4)-(6.2.13). The sample in centre is lighter than those on either side, implying it is warmer, which demonstrates that sample arrangement does affect cooling.	200
6.9	The reciprocal of the automatically-determined step size for the three sample ‘Line’ arrangement shown in Figure 6.8. The step size converges, implying stability. Similar convergence is observed for all arrangements, which suggests the model is stable.	201
6.10	Two merged droplets rendered in COMSOL Multiphysics 5.5. The same set-up is displayed from two different angles. The droplet is formed from the union of the two separate droplets accompanying each sample.	202
6.11	Top-down visualisations of the six arrangements our model. The ‘Line’ arrangement is considered in three permutations with variable sample number, named ‘Line-1’, ‘Line-2’ and ‘Line-3’ according to the number of samples considered.	204
6.12	The average temperature (Equation (6.2.14)) in samples arranged in a straight line on the Open Cryotop [®] device, namely the ‘Line’ arrangement (Figure 6.11a) with one to three samples, and the ‘Cylinder’ arrangement (Figure 6.11b), which is equal to the ‘3-Line’ arrangement, with a semicircular droplet cross-section. Temperature is determined by numerically resolving the PDE system given by Equations (6.2.4)-(6.2.12) in COMSOL Multiphysics 5.5. Each sub-figure uses a different choice of thermophysical parameters for the CPA solution, with the constant parameters for ethylene glycol and and vitreous water given in Table 6.2, and the temperature-dependent thermophysical parameters for water given in Table 6.3.	205

6.13	The average temperature (Equation (6.2.14)) in samples in the ‘Cylinder’, ‘Triangle’, ‘Square’, and ‘Cross’ arrangements (Figures 6.11b, 6.11c, 6.11d and 6.11e respectively), as well as the ‘Worst case’ arrangement (Figure 6.11f) that considers a much larger droplet. Temperature is determined by numerically resolving the PDE system given by Equations (6.2.4)-(6.2.12) in COMSOL Multiphysics 5.5. Each of the figures uses a different choice of thermophysical parameters for the CPA solution, with the constant parameters for ethylene glycol and vitreous water given in Table 6.2, and the temperature-dependent thermophysical parameters for water given in Table 6.3.	206
7.1	A reproduction of Figure 1.2 in Chapter 1, the re-expansion of a single blastocyst is shown. The images are approximately 25 minutes apart, showing the initially compacted and shrunken blastocyst re-expanding to its pre-vitrified state.	215
7.2	Visualisation of the outcomes of the 147 embryos from LWC data. We excluded 23 embryos from the study; 22 were missing corresponding movies, and one did not survive cryopreservation and was not transferred. Of the remaining 124 embryos, 67 resulted in a positive pregnancy test (54.0%), 53 in a clinical pregnancy (41.7%), and 37 in live birth (29.1%). These rates are in line with expected European live birth rates [85]. The two embryos listed as ‘?’ had no recorded live birth outcome data.	216
7.3	A single frame from one time lapse video, at the start of the re-expansion process when the blastocyst is compacted. The two major features of interest are labelled: the border of the blastocyst (red solid line) and the inner edge of the Zona Pellucida (ZP, black dashed line), labelled by hand. The blastocyst area is contained within the red contour, and the perivitelline space is the area contained between the red and black dashed contours.	218
7.4	ResNet architectures in machine-learning models differ from standard neural networks due to the inclusion of ‘shortcuts’ that can skip layers. Adapted from [388].	219
7.5	A single frame from an embryo expansion movie provided by LWC, manually labelled using the MATLAB 2022b image labelling app.	220

7.6	The blastocyst area in a single embryo, plotted over time, in two different representations. In Figure 7.6a, the area is plotted as a function of time, in frames (where one frame is taken every 5 minutes), whilst in 7.6b, we plot the blastocyst area in the frame $t + 1$ against the blastocyst at frame t	221
7.7	Scatter plot of the blastocyst area, at time $t + 1$ against the area at time t , for the 124 embryos, provided by LWC. Each point represents a single frame of an embryo, with crosses marking embryos which did not result in a clinical pregnancy, and open circles marking those which did. In general, the points farthest from the $y = x$ curve tend to be originate from embryos which did not result in a clinical pregnancy. Colour denotes which point originated from which embryo.	222
7.8	The same embryo data plotted in Figure 7.6, with bounding lines given by Equation (7.3.1). When $\Delta\mathcal{A} = 1\%$, the embryo remains inside these bounds since the biggest change in the blastocyst area is smaller than this value, but for $\Delta\mathcal{A} = 0.5\%$, a number of points (filled red circles) sit outside of the boundary. These points would cause the embryo to be classified as ‘OB’	223
7.9	(7.9a) The number of embryos labelled IB or OB as a function of $\Delta\mathcal{A}$. (7.9b) the percentage of embryos that led to clinical pregnancy for embryos classified as IB and OB, as a function of $\Delta\mathcal{A}$. When $\Delta\mathcal{A} > 0.42\%$, IB embryos are more likely to generate a clinical pregnancy than OB embryos, and beyond $\Delta\mathcal{A} = 2\%$, all OB embryos had a negative clinical pregnancy outcome.	224
7.10	The χ^2 test statistic value as $\Delta\mathcal{A}$ varies from 0 to 3.5%, compared to the critical value 3.84 corresponding to the p-value $p = 0.05$. Any test statistic value over the critical value 3.84 is considered statistically significant.	226

7.11	Histogram exhibiting the Probability Density Function (PDF) of the changes in the blastocyst area between each time frame in the time-lapse movies; plotted separately for embryos with a positive and negative clinical pregnancy outcome. Black dotted vertical lines denote the upper and lower bounds on the blastocyst area change for embryos with a positive clinical outcome, and red dashed lines show the upper and lower limits on the blastocyst area change for negative clinical outcome. Embryos resulting in negative clinical pregnancy have wider tails in PDF, and are skewed negatively compared to the embryos resulting in positive clinical pregnancy.	228
B.1	An example of linking error in SPT, where two particles have moved from the grey to the red location over a fixed time step. Based on only observation of the initial and final state of the system, the trajectory of each particle is uncertain.	246
C.1	Fitting results performing DDM analysis for brightfield imaged mouse oocytes, under control conditions, where the Brownian fitting model is used.	249
C.2	Fitting results performing DDM analysis for brightfield imaged mouse oocytes, in the FCCP condition, where the Brownian fitting model is used.	250
C.3	Fitting results performing DDM analysis for brightfield imaged mouse oocytes, under control conditions, where the bacteria fitting model is used.	251
C.4	Fitting results performing DDM analysis for bright-field imaged mouse oocytes, under FCCP conditions, where the bacteria fitting model is used.	252

C.5 A sketch of the principle of operation of a DIC microscope with de-Sénarmont compensation. After adjusting a transmission microscope for Köhler illumination, a linear polariser at $45^\circ + \theta$ to the sample x (horizontal) and y (out of plane) directions creates a linear polarisation (see electric field sketches on the right), which is converted by a quarter wave plate to a relative phase shift of 2θ of the x and y polarised field components E_x and E_y of equal amplitude (see sketch of temporal field oscillations on the right). A Wollaston prism splits the two fields, creating two focal points separated by the shear \mathbf{s} . The transmitted fields, which experienced a relative phase shift $\delta(\mathbf{r})$ as function of sample position $\mathbf{r} = (x, y)$, are collimated by the objective, recombined in direction by a second Wollaston prism, and projected by the analyser of -45° orientation, providing an intensity proportional to $1 - \cos(\psi - \delta(\mathbf{r}))$. Using $\psi = 90^\circ$, a nearly linear dependence of the intensity on the phase shift $\delta(\mathbf{r})$ is obtained, leading to the relief-type contrast. . . 254

List of Tables

- 1.1 IVF success rates reported across world-wide registries. No registry exists for Asia to date [84]. Additionally, reporting differences such as the choice of success metric and patient demographics may account for variation between reporting bodies. 8

- 3.1 Parameter selection for generation of synthetic image data of the form given by Equation (3.1.2), which has been shown to yield well-conditioned images for DDM application. 90

- 6.1 Reported cooling rates of the Open Cryotop[®] device. In particular, [224] is used by the manufacturer, Kitazato[®] as the official reported device cooling rate [350]. The interval over which cooling rate is defined varies, which by the exponential temperature decay predicted in Equation (6.2.16), may affect reported cooling rates. The variable cooling rates reported in [240] correspond to different definitions of the heat transfer coefficient, see Section 6.2.1.3. Cooling rates are equivalent in K/min. 188
- 6.2 The thermophysical properties for the materials involved in the Open Cryotop[®] system. 194
- 6.3 Thermophysical properties of water [368]. Two values are reported at 273.15K for ice or water; with linear interpolation, we determine the values of ice at 273.14K and the value for water at 273.16K. . . 194
- 6.4 Time taken (seconds) for the average temperatures inside the samples, plotted in Figures 6.12 and 6.13, to reach the target vitrification temperature, 143.15K. Each column denotes a separate case of the choice of thermophysical parameters, with constant parameters for ethylene glycol (EG) and vitreous water (vitW) available in Table 6.2, as well variable thermophysical parameters for water (W) in Table 6.3. 207

6.5	Cooling rates from 310.15K to 143.15K in the simulations, calculated by dividing the temperature difference by the cooling times from Table 6.4. Results shown for the three choices of thermophysical parameters,: ethylene glycol (EG), vitreous water (vitW) or water with variable thermophysical parameters (W).	207
-----	---	-----

Chapter 1

Introduction

Assisted reproductive technologies (ART) are treatments and procedures that include the in vitro handling and manipulation of oocytes (egg cells), sperm or embryos, with the purpose of achieving a pregnancy [1, 2]. One of the most well known forms of ART is In Vitro Fertilisation (IVF) [3], a procedure in which oocytes are fertilised with sperm in a laboratory to create an embryo, which is then transferred to the uterus of the patient [4]. Since the first live birth resulting from IVF treatment in 1978 [5], there have been over 8 million IVF babies worldwide [6]. In the United Kingdom (UK) alone, 595,000 patients have received IVF treatment since 1991, totalling over 1.3 million IVF treatments that have resulted in 390,000 babies born [7]. ART is becoming increasingly common [8], with medically assisted births accounting for 2.9% of all live births in the UK, and as much as 5% of live births in countries such as Denmark, Austria and Slovenia [3, 6, 9].

The objective of this thesis is to address some of the mathematical challenges involved in the pursuit of improving IVF. This thesis is linked to a project led by the School of Mathematics, Cardiff University (headed by Katerina Kaouri and Thomas Woolley), and is funded by a Knowledge Economy Skills Scholarship (KESS2). The project has been co-funded by, and was in collaboration with, the London Women's Clinic (LWC), who operate 14 fertility clinics across the UK [10]. In particular, supervision was provided by Helen Priddle, Andrew Thomson and Giles Palmer at LWC Wales and Bristol branch, based in Cardiff. The project is additionally co-supervised by Professor Karl Swann in the School of Biosciences, Cardiff University, who leads an experimental lab on oocyte and fertilisation research. The project also involved collaboration with the School of Physics and Astronomy, Cardiff University (more detail can be found in Chapter 5). We combine mathematical modelling and data analysis, the clinical expertise of LWC, experiments at the Swann lab, and imaging experiments to tackle some of the challenges limiting IVF success.

We focus on three particular aspects of IVF; first, we introduce a set of new methods for analysing time-lapse images of oocytes, in order to facilitate non-invasive oocyte health assessment. Second, we construct computational models to simulate the process of cryopreservation, testing how variations in oocyte and embryo freezing protocol may affect the survival rates of preserved cells. Finally, we explore the use of time-lapse imaging of thawing embryos after removal from storage, identifying a marker which may predict poor IVF potential. Before tackling these challenges, we will first introduce IVF in Section 1.1, alongside the related technologies and practices of interest. In Section 1.2, we will discuss some of the mathematical pursuits being explored in IVF research. Finally, in Section 1.3, we narrow our focus to a select number of challenges, providing an outline of this thesis.

1.1 Overview of In Vitro Fertilisation (IVF)

A typical IVF cycle is described in Figure 1.1, using [4] as a guide. First, the patient is given hormonal medication to encourage the development of multiple oocytes, in a process called controlled ovarian stimulation (COS) [11]. Oocytes are then retrieved from the ovary and fertilised in the laboratory, either by mixing of the oocytes with sperm, or by intra-cytoplasmic sperm injection (ICSI), a process in which a single sperm cell is injected directly into the cytoplasm of the oocyte, usually when problems relating to the sperm prevent fertilisation [12]. After culturing the developing embryo for 3 or 5 days [13], one or more embryos are transferred to the uterus of the patient (embryo transfer) in the hope of establishing a pregnancy.

IVF is often used as a treatment for infertility, defined as a disease of the male or female reproductive system, indicated by failure to achieve pregnancy after 12 or more months of regular unprotected sexual intercourse [14–16]. Infertility is highly distressing [17], and is a growing, worldwide challenge [18], currently estimated to affect one in six people of reproductive age [16]. Infertility can be caused by a number of health conditions [19, 20], including sexually transmitted infections [21], polycystic ovarian syndrome (PCOS) [22] and damage to reproductive function resulting from cancer treatment [23]. IVF is also being increasingly used by single women [24, 25] and homosexual couples [26–28].

Fertility has become a booming global industry, estimated to be worth around £19.8 billion worldwide in 2022, projected to increase to £32.4 billion by 2026 [29]. The number of clinics offering IVF is also increasing; in Europe in 2017, there were 1382 clinics registered with the European Society of Human Reproduction

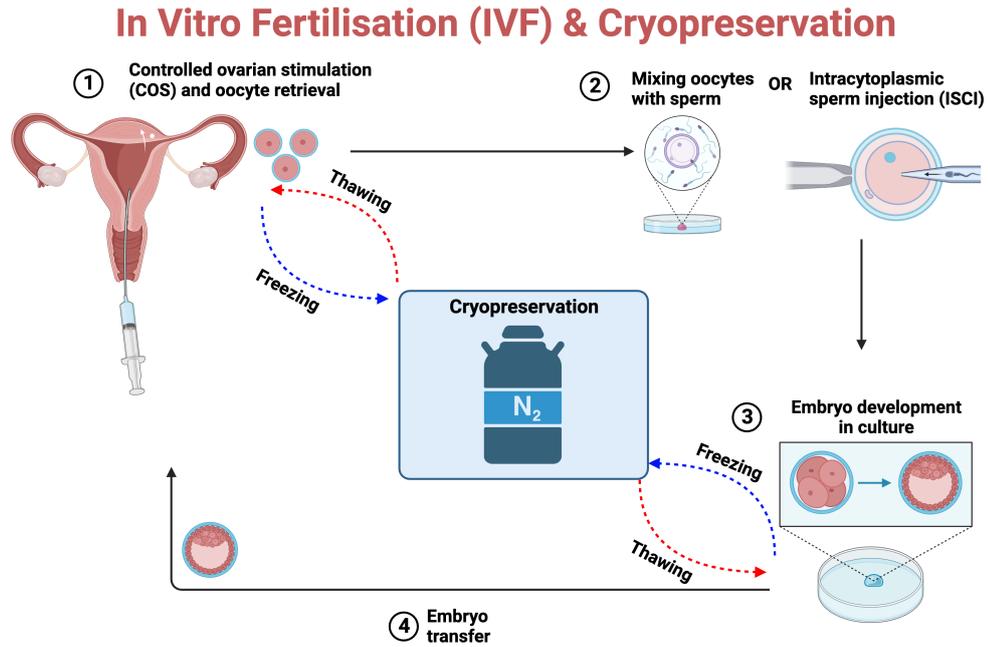


Figure 1.1: Diagrammatic representation of a general IVF process, as described in [4]. Freezing and thawing of oocytes and embryos are optional stages, with cycles using cryopreserved embryos are known as frozen embryo transfers (FET). Created with BioRender.com.

and Embryology (ESHRE) [30], increasing to 1552 by 2018 [9], of which 86 were based in the UK [9].

As part of the modern IVF package, additional technologies and treatments are frequently used to improve or augment the treatment pathway. Of these, we wish to highlight cryopreservation and time-lapse microscopy.

1.1.1 Cryopreservation (oocyte or embryo freezing)

The objective of cryopreservation (see Figure 1.1) is to maintain the health of the oocytes or embryos beyond their typical lifespan by storing them in extremely cold conditions, achieved through the application of liquid nitrogen at -196°C [31] (other forms of cryopreservation may alternatively place samples in a deep freezer at around -80°C). IVF cycles using previously cryopreserved embryos are referred to as frozen embryo transfers (FET), compared to ‘fresh’ IVF cycles. Cryopreservation of embryos has been instrumental in maximising the use of a single oocyte retrieval, by preserving ‘surplus’ embryos from fresh cycles to give patients more attempts at pregnancy [32, 33]. Cryopreservation of oocytes and sperm, meanwhile, is sometimes undertaken prior to treatments which might damage fertility, such as chemotherapy or gender-affirming surgery. It has also enabled the storage and transfer of donor oocytes and sperm, in cases where the

those provided by the patient are not viable for IVF [34, 35], and for patients with no sperm or egg provider, including single women [24, 25], homosexual couples and people with transgender partners [24, 25, 36].

The importance of cryopreservation is growing with an increase in age-related infertility, as more people in the developed world postpone childbearing until later in life [37, 38]. Fuelled by a combination of changing economic status, education levels and healthcare availability, more people are delaying the decision to have children, resulting in the average age of parents at the time of birth increasing [39]. For example, in England and Wales, the mean age of parents at the time of birth has been steadily increasing since the 1970s, from 26.4 to 30.9 years of age for women, and 29.4 to 33.7 for men [40]. This trend is also seen across the rest of the developed world, including Spain, Japan, Russia and the United States [3, 39]. As both maternal and paternal age increase, the quality of sperm and oocytes decreases, resulting in a declining ability to conceive [37, 38, 41, 42].

In the UK, the number of oocyte freezing cycles has approximately doubled every 3 years between 2010 and 2016, with 10253 oocytes frozen in 2016 as a result of over 1300 oocyte freezing cycles [43]. The number of IVF cycles using previously cryopreserved embryos has increased by 86% from 2014 to 2019, accounting for 41% of all IVF cycles [7]. Embryo storage was the only form of ART to increase during the Covid-19 pandemic [44], with freezing being undertaken to avoid the risks of ovarian hyper-stimulation during fresh cycles, which would add additional burden to the NHS at a time when it was already under enormous strain [36]. The UK recently changed legislation to increase the maximum storage period for cryopreserved oocytes, sperm and embryos from 10 to 55 years [45, 46], and many major companies like Apple and Facebook (now Meta) have offered to fund oocyte freezing to attract more women to the workforce [47].

1.1.2 Time-lapse imaging

Many fertility clinics now offer, for an additional cost, several ‘add-on’ options that augment the typical IVF process outlined in Section 1.1 [48–52]. These treatments are made available on the grounds that they may increase the likelihood that treatment will result in pregnancy, although there is often little or conflicting evidence for their efficacy [48–52], but the high emotional stakes at play in infertility may lead to nevertheless proceeding with such options [50, 51]. In this thesis, we will explore one of these add-ons, known as time-lapse microscopy.

‘Time-lapse’ refers to taking multiple images of a biological sample over time, to provide data on how the sample changes [53]. Conventional time-lapse approaches in the context of IVF involve removing embryos from the incubator at fixed time points to make static observations. Although such disturbance of

the embryo from its culture medium is considered unsafe [54–56], the usefulness of the information gathered from such observations is considered to be a worthwhile trade-off [36]. More recently, however, time-lapse has become more common due to the commercial availability of new specialist incubators, such as the Embryoscope[®][57, 58] and Geri[®][54, 58]. The benefit of these incubators is that cameras and microscopes are integrated into the incubator, so imaging does not require the removal of the oocyte or embryo, and is therefore more safe [54–56]. The availability of more data may help guide evidence-based clinical decision making [54, 56, 57, 59]. Currently, time-lapse analysis is mostly used to assess embryo health by statistically correlating the timings of milestone events, such as cell division or changes in appearance, with the likelihood of clinical pregnancy or live birth (more detail is given in Sections 1.2.2.1 and 1.2.4 [57, 59–62]).

We display a number of frames from a time-lapse image stack in Figure 1.2, provided by LWC. These frames, captured using a Geri[®]incubator, display the re-expansion of a human embryo after thawing from cryopreservation. More detail on this image can be found in Chapter 7.

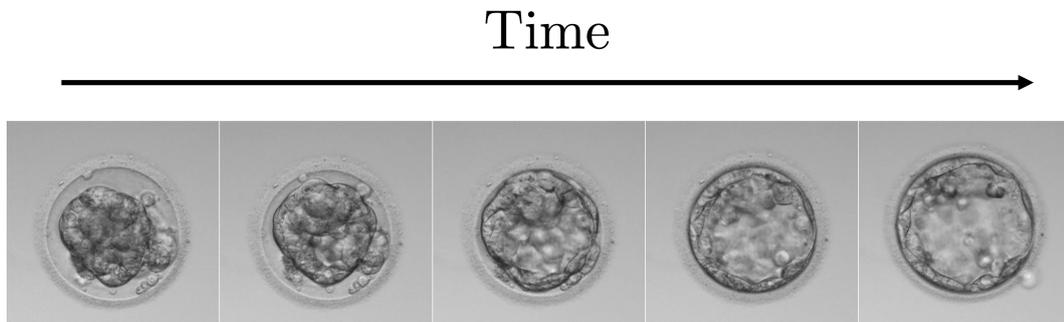


Figure 1.2: A collection of frames from a time-lapse movie obtained by the London Women’s Clinic, Wales and Bristol. A single human embryo is depicted, which re-expands after being thawed from cryopreservation. More detail on this image stack, and analysis of similar movies, can be found in Chapter 7

1.1.3 The need for improving IVF

The goal of IVF is that every treatment results in the live birth of a healthy child [63], but in reality, this is not the case. Evaluating the effectiveness of IVF

treatments is made difficult, however, by the wide variety of ways that outcomes can be reported. The multi-stage nature of IVF means that there are several milestones which must be achieved before a healthy live birth is possible [64], summarised in Figure 1.3, which is based on outcome reports from LWC (and is a more general form of Figure 7.2 in Chapter 7). We will briefly elaborate on the possible definitions of success in IVF, so that we may evaluate the current state of treatments.

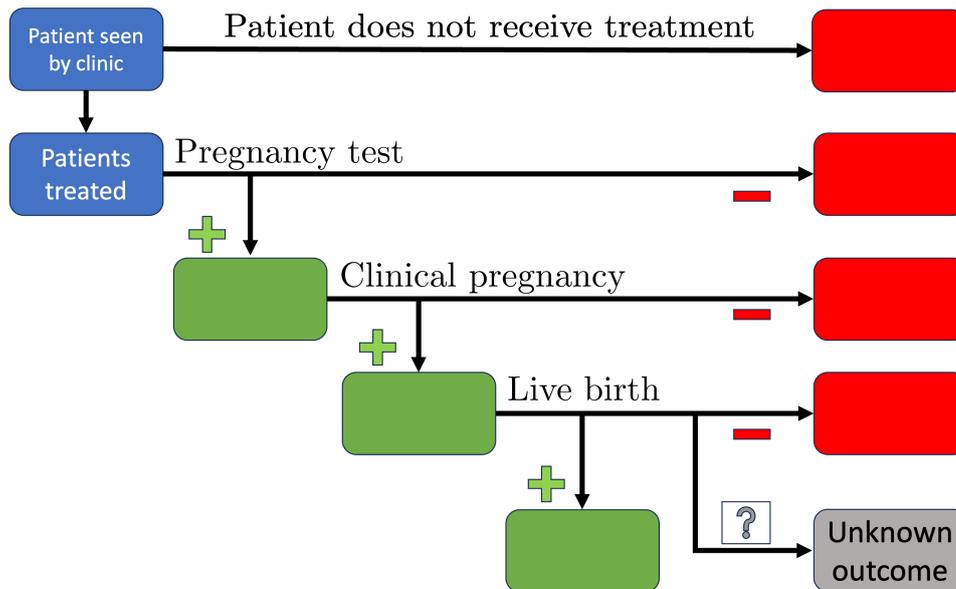


Figure 1.3: Flowchart of the reportable outcomes of IVF treatment. Whilst the overall objective of IVF is live birth, success can be reported as reaching any of the positive milestone stages that are pre-requisite for live birth, and statistics on each stage can be complicated by the exclusion of patients who did not reach the previous milestone.

Of the patients who receive an embryo transfer, a pregnancy test is administered to confirm pregnancy. Those who have a positive pregnancy test will undergo an ultrasound scan to determine the existence of a clinical pregnancy, defined by the established presence of a fetal heartbeat via ultrasound [65]. Finally, not all patients with a clinical pregnancy will result in a live birth. Success has been defined by a number of different metrics in literature, including clinical pregnancy rate [66–70], live birth rate [8, 71–73], live birth per embryo transferred [7, 9], and cumulative live birth rate [74], defined by the number of patients with at least one live birth over a course of repeated IVF cycles.

An additional measure commonly used to record success in IVF is ‘implantation rate’, defined by the number of gestational sacs confirmed by vaginal ultrasound 3-5 weeks after transfer, divided by the number of transferred embryos [75]. This is a measurement more commonly made in controlled studies than in

retrospective datasets; the NHS only administers scans at weeks 10-14 and 18-21, for example [76], so this metric cannot be used to measure success in these patients. Additionally, implantation rate has been reported to be problematic for comparing groups, since it is often distorted by non-equal numbers of transferred embryos [75].

The choice of success metric may depend on the aspect of IVF we wish to study. For example, oocyte and embryo assessment (discussed further in Sections 1.2.2.1 and 1.2.4) aim to characterise the potential of the embryo to implant in the uterus, so the success of these techniques is more commonly measured using clinical pregnancy, as opposed to live birth [77, 78]. In this thesis, we will use clinical pregnancy as the definition of success, taking care to specify when other outcome measurements are considered.

Despite the considerable use and development of IVF in the last 40 years, the majority of treatments do not result in clinical pregnancy, and improvements in live birth rates have begun to stagnate at around 30% in Europe [79]. The live birth rate per embryo transferred in the UK, plotted in Figure 1.4, has been increasing steadily over the years for women using their own oocytes. However, this rate is still only 32% for patients under 35, a figure which falls drastically as patient age increases, down to only 3.7% in patients aged 45-50 [7]. Whilst using donor eggs increases live birth rates to between 31-34% across all age ranges [7], the majority of IVF treatments do not result in live birth.

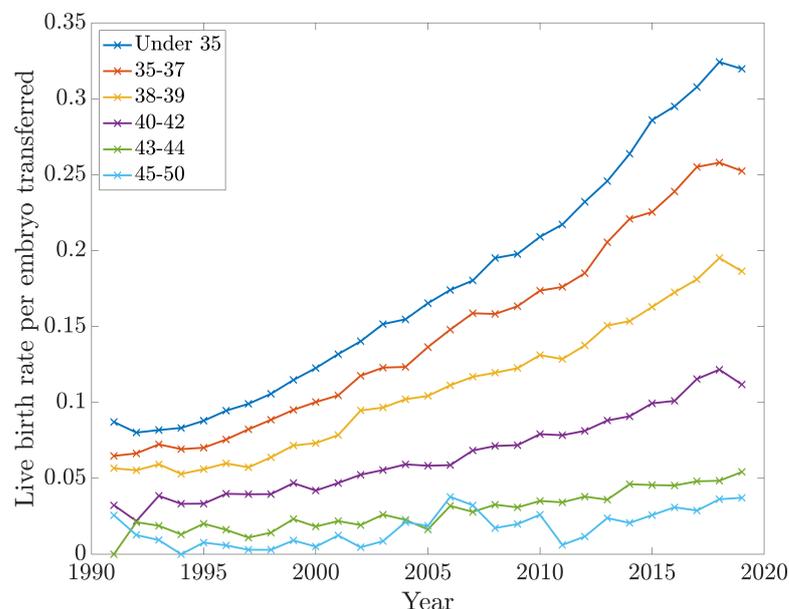


Figure 1.4: The live birth per embryo transferred rate, reported by the Human Fertility and Embryology Authority (UK) between 1991 and 2019. Reproduced from [7].

Outside of the UK, some international IVF outcome statistics are collected in

Table 1.1, which suggests that only around 19% of embryo transfers worldwide result in a live birth, meaning patients are unlikely to get pregnant at their first IVF treatment. Many patients undergo repeated IVF cycles with the hope to get pregnant, but evidence suggests that the likelihood of clinical pregnancy decreases with each unsuccessful cycle [80], implying many patients may never get pregnant. Going through IVF can be incredibly difficult, both physically and psychologically [81], and failed treatments can result in diminished self-esteem and depression [82].

Location, year (reference)	Outcome	Notes
Africa, 2019 [83].	38.4% clinical pregnancy rate per embryo transferred.	Under 35s make up higher demographic in Africa, inflating outcome [84].
Europe, 2022 [85].	34.1% birth rate per embryo transferred.	49.6% for donor embryos.
Latin America, 2020 [86].	6.9-36.9.% delivery rate per embryo transferred, depending on patient age.	
USA, 2020 [87].	23.0% live birth rate per cycle.	More specific reporting available by state.
World, 2022. [85]	19% birth rate per embryo transferred.	ESHRE estimate, older reports suggest 5-29% live birth per fresh cycle [88].

Table 1.1: IVF success rates reported across world-wide registries. No registry exists for Asia to date [84]. Additionally, reporting differences such as the choice of success metric and patient demographics may account for variation between reporting bodies.

Another key challenge in IVF, compounded by the expectation that multiple treatments are necessary, is the considerable expense incurred by each patient. Treatment is often prohibitively expensive for many patients [89]. There is considerable diversity in IVF funding models and policies across the world, motivated by a combination of safety concerns, legal considerations and rationale on the cost-effectiveness of treatments [90]. State-funded options exist, but do not cover demand; in the UK, for example, the National Institute for Health and Care Excellence (NICE) restricts state-funded IVF based on patient age and lack of success in alternative treatments, and allows a maximum of 3 cycles [91]. The proportion of publicly funded treatments in the UK was around 28% in 2020 [44], although this value varies significantly by country, from 62% in Scotland to only 20% in some parts of England [7]. Even when public funding is available, long waiting lists often lead patients to seek IVF treatment in private clinics [92–94]. Private funding thus forms the majority of treatments in the UK, as well as in Europe [95] and the USA [29].

The cost of a single IVF treatment varies depending on the clinic and country. In 2010, ESHRE estimated the costs of a single fresh IVF cycle to range from £2,114 in Belgium to over £8,591 in the USA [96]. These prices are increasing; the average cost of an IVF cycle in the USA jumped from £9,781 in 2014 [97] to £15,145 in 2021 [98]. In the UK, the HFEA estimates that a single IVF cycle costs £5,000 [99]. The considerable cost of IVF treatments is a source of socio-economic inequality between patients who can afford to treat their infertility, and those who cannot [89].

Improving IVF success rates lies at the heart of ART research, with benefits for patients and clinics. Understanding the cause of the embryo failing to implant in IVF, and improving treatment for women who respond poorly to ovarian stimulation during IVF, are listed as the top two priorities for research in ART [100]. ESHRE announced in 2022 that interdisciplinary collaborations involving the use of mathematics are key to innovation in reproductive medicine [101]. Having described some of the challenges faced in IVF, we will now present some mathematical approaches to improving success rates in IVF.

1.2 Mathematical challenges in IVF

Attempting to discuss every mathematical challenge in IVF is beyond the scope of this work, but equally, we aim to demonstrate the wide variety of challenges faced in the pursuit of improving clinical pregnancy rates. Many of the mathematical challenges in IVF focus on a specific stage of the IVF process (see Figure 1.1), so we order our review chronologically by the stage of IVF being targeted. Of these challenges, we highlight oocyte selection (Section 1.2.2.1), optimising cryopreservation protocol (Section 1.1.1) and post-thaw embryo selection (Section 1.2.6) to be relevant to the work undertaken in this thesis. To provide context on additional challenges beyond the scope of this work, additional information is presented on other topics, namely optimal COS strategy (Section 1.2.1), sperm selection (Section 1.2.3), embryo selection (Section 1.2.4) and logistics and operations optimisation (Section 1.2.7).

1.2.1 Personalising controlled ovarian stimulation (COS)

Human females are born with all the oocytes capable of participating in reproduction during their life; this is referred to as the ovarian reserve, which decreases with age [102]. These oocytes are found in the ovaries, in a stage of arrested development and surrounded by a layer of pregranulosa cells in a structure called the primordial follicle [103]. In a normal menstrual cycle, a follicle is recruited from the reserve as a result of complex hormonal signalling [102–104], resulting

in a single oocyte being released from the ovary. In IVF, however, a large number of retrieved oocytes are required to maximise the success rate. The number of retrieved oocytes is increased by COS, using hormonal medication that encourages the release of multiple oocytes. Although COS has potentially detrimental effects on oogenesis, embryo quality, endometrial receptivity and perinatal outcomes [105], the need to retrieve multiple oocytes necessitates its implementation. The comparison of oocytes retrieved following COS, with those released in natural cycles, is not relevant to the objectives of this thesis, as COS is a core aspect of modern IVF treatment [105].

Good response to COS is a key factor for IVF success. Predicted pregnancy rate increases when some optimal number of oocytes is retrieved, a value which depends on the stimulation protocol used; for mild stimulation, the optimum is 5 oocytes, compared to 10 oocytes being optimal in conventional protocols [106]. Overstimulation results in ovarian hyper-stimulation syndrome (OHSS), which lowers the number of good embryos available to transfer and, hence, reduces the clinical pregnancy rate [106]. A number of different COS protocols have been described as a result of varying the dosage of medication, the strategy used, and even the production method of the hormonal medication [107, 108]. Achieving the optimal response is difficult, because patients are heterogeneous and respond differently to stimulation protocols [11].

A ‘one size fits all’ approach is not sufficient to increase IVF success rates for all patients, so a key challenge in the clinic is developing personalised COS protocols [107]. This has been cast as a mathematical optimisation problem, and has been well described in [109]. COS protocols which are easy to manage, and can ensure a sufficient number of oocytes, are desirable. These objectives must be balanced by offering ‘patient friendly’ protocols that minimise side effects, encourage compliance, and minimise the risk of OHSS.

Mathematical exploration of this challenge can be partitioned into two distinct approaches. The first approach is statistical [11, 110–113]; the success of different strategies is evaluated in patients grouped by characteristics [11, 111] or genetics [113]. A patient would then be assigned the strategy which has historically performed best in patients with similar characteristics. Statistical approaches are limited, however, because the exact molecular mechanisms at play are poorly understood, and the complex statistical interaction between a high number of variables that affect COS response makes it difficult to quantify the significance of each variable [113].

A second approach to optimise COS uses ordinary and delay differential equations to model hormone levels in patients [114–118]. These models can be used to

experiment with different COS protocols, but current models do not match experimental data [118], which could be attributed to an incomplete model, poorly selected parameters or inability to fully simulate realistic protocols [118]. In both the statistical and modelling approaches, there is important and interesting work to be done in order to maximise the potential of COS.

1.2.2 Health assessment and selection in IVF

In every IVF cycle, embryologists are required to select which oocytes, sperm and embryos should be used. These decisions, which are fundamental to the success of IVF [100], are all motivated by the assessment of quality, requiring knowledge of which factors indicate ‘good quality’, and the correct tools to measure these indicators. It is a great mathematical challenge to develop new tools with which we can accurately assess the quality of oocytes, sperm and embryos.

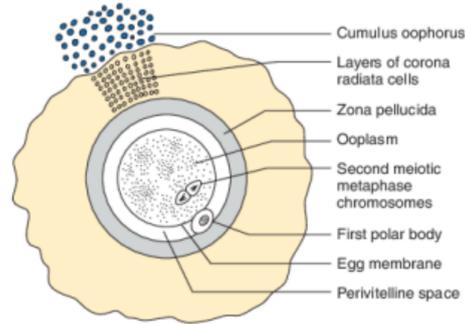
Before exploring selection challenges in more detail, we highlight the important requirement that clinical assessment methods in IVF must be non-invasive, meaning they have zero detrimental impact on the oocyte, fertilisation, subsequent embryo development, ability to establish a pregnancy or the expected lifespan and health of the offspring compared to non-ART conceived individuals [119]. With this in mind, we will first review mathematical works in oocyte health assessment.

1.2.2.1 Oocyte assessment

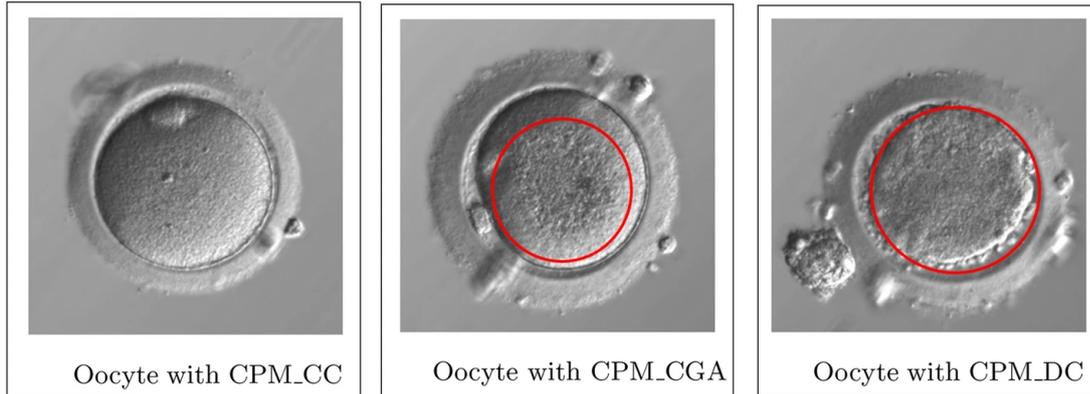
Following a successful COS protocol, a cohort of oocytes is available for fertilisation. Whilst fertilisation requires both sperm and oocyte, there is evidence suggesting that the quality of the oocyte primarily determines the viability of the resulting embryo [119–122]; this means oocytes must be carefully assessed in order to identify those most likely to develop when fertilised, and information on the quality of each oocyte may be useful in embryo selection later during treatment. Identifying poor quality oocytes could also be a mechanism for avoiding fertilising those which have a high risk of failure, reducing the number of excess embryos and, therefore, reducing wasted resources [109].

A popular non-invasive method of oocyte assessment is through describing the oocyte morphology, the physical appearance of the oocyte under a microscope [109, 119, 122]. A schematic diagram of the human oocyte, taken from [123], is given in Figure 1.5a, and microscope images of human oocytes in Figure 1.5b taken from [124].

The visual identification of morphological features and anomalies is common practice in the clinic [122, 125–127]. For example, grading of the cytoplasm is



(a) Diagram of a human oocyte, taken from [123].



(b) Microscope images of human oocytes which have been stripped of the cumulus and corona cells, taken from [124]. CPM is used as shorthand for cytoplasm, with CPM_CC on the left representing a healthy oocyte with a smooth, homogenous cytoplasm surface. The middle and right images represent oocytes with morphological abnormality, where CPM_CGA demonstrates cytoplasmic granularity and CPM_DC has a darkened cytoplasm.

Figure 1.5: The human oocyte: (a) schematic diagram. (b) microscope images.

often considered [122, 124], as shown in Figure 1.5b, where the colour and texture of the cytoplasm are assessed. Both [109] and [122] are comprehensive reviews of the use of morphological oocyte assessment, which agree that studies reporting on the correlation between abnormalities and outcome are often contradictory or unclear. With the exception of very severe defects, abnormalities often do not predict outcome when considered in isolation [109, 122]. The weak predictive power of abnormalities is partly attributed to their high incidence rate, with one study reporting at least one form of abnormality in up to 70% of retrieved oocytes [125]. The clinical significance of these abnormalities remains, thus, uncertain; studies of morphological assessment of oocytes seem to yield conflicting conclusions on their relation to embryo viability [128–132]. For example, darkness and discolouration of the oocyte have been reported as both detrimental [129] and unrelated [130] to outcome. Furthermore, many oocytes with ‘normal’ morphology also do not result in pregnancy, despite leading to healthy-looking embryo in the lab following fertilisation [125].

Current oocyte morphology assessment methods appear to prioritise identifying singular features which deviate from the expected morphology of a healthy oocyte, before testing whether their correlation with clinical pregnancy rates. Instead, we propose tackling this problem in reverse; statistical exploration of large, labelled datasets to identify combinations of morphokinetic features which correlate with improved, or diminished, clinical pregnancy outcomes. Such an exploration is a form of pattern recognition, and would hence be well suited to machine learning.

In general, despite morphological parameters remaining the main criteria by which oocyte health is assessed, they are limited by being subjective and dependent on the judgement and expertise of the embryologist [122]. A proposed, albeit invasive, alternative is chemical analysis of both the contents of the cytoplasm, and the material excreted by the oocyte into the surrounding media [119, 125], which have a more firmly established relationship with the oocyte health. Another alternative, non-invasive approach is to use specialist imaging techniques, such as fluorescent lifetime imaging [133] or Raman spectroscopy [134], which can probe chemical markers inside the cytoplasm that may in turn be correlated with oocyte health. Unfortunately, these chemical and imaging approaches have not been implemented in the clinic thus far, since they either lack human trials, use technology which is not available in the clinic, are too expensive or time-consuming to incorporate into regular treatment protocol, or are invasive [122].

There is a gap within current oocyte assessment techniques, since no technique proposed thus far is non-invasive, objective, reproducible and viable in the clinic [122]. A new non-invasive approach proposed in [135] bridges this gap, however, by correlating the rate of visible movements in the cytoplasm of embryos immediately after fertilisation, imaged using high speed cameras, with known metabolic indicators of viability, which are not visible. Understanding this work first requires that we introduce some biology of the cytoplasm.

The cytoplasm describes the contents inside the cell wall around the nucleus [136], and is made up of membrane-enclosed organelles, such as the mitochondria and endoplasmic reticulum, in addition to a concentrated aqueous gel known as the cytosol. The cytosol is made up of a crowded system of particles of varying size [136], from small molecules such as carbon compounds and monomers, to larger particles such as proteins, nucleic acids and polysaccharides [136]. The cytoplasm additionally contains a network of protein filaments called the cytoskeleton, which are crucial to the shape, organisation and mechanical properties of the cell [136].

In oocytes, sperm entry triggers rhythmic spasms in the cytoplasm [135], which can be observed under the microscope [135]. It was shown in experiments

[135] that the pattern and speed of these movements are correlated with Ca^{2+} waves in the cytoplasm, which are indicative of embryo viability. Additionally, the movement of the cytoplasm was related to cytoskeleton health in [135], through the use of invasive experiments in which oocytes were injected with messenger RNAs, which encode for two proteins that bind either to myosin or actin, which together form a crucial element of the cytoskeleton. The proteins encoded by the messenger RNAs were tagged with forms of Green Fluorescent Protein (GFP), such that increased actin or myosin activity could be measured as an increase in the intensity of fluorescence. Spikes in fluorescence were correlated with spikes in the cytoplasmic movement speed, indicating that cytoplasmic movement could be used as an alternative method to measure actin and myosin activity. Crucially, direct measurements of Ca^{2+} waves or cytoskeletal properties usually requires invasive techniques, such as potentially harmful dyes and excitation light [135]. The latter work is, thus, an exciting development, because imaging of the cytoplasm can be achieved using red light, which is known not to affect development of the embryo and therefore is much less invasive than other methods of studying properties of the cytoskeleton.

The movement of the cytoplasm therefore has the potential to be used as a proxy variable for biological health markers in zygotes, with a proven relationship to embryo viability. Mathematical modelling of the mechanochemical coupling between the Ca^{2+} dynamics and the observed spasms in cytoplasm of the zygote is therefore of interest, since understanding this coupling could lead to the development of a day 1 assessment tool for IVF treatments. Whilst [135] focusses on very early zygotes, immediately after fertilisation, even earlier analysis of the movement in the cytoplasm of oocytes, prior to fertilisation, could be used to select the highest quality oocyte for treatments.

For example, cytoplasmic movement profiles of mouse oocytes generated using Particle Image Velocimetry (PIV) have been used to identify the absence of a ring of Hoechst-positive chromatin surrounding the nucleolus, which leads to failure in development of the resulting embryo [137]. Whilst this work demonstrates that the movement of the oocyte cytoplasm contains predictive information which can be used to identify poor quality embryos, we propose to extend this analysis to explore whether we can rank-order oocytes by quality.

The cytoplasm of eukaryotic cells is in constant motion, held in a state of non-equilibrium by active processes [138], where the term ‘active’ implies the expenditure of energy by the cell. One of these processes is *active diffusion* [138], a biological analogue of the random movement, known as Brownian motion, which derives from thermal motion [139]. The term *active diffusion* has been coined because particles in the cytoplasm move in a random directional manner, which

is characteristic of diffusion, but over much larger distances than Brownian motion would cause in a viscoelastic medium of such high viscosity [140, 141]. Active diffusion primarily acts through the actin-myosin interactions in the cytoskeleton [140, 142], as so called ‘myosin motors’ interact with actin filaments to create non-thermal pulses which pull the cytoplasm, and the particles contained therein, along with it. Detail of the exact mechanics of interaction between actin and myosin is beyond the necessary information required for this work; we need only be aware that as the filaments interact, motor molecules enact a directed linear movement along the filament they interact with, resulting in cytoplasmic drag which follows this movement.

There is a complex relationship between thermal and mechanical effects within the definition of active diffusion, with thermal movement driving the likelihood of interaction between actin and myosin filaments which, in turn, create mechanical movement [140]. Active diffusion also exhibits a complex relationship with the viscosity of the cytoplasm, which controls the response of the cytoplasm to the short non-thermal pulses enacted by actin-myosin interactions. Elucidating the exact nature of these interactions is beyond the scope of this thesis; we need only know that considered over the scale of the entire cell, active diffusion is characteristically equivalent to Brownian diffusion on a faster timescale. The presence of active diffusion in the oocytes of various species is well documented, including mice [143], *Drosophila* (fruit flies) [142, 144] and frogs [145].

A second process of interest is *cytoplasmic streaming*, in which a large-scale net flow results in the mixing of the cytoplasm [146, 147]. It is a hypothesis that cytoplasmic streaming is important to cell metabolism [148]. Streaming is observed commonly in plants and algae [148, 149], as well as in the oocytes of *Drosophila* [142, 150], mice [151] and humans [152]. Whilst the presence of streaming in human oocytes is clearly significant in the context of IVF, we note that understanding the streaming observed in mouse oocytes is also important, since mouse models are used to inform human studies [153].

There are several types of cytoplasmic streaming [154], as shown in Figure 1.6; *Drosophila* oocytes, for example, demonstrate a semi-random pattern where multiple flows in different directions can appear across the cell, although the speed of the flows are faster at the anterior than the posterior of the cell, moving down a gradient [142]. In mouse oocytes, the streaming takes on a bi-directional circulatory flow [151], where the cytoplasm streams upwards through the centre of the cell from some ‘drain’ pole to a ‘source’ pole, flowing down the sides of the cell back to the drain again. Different types of streaming are seen in plant cells, such as ‘spiralling’ flows in algae [147]. All three of these types of streaming are

visualised in Figure 1.6, although of the three, we are primarily interested in the mouse oocyte due to its applications in IVF research.

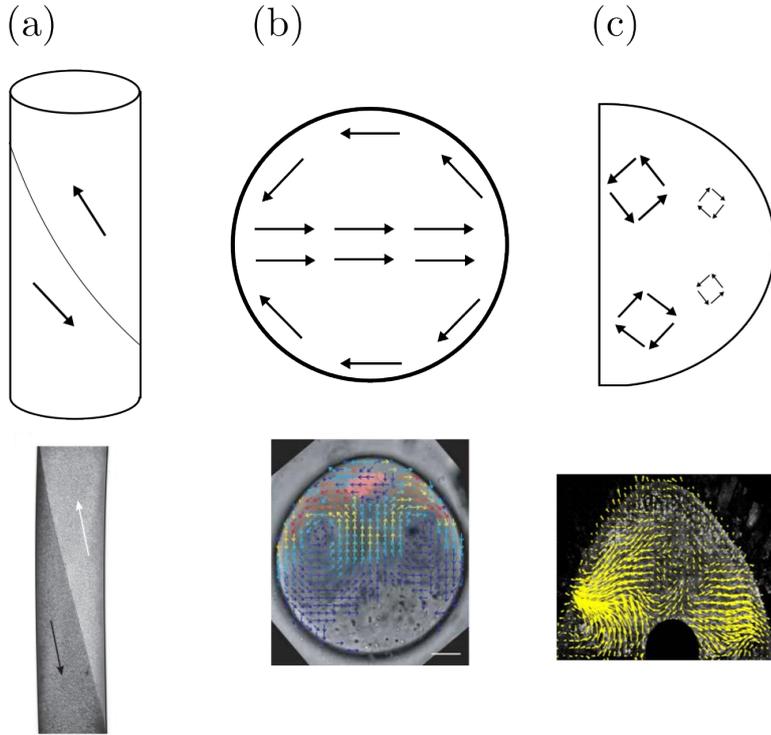


Figure 1.6: Schematics of the cytoplasmic streaming behaviour in various cells, alongside annotated microscope images of each showing the direction of the local velocity field. (a) *Chara corallina* algae, schematic adapted from [154] and microscope image taken from [154], (b) mouse oocyte, schematic adapted from [151] and microscope image taken from [155] (c) *Drosophila* (fruit fly) oocyte, schematic adapted from [142, 154] and microscope image taken from [150].

Together, active diffusion and cytoplasmic streaming represent the combined effects of diffusion and advection. Hence, advection-diffusion equations are used to describe cytoplasmic streaming [147, 150, 156, 157],

$$\frac{\partial C}{\partial t} + \mathbf{v} \cdot \nabla C = \nabla \cdot (D \nabla C), \quad (1.2.1)$$

where C is the concentration of some chemical component of the cytoplasm, D ($\mu\text{m}^2/\text{s}$) is an unknown (not necessarily constant) diffusion coefficient and \mathbf{v} ($\mu\text{m}/\text{s}$) is the velocity field. The cytoplasm is frequently modelled as an incompressible Newtonian fluid [147, 150, 154], which allows the velocity field to be modelled by the Navier-Stokes equations. The validity of this assumption has been studied in depth; despite the cytoplasm being densely packed with polymers, non-Newtonian behaviour is only observed in the cytoplasm of *drosophila* oocytes at timescales less than 0.4 seconds. Whilst in depth study of conditions precluding incompressible or compressible behaviour in oocytes is available [158],

simulations based on the Navier-Stokes equations have accurately reproduced the flows seen in the cytoplasm of the embryos of *C. elegans*.

In addition to the assumption that the cytoplasm is an incompressible Newtonian fluid, steady state flow is also assumed, and the system is initially considered to be at rest [159]; boundary conditions are dependent on the cell type being modelled [147]. A common and appropriate assumption is that cytoplasmic flows have low Reynolds number. The Navier-Stokes equations then reduce to the (linear) Stokes flow [147, 154], i.e

$$\nabla \cdot \mathbf{v} = 0, \quad \eta \nabla^2 \mathbf{v} = \nabla p, \quad (1.2.2)$$

where η is the dynamic viscosity and p the pressure of the fluid. Given estimates for D , Equations (1.2.1) and (1.2.2) can be numerically solved to determine \mathbf{v} and C , and hence the movement of the cytoplasm in an oocyte [151]. However, we are interested in the inverse problem. Since active diffusion and cytoplasmic streaming both require the expenditure of energy, it is possible that oocytes with more energy display faster cytoplasmic movements. Hence, a new non-invasive oocyte assessment method could be developed by measuring D and \mathbf{v} from a given time-lapse movie of an oocyte, and correlating these values with some measure of metabolic activity, such as ATP concentration, with a known relationship to embryo potential.

For this assessment method to be implemented in the clinic, the potentially non-monotonic relationship between movement parameters and oocyte quality should be elucidated, but investigating this correlation is only possible if we have appropriate tools to extract movement parameters from the cytoplasm. Equations (1.2.1) and (1.2.2) provide a model for the streaming behaviour we expect to see in the cytoplasm, which can be compared with live imaging data which visualises streaming. However, new mathematical tools are required to understand the relationship between cytoplasmic movement and the images such movement will generate. These tools, which will combine statistical-mechanical modelling of cytoplasmic movement with mathematical descriptions of image formation and image analysis techniques, may be used to work backwards and determine the parameters which resulted in the observed data. More detailed exploration of the importance of mathematical descriptions of image formation in the pursuit of oocyte health assessment can be found in Chapter 5.

A number of image analysis methods have been previously used to track intracellular movement in oocytes, namely single particle tracking (SPT) [160, 161], particle image velocimetry (PIV) [135, 150, 152, 162], and more recently, differential dynamic microscopy (DDM) [142]. More detailed descriptions of these methods are presented in Chapter 2. Of these techniques, we are interested in

DDM, which has performed well in classifying the velocity of advection-diffusion processes [142, 163, 164]. In particular, DDM has been successfully used to quantify the rates of active diffusion and cytoplasmic streaming in *Drosophila* oocytes [142]. However, to date, the statistics reported from DDM have not been used to predict oocyte health; this poses another exciting mathematical challenge, to probe the relationship between auto-correlative statistics returned by DDM analysis, and measurements of oocyte health and viability. In Chapter 3, we introduce methods for simulating images, which are used in Chapter 4 to explore the relationship between the statistics determined in DDM analysis, and the parameters driving the simulated movements.

1.2.3 Sperm selection

The health of the sperm is also assessed in IVF [165]. The sperm cell is a highly adapted DNA delivery system, essentially a ‘stripped down’ cell designed to propel itself through an aqueous medium using its flagellum [166]. These are the smallest cells in the human body [167], lacking many of the typical cytoplasmic organelles found in eukaryotic cells, such as ribosomes or endoplasmic reticulum. Once the sperm has reached an oocyte, its purpose is to bind to the zona pellucida and penetrate it [166, 168]. Sperm research could, thus, be broken down into two stages; reaching the oocyte and penetrating the oocyte.

Mathematical modelling has been extensively used to study sperm motility. Highly complex fluid-solid interactions dictate the coupling of biomechanics of the sperm to the surrounding fluid [169–172]. We will not review these models in this work; we will instead focus on measurements from the clinic used to assess sperm quality. These include semen volume, concentration, total sperm count, motility, vitality (whether non-motile cells are alive or dead), pH and morphology [165]. Together, these measurements influence whether or not to use ICSI (intracytoplasmic sperm injection), by identifying cases where poor motility or inability to bind with the zona pellucida would prevent natural fertilisation. Mathematics also plays a role in this sperm classification process, through computer-aided semen analysis [165], which uses automated cameras and software to analyse microscopy images of semen, measuring parameters much faster than humans [173]. However, computer-aided analysis is imperfect, with unreliable results being reported in samples which have higher viscosity, are debris-filled or clumped together [173]. Additionally, morphology appears difficult to classify through automated means [173]. The solution to these challenges may be novel AI approaches; despite these showing promising results [174, 175], they could still benefit from further development. Additionally, coupling sperm microscopy with

the available biomechanical models is not a current feature of semen analysis in the clinic [165], but could provide new insight into semen analysis.

1.2.4 Embryo selection

Having assessed the quality of the oocytes and sperm, fertilisation can be undertaken either by mixing of oocytes and sperm, or through ICSI. The average number of oocytes retrieved per patient is between 10-12 [176, 177], and around 70% of oocytes are expected to fertilise successfully [178], which means a cohort of 7-9 embryos can be expected. However, the transfer of more than one embryo per cycle is strongly correlated with a significantly increased risk of multiple pregnancy, with a reported 33.1% multiple birth rate in double embryo transfer patients compared to 0.8% in single embryo transfer patients [179]. Multiple pregnancy is associated with increased risk of serious medical conditions [180] and increased mortality, particularly in older patients [181]. Subsequently, the HFEA has set a directive that less than 10% of IVF pregnancies should result in multiple births, leading to the prevalence of elective Single Embryo Transfer (eSET) [182].

In reality, whilst eSET is strongly encouraged, multiple embryo transfer can be undertaken when deemed appropriate, depending on factors such as patient age and the quality of the available embryos [36]. We focus on the encouraged practice of eSET, however, in which only one embryo is transferred in the first IVF cycle, regardless of whether that cycle is fresh or frozen. The unselected embryos can be cryopreserved if they are of good quality, and used in future cycles, as discussed in Section 1.1.1. However, by improving the embryo selection criteria, we may be able to identify the ‘best’ embryo, the one with the greatest potential to implant. This may reduce the need for many IVF cycles, and improve clinical pregnancy rates [183, 184].

In many ways, assessment of embryos resembles that of oocytes, since morphological features are studied for both [184–186]. The typical morphology of an embryo at the blastocyst stage is presented in Figure 1.7. The key features of a blastocyst are [187, 188]:

- The inner cell mass (ICM), a dense ball of cells inside the trophectoderm which will develop the embryo itself [187].
- The zona pellucida, a thick extracellular matrix that surrounds the embryo [188].
- The trophectoderm, a thin layer of cells that will eventually contribute to the placenta [187].

- The blastocoel, a fluid filled cavity inside the trophectoderm [187].
- The perivitelline space, not shown in Figure 1.7, which describes the space between the trophectoderm and the zona pellucida. Although the space itself is not considered during embryo quality assessment [36], the abnormal presence of extracellular cytoplasm in the perivitelline space has been evaluated as a measurement of embryo quality [189].

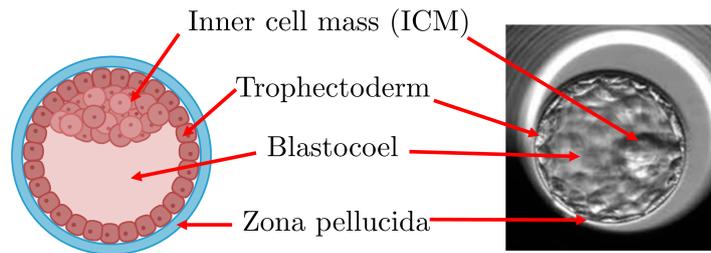


Figure 1.7: The morphology of a typical embryo at the blastocyst stage. The cells of the embryo are contained inside the zona pellucida, with the trophectoderm surrounding the inner cell mass (ICM) and fluid-filled blastocoel. (Left) Diagram produced using BioRender. (Right) Image of an embryo graded as 3AA using the Veeck and Zaninovic system, a slightly modified form of the Gardner system (using which, this embryo would be classified as 4AA [36]) . Taken from [190].

Embryo classification often takes the form of a standardised, categoric grading system, which consists of a set of quality labels. The Gardner system is the most well known classification scheme [77, 78, 191]; size and development stage are assigned an integer from 1 to 6 (least to most developed), whilst embryos of grade 3 or above receive two additional classifications, for the ICM and trophectoderm respectively. These grades, either A, B or C, reflect the number and arrangement of cells, from highly dense to sparse [192]. Hence, the ‘top scoring embryos’ are 3AA and above. Although embryo grading is widely used in the clinic [193], it is highly subjective. Studies that involved classifying the same embryos in different clinics have highlighted this limitation [194, 195], which considerably weakens the predictive power of assessment.

A further drawback of embryo grading is that the grade is based on morphology at a single time point [185], despite the growth of an embryo being a dynamic process with dramatic changes observed during the 5-6 days of development before transfer or freezing [196, 197]. Research has, thus, moved away from studying morphology at one time point, and focuses on an emerging set of non-invasive

assessment techniques using time-lapse microscopy. Where morphology described the structure of the embryo in a static image, the term ‘morphokinetics’ is used to describe these time-dependent morphological observations and measurements [57, 198, 199].

Morphokinetic studies have identified a number of potential indicators of embryo viability, resulting in several hierarchical classifications of embryos which propose to predict viability. Three implementations of time-lapse morphokinetic analysis stand out. Chavez et al. [200] collected several morphokinetic parameters, including the time between mitosis (cell division) stages and the duration of first cytokinesis (cell cleavage during cell division), to identify combinations of measurements that correlate with genetic abnormality in the embryo, which is negatively correlated with viability [200]. Additionally, automated boundary detection was used to measure and categorise the level of cell fragmentation, the appearance of membrane-bound extracellular cytoplasm in the perivitelline space [189], which is considered a predictor of negative outcome.

Secondly, Meseguer et al. [57] proposed that a smaller variance in the timing of key embryo developmental stages is an indicator of embryos more likely to implant successfully, and provided a hierarchical classification that could be trialled as a selection protocol in the clinic.

Similarly, Wong et al. [61] identified the duration of the first cytokinesis, in addition to the timings of the P2 (time between 1st and 2nd mitoses) and P3 (time between 2nd and 3rd mitosis) events, as key development markers, using automated tracking algorithms to determine the timing of the events [201, 202]. They concluded that normal embryos obeyed strict timing in all three events; large deviations from these timings correlated with developmental failure. These studies demonstrate that time-lapse imaging can be used as an extension of the morphological assessment currently undertaken in IVF clinics, adding event timings to the collection of available methods used to select embryos.

Reviews of time-lapse incubation [55, 56, 60, 198, 203–205] paint a mixed picture of current practice. Integrating time-lapse incubators into the clinic has not yielded immediate improvements to clinical outcomes, but whilst many identify potential for time-lapse imaging to provide additional information for embryo selection, the general consensus is that we lack the tools to currently make full use of these data [56, 60, 198, 203–205]. One barrier to utilising time-lapse imaging is a lack of randomised control trials, which are difficult to undertake due to stringent ethical restrictions and lack of available human embryos, compared to retrospective studies. Additionally, a significant number of confounding variables such as patient age, the number of previous cycles and the variation in COS protocols additionally determine clinical outcomes, making it more difficult to

evaluate the effect of embryo selection. Finally, there is a shortage of reliable, quantifiable measurements, related to embryo quality by a sound scientific basis, that could be used to inform embryo selection [55, 56, 60, 203–205].

In the grading of both embryos and oocytes, there is emerging interest in automation. One commercially available system, EEVA [206], claims to predict viability from time-lapse imaging of Day 3 embryos, by measuring the timing of P2 and P3 events [206, 207]. There is also great interest in AI-based tools for embryo selection, summarised in [208], to perform a number of tasks including embryo quality assessment [190] and counting the number of cells in the embryo [209], as well as predicting genetic abnormality [210], implantation potential [211] and pregnancy outcome [70].

Whilst these appear promising, AI implementation in the IVF clinic faces several limitations. Given the success rate for IVF treatments is no greater than 30%, training data frequently has significant class imbalance (where the negative outcomes significantly outweigh the positives), which can result in overestimation of negative outcome by an AI model [212]. Furthermore, where different studies use a variety of network architectures and evaluation metrics, comparing studies can be challenging, so it is difficult to determine which AI approaches show the greatest performance [212]. Many AI approaches, though apparently successful in the data they are trained in, fail to perform well when generalised to other datasets [213]. Overcoming the limitations of current AI approaches is an important and interesting mathematical challenge; key hurdles include finding effective methods for overcoming class imbalance, beyond simple downsampling of the dominant class, and implementation of metrics to make AI decisions more ‘explainable’. Explainability is a crucial characteristic of AI models in medical settings, which can highlight important flaws or bias in the AI decision making process and improve reliability of trained networks. Many of the studies discussed here lack consideration of potential bias in the training dataset; these approaches would benefit from even very simple analysis of the accuracy of the network within subpopulations, to guard against discrimination. Furthermore, mathematical understanding of the design and architecture of AI models may yield new AI models which are robust to highly variable time-lapse image data available to clinics.

Although these statistical challenges are beyond the scope of this thesis, we explore the utility of AI in embryo selection in Chapter 7 in a different way. Rather than using AI as a decision making tool, we instead train a neural network to perform automated segmentation of time-lapse images, yielding measurements which we correlate statistically with clinical pregnancy outcome to identify a new potential metric of embryo viability.

1.2.5 Optimal cryopreservation protocols

In this section, we will examine cryopreservation protocols in more detail. Modern cryopreservation practice has a survival rate of around 90% for oocytes [214], and over 95% for embryos [215, 216]. Though the risk is small, some oocytes and embryos may not survive the cryopreservation process. Freezing is lethal to most living systems [217], with a variety of mechanical and chemical factors resulting in cell death at low temperatures [218, 219]. A key contributor to cell death during cryopreservation is the presence of large ice crystals forming in the cytoplasm [220], which expand during their formation and result in mechanical damage to the cell. Cryopreservation techniques in IVF have therefore been designed to minimise the formation of such crystals. There are two fundamentally contrasting cryopreservation techniques, slow freezing and vitrification [31]. Both techniques involve some form of Cryoprotective Agent (CPA). Slow freezing relies on controlled cooling and the dehydration of cells over a large time, around 10 minutes [221], to minimise ice formation [31, 222], whereas vitrification uses high cooling rates to cause the tissue to enter a glass-like state of thermodynamic in-equilibrium [31, 222, 223], over a timespan of seconds [224]. Vitrification has been used for cryopreservation of embryos and gametes for decades [225], and is now the dominant method of cryopreservation in IVF [226].

The key factors determining the success of vitrification are high cooling rates [222, 224, 227], high rates of warming during oocyte/embryo thawing after storage [228], the application of high concentration permeating CPAs [223, 229] in low volumes (typically less than $1\mu\text{l}$) [224, 230, 231], and minimising the distance between the cells and the liquid nitrogen [232]. A number of different devices and techniques have been developed to facilitate vitrification [233–235], but the device we will focus on in this work is the Open Cryotop[®] [224], which is used in over 2200 clinics distributed over more than 100 countries [236]. Further detail on the Open Cryotop[®] device can be found in Chapter 6.

Whilst the Open Cryotop[®] protocol is very detailed and specific about the timing of each vitrification stage, it allows flexibility on the number of embryos or oocytes cryopreserved on a single device, with up to four being permitted [237]. The optimal number of embryos or oocytes vitrified at once may therefore be questioned. A retrospective study [238] reports that the post-thaw survival rate of embryos is affected by the number of embryos vitrified simultaneously, and states that further work is required to determine the optimum number.

When more than one embryo or oocyte is vitrified simultaneously, the relative position becomes important. Hence, it is not only the number of embryos or oocytes, but also their arrangement that may affect the cooling process. Embryos

and oocytes are likely to affect the cooling rates of their near-neighbours more so than those spaced far apart. Hence, it is possible that even with a constant number of embryos or oocytes, there exists an optimal arrangement the embryologist should aim for. Some examples of arrangements of four oocytes, based on real clinical observations, are given in Figure 1.8.

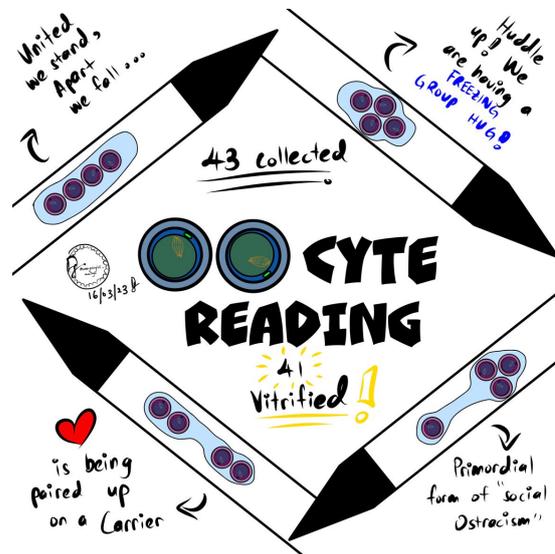


Figure 1.8: An illustration by senior embryologist, Diana Tain, depicting some of the observable patterns which groups of four oocytes have formed when vitrified together on the same Cryotop[®] device. Taken with permission from [239].

Improving cryopreservation requires a multidisciplinary approach, drawing on knowledge from Engineering, Biology, Chemistry, Physics and Mathematics. The main mathematical approach to this date has been the development of computational PDE models to predict the cooling and warming rates of different cryopreservation devices [240–245]. Experimental work [217, 228, 229, 246] allows cooling and warming rates to be related to the relative risk of damage to the vitrified embryos and oocytes. Combining these experiments with mathematical models allows for the design of in silico experiments, which can predict the outcomes we would observe in real experiments. In Chapter 6, we use a computationally-resolved mathematical model to demonstrate that there is no optimal number or arrangement of oocytes or embryos on the Cryotop[®] device, with respect to ensuring a sufficiently high cooling rate.

1.2.6 Post-thaw embryo assessment

In FET, the embryos used are often referred to as ‘frozen-thawed’ [53], which reflects the process by which embryos are taken out of storage and warmed before transfer [73, 247]. The warming protocol for the Open Cryotop[®] resembles the vitrification process in reverse [248]. The embryos are removed from storage and

warmed in a series of different chemical solutions, to resume their physical and chemical processes. Thawed embryos are then transferred to the uterus in the same way as fresh ones. So far, we have considered assessment of oocytes and embryos, in order to choose the best candidate from a cohort. The selection process in thawed embryos is subtly different, because when a number of frozen embryos are available, clinicians will thaw only a single embryo for transfer. In the event of cell death being identified by the embryologist following cryopreservation, another embryo is thawed only when the first is damaged. The average clinical pregnancy rates of FET in Europe is 22.3% [216], which is comparable with the overall clinical pregnancy rates for all treatments (see Table 1.1); this suggests that cryopreservation is not detrimental to success rates, and further suggests that embryos which do not survive cryopreservation are identified correctly by embryologists, since their use would reduce success rates.

Although an embryo may survive vitrification, this does not guarantee its success in IVF treatments. Assessment of embryos following thawing from cryopreservation is therefore an interesting and important challenge, with the potential to improve clinical pregnancy rates by identifying and discarding poor quality embryos in favour of a second candidate. This selection process has different implications than gamete or embryo selection, however; each embryo is the result of considerable financial, psychological and time investment for the patient, so the unnecessary waste of good quality candidates is an important factor to be considered. Therefore, designing reliable metrics that can measure the implantation potential of a thawed embryo is critical to make the right decision at this stage.

Since FET occurs after oocytes and embryos have been assessed prior to cryopreservation, their respective grades at earlier stages can be used to choose which embryo to thaw first. However, the assessment prior to vitrification cannot fully determine the viability of the embryo post-thaw [66], due to the damage which can be incurred through vitrification. During cooling, cells undergo significant volume reduction as water is replaced with cryoprotectant; during warming, they re-expand through rehydration [78, 249]. Shrinking, expansion or their combination can result in damage [78, 249, 250]. Therefore, we need to develop assessment methods for accurately determining embryo viability after thawing.

Some efforts in this direction have been undertaken [67, 69, 71–73, 78, 247, 249–254]. These works utilise the same machinery involved in oocyte and embryo assessment prior to cryopreservation, namely time-lapse or single-frame microscopy. Time-lapse imaging generates complex spatio-temporal data, so many studies tend to make simplifying assumptions by restricting the number of features of interest, such as the formation of the blastocoel [67, 71, 73] (see Figure 1.7), as well as the complexity of each feature, and the number of time points

sampled at. Below, we classify the literature by the method used to simplify and analyse the time-lapse (or single frame) image data available in the clinic.

The most restrictive approach involves reducing the measurements on each embryo down to only a small set of discrete classifiers, thus removing the dependence on continuous measurements. These approaches wait until full re-expansion stops, labelling embryos depending on whether they have expanded to reach their original, pre-vitrification size within two to four hours [249, 252]. Failure to re-expand within this time-frame is considered a negative indicator of the potential to implant [69]. A slightly more complex method is adopted in [253], where a hierarchical structure is defined based on a combination of two binary classifiers, namely whether the embryo has re-expanded back to its original size at a fixed cut-off time of 5-6 hours, and whether the blastocoel has formed or not. Considering ‘good’ embryos to be those which developed to the hatching stage, the four groups defined by the binary classifiers were shown to exhibit different rates of development, and it was implied that this classification could be used to distinguish viable embryos from those which have shrunk [253].

In other approaches, pairs of time points are sampled. In [251], the initial cell area and zona pellucida thickness were measured, as well as the minimum zona pellucida thickness, the maximal cell area and the time taken to begin re-expansion. Each of these measurements was assigned a threshold (i.e. the zona pellucida being less than, or greater than, $14\mu\text{m}$ in thickness), such that the authors could compare implantation and survival rates of the embryos based on each measurement. The results tended to suggest that embryos with thinner zona pellucida and higher initial area implant more frequently, with the timing of the re-expansion of the embryo not being a statistically significant indicator of implantation rate. This work was further developed in another study [250], which recorded the time after warming at which re-expansion begins, as well as the initial and maximal area of the blastocyst during re-expansion. The key development of this work was the creation of a scorecard, similar in style to the Gardner criteria described in Section 1.2.4. A single grade from A to D is assigned to each embryo, depending on a mixture of measurements of the ICM area and trophectoderm morphology. It has been concluded in [250] that the assigned label, and the time at which re-expansion begun, does not predict implantation rate.

Another common approach using two time points is to consider either the ‘mean expansion rate’, by dividing the change in size of features the time elapsed [247, 255], or an absolute change in embryo size over a fixed time [78]. In these works, the cross-sectional area of the ICM is measured by treating the ICM as an ellipsoid with perpendicular major and minor axes. These measurements

give an approximation of the rate of re-expansion of each embryo. All three works [78, 247, 255] suggest that faster re-expansion rates correlate with improved implantation rate. A similar metric is the time taken for complete re-expansion [69], which, unlike the previous expansion rate measurements, is independent of initial and final embryo size. The study of 144 women suggested that shorter expansion times correlate with increased clinical success, with a mean time to re-expansion of 0.6 ± 0.42 hours for live births, compared to 1.3 ± 1.19 for non-pregnancy [69]. We note, however, that the high standard deviations of these reported outcomes mean we cannot necessarily predict the outcome of an embryo based on its re-expansion time .

Notably, by only considering the start and end time points for re-expansion, the approaches in [69, 78, 255] cannot observe non-linear expansion rates. Examples of observations of non-linear behaviour are found in both [247] and [250], which count the number of contraction events, where the blastocyst cross-section decreases in size and pulls away from the zona-pellucida following re-expansion. The final class of literature therefore considers *continuous* measurements of embryo features during re-expansion, such as contractions. More frequent contractions were shown in [247] to correlate to poor embryo quality (organelle and genetic abnormality). In contrast, [250] only measured whether one or more contractions had occurred, concluding that the presence of a contraction does not predict implantation rate.

In summary, much of the existing work on the morphology of post-thaw embryos either does not consider morphokinetics during re-expansion [67, 69, 71, 73, 249, 250], or takes into account only a few time points [78, 247, 250, 255]. The only feature considered over multiple time points was contraction during the re-expansion process [247]. The current state of time-lapse imaging in post-thaw embryos shows that there is scope for more mathematical and statistical analysis of post-thaw embryo behaviour *during* re-expansion, rather than *after* re-expansion.

1.2.7 IVF logistics, planning and optimisation

So far, we have discussed a number of mathematical challenges related to the protocols and decision making processes in IVF, motivated by the underlying biological, chemical and physical forces that determine each stage. The IVF clinic itself, however, is a complex organisation, requiring balance and synchronisation of patient interaction and treatment, laboratory management and collaboration with egg and/or sperm banks when donor cells are used. Logistics are important to consider, in order to maximise efficiency and reduce costs.

For example, the wait times involved in COS have been studied in [256]. Long wait times are detrimental to the patient experience, and reduce the number of patients who can be seen by the clinic. They may also increase the risk of healthcare-associated infections [256]. A ‘lean management’ strategy is considered in [256], which allows a clinic to identify various causes of operational inefficiency, such as repeated and unnecessary registration steps, poor utilisation of nurses and even poorly designed office space, leading to reduced productivity. Solving some of these challenges cut total waiting times by more than half in the Yasmin Fertility Clinic in Jakarta, Indonesia, demonstrating the power of logistics planning in the clinic [256].

Another important logistics consideration for the fertility industry is the transport and storage of donor oocytes and sperm, addressed in [257]. Cells which have been vitrified are highly sensitive to temperature changes, making it critical to ensure temperatures never rise above a threshold of -180°C . Despite stringent safety steps being obeyed, a large difference in survival rates was observed in [257] for oocytes transported using different means. The data from this study came from real multi-centre oocyte transports, but future investigations would benefit from predictive mathematical models that facilitate safer, cheaper *in silico* experiments before moving into trials with human cells.

Whilst operations research for IVF clinics are few, other healthcare settings have already been studied in detail [258–260]. Each healthcare setting is different, however, so approaches considered in one field may not be suitable in another; still, we do not need to reinvent the wheel. Hence, collaboration between operations research experts (mathematical modellers) and IVF clinics is an open and exciting avenue in this field.

An increasingly important consideration is the environmental impact of IVF. It is interesting to note that literature on the detrimental impact of pollution on IVF success rates is well documented [68, 261–264], but the impact of IVF treatments on the environment is less understood. Although we could not find any published research on the impact of IVF on the environment, a lecture given as part of the International IVF Initiative (i3) by the embryologist Roisin O’Raghallaigh identified single use plastics, fuel emissions from transport, energy usage by the clinic and more as factors to consider [265]. With the increased popularity of IVF, the environmental impact of IVF will continue to grow, so there is an increasing need to make treatment practices more sustainable.

1.3 Thesis outline

In this thesis, we cannot hope to address every mathematical challenge discussed above; we focus our efforts on a few key challenges. The thesis is divided into two parts.

In the first part, we explore the use of Differential Dynamic Microscopy (DDM) as a non-invasive oocyte health assessment tool. In Chapter 2, we define the DDM algorithm, and apply it to a sample oocyte dataset in 2.3. Based on the results of this analysis, we introduce simulation-driven methods of enhancing the existing DDM algorithm to improve its performance in oocyte data. In Chapter 3, we introduce synthetic image data in the context of DDM, expanding on parameter choices for the appearance of simulated images which ensure the robustness of DDM analysis. Then, in Chapter 4, we use simulated data to test the limitations of current DDM analysis approaches, and propose new simulation-driven fitting regimes which relax some of the assumptions required for DDM analysis, overcoming some of the key challenges faced when applying DDM to oocyte data. Finally, in Chapter 5, we explore a previously unknown interaction between the output of DDM analysis, and image shadowing from the phase contrast microscopy used to generate images in commercial time-lapse incubators. We derive an analytic expression which describes the effect of phase contrast on DDM, from which we conclude that shadowing does not affect parameter fitting when the studied motion is isotropic. We further expand our analysis to consider anisotropic advection-diffusion, highlighting potential challenges which could be faced when implementing DDM in clinical oocyte data.

The second part of this thesis moves away from oocyte selection. In Chapter 6, we develop a computational PDE model (heat equation) of the Open Cryotop[®] device, to predict cooling rates and, hence, show that there is no variable risk of cryogenic injury to vitrified embryos or oocytes in different numbers and arrangements. Finally, in Chapter 7, we tackle the challenge of predicting embryo viability from time-lapse images depicting the re-expansion process after thawing from cryopreservation.

We summarise and discuss the work in this thesis, along with future directions, in Chapter 8.

Chapter 2

Differential Dynamic Microscopy (DDM)

2.1 Introduction

Chapters 2-5 set out to evaluate and further develop Differential Dynamic Microscopy (DDM), applied to time-lapse microscopy data, as a non-invasive oocyte health assessment technique. This chapter has two objectives: to provide a formal mathematical introduction to the theory of DDM, and to evaluate the validity of current DDM analysis approaches when applied to time-lapse oocyte data. First, we will introduce the time-lapse images we analyse in Section 2.1.1. Understanding the data to be analysed is critical to motivating the selection of the analysis technique; we discuss current available image-analysis techniques in Section 2.1.2, drawing attention to their strengths and weaknesses, which justifies using DDM on the available data.

We then provide a review of the theory of DDM in Section 2.2. Whilst we are not the first to review DDM [164, 266, 267], justification of the steps taken when handling equations is often omitted, and many of the important assumptions which must be considered when applying the technique are not discussed. Therefore, our review of DDM aims to be the most comprehensive to date, beginning in Section 2.2.1 with the relevant theory that relates the positions of moving particles to the images that depict them. We then introduce the DDM method in Section 2.2.2, with Sections 2.2.3-2.2.6 focussing on the analytic theory of constructing the so-called ‘intermediate scattering function’ (ISF), whilst Sections 2.2.7-2.2.8.4 discuss algorithmic considerations such as parameter fitting. We briefly summarise the important assumptions made during DDM in Section 2.2.9.

Finally, we apply the DDM technique to the available time-lapse image data in Section 2.3. From this analysis, we conclude that whilst DDM analysis has potential to provide new information related to oocyte health, there are a number of

limitations to overcome, summarised in Section 2.3.5. Tackling these assumptions forms the basis of the work in Chapters 3 and 4.

2.1.1 Time-lapse images of mouse oocytes

From our collaborators Emily Lewis and Karl Swann in the School of Biosciences, Cardiff University, we have access to bright-field images of 11 mouse oocytes. Details regarding experimental conditions for the retrieval of the oocytes, and ethics related to animal work, can be found in Appendix A. Bright-field microscopy passes light through a specimen, with the formation of the image depending on the level of absorption of light by the sample [268]. These images, which are taken once every second, are 2688×2192 pixels in size, where the length of a single pixel is 230nm. A single frame depicting the oocytes is shown in Figure 2.1.

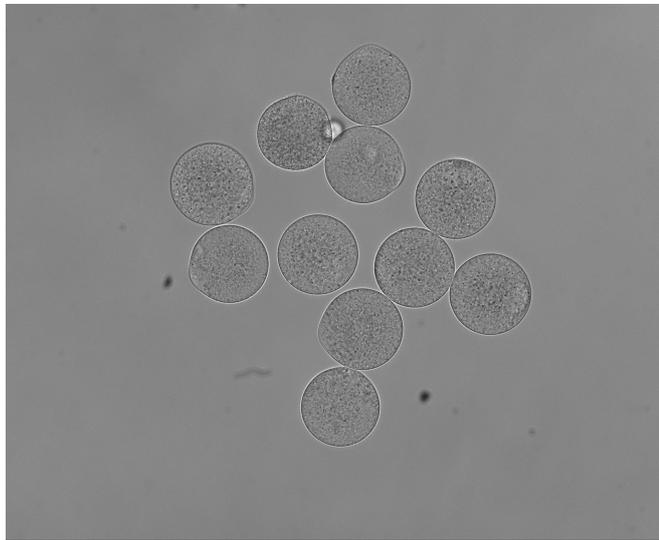


Figure 2.1: The first frame of a bright-field time-lapse movie depicting eleven mouse oocytes. Analysis of this data is presented in Section 2.3. Images provided by Emily Lewis and Karl Swann in the School of Biosciences, Cardiff University.

These oocytes are imaged under two different experimental conditions. The first condition acts as a control experiment, depicting healthy oocytes under standard incubation conditions. The second set of images show the same oocytes 30 minutes after the application of FCCP, a metabolic poison that inhibits the production of ATP in the mitochondria [166]. Without ATP, neither active diffusion or cytoplasmic streaming, described in Section 1.2.2.1, can take place. Hence, cytoplasmic movement will no longer be described by the advection-diffusion equation (Equation (1.2.1)). We therefore expect that the rate of movement in the poisoned oocytes will be slower than in the control oocytes. Determining whether DDM analysis applied to both datasets matches this expectation is a useful way of evaluating the performance of current DDM approaches.

Having introduced the context for the analysis we wish to undertake, we now briefly review image analysis algorithms in Section 2.1.2.

2.1.2 Image analysis algorithms for quantifying movement parameters

First introduced in 2008 [266], DDM probes the spatial cross correlation of multiple particles over time. Applications of DDM include characterising diffusion coefficients in colloidal dispersions [164, 266], measuring advection-diffusion speeds of bacteria [163, 269] and sperm [270], as well as a host of soft matter applications [271–273]. DDM is just one example of a set of image analysis algorithms designed to extract statistical information about the movement of a collection of particles or a fluid, depicted in a series of time-lapse images. When choosing which algorithm to employ, one must consider the set of strengths, weaknesses and conditions accompanying each approach, against the nature of the data being considered. We discuss the two most commonly used alternatives to DDM, single particle tracking (SPT) [274, 275] and particle image velocimetry (PIV) [276, 277], which have both been successfully applied to extract the speed of movement in the cytoplasm of oocytes [135, 150, 152, 160, 162]. By discussing these techniques, we identify the niche that DDM can occupy in the classification of oocyte health. Whilst only a brief account of SPT and PIV is provided here, additional details are given in Appendix B.

2.1.2.1 SPT

SPT assigns a position to each particle in the image in every frame [274, 277]. Each particle is also assigned a unique identity, which means the position from each frame can be collected to form a trajectory for that particle over time. Knowing the position of every particle at every time is the most amount of information that can be extracted from a time-lapse movie, but this is only possible when individual trajectories are identifiable. This first requires that the position of all particles in the frame can be determined, known as *localisation*, excluding systems where particles can move out of the focal plane, or are sub-pixel in size and bulk movement is analysed.

Assuming particle localisation is possible, SPT further requires correct linkage of locations to reconstruct the ‘true’ trajectory of each analysed particle. The certainty with which we may correctly link particles across frames is a function of the optical density, which is the distance between each particle, and the size of the displacements between frames [274], with the exact trade-off being subject

to the type of movement exhibited. This means SPT ceases to be appropriate in optically dense, fast-moving systems.

SPT has been applied to oocyte images to track large endogenous particles in the cytoplasm [160, 161], but the statistical validity of such analysis will be limited by the number of observable particles, with around 30 accessible in a single *Xenopus* (frog) oocyte using epifluorescence microscopy [160]. This number may be reduced further when considering non-invasive imaging techniques; it is noted in [160] that when using Differential Interference Contrast (DIC) microscopy, endogenous probes may be required to perform particle tracking. That means that we are not extracting movement information from the majority of the cytoplasm, wasting much of the available data. Additionally, in the bright-field images of mouse oocytes shown in Figure 2.1, endogenous particles are not well visually contrasted, so SPT will be unable to identify the position of each particle in the frame. Hence, SPT is not suitable for our data.

2.1.2.2 PIV

PIV tracks local velocity fields which create bulk movement throughout the time-lapse movie [277]. Each image is partitioned into a regular grid of ‘interrogation windows’. If particles are subject to a locally coherent velocity field, then particles in the same interrogation window will have approximately the same displacement between consecutive frame. These displacements are estimated by sliding each interrogation window from the first frame over the second frame, until the positions of the particles contained match up.

If SPT yields very specific information, following the trajectory of every particle individually, PIV can be thought of as a more general technique, since it determines the average velocity over a grid on the image. The trade-off for losing such specific information is the relaxation of the sparsity requirement of SPT [277]. PIV is a very popular technique for studying cytoplasmic streaming [135, 150, 152, 162], and is highly appropriate for measuring flows, seeing extensive use in fluid mechanics; a summary of such applications may be found in [278]. Whilst one of the strengths of PIV is its ability to describe vector fields which are spatially anisotropic across the whole image, this still requires the assumption that movement is *locally* coherent. This assumption is not true of Brownian motion, and therefore also not true of active diffusion, as particles move independently. Therefore, PIV may not be suitable for studying isotropic movement of independently moving particles, particularly diffusion [279].

2.1.2.3 DDM

SPT and PIV are both powerful techniques which offer unique insights into cytoplasmic movement and flows. However, features of the oocyte data we aim to analyse violate key requirements of both SPT and PIV. SPT requires a sparse distribution of particles, but the cytoplasm is a dense, crowded environment [136], with no uniquely identifiable particles to act as tracers (see Figure 2.1). PIV relaxes the assumption on low particle density, but in turn requires a locally coherent velocity field, violated by the active diffusion observed in the cytoplasm [140–143]. Hence, we turn to DDM, which does not require low particle density, and has been shown to perform well at characterising diffusive processes [164, 266].

Compared to the locally averaged information provided by PIV, DDM yields movement statistics which are globally averaged across the whole image. Losing this local information is necessary in order to relax the assumption of local spatial coherence, even with particle densities too high for SPT. These averaging processes utilised in DDM have associated conditions on the movement studied, namely ergodicity [164, 266] and isotropy [266, 267] (although some work has been undertaken to adapt DDM to specific anisotropic behaviours [280]). Of interest to us is the application of DDM to classifying cytoplasmic advection-diffusion in *Drosophila* oocytes [142], yielding comparable results to PIV analysis of the oocyte [150]. Crucially, the strengths highlighted in this work are that DDM is robust to noise, can be linked with prior understanding of the system being modelled, and is capable of characterising both advection and diffusion simultaneously. We note, however, that DDM does not yield movement statistics directly like SPT and PIV; it instead algorithmically determines the DDM matrix, \mathcal{D}_R [266], which must be interpreted by either drawing inference from its shape, or comparing its structure to some model determined by the distribution from which displacements are drawn from [142, 164, 280–282].

High-quality, open source implementations of DDM are available, allowing the user to convert any movie into a DDM matrix [164]. It is always possible to apply parameter fitting to this matrix to determine some collection of movement parameters. However, there are many potential sources of error. If the incorrect model choice is made, these fitted parameters will not correspond to any observable phenomena in the movie. Also, even with the correct model, experimental design conditions such as the image sampling rate, image size and relative magnitudes of the underlying types of movement could affect the accuracy of the parameter fitting. Heuristic arguments to combat these sources of error have been developed [164, 266], but ultimately, the question of whether DDM can be

used in a given dataset to produce accurately determined, meaningful movement statistics remains open.

Before exploring whether DDM can be used in non-invasive oocyte health assessment, we first provide a detailed overview of the theory of DDM, drawing particular attention to the assumptions and requirements of DDM in current practice.

2.2 DDM Theory

Before presenting our own review of the theory of DDM, we highlight some existing reviews which were instrumental in developing this work. A technical overview of DDM is found in [267], whilst a more ‘hands-on’ pedagogical application is given by [164]. Finally, [279] gives a thorough examination of recent developments in the theory and application of DDM prior to 2017. Whilst these works allow a practical understanding of DDM application, fundamental details which justify the steps involved are often lacking or not widely discussed, and many assumptions remain hidden or unjustified. Hence, this section will present the most comprehensive summary of DDM to date. The steps presented in this section are not novel, but rather represent the union of several source documents to provide a single coherent explanation of the origin of DDM. We provide both the theory which justifies the technique, including explicit statement of the assumptions involved, and practical tips and advice for improving the performance of DDM.

2.2.1 Inferring particle movement from image data

The input data for DDM is always a time-series image stack, which depicts a collection of moving particles (or, if the particle number is sufficiently high, a fluid). The displacements of the particles are assumed to be described by a probability density function whose general form is known, but with undetermined parameters we wish to quantify. We make two assumptions about the system we study, drawing attention to why these assumptions are important as they become relevant. The first assumption is that movement is ergodic [266], which means a sufficiently large collection of particles can represent the average statistical properties of the entire process [283]. Alternatively, ergodicity can be interpreted as the property that a single particle, observed for a sufficiently long time, will be sufficient to describe the behaviour of the whole system. The second assumption is that the system is stationary [266], which means the laws which define movement in the system are invariant in time [284], such that particle displacement is a function only of the time between observations, and not the time of the initial

observation. The importance of these assumptions will be highlighted as we progress through the theory of DDM.

We first establish the relationship between a collection of particles, and the corresponding images that depict them. Consider a collection of N_p particles which are assumed to be homogeneous, meaning they have the same size, shape and visual appearance. Let $\mathbf{r} = (x, y)$ denote the 2D Cartesian coordinate system parallel to the imaging plane, and let z denote the position in the axis orthogonal to the imaging plane. We may describe the positions of a finite set of particles at time t by the sample density function, [285]

$$c(\mathbf{r}, z, t) = \sum_{n=1}^{N_p} \delta(\mathbf{r} - \mathbf{r}_n(t))\delta(z - z_n(t)), \quad (2.2.1)$$

where $\delta(\mathbf{r})$ is the Dirac delta function, whilst \mathbf{r}_n and z_n denote the position of the center of the n^{th} particle in the imaging plane and orthogonal axis respectively. Alternatively, if N_p is sufficiently large and we divide by the domain volume, we could consider the concentration of a chemical or compound within the fluid, which we denoted by $C(\mathbf{r}, z, t)$ in Equation (1.2.1).

The input image stack, labelled $I(\mathbf{r}, t)$, consists of N_t digital images that have $L \times W$ pixels, where L is the length and W is the width of the image. We define the set of times at which each frame is generated to be $T = \{t_1, t_2, \dots, t_{N_t}\}$, where without loss of generality, $t_1 = 0$. We will assume for simplicity that images are regularly spaced in time, such that $t_{n+1} - t_n = \Delta t_{\text{min}}$ for some constant Δt_{min} . We assume the relationship between the time-dependent image intensity distribution $I(\mathbf{r}, t)$, and particle positions $c(\mathbf{r}, z, t)$, to be Linear Space-Invariant (LSI) [286], i.e

$$I(\mathbf{r}, t) = i_0 + \int \int K(\mathbf{r} - \mathbf{r}', z) c(\mathbf{r}, z, t) d^2 \mathbf{r}' dz, \quad (2.2.2)$$

where i_0 is an approximately constant background illumination and $K(\mathbf{r} - \mathbf{r}', z)$ is known as the *Point Spread Function* (PSF) [267], which defines the appearance of the image of a particle located at (\mathbf{r}, z) . We can break down the LSI relationship into two parts, based on the function of each integral. The innermost integral over \mathbf{r}' represents a convolution of the particle density $c(\mathbf{r}, z, t)$ with the PSF, where the convolution operator \star on two functions f and g is defined by [287, 288]

$$f \star g = \int f(\mathbf{r}) g(\mathbf{r} - \mathbf{r}') d\mathbf{r}' \quad (2.2.3)$$

Convolution pairs the visual appearance of a particle defined by the PSF, with the position of the particles defined by $c(\mathbf{r}, z, t)$. The second integral in Equation (2.2.2), over z , represents a flattening or projection of the 3D appearance of particles into the 2D imaging plane, by taking the sum of the contributions from the convolution of the PSF and particle density distribution through each thin slice of the 2D sample, each of which has some corresponding depth z attenuating the contribution to the resulting image. If we generate images which are themselves only slices of the 3D appearance, then the domain of integration over z is only a single point, corresponding to the height at which the slice is taken. We note that in the LSI definition, the contribution of a particle at some given position (\mathbf{r}, z_1) is unaffected by the existence of a particle located at (\mathbf{r}, z_2) , implying no shadowing or obscuring effects are seen. A graphic representation of the creation of an LSI image is given in Figure 2.2.

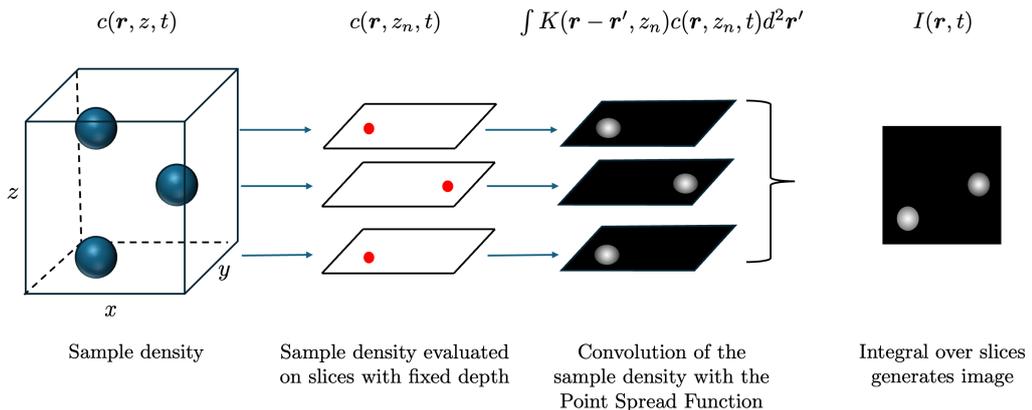


Figure 2.2: Example of the construction of a Linear Space Invariant (LSI) image. By definition, these images would exhibit no shadowing or obstruction when objects appear above or below one another in the z axis. Reproduced from [267].

It is relevant to note here that the LSI definition is continuous, but images are clearly discrete. This discretisation is typically performed by a Charge-Coupled Device (CCD), which converts brightness to charge across small regions corresponding to each pixel [289]. If the area denoted by a pixel is Ω , then one might consider the brightness of the pixel to be proportionate to the integral of $I(\mathbf{r}, t)$ over Ω . For our application, we will typically make the assumption that the brightness is approximately constant over the area defined by each pixel, an assumption which gets stronger when the pixel size is much smaller than the particle

size. Clearly, as pixels become smaller and more numerous, we have a better *resolution*, and are able to describe the exact location and appearance of the particle in further detail, which is advantageous for DDM.

When particles are much smaller than the pixel size (where particle size is encoded within the PSF in this case, since the particles are treated as point particles by the sample density function), we typically lose the ability to localise particles, knowing only that they must be inside the pixel somewhere. Small particle size can be offset by an increased number of particles, such that the brightness of a pixel corresponds to the number of particles contained therein; we refer to previous discussion of how Equation (2.2.1) can be converted to consider particle concentration, rather than position. We have omitted a great deal of complexity here regarding the transition from continuous to discrete by making the assumption that information is not lost in the capture of a digital image; additional insight into this topic is contained in Chapter 3, where we explore conditions on the relationship between particle size and pixel size in order to ensure the validity of DDM analysis.

The LSI condition is important to the context of DDM compared to imaging alternatives that use laser-based inference to access the same information we seek from images [267], although this is not relevant to our application. The importance of the LSI condition here is that it is sufficient to justify that the images contain information about the position of the particles we wish to study, although whether that information is accessible or not is a separate consideration that features prominently in our work.

Having defined the position of the studied particles, and the images which depict them, we must also describe how the particles move. A convenient way to do so is the van Hove function, derived from the Born approximation scattering formula and defined by [290]

$$G(\Delta\mathbf{r}, \Delta t) = \frac{1}{N_p} \left\langle \sum_{i,j=1}^{N_p} \int \delta(\Delta\mathbf{r} + \mathbf{r} - \mathbf{r}_j(\Delta t)) \delta(\mathbf{r} - \mathbf{r}_i(0)) d\mathbf{r} \right\rangle, \quad (2.2.4)$$

where $\langle \cdot \rangle$ denotes the ensemble average or temporal average, which are equivalent for ergodic systems [285]. The van Hove function is a temporal extension of the pair correlation function [290], which describes the distribution of objects throughout a medium. The van Hove function is defined in 3D, but the direction orthogonal to the imaging plane is not important to the next few steps, so we omit the z -dependent features here for simplicity.

The function within the ensemble average queries the number of particles separated in space by $\Delta\mathbf{r}$, and in time Δt . Whilst for any given finite sample this is essentially 0 everywhere, we draw attention here to the assumption of

ergodicity, such that an ensemble average is sufficient to reproduce the average behaviour of the system. The van Hove function therefore represents the average correlation of particles in space and time. Another physical interpretation of the van Hove function is that $G(\Delta\mathbf{r}, \Delta t)d\mathbf{r}$ is equal to the expected number of particles j located within the region $d\mathbf{r}$ about a point $\Delta\mathbf{r}$ at time t , given a particle i is located at the origin at time $t = 0$ [285, 291]. The van Hove function naturally separates into the sum of the ‘self-correlating’ and ‘distinct’ parts, $G_s(\Delta\mathbf{r}, \Delta t)$ and $G_d(\Delta\mathbf{r}, \Delta t)$, for $i = j$ or $i \neq j$ respectively, as follows [290]:

$$G(\Delta\mathbf{r}, \Delta t) = G_s(\Delta\mathbf{r}, \Delta t) + G_d(\Delta\mathbf{r}, \Delta t),$$

where

$$G_s(\Delta\mathbf{r}, \Delta t) = \frac{1}{N_p} \left\langle \sum_{i=1}^{N_p} \int \delta(\Delta\mathbf{r} + \mathbf{r} - \mathbf{r}_i(\Delta t)) \delta(\mathbf{r} - \mathbf{r}_i(0)) d\mathbf{r} \right\rangle,$$

and

$$G_d(\Delta\mathbf{r}, \Delta t) = \frac{1}{N_p} \left\langle \sum_{j \neq i=1}^{N_p} \int \delta(\Delta\mathbf{r} + \mathbf{r} - \mathbf{r}_j(\Delta t)) \delta(\mathbf{r} - \mathbf{r}_i(0)) d\mathbf{r} \right\rangle,$$

The self-correlating term is normalised, describing the probability that any given particle undergoes a displacement $\Delta\mathbf{r}$ during the time period Δt , which is the main property we wish to extract from the images. However, the distinct term is more complex to tackle, describing the number of particles we expect to find separated by \mathbf{r} in space and Δt in time where the two particles considered are different; when $\Delta t = 0$, this is equal to the pair correlation function, $g(\mathbf{r})$ [290]. G_d depends on both the particle movement and the manner in which particles are distributed throughout the domain. For example, particles initially distributed on a regular lattice will have a strong distinct correlation over distances equal to the separation within the lattice than particles uniformly randomly distributed throughout the system.

When particle trajectories are statistically independent, such that the displacement of a particle is not affected by the position or displacement of any other particle, the distinct part is dominated by the self-correlating part for all $\Delta\mathbf{r}$ since there is no correlation between the positions of separate particles. We therefore make the assumption throughout this work that $G_d(\Delta\mathbf{r}, \Delta t) = 0$, such that we need only consider the self-correlating term, G_s , which describes the probability that any given particle undergoes a displacement $\Delta\mathbf{r}$ during the time period Δt . Under this assumption, we can recast the aim of DDM analysis to be characterising of the self part of the van Hove function. For more thorough

insight into the properties of the van Hove function when the distinct part is non-zero, see [285] and [290].

The reason for introducing the van Hove function is that the relationship between the particle density, $c(\mathbf{r}, t)$, and the van Hove function, $G(\Delta\mathbf{r}, \Delta t)$, is given by [285]

$$\begin{aligned} G(\Delta\mathbf{r}, \Delta t) &= \frac{1}{N_p} \left\langle \sum_{i,j=1}^{N_p} \int \delta(\Delta\mathbf{r} + \mathbf{r} - \mathbf{r}_i(\Delta t)) \delta(\mathbf{r} - \mathbf{r}_j(0)) d\mathbf{r}' \right\rangle \\ &= \frac{1}{N_p} \left\langle \int c(\Delta\mathbf{r} + \mathbf{r}, \Delta t) c(\mathbf{r}, 0) d\mathbf{r} \right\rangle \\ &= \frac{1}{\rho} \langle c(\Delta\mathbf{r}, \Delta t) c(0, 0) \rangle, \end{aligned} \quad (2.2.5)$$

where $\rho = N_p/V$ is the particle density describing the expected number of particles found in a given volume V . The final step in Equation (2.2.5) requires the assumption of ergodicity, which gives translational invariance, and stationarity, which means we no longer depend on the time of the initial observation. We note that the relationship in Equation (2.2.5) includes the distinct and self correlating parts of the van Hove function; this underpins a critical element of DDM analysis, that we cannot distinguish between the contributions of G_s and G_d in any subsequent analysis of the image stack, explaining the need for dominance of the self-correlating part over the distinct part.

If we could determine the location of particles in the image stack, Equation (2.2.5) would immediately give us the van Hove function. However, when dealing with optically dense samples, the particle density may not be immediately available from the images [292]. Instead, we can extract information about the van Hove function from images in the Fourier space, by introducing a structure known as the normalised image intermediate scattering function (ISF) which acts as a bridge between the van Hove function and image stack.

Let the spatial Fourier transform \mathcal{F} of a given function $f(x)$ be denoted by $\tilde{f}(q)$, where q represents the angular frequency variable, such that [293]

$$\begin{aligned} \mathcal{F}(f(x)) &= \int_{-\infty}^{\infty} f(x) e^{-ixq} dx \\ &= \tilde{f}(q), \end{aligned}$$

where we note the absence of a consideration of the normalisation factor. Denoting the complex conjugate by \cdot^* , element-wise multiplication (Hadamard product) by \odot , and an average over time by $\langle \rangle_t$, we define the normalised image ISF for LSI images by [267, 285]

$$F(\mathbf{q}, \Delta t) = \frac{\langle \tilde{I}^*(\mathbf{q}, t + \Delta t) \odot \tilde{I}(\mathbf{q}, t) \rangle_t}{\langle |\tilde{I}(\mathbf{q}, t)|^2 \rangle_t},$$

which is a normalised Fourier transform of the autocorrelation of the images. We can redefine the normalised ISF in terms of the particle positions, $c(\mathbf{r}, z, t)$, and PSF, $K(\mathbf{r} - \mathbf{r}', z)$, by substituting the definition of an LSI image from Equation (2.2.2) into the normalised ISF. Before doing so, it is useful to present the convolution theorem. This states that for two functions, f and g , the Fourier transform of their convolution given by $f \star g$ is equal to the Hadamard product of Fourier transforms of f and g respectively. Formally, [294]

$$\mathcal{F}(f \star g) = \mathcal{F}(f) \odot \mathcal{F}(g). \quad (2.2.6)$$

We previously identified that the LSI definition in Equation (2.2.2) can be expressed as convolution of the PSF and particle density, implying that by the convolution theorem (Equation (2.2.6)),

$$\tilde{I}(\mathbf{q}, t) = \int \tilde{K}(\mathbf{q}, q_z) \odot \tilde{c}(\mathbf{q}, q_z, t) dq_z. \quad (2.2.7)$$

Hence,

$$\begin{aligned} F(\mathbf{q}, \Delta t) &= \frac{\langle \tilde{I}^*(\mathbf{q}, t + \Delta t) \odot \tilde{I}(\mathbf{q}, t) \rangle_t}{\langle |\tilde{I}(\mathbf{q}, t)|^2 \rangle_t}, \\ &= \frac{\int |\tilde{K}(\mathbf{q}, q_z)|^2 \odot \langle \tilde{c}(\mathbf{q}, q_z, t + \Delta t)^* \odot \tilde{c}(\mathbf{q}, q_z, t) \rangle_t dq_z}{\int |\tilde{K}(\mathbf{q}, q_z)|^2 \odot \langle |\tilde{c}(\mathbf{q}, q_z, t)|^2 \rangle_t dq_z}. \end{aligned} \quad (2.2.8)$$

Unless otherwise stated, throughout this work, matrix multiplication of Fourier transformed values will be element wise, so we drop the explicit \odot notation. Assuming the medium is homogeneous, we expect a uniform distribution of particles throughout the system, and therefore the value of $|\tilde{c}(\mathbf{q}, q_z, t)|^2$ should be time invariant; we can set $t = 0$ without loss of generality. When motion is confined to a single 2D plane orthogonal to the z direction, or the PSF does not depend strongly on z , Equation (2.2.8) simplifies to

$$F(\mathbf{q}, \Delta t) = \frac{\langle \tilde{c}(\mathbf{q}, t + \Delta t)^* \tilde{c}(\mathbf{q}, t) \rangle_t}{|\tilde{c}(\mathbf{q}, 0)|^2}. \quad (2.2.9)$$

Assuming truly 2D dynamics is not a particularly realistic assumption, but fortunately in [286] it has been shown that in many practical cases, relationship (2.2.9) can be recovered for a range of q , under the condition that movement in the z direction is small relative to displacement within the xy plane. Whether this assumption is valid in oocytes is not necessarily known; whilst a practical implementation of DDM in oocytes exists [142], for directed streams as shown in Figure 1.6 it is possible that the stream is oriented in the axis orthogonal to the imaging plane, potentially invalidating the assumption.

We now observe that the right hand side of Equation (2.2.5) is the autocorrelation of c , a specific case of convolution where both functions are the same. Applying the convolution theorem in Equation (2.2.6), it follows that the Fourier transform of Equation (2.2.5) yields

$$\frac{1}{\rho} \langle \tilde{c}(\mathbf{q}, \Delta t) \tilde{c}^*(\mathbf{q}, 0) \rangle = \tilde{G}(\mathbf{q}, \Delta t), \quad (2.2.10)$$

Combining Equations (2.2.9) and (2.2.10), we derive the key relationship between the van Hove function and the normalised image ISF, [267],

$$F(\mathbf{q}, \Delta t) = \frac{\tilde{G}(\mathbf{q}, \Delta t)}{\tilde{G}(\mathbf{q}, 0)}.$$

We may further simplify this relationship when particle trajectories are independent; In this case, $G(\mathbf{q}, 0) = 1$ [285], and we deduce

$$F(\mathbf{q}, \Delta t) = \tilde{G}(\mathbf{q}, \Delta t). \quad (2.2.11)$$

Thus, the normalised image ISF defined on the image stack $I(\mathbf{r}, \Delta t)$ is equal to the Fourier transformed van Hove function $G(\Delta \mathbf{r}, \Delta t)$. This relationship is the key motivation for DDM analysis, which is described in section 2.2.2.

2.2.2 The DDM tensor, $\mathcal{D}(\mathbf{q}, \Delta t)$

The objective of DDM analysis is to numerically determine $F(\mathbf{q}, \Delta t)$ from our image stack $I(\mathbf{r}, \Delta t)$, so that we may use Equation (2.2.11) to deduce unknown constant parameters in $\tilde{G}(\mathbf{q}, \Delta t)$ if an analytic form of the van Hove function is known. However, $F(\mathbf{q}, \Delta t)$ often may not be calculated from $I(\mathbf{r}, \Delta t)$ using the relationship in Equation (2.2.8), due to the possible influence of static artefacts. Stationary objects are clearly not the objective of our analysis, but will very strongly correlate with themselves over time at small distances, and thus dominate the calculated signal in $F(\mathbf{q}, \Delta t)$. For this reason, DDM analysis applies a series of processing steps to convert the image stack into a useable form known as the DDM matrix. An overview of this process is presented in Figure 2.3, with each step fully explained below.

Static artefacts are removed in DDM analysis by using image differences (step (I) in Figure 2.3), defined by

$$\Delta I(\mathbf{r}, t, \Delta t) = I(\mathbf{r}, t + \Delta t) - I(\mathbf{r}, t),$$

where $t \in T$ is the time of the first frame and $t + \Delta t \in T$ is the time of the subsequent frame. Next, the image differences are converted into a multi-dimensional array which approximates the function we refer to as the DDM tensor, also known

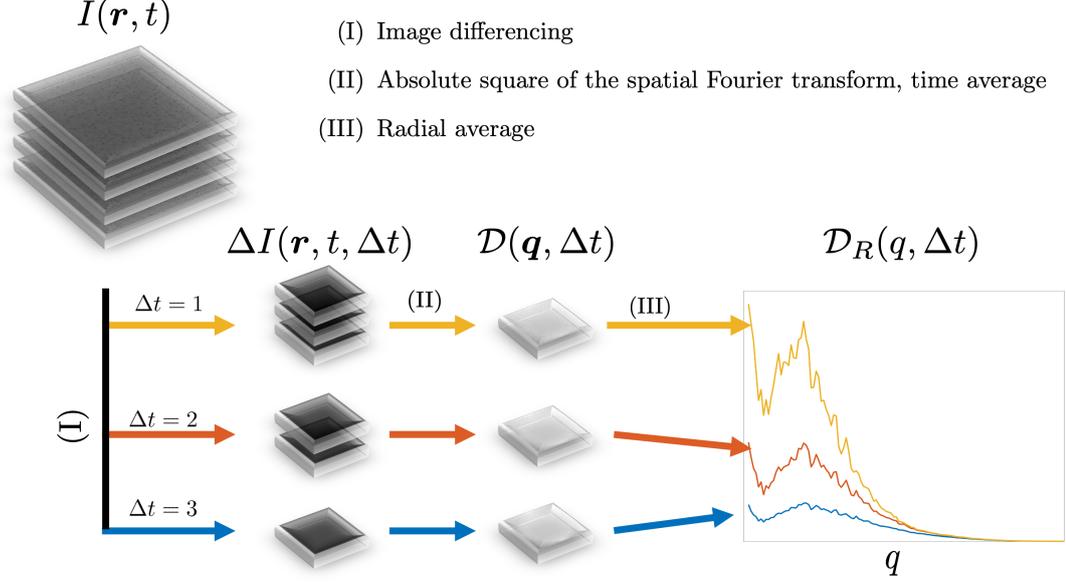


Figure 2.3: An overview of the algorithm which converts the image stack $I(\mathbf{r}, t)$ into the DDM matrix $\mathcal{D}_R(q, \Delta t)$, from which movement parameters in the van Hove function $G(\Delta \mathbf{r}, \Delta t)$ may be determined. Step (I) constructs image differences over all pairs of images, which can be grouped by time lag. The time-averaged absolute square of the spatial Fourier transform (Equation (2.2.12)) constructs the DDM tensor $\mathcal{D}(\mathbf{q}, \Delta t)$, which is reduced to a matrix by radial averaging (Equation (2.2.16)).

as the image structure function [295] (Step (II) in Figure 2.3), defined by the formula

$$\mathcal{D}(\mathbf{q}, \Delta t) = \left\langle |\Delta \tilde{I}|^2 \right\rangle_t = \frac{1}{M} \sum_{k=1}^M \left| \Delta \tilde{I}(\mathbf{q}, t_k, \Delta t) \right|^2, \quad (2.2.12)$$

where M is the number of image pairs we choose to sample from our image stack, satisfying the following relationship

$$0 < M \leq \frac{\max(T) - \Delta t}{\Delta t_{\min}}.$$

The DDM tensor can be expanded and written as [164, 266]

$$\begin{aligned} \mathcal{D}(\mathbf{q}, \Delta t) &= \left\langle \Delta \tilde{I} \Delta \tilde{I}^* \right\rangle_t \\ &= \left\langle \left(\tilde{I}(\mathbf{q}, t + \Delta t) - \tilde{I}(\mathbf{q}, t) \right) \left(\tilde{I}^*(\mathbf{q}, t + \Delta t) - \tilde{I}^*(\mathbf{q}, t) \right) \right\rangle_t \\ &= 2 \left\langle \left| \tilde{I}(\mathbf{q}, t) \right|^2 \right\rangle_t \left(1 - \frac{\Re \left(\left\langle \tilde{I}^*(\mathbf{q}, t + \Delta t) \tilde{I}(\mathbf{q}, t) \right\rangle_t \right)}{\left\langle \left| \tilde{I}(\mathbf{q}, t) \right|^2 \right\rangle_t} \right) \\ &= 2 \left\langle \left| \tilde{I}(\mathbf{q}, t) \right|^2 \right\rangle_t \left(1 - \Re \left(\tilde{G}(\mathbf{q}, \Delta t) \right) \right), \end{aligned} \quad (2.2.13)$$

where Equations (2.2.8) and (2.2.11), allow us to replace the term in the rightmost bracket with the Fourier transformed van Hove function, and \Re denotes the real part. Equation (2.2.13) is most commonly expressed in the form

$$\mathcal{D}(\mathbf{q}, \Delta t) = A(\mathbf{q})(1 - f(\mathbf{q}, \Delta t)), \quad (2.2.14)$$

$$\text{where } f(\mathbf{q}, \Delta t) = \Re\left(\tilde{G}(\mathbf{q}, \Delta t)\right), \quad (2.2.15)$$

and $A(\mathbf{q}) = 2 \left\langle |\tilde{I}(\mathbf{q}, t)|^2 \right\rangle_t$ is a scaling parameter that depends only on the Fourier transform of individual frames in the image, and hence is defined by the PSF [266]. In addition to knowing the PSF, deriving an accurate analytical approximation of $A(\mathbf{q})$ would require knowing the concentration of particles within the frame, as well having zero noise and static artefacts. For this reason, approximation of $A(\mathbf{q})$ from first principles is limited to synthetic image data [286, 292], and in practice, numerical calculation of the scaling term is necessary.

The term $f(\mathbf{q}, \Delta t)$ is referred to in the literature interchangeably as the auto-correlation [164, 292], correlation function [296] or normalized ISF [142, 270, 297]. We will adopt the term ISF. The ISF is the real part of the spatial Fourier transform of the autocorrelation of the image differences [267, 286]. As such, it is a measure of the similarity of two images separated by a time lag Δt . Equation (2.2.14) is defined by the negative of the ISF, and hence shows that $\mathcal{D}(\mathbf{q}, \Delta t)$ is a measure of the decorrelation between two images taken with time lag Δt . The more different the image pairs are, the smaller the ISF and the larger the DDM tensor.

2.2.3 The radial average of the DDM tensor

In most applications of DDM, a radial average (Step (III) in Figure 2.3) is taken on $\mathcal{D}(\mathbf{q}, \Delta t)$ [142, 163, 270]. This operation leads to dimensional reduction and improved statistical sampling [266] (see Figure 2.4). The analytic radial average is evaluated as the integral of $\mathcal{D}(\mathbf{q}, \Delta t)$ along a circular contour Γ centred at $\mathbf{q} = (0, 0)$, divided by the contour length $||\Gamma||$, i.e,

$$\mathcal{D}_R(q, \Delta t) = \frac{1}{||\Gamma||} \oint_{\Gamma} \mathcal{D}(\mathbf{q}, \Delta t) d\Gamma = \frac{1}{2\pi q} \int_{-\pi}^{\pi} q \mathcal{D}(q, \theta, \Delta t) d\theta, \quad (2.2.16)$$

where

$$\Gamma = q(\cos(\theta), \sin(\theta)), \quad \theta \in [0, 2\pi).$$

It is common in the literature [142, 164, 267] to distinguish between the multi-dimensional array $\mathcal{D}(\mathbf{q}, \Delta t)$, and the matrix formed by its radial average, $\mathcal{D}(q, \Delta t)$, by dropping the vector notation on the frequency \mathbf{q} . In order to avoid confusion, however, we use a subscript R to denote the radial average.

For image data, the discrete nature of $I(\mathbf{r}, t)$, and therefore $\mathcal{D}(\mathbf{q}, \Delta t)$, requires the radial average to be evaluated numerically by averaging over values of $\mathcal{D}(\mathbf{q}, \Delta t)$ contained within an annulus, with an inner radius $q - \Delta q/2$ and outer radius $q + \Delta q/2$, where Δq is defined by the number of ‘bins’ we wish to discretise over [164], i.e

$$\mathcal{D}_R(q, \Delta t) = \frac{\sum_{\mathbf{q}} \zeta(\mathbf{q}) \mathcal{D}(\mathbf{q}, \Delta t)}{\sum_{\mathbf{q}} \zeta(\mathbf{q})},$$

where

$$\zeta(\mathbf{q}) = \begin{cases} 1, & \text{if } q - \frac{\Delta q}{2} < |\mathbf{q}| \leq q + \frac{\Delta q}{2}, \\ 0 & \text{else.} \end{cases}$$

A pixel thus belongs to a bin if its centre lies within the annulus of width Δq , with inner radius $q - \Delta q/2$. For an image which is L pixels long and has real pixel length ν , the smallest bin size is

$$q_{\min} = \frac{2\pi}{L\nu}, \quad (2.2.17)$$

We visually demonstrate the radial average on a slice of a DDM tensor in Figure 2.4, derived from simulated data described in Section 3.

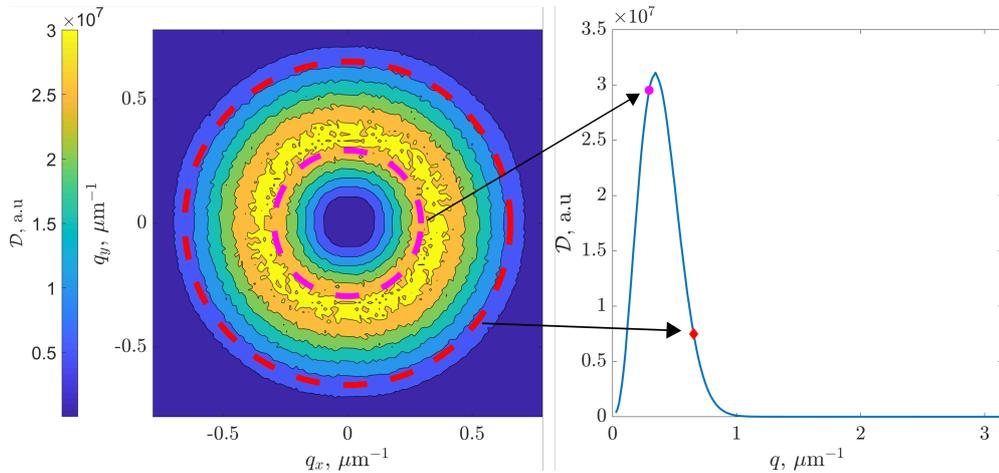


Figure 2.4: (Left) The contour plot of $\mathcal{D}(\mathbf{q}, \Delta t)$ for a fixed Δt , for simulated Brownian motion data, as described in Chapter 3. Each dashed circle represents an annulus over which the radial average is applied. (Right) The curve generated by the radial average, where the average over the smaller purple annulus produces the circular point in the right hand figure, and the average over the larger red annulus produces the diamond point.

The numerical azimuthal average is only truly accurate when $\Delta q \gg q_{\min}$, but the accuracy of the average can be improved by ‘oversampling’. This involves splitting each pixel into a grid of smaller pixels of equal brightness, which increases

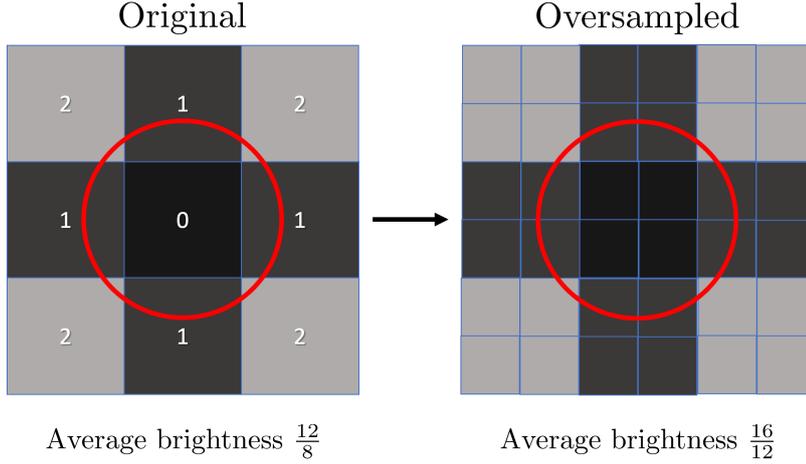


Figure 2.5: Example of oversampling of pixels to improve the accuracy of the radial average. (Left) The example image is made up of pixels of brightness 0, 1 or 2. In the original image, the 2 brightness pixels intersect the red circular contour over a shorter distance than the 1 brightness pixels, but have equal weight in the calculation of the average. (Right) Oversampling splits pixels into smaller grids of equal brightness, allowing us to weight contributions towards the radial average based on the length of the contour intersecting the original pixel. The 1 brightness pixels now have a greater contribution to the average than the 2 brightness pixels, resulting in a drop in the average brightness.

the level of refinement and makes the numerical average a better approximation to the contour integral. We visualise this concept in Figure 2.5.

When applying the radial average to Equation (2.2.14), we notice that the product of the ISF and scaling term can create non-trivial integrals. In literature, this concern is mitigated by always assuming that at least one of these terms does not depend on the orientation of \mathbf{q} [142, 163, 164, 280]. The ISF, $f(\mathbf{q}, \Delta t)$, is independent of orientation when the observed movement is isotropic [142, 164, 267], whilst the scaling term $A(\mathbf{q})$ is independent of orientation when the underlying images are radially symmetric, meaning there is no illumination gradient and particles are uniformly distributed throughout the system. We can study DDM when either the scaling term [292] or ISF [280] are radially asymmetric, but there are no studies considering both.

We define the radial average of $A(\mathbf{q})$ and $f(\mathbf{q}, \Delta t)$, respectively, as

$$A(q) = \frac{1}{\|\Gamma\|} \oint_{\Gamma} A(\mathbf{q}) d\Gamma,$$

$$f(q, \Delta t) = \frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) d\Gamma.$$

Assuming at least one of these terms is independent of the orientation of \mathbf{q} , the

radial average applied to (2.2.14) gives

$$\mathcal{D}_R(q, \Delta t) = A(q)(1 - f(q, \Delta t)), \quad (2.2.18)$$

In most applications of DDM, this relationship is presented with additional term $B(q)$ that accounts for imaging noise [142, 163, 164, 267], such that the DDM matrix is described by

$$\mathcal{D}_R(q, \Delta t) = A(q)(1 - f(q, \Delta t)) + B(q), \quad (2.2.19)$$

If an analytic form of $f(q, \Delta t)$ is known, we can evaluate $\mathcal{D}_R(q, \Delta t)$ numerically by iterating Equation (2.2.12) over the image stack, and then employ numerical fitting approaches to Equation (2.2.19) to determine the parameters $A(q)$, $B(q)$ and those within $f(q, \Delta t)$.

The standard fitting approach is to split the matrix $\mathcal{D}_R(q, \Delta t)$ into vectors, in which q is fixed and Δt is variable. On each vector, $A(q)$ and $B(q)$ will be constant, so we can probe the dependency of $\mathcal{D}_R(q, \Delta t)$ on $f(q, \Delta t)$ with respect to only the time lag [142, 164]. The exact fitting approach will depend on the type of motion expected within the image stack, however. We will outline the approach taken when studying Brownian motion as an example.

2.2.4 Example: Brownian motion

Brownian motion is a form of random motion caused by collisions between the molecules constituting a fluid and particles suspended in the fluid [139, 298, 299]. The verification and characterisation of Brownian motion is considered strong evidence for the atomic description of matter [298]. It was shown by Einstein that the van Hove function that describes the displacement caused by Brownian motion, $G(\mathbf{r}, \Delta t)$, satisfies the diffusion equation on an unbounded domain [299], as follows

$$\begin{aligned} \frac{\partial G(\Delta \mathbf{r}, \Delta t)}{\partial \Delta t} &= D \nabla^2 G(\Delta \mathbf{r}, \Delta t), \\ G(\Delta \mathbf{r}, 0) &= \delta(\Delta \mathbf{r}), \end{aligned} \quad (2.2.20)$$

where D ($\mu\text{m}^2\text{s}^{-1}$) is the (constant) diffusion coefficient. The initial condition stipulates that all objects have zero displacement when $\Delta t = 0$. In the one-dimensional case, the solution of Equation (2.2.20) under this initial condition is

$$G(\Delta x, \Delta t) = \mathcal{N}(0, 2D\Delta t) = \frac{1}{\sqrt{4\pi D\Delta t}} e^{-\frac{\Delta x^2}{4D\Delta t}}, \quad (2.2.21)$$

where \mathcal{N} is the Normal probability density function. Higher dimensional solutions are immediately available from Equation (2.2.21) under the observation that

Brownian motion acts in each dimension independently. Alternatively, we may solve this PDE by taking the continuous spatial Fourier transform,

$$\begin{aligned}\frac{\partial \tilde{G}(\mathbf{q}, \Delta t)}{\partial t} &= -\mathbf{q}^2 D \tilde{G}(\mathbf{q}, \Delta t), \\ \tilde{G}(\mathbf{q}, 0) &= 1,\end{aligned}$$

which has the solution

$$\tilde{G}(\mathbf{q}, \Delta t) = e^{-\mathbf{q}^2 D \Delta t}.$$

This solution is radially symmetric, so under the radial average, we can reduce the solution to the one dimensional case,

$$\tilde{G}(q, \Delta t) = e^{-q^2 D \Delta t}.$$

Since \tilde{G} is already real valued, Equation (2.2.15) gives us immediately that

$$f(q, \Delta t) = e^{-q^2 D \Delta t}, \quad (2.2.22)$$

which may be substituted directly into Equation (2.2.19) in order to determine an estimate of the diffusion coefficient, D (see Section 2.2.7).

2.2.5 Example: Bacterial advection-diffusion

In addition to experiencing Brownian motion, bacteria have flagella that enable self-propulsion. We will describe the ISF for a collection of active swimmers, ignoring Brownian motion for now. A single swimmer is considered to have velocity given by an incompressible velocity field $\mathbf{v} \in \mathbb{R}^3$, such that it acts as a ballistic projectile. The position of such a projectile is deterministic; the van Hove function G_v for this single swimmer is given by the indicator function,

$$G_v(\Delta \mathbf{r}, \Delta t) = \delta(\Delta \mathbf{r} - \mathbf{v} \Delta t).$$

This satisfies the continuity equation for an incompressible velocity field \mathbf{v} ,

$$\frac{\partial G_v(\mathbf{r}, \Delta t)}{\partial \Delta t} = \mathbf{v} \cdot \nabla G_v(\mathbf{r}, \Delta t),$$

from which we derive the solution in Fourier space

$$\begin{aligned}\frac{\partial \tilde{G}_v(\mathbf{q}, \Delta t)}{\partial \Delta t} &= i\mathbf{q} \cdot \mathbf{v} \tilde{G}_v(\mathbf{q}, \Delta t), \\ \implies \tilde{G}_v(\mathbf{q}, \Delta t) &= e^{i\mathbf{q} \cdot \mathbf{v} \Delta t}.\end{aligned}$$

Hence,

$$f_v(\mathbf{q}, \Delta t) = \cos(\mathbf{q} \cdot \mathbf{v} \Delta t).$$

Although f_v is the ISF for a single population of swimmers, in reality not all bacteria will have the same velocity magnitude or direction. A common assumption is that the speeds of the collective swimmers follow a Schulz distribution, plotted in Figure 2.6 and given by [163, 164, 281, 300]

$$\text{Sch}(v) = \frac{v^Z}{Z!} \left(\frac{Z+1}{\bar{v}} \right)^{Z+1} e^{-v(Z+1)/\bar{v}},$$

where \bar{v} is the mean velocity and Z is a measure of spread, related to the standard deviation σ by

$$Z = \left(\frac{\bar{v}}{\sigma} \right)^2 - 1.$$

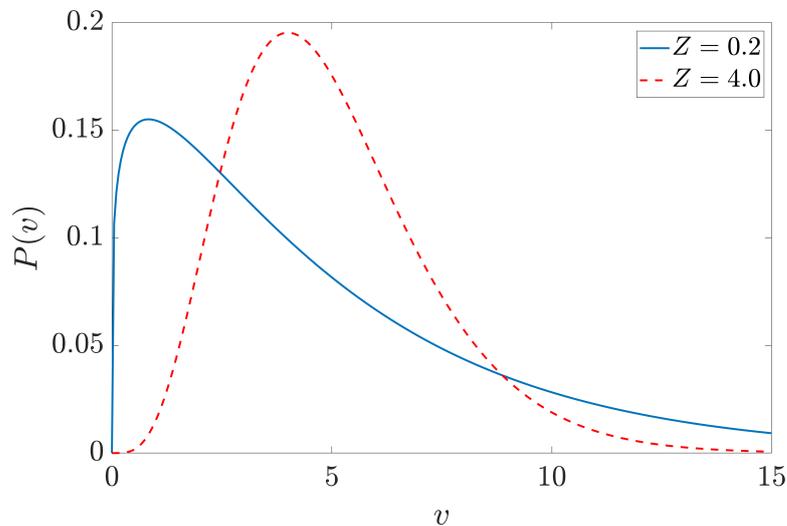


Figure 2.6: A plot of the Schulz distribution for $\bar{v} = 5$, when $Z = 0.2$ (blue solid line) and 4 (red dashed line).

To calculate the ISF for a collection of ballistic objects, which we denote by $f_{\text{bact}}(\mathbf{q}, \Delta t)$, we must consider the average of $f_v(\mathbf{q}, \Delta t)$ over all possible velocities, [300]

$$f_{\text{bact}}(\mathbf{q}, \Delta t) = \int \text{Sch}(v) e^{i\mathbf{q}\cdot\mathbf{v}\Delta t} d\mathbf{v}, \quad (2.2.23)$$

where we integrate over all three-dimensional $\mathbf{v} \in \mathbb{R}^3$. The consideration of the third dimension is a notable difference compared to the approach taken for Brownian motion. We assumed in Equation (2.2.9) either that movement in the axis orthogonal to the imaging plane is negligible, or the PSF $K(\mathbf{r} - \mathbf{r}', z)$ does not depend on z . When making either of these assumptions, we cannot observe movement in the orthogonal axis. Instead, we observe a projection of the particle locations onto into the imaging plane, from three to two dimensions. This has no effect when generating the ISF for Brownian motion, since the displacement in each Cartesian axis is independent, so the orthogonal axis is not needed to

accurately describe the diffusion coefficient. For bacterial motion, however, the projected displacement observed depends not only on the magnitude of the velocity vector, but also on its orientation.

To tackle the integral in Equation (2.2.23), we can convert to polar coordinates, setting $\mathbf{v} = v(\cos(\theta) \sin(\phi), \sin(\theta) \sin(\phi), \cos(\phi))$ where $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. We assume that the velocity is isotropic, such that the direction of the velocity is uniformly distributed in ϕ and θ . The integral is therefore independent of the orientation of \mathbf{q} , and as before in the Brownian case, we can reduce the solution to the one dimensional case for spatial frequency magnitude, such that the ISF is given by [300]

$$\begin{aligned} f_{\text{bact}}(\mathbf{q}, \Delta t) &= \int_0^{2\pi} \int_0^\pi \int_0^\infty \sin(\theta) v^2 \text{Sch}(v) e^{iqv\Delta t \cos(\theta)} dv d\theta d\phi \\ &= 4\pi \int_0^\infty v^2 \text{Sch}(v) \frac{\sin(qv\Delta t)}{qv\Delta t} dv, \\ &= \frac{\sin(Z \tan^{-1}(\eta))}{Z\eta(1 + \eta^2)^{Z/2}}, \end{aligned}$$

where

$$\eta = \frac{q\bar{v}\Delta t}{Z + 1}.$$

This is purely real valued and independent of the orientation of \mathbf{q} , so therefore

$$f_{\text{bact}}(q, \Delta t) = \frac{\sin(Z \tan^{-1}(\eta))}{Z\eta(1 + \eta^2)^{Z/2}}. \quad (2.2.24)$$

At the start of this example, we noted that Brownian motion was momentarily neglected. In Section 2.2.6, we introduce some helpful properties of the ISF that allow us to construct an ISF for behaviours composed of several different types of independent motion, allowing us to reintroduce Brownian motion.

2.2.6 Constructing ISFs for multiple independent processes

Particle movement often originates from more than one independent process. For example, bacteria and sperm move by beating their flagella, but are also small enough that random Brownian fluctuations within a fluid influence their trajectory [163, 270]. Bacteria are an example of a system in which there is more than one van Hove function determining the displacement of observed particles, but in Equation (2.2.19), the DDM matrix is described with only a single ISF. The following lemma allows us to define a single ISF for particles affected by more than one process.

Lemma 2.2.1. *Let the displacements $\Delta\mathbf{r} \in \mathbb{R}^3$ of some collection of N particles be the sum of two independent statistical processes A and B , with van Hove functions $G_A(\Delta\mathbf{r}, \Delta t)$ and $G_B(\Delta\mathbf{r}, \Delta t)$ respectively. If the corresponding ISFs for processes A and B are $f_A(q, \Delta t)$ and $f_B(q, \Delta t)$, respectively, then the ISF for the sum of the two processes is given by*

$$f_{A+B}(q, \Delta t) = \Re(\mathcal{F}(G_A(\Delta\mathbf{r}, \Delta t))\mathcal{F}(G_B(\Delta\mathbf{r}, \Delta t)))$$

If at least one of either process A or B is radially symmetric, then this simplifies further to

$$f_{A+B}(q, \Delta t) = f_A(q, \Delta t)f_B(q, \Delta t).$$

Proof. To describe f_{A+B} , we need to describe the joint probability distribution given by the van Hove function $G_{A+B}(\Delta\mathbf{r}, \Delta t)$. To do so, consider a single particle, which has made a displacement $\Delta\mathbf{r}$ during the time lag Δt . The displacement $\Delta\mathbf{r}$ is achieved when the displacement due to the distribution A is equal to \mathbf{r}_A , and the corresponding displacement under B is $\Delta\mathbf{r} - \mathbf{r}_A$. This particular trajectory occurs with probability $G_A(\mathbf{r}_A, \Delta t)G_B(\Delta\mathbf{r} - \mathbf{r}_A, \Delta t)$. Integrating over all possible \mathbf{r}_A yields the joint probability distribution, namely

$$\begin{aligned} G_{A+B}(\Delta\mathbf{r}, \Delta t) &= \int_{\mathbb{R}^3} G_A(\mathbf{r}_A, \Delta t)G_B(\Delta\mathbf{r} - \mathbf{r}_A, \Delta t)d\mathbf{r}_A \\ &= G_A(\Delta\mathbf{r}, \Delta t) \star G_B(\Delta\mathbf{r}, \Delta t) \end{aligned} \quad (2.2.25)$$

where \star denotes the convolution operator (see Equation (2.2.3)). Recalling the convolution theorem in Equation (2.2.6),

$$\mathcal{F}(f \star g) = \mathcal{F}(f)\mathcal{F}(g),$$

and additionally from Equation (2.2.11) that the Fourier transform of the van Hove function gives the ISF, the Fourier transform of the Equation (2.2.25) gives us immediately that

$$G_{A+B}(q, \Delta t) = G_A(q, \Delta t)G_B(q, \Delta t).$$

This property holds, inductively, for any finite number of independent van Hove functions. Taking the real part of the Fourier transform of this joint van Hove function returns the desired result immediately. When A and B are not both anisotropic, we expect radial symmetry in the ISF in one of these functions, which implies only one ISF contains a complex part. In this case, the product of the ISFs of the van Hove functions will be equal to ISF function of the joint process. \square

In addition to particles which are acted on by multiple independent sources of motion, we also consider cases where the displacement of different particles may not be described by the same van Hove function. For example, in bacteria, a subset of the imaged cells may perish and thus no longer be subject to ballistic motion [163, 164]. Bacteria are therefore an interesting population to study, because they can be split into an active and inactive population, each of which is described by a different ISF. For populations which can be partitioned into a finite number of subgroups, where each group is associated with a different van Hove function, the linearity of the Fourier transform gives us the following rule to construct the ISF.

Lemma 2.2.2. *Let a set of particles be divided into $N \in \mathbb{N}$ partitions, such that the probability of a particle belonging to partition $i \in N$ is α_i . Let each element of the partition have an associated van Hove function, denoted $G_i(\Delta\mathbf{r}, \Delta t)$, with corresponding ISF $f_i(q, \Delta t)$. Then, the ISF for the collection of all particles is given by*

$$f(q, \Delta t) = \sum_{i=1}^N \alpha_i f_i(q, \Delta t).$$

Proof. Bayes theorem gives us immediately that the probability of displacing by $\Delta\mathbf{r}$ over time lag Δt is the weighted sum of the probability of displacing by $\Delta\mathbf{r}$ given the particle belongs to partition $i \in N$, where the weights are the respective probabilities of belonging to partition i , i.e.

$$G(\Delta\mathbf{r}, \Delta t) = \sum_{i=1}^n \alpha_i G_i(\Delta\mathbf{r}, \Delta t).$$

Taking the Fourier transform of both sides, the linearity property immediately yields the desired result. \square

Combining Lemmas 2.2.1 and 2.2.2 allows us to construct a single ISF for a behaviour that can be described as the union of a set of independent processes. For example, when studying bacteria, all particles undergo Brownian motion, but only those which are active will undergo ballistic motion. The ISF of the active bacterial population is defined by the product of the ballistic ISF defined in Equation (2.2.24) and the ISF for Brownian motion, whilst the inactive population is described by the Brownian ISF alone. If the motility fraction is equal to α , the ISF for a population of bacteria will be given by [163, 164]

$$\begin{aligned} f(q, \Delta t) &= \alpha e^{-q^2 D \Delta t} f_{\text{bact}}(q, \Delta t) + (1 - \alpha) e^{-q^2 D \Delta t}, \\ &= e^{-q^2 D \Delta t} (\alpha f_{\text{bact}}(q, \Delta t) + (1 - \alpha)). \end{aligned} \quad (2.2.26)$$

Now that we can describe the ISF for a given system, we need some way of comparing the expected ISF to the DDM matrix, in order to quantify movement statistics. In Section 2.2.7, we review literature approaches to parameter fitting, including optimisation steps taken to increase accuracy and efficiency.

2.2.7 Parameter fitting

Once we have numerically generated the DDM matrix $\mathcal{D}_R(q, \Delta t)$, we wish to interpret its shape by comparing it to some model of the ISF, $f(q, \Delta t)$. The relationship between the two is given in Equation (2.2.19), reproduced here for clarity.

$$\mathcal{D}_R(q, \Delta t) = A(q) (1 - f(q, \Delta t)) + B(q).$$

There are several unknowns in this relationship, namely the scaling term $A(q)$, the noise term $B(q)$, and some set of parameters inside of $f(q, \Delta t)$ that correspond to rates of movement within the observed system. We find these unknowns through parameter fitting, finding values which minimise the distance between the numerically determined left hand side of Equation (2.2.19), and the model function on the right.

It is important to note that although there is a correct choice for $f(q, \Delta t)$ based on the behaviour being imaged, we can always determine some set of parameters that minimises the distance between these curves, regardless of our choice of $f(q, \Delta t)$. If we choose a model equation for $f(q, \Delta t)$ that does not match the imaged behaviour, the parameters we determine will not correspond to any of the observed phenomena we are trying to study. For example, if we use the ISF for Brownian motion to try and describe the movement of bacteria, we will successfully fit a numerical value for the diffusion coefficient. However, this value would not reflect the real rate of Brownian motion observed in the bacteria, since the bacterial motility will also contribute to the fitting of D . Hence, we currently have to work under the assumption that we have chosen the correct model function to represent the ISF. Later, in Chapter 4, we will show that it is possible to conduct parameter fitting without a model for the ISF if we can describe the behaviour of individual particles. Furthermore, in Chapter 5, we will also show that for some behaviours, it is possible to find different models for the ISF which will converge to the same, or similar, parameter fits.

Assuming, for now, that we have the correct model for the ISF, we can progress with parameter fitting. This fitting approach in DDM depends on the type of movement we are trying to study, but in general, is performed in two stages [142, 164, 266]. The first stage determines $A(q)$, $B(q)$ and some frequency-dependent parameters inside of $f(q, \Delta t)$, by fixing the frequency as a constant and fitting

curves over variable time lag. The second stage extracts constant movement parameters from the frequency-dependent results of the first stage. It is easier to understand why this approach is necessary once we have seen an example; we will present the fitting approach taken for Brownian Motion, where the ISF was given in Equation (2.2.22) as

$$f(q, \Delta t) = e^{-q^2 D \Delta t}.$$

The parameter we want to extract is the diffusion coefficient, D , but we also need to classify the scale and noise terms to satisfy Equation (2.2.19). Trying to determine D immediately from this relationship is difficult, because there will exist some large and small frequencies for which Equation (2.2.19) does not hold. At high frequencies, the signal-to-noise ratio in the DDM matrix becomes poor [164, 266]. For very small frequencies, which correspond to large distances, observed displacements may be too small to yield any information at these frequencies [164]. Additionally, at small q , the radial average is performed over a small number of pixels and becomes unstable [164] (see Figure 2.5). As such, there is some optimal frequency region to perform fitting over, which can be identified when fitting is performed over each frequency independently first.

In Brownian motion, the first stage of fitting normally attempts to fit three unknowns by fitting the following model function to $\mathcal{D}_R(q, \Delta t)$,

$$a_1(q) (1 - e^{-\Delta t / \tau_D(q)}) + a_2(q),$$

where we expect to fit

$$\begin{aligned} a_1(q) &= A(q), \\ a_2(q) &= B(q), \\ \tau_D(q) &= \frac{1}{q^2 D}. \end{aligned}$$

Fitting $1/\tau_D(q)$ instead of the reciprocal is common practice [142, 164], and has numerical stability benefits, in our experience. It is also often more numerically stable to fit the log of the DDM matrix, $\log(\mathcal{D}_R(q, \Delta t))$. This is permissible since, by Equation (2.2.12), $\mathcal{D}(q, \Delta t)$ is non-negative and the log is always well defined. For Brownian motion, since the shape of the DDM matrix is exponential (see Equation (2.2.22)), fitting the logarithm reduces to a linear fitting regime.

The second stage of the fitting algorithm takes the parameters determined from the first stage, and explores their dependency on the frequency q in order to extract some set of constant scalar values describing the rate of motion in the images. For example, in Brownian motion, we wish to extract the diffusion

coefficient, D , from $\tau_D(q)$. This is most easily achieved by rewriting τ_D as the linear regression,

$$\log \tau_D = -2 \log q - \log D. \quad (2.2.27)$$

For motile micro-organisms such as bacteria, the process is similar. We fit the DDM matrix to the ISF in Equation (2.2.26), where f_{bact} is given in Equation (2.2.24). In addition to fitting the motility fraction α and Schulz distribution variance Z at each frequency, we also seek to determine the value of the parameter group $\tau_v = \frac{1}{qv}$ for the average velocity v . Hence, the average velocity can be extracted by fitting the linear regression

$$\log \tau_v = -\log q - \log v. \quad (2.2.28)$$

A typical diagram we observe during this fitting process is given in Figure 2.7, taken from [164]. In this example, the diffusion coefficient is fit by plotting τ_D (called τ_d in the original work) on a log-log scale against spatial frequency, identifying the region where the plot is linear with gradient -2 (region (2)) and finding the vertical intercept of the curve. Often, the fitted scaling and noise terms are displayed below. These plots are useful because they allow us to quickly identify the optimal fitting region based on the shape of τ_D and the signal-to-noise ratio, and to highlight any problematic fitting results. For example, in region (3), the scale term $A(q)$ (red dashed line) is of the same order of magnitude as the noise term (blue dotted line), implying a poor signal to noise ratio.

We now outline some practical considerations for improving the speed and accuracy of DDM.

2.2.8 Algorithm optimisation

Practical applications of DDM frequently require the inclusion of additional steps which are not derived from the well-established underlying theory of the technique. Instead, these considerations are computational in nature, relating to optimising the various algorithms included in DDM analysis, with respect to both run-time and numerical accuracy. Whilst not related to the theory of DDM, these considerations have significant consequences in any analysis undertaken of the DDM matrix, and will therefore be detailed in this section.

2.2.8.1 Time lag sampling

Consider an image stack with N_t frames, which are uniformly spaced in time such that the time lag between consecutive frames is equal to Δt_{min} . In such movies, the time lag between any two frames is equal to $m\Delta t_{\text{min}}$, where $m \in \mathbb{N}$ with $m <$

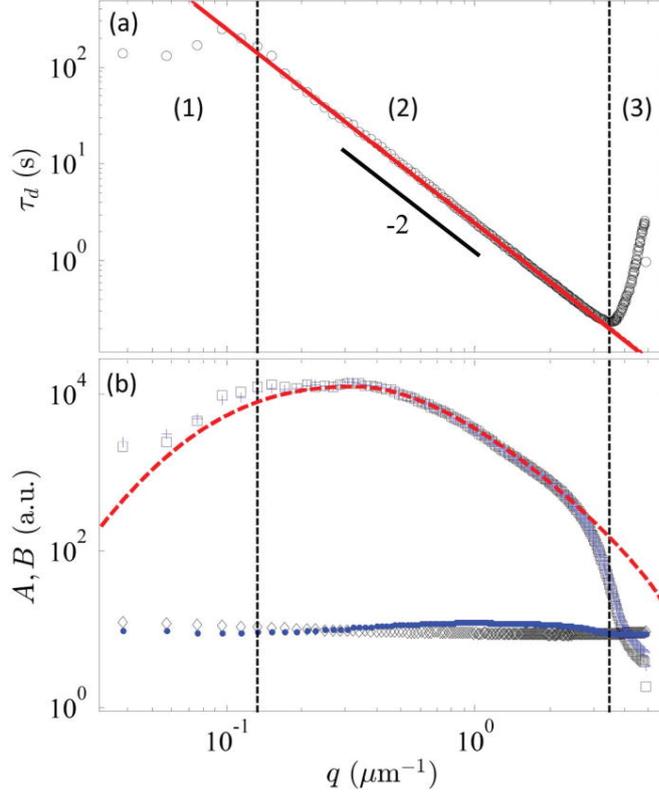


Figure 2.7: A typical parameter fitting diagram, taken from [164], which shows the fitting results obtained when attempting to classify Brownian motion in a colloidal dispersion. (a) The parameter $\tau_D(q)$ is plotted as a function of spatial frequency. Fitting is performed over region (2), where the log-log plot of τ_D is linear with gradient -2. (b) The fitted values of $A(q)$ (red dashed line) and $B(q)$ (blue dotted line) are plotted, alongside estimates of these parameters obtained by taking short and long time behaviours of $\mathcal{D}_R(q, \Delta t)$.

N_t . Hence, a bijective map exists between the actual time lag between frames, Δt , and the number of frames m . Although we can determine the DDM matrix for all time lags $m < N_t$, this is unnecessary for two reasons. First, the long-time dynamics of the particles in the frame may be completely spatially decorrelated, such that we cannot extract useful information from the DDM matrix for large Δt . In Brownian motion, where the ISF is $e^{-q^2 D \Delta t}$, there is a negative exponential decay with respect to time, so the ISF decays quickly and large time lags are not needed to extract the diffusion coefficient. Secondly, as m increases, the number of image pairs separated by m frames decreases. The DDM matrix is an average over the contribution from each pair of images, which mean for large times and, hence, small values of m , our statistical sample may not be a good representation of the underlying distribution of particle displacements.

To deal with both issues, we sample over a non-uniformly spaced subset of all available time lags, such that small Δt are sampled more frequently. We adopt

the scheme proposed by [164], namely

$$m \in \left\{ \left\lfloor 10^{k/b} \right\rfloor \right\}, \quad k \in \left\{ 1, \dots, \frac{\log(N_t)}{\log(10^{1/b})} \right\}, \quad (2.2.29)$$

for some suitable choice of $b \in \mathbb{N}$ [164] (for example, $b = 12$ in [164], which would reduce an image stack of $N_t = 900$ to a sample of only 31 unique time lags). The choice of b balances sampling rate with computational speed: a larger b means more time points are sampled, but as this is more computationally expensive, so the complexity of the ISF with respect to the time lag should be considered when deciding which time points to sample. We plot the sampling rate for b up to 20 in a movie of length $M = 900$ in Figure 2.8, which shows that with increasing b , the largest sampled m increases, as does the density of sampling for smaller m . In this work, we choose $b = 12$ to be consistent with literature [164]. A consequence of such time lag sampling is that the model function given in Equation (2.2.18) may over-fit to the DDM matrix at smaller time lags.

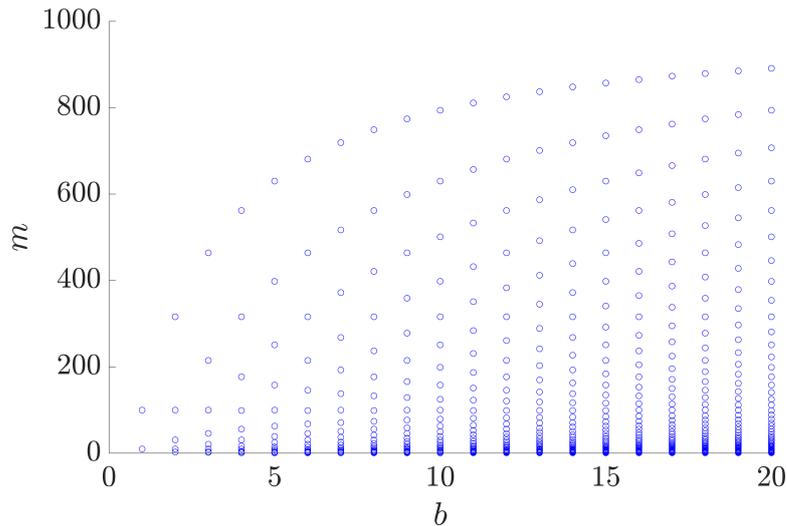


Figure 2.8: In movies with equally spaced frames, such that the time lag can be mapped to a natural number m denoting frame difference, we plot the sampling rate of time lags when using the logarithmic spacing given in Equation (2.2.29). As b increases, both the maximum sampled time and the density of time sampling increase.

2.2.8.2 Choosing a spatial frequency interval for fitting

An additional consideration in DDM analysis is that the quality of the first fitting step in Section 2.2.7 is known to be poor for very small and very large frequencies [266]. Hence, in general, there exists some frequency interval $[q_l, q_u]$ on which the second fitting step is applied. In Figure 2.7, for example, the vertical dashed line dividing regions 1 and 2 denotes q_l , and the same line between regions 2

and 3 denotes q_u . This frequency band is not particularly well characterised, so it is often subjectively chosen by visualising the fit of parameters, e.g. τ_D , and manually selecting an interval where the fit is good. The optimal selection is the largest interval for which the residuals remain small, so automating interval selection is possible if we define some cost function to quantify this heuristic. For example, a proposed cost function to optimise the fitting region when tackling Brownian motion is given by [301]

$$\text{loss}(q_l, q_u) = \frac{\sum_{q_l}^{q_u} \left(\frac{\partial \log \tau_D}{\partial q} - (-2) \right)^2}{q_{\min}},$$

where q_{\min} is the smallest frequency available. The numerator is minimised by finding the interval which gives the closest gradient to -2 , as expected by Equation (2.2.27), whilst the denominator is maximised by choosing larger intervals. Hence, the proposed cost function balances using as much data as possible, without including data which does not conform to the expected fitting pattern. By iterating over possible combinations of q_l and q_u , we aim to find an appropriate interval, although care must be taken to prescribe a minimum number of points to prevent selection of excessively small fitting intervals for which the loss is low. The choice of fitting interval has direct consequences on the fit of the key output parameters we seek to determine during DDM, so choosing a sensible interval is highly important.

2.2.8.3 Image windowing

Image windowing is the practice of multiplying the image by some filter, in order to force the edges of the image to gradually fade to zero, whilst preserving the centre of the image. We illustrate the benefits of windowing in Figure 2.9, using diagrams reproduced from [295], who first showed that windowing is good practice in DDM.

When particles are located at the image boundary (top left of Figure 2.9), the spatial Fourier transform performs poorly when trying to approximate discontinuities, resulting in a characteristic ‘smearing’ in the direction of the crossed boundary (top right). To mitigate this, we apply a so-called ‘window’ function, a matrix with the same shape as the image, which has brightness 1 at the centre and tapers off towards the edges of the frame. Element-wise multiplication of the image differences ΔI with the window function reduces the contribution of particle displacement at the edge of the frame to zero, whilst preserving the dynamics closer to the centre of the frame (bottom left of Figure 2.9). When the Fourier transform is taken of the windowed function, smearing is reduced (bottom

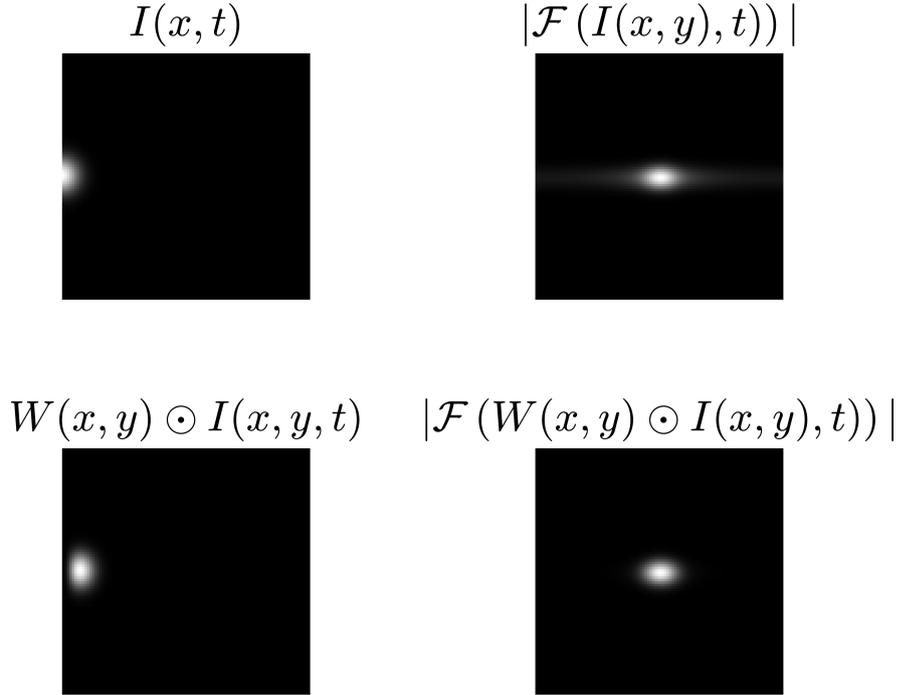


Figure 2.9: Visualising the benefit of image windowing, reproduced from [267]. When particles touch the boundary of the image (top left), the discontinuity in brightness causes smearing in the Fourier transform (top right). If we premultiply by a window function such Equation (2.2.30) (Bottom left), smearing is reduced (bottom right).

right). In this work, we use a window function $W(x, y) \in \mathbb{R}^{L \times W}$ derived from the Blackman filter, given in MATLAB by

$$W(x, y) = w(x)w(y), \quad (2.2.30)$$

$$\text{where } w(x) = \left(0.42 - 0.5 \cos\left(\frac{2\pi x}{L-1}\right) + 0.08 \cos\left(\frac{4\pi x}{L-1}\right) \right),$$

We visualise the brightness of the filter when $L = 100$ in figure 2.10. Whilst windowing has been shown to improve the performance of DDM [295], taking a window clearly affects the brightness of the images considered, and will therefore affect the scale term $A(q)$. Hence, any analysis of the scale term should account for the effect of windowing; for example, if the image intensity for a single particle is given by $\psi(\mathbf{r}, z)$, the scale term for a synthetic image dataset containing an average \bar{N}_p particles per frame, using an image window $W(\mathbf{r})$, is given by [295]

$$A(q) = 2\bar{N}_p |\tilde{W}(\mathbf{q})|^2 \star \left\langle \left| \tilde{\psi}(\mathbf{q}, z) \right| \right\rangle.$$

2.2.8.4 Assuming the absence of boundaries

Consider the ISF presented for Brownian motion in Equation (2.2.22). We observe that the partial differential equation we solved to determine the ISF has

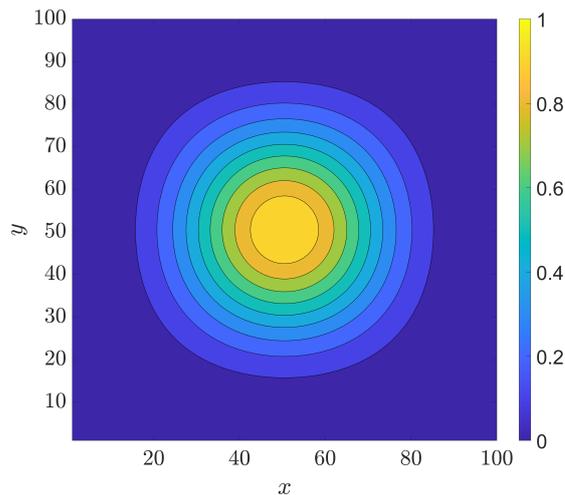


Figure 2.10: Contour plot of the Blackman-Harris window given in Equation (2.2.30), for a square image of length $L = 100$. The brightness is 1 at the centre and 0 at the edges, with a continuous slope between.

an associated initial condition, but no boundary condition. This highlights an important, though often unstated, assumption of DDM that there are no boundary effects, or that the boundaries of the system are sufficiently far away from the imaged domain that they have little to no effect on the observed displacements. To validate this assumption, DDM is often applied only on a subsection of the available images [142, 280, 302]. For example, in an application of DDM to *Drosophila* oocytes [142], images were cropped to different sections inside of the cytoplasm which do not include the cell wall, as shown in Figure 2.11. In the application of DDM to oocytes, this means we must discard image data close to the cell wall, which means we limit the amount of useable data available. This concern is addressed in more detail in Chapter 4.

2.2.9 Summary of assumptions

To summarise, we have now introduced the DDM technique, a method of numerically quantifying the rate of movement within time series image data. DDM is a powerful tool, which has performed well in a variety of applications [142, 280].

By providing the mathematical theory behind DDM, we have also highlighted a number of assumptions we needed to justify particular steps. In brief, these were:

- The motion being studied must be distributed by some probability density function which is stationary, to justify taking time averages.
- This distribution must also be ergodic, so that the study of an ensemble can be used to reconstruct the original distribution.

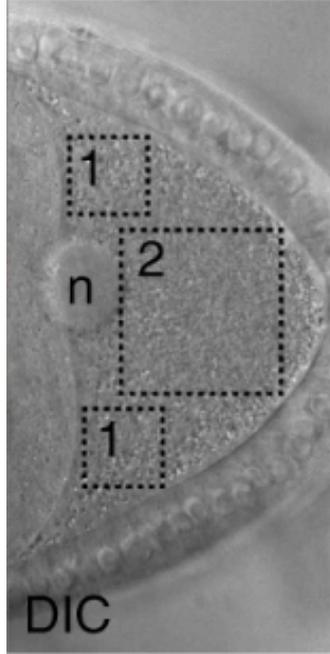


Figure 2.11: Subdomains of the cytoplasm of a *Drosophila* oocyte, over which DDM was applied in [142], in order to avoid imaging the nucleus (labelled n) and the cell walls. Taken from [142].

- We must be able to infer the position of particles in the frame from their appearance in the image stack. The LSI condition in Equation (2.2.2) is sufficient, but not necessary, to meet this assumption.
- The ISF, $f(\mathbf{q}, \Delta t)$, and the scaling term, $A(\mathbf{q})$, cannot both be anisotropic, in order to prevent a product integral forming in the radial average.
- Movement in and out of the frame through the direction orthogonal to the imaging plane is limited. This prevents images changing their appearance over time, enforcing the time-independence of $A(q)$.
- The images do not contain boundaries with which the particles might interact. This assumption prevents displacement in the frame from depending on the initial position of the particle.
- We need to know the ISF of the underlying behaviour before we can attempt fitting.

Some of these assumptions are highly limiting in reality, and in particular, may not be satisfied when dealing with images of oocytes. Two features of oocytes which are particularly challenging in this context are that the vast majority of molecules and particles imaged are sub-pixel in size (see Figure 2.1), and that the high concentration of these molecules and particles means that inter-particle interaction is highly likely. Overcoming these features relies on replacing the sample

density function in Equation (2.2.1) with the sample concentration, by dividing by domain volume for sufficiently high particle number. This means that we are not tracking individual particles; we are instead tracking visible shifts in the brightness of the cytoplasm, driven by bulk movements in molecules. This allows us to overcome the lack of localisation of particles, and makes the assumption that particle interactions have minimal effect on the bulk movements described by the combination of active diffusion and cytoplasmic streaming.

Additionally, without validation steps and confidence that the DDM analysis output matches the observable input, statistics determined from DDM will be no better than existing morphokinetic parameters in the IVF clinic. However, designing reliable, reproducible validations is practically very challenging. Without seeing a dataset, it is difficult to decide if the unique combination of biology and microscopy involved in its generation is conducive to accurate analysis. Therefore the first step in our exploration of the problem, in Section 2.3, is to implement DDM analysis in an example dataset, in order to determine exactly the kinds of problem we might face in the pursuit of non-invasive oocyte health assessment through DDM.

2.3 Preliminary exploration of DDM in oocyte data

In this section, we outline the challenges a user in the lab might face when performing DDM analysis on oocyte data, highlighting whether the particular challenges faced arise from features of the biology being studied, or are inherent in the theory of DDM. We approach this task as a typical new user of DDM would, applying pre-packaged code to an unconditioned dataset, so that we might emulate the same problems and outcomes expected within the lab. The data we analyse was introduced at the beginning of this chapter in Figure 2.1.

Pre-packaged DDM code, for both generation of the DDM matrix and subsequent analysis, is taken from a repository made available by Germain et al. [164]. This code was chosen because it was designed to teach new users how to implement DDM, and hence is well documented and accessible for non-expert users. Alongside code to generate the DDM matrix, the parameter fitting code available can consider one of two behaviours, each with its own corresponding ISF:

- Brownian motion, with the ISF given in Equation (2.2.22).
- Bacterial ‘run and tumble’ dynamics, consisting of a mixture of inactive and active swimmers, with velocities drawn from a Schulz distribution. All

particles, regardless of their active status, are additionally subject to random Brownian motion. The ISF for the collection of particles is given in Equation (2.2.26), with f_{bact} in Equation (2.2.24) describing the ISF for active swimmers.

Neither of these behaviours matches the proposed models for cytoplasmic streaming given in Section 1.2.2.1, but the bacterial model has been used to fit cytoplasmic streaming and active diffusion in *Drosophila* oocytes [142], reasonably suggesting we could use the same approach when studying the control oocytes. In the poisoned oocytes, however, since active processes are inhibited, we expect movement to be attributed exclusively to Brownian motion, implying the Brownian ISF should be suitable. Regardless of which behaviours we think are appropriate, once we have generated a DDM matrix, it is trivial to perform analysis using both choices of ISF, allowing us to compare how fitting of the same matrix looks given two different models.

DDM is applied to each oocyte individually, rather than over the whole image, allowing us to compare the variation between individual oocytes. The image stack is thus divided into a series of 11 smaller image stacks, each defined by a square subdomain centred on each oocyte. The length of the subdomain, $L = 200$ pixels, is chosen so that each square is contained within the cytoplasm, and does not include the boundaries of the oocyte, to respect the assumption outlined in Section 2.2.8.4 that no boundary effects exist within the imaged domain. We will consider this assumption further in Chapter 4. The approach of taking a subdomain is consistent with the previous application of DDM to oocyte data [142].

A Blackman filter is also applied, as defined in Equation (2.2.30), to remain consistent with literature approaches to DDM [295]. One might observe that there is some interaction between cropping of the image and application of a filter; a filter could be applied in such a way that the image brightness goes to zero at the cell boundary. However, considering a cropped image reduces the size and, therefore, computational cost associated with performing Fourier transforms on the image data. The interaction between image windowing and cropping of the image is further explored in Section 4.2, but for now, we consider the standard approach of applying both windowing and cropping of the image.

Subsequent analysis will generate, for each oocyte and experimental condition, two graphs. The first graph, an example of which is given by the top plot in Figure 2.12a, visualises the fitted characteristic timescales $\tau_D = 1/(q^2D)$ to describe Brownian motion, and $\tau_v = 1/(qv)$ to describe the mean velocity of advection when using the bacteria model. To interpret these characteristic timescales, we

observe that by Equations (2.2.27) and (2.2.28), both τ_D and τ_v are linear when plotted against the frequency q on log-log scales, with gradient -2 for τ_D and -1 for τ_v . In these equations, D and v can be found by linear regression of $\log(\tau_D)$ and $\log(\tau_v)$ against $\log(q)$ to find the intercept term, the exponential of which returns the diffusion coefficient and mean velocity, respectively. The interval on q for which we undertake fitting is given by the two dashed vertical lines in each plot; the same fitting region $3 < q < 6 \mu\text{m}^{-1}$ is chosen for all oocytes for fair comparison, where we estimate the upper and lower bounds by inspection for this dataset.

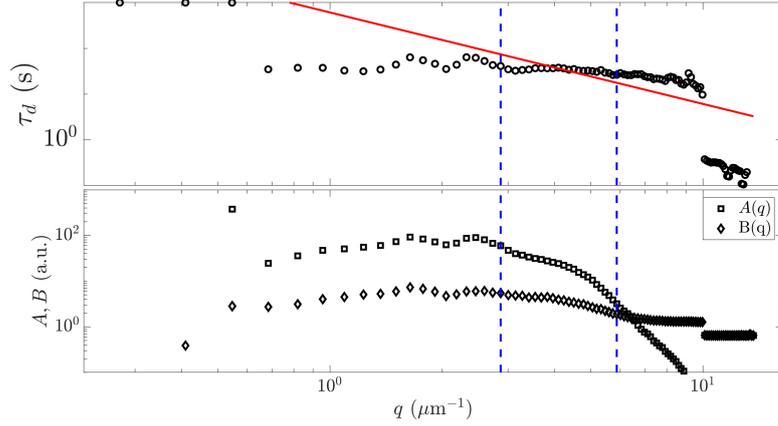
The second graph, for example the bottom plot in Figure 2.12a, shows the fitted scaling term, $A(q)$, and the noise term, $B(q)$, allowing us to estimate the signal to noise ratio in the DDM matrix. Since there are two experimental conditions, two choices of ISF and 11 oocytes, there are a total of 44 graphs, all of which are given in Appendix C. Within this chapter, however, we will present only a small selection to illustrate the findings from our analysis, summarised below.

2.3.1 Brownian motion describes cytoplasmic movement in FCCP treated oocytes, but active diffusion is insufficient to describe healthy oocytes

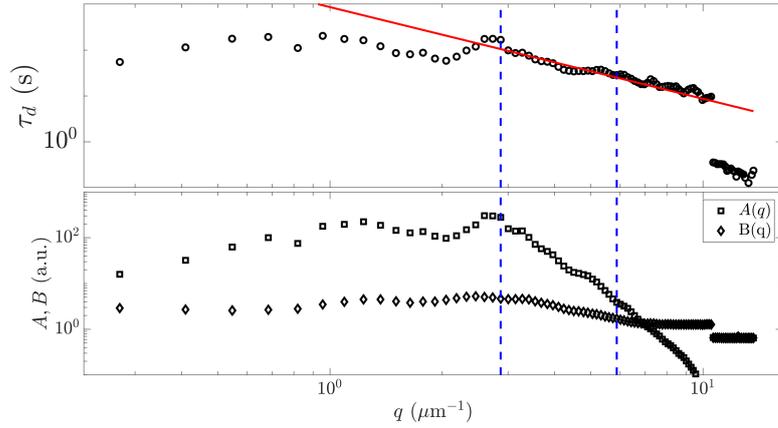
In Figure 2.12, using the Brownian model for the ISF, we plot the parameter fitting output for a single oocyte (numbered 7 in Appendix C), for both the control and FCCP experiment. As both image stacks depict the same oocyte, we expect equivalent fitted scale and noise parameters, $A(q)$ and $B(q)$ respectively; this expectation is met within the fitting region denoted by vertical dashed blue lines. The fit of these terms is dependent on the quality of the fit of τ_d , however; under-estimation of the diffusion timescale can be somewhat accounted for by under-estimation of the noise and scale terms. This is evidenced in 2.12 by small local fluctuations in all three parameters; in particular, just below the lower bound of the fitting region in Figure 2.12b. Hence, in depth comparison of the scale and noise terms will not be informative, and we continue to the comparison of the fitted values of τ_d in each case.

For the FCCP poisoned oocyte, the parameter τ_D shown in 2.12b follows the linear relationship given in Equation (2.2.27). This matches our expectation that, in the absence of any active processes, the dominant driving force is Brownian motion. However, in the control experiment in the same egg prior to poisoning, Figure 2.12a shows τ_D to be fit with the incorrect gradient.

The patterns observed in Figure 2.12 can be seen across the rest of the oocytes (see Figures C.1 and C.2 in Appendix C). The poor quality of the fitting



(a) Control, Brownian model.



(b) FCCP, Brownian model.

Figure 2.12: Fitting results for a given oocyte (7 in Appendix C), using the Brownian model in both the control condition 2.12a and the FCCP condition 2.12b. The FCCP condition yields better fitting quality under the Brownian model than in control conditions.

parameters in the healthy oocytes could be attributed either to the incorrect ISF being chosen, or poor conditioning of the studied images leading to a breakdown of the relationship between the image appearance and the particle density in the frame. For example, we previously highlighted the LSI condition in Equation (2.2.2) as a sufficient, but not necessary, condition for DDM to be viable in a dataset. Under this condition, the image can be represented in terms of a convolution of the particle locations $c(\mathbf{r}, z, t)$ and the PSF $K(\mathbf{r} - \mathbf{r}', z)$. If no equivalent condition holds in our dataset, the relationship between DDM and the ISF would no longer hold, and parameter fitting would fail. However, we were able to fit Brownian motion to the poisoned oocytes, which tells us the relationship between the particle position and image brightness holds. Therefore, the incorrect choice of the ISF is the cause for the poor performance of the Brownian model in the healthy oocyte data.

If cytoplasmic movement cannot be described solely by active diffusion, then

we require a new model for the ISF. We next consider the application of the bacteria model to the oocyte data.

2.3.2 The bacteria model performs poorly in both control and FCCP oocytes.

Replacing the Brownian fitting model with the bacterial ISF, we can replot the curves in Figure 2.12 under our new fitting regime. The same oocyte is considered to allow comparison between the Brownian and bacteria models. We plot the new fitted parameters in Figure 2.13.

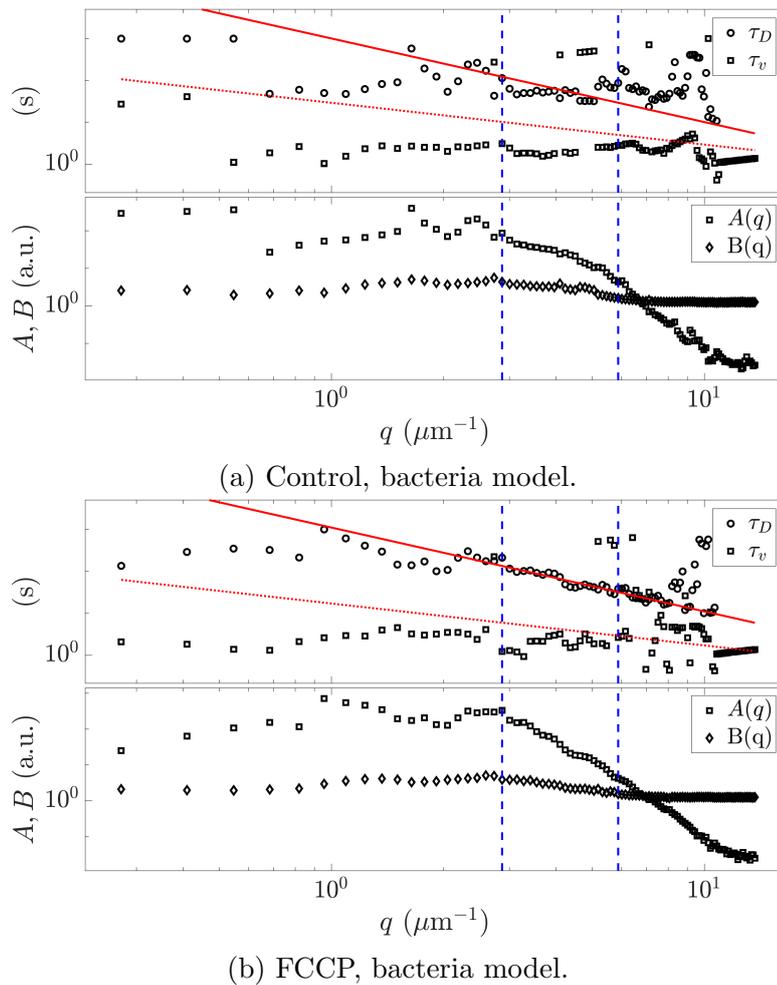


Figure 2.13: Fitting results for the an oocyte (7 in Appendix C), using the bacterial model in both the control experiment (2.13a) and the FCCP condition 2.13b. The solid and dashed red lines represent the curves fit to τ_d and τ_v respectively, generated by linear regression as presented in Equations (2.2.27) and (2.2.28). The performance of the bacterial model is poor in both cases.

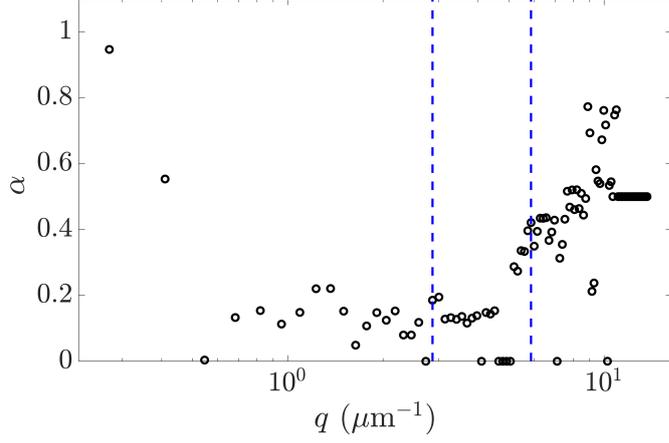
In the control oocyte, Figure 2.13a, the fit of both parameters is poor, with both τ_D and τ_v being fit as flat curves with high amounts of noise. This suggests that the model is a poor fit for the observed movement of the cytoplasm. In the

FCCP dataset, Figure 2.13b, the diffusive term τ_D is well fit, but the velocity term τ_v is not.

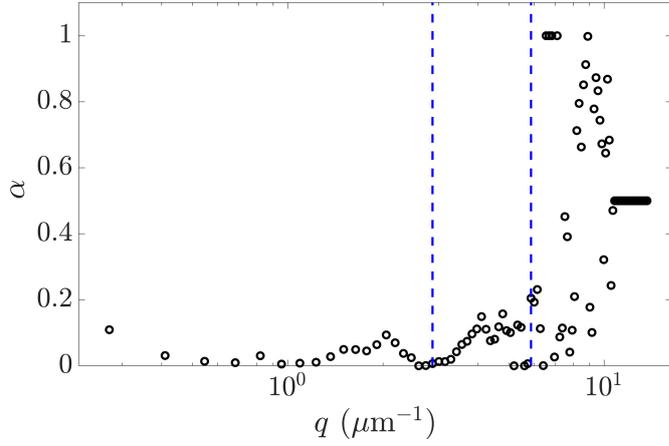
When discussing the fit of the velocity term, we highlight that the bacteria model has a variable parameter α controlling the motile fraction of the moving particles, where $\alpha = 0$ means no particles in the frame are subject to advection. Hence, when α is fitted to be approximately zero over the fitting interval, the parameter τ_v can be ignored, because it no longer effects parameter fitting algorithms. In Figure 2.14, we show that the fitted motility fraction for this oocyte in both the FCCP and control condition is approximately 20% for both the control and FCCP cases. This means a possible suggestion for the poor quality fit of τ_v is that advection may affect only a small fraction of the objects within the cytoplasm, and the DDM matrix is dominated by contributions from diffusive movement. An alternative explanation is that, since α is not fit as a constant, the motility fraction has no true biological meaning for the observed movement, and therefore we still do not have the correct model for the ISF.

Referring to Appendix C, the patterns observed in the fitting of τ_D under the bacteria model vary significantly. Considering healthy oocytes, τ_D is well fit in some cases, (Figures C.3b and C.3j), in others the general linear trend is observed but with a large amount of noise (Figures C.3c and C.3g), and in some oocytes, the diffusion fit term demonstrates no correlation at all with the frequency q (Figures C.3d, C.3f, C.3h and C.3j). The fit of τ_v likewise varies from good (Figures C.3e and C.3k), to noisy (Figures C.3i and C.3j) to completely decorrelated with no observable relationship between frequency and τ_v (Figures C.3b and C.3f). This inconsistency in the type of error we see between oocytes makes it difficult to identify exactly why our model ISF performs poorly. It may even be that different oocytes are subject to different ISFs; any defect or alteration in the quality of the cytoskeleton could generate variations on the expected movement from within the cytoplasm. Such differences may be difficult to detect by eye due to the optical density of the cytoplasm, but would cause our chosen model to perform differently across the oocyte cohort.

If the ISF itself can vary from oocyte to oocyte, we would need to construct a ‘global’ ISF as the weighted sum of each possible choice. In theory this is possible, treating parameters like the motility fraction α as binary variables which switch each type of behaviour on or off. However, in practice, numerical constraints mean that we will likely not fit these weights to be exactly one or zero to identify which behaviour we observe. This can cause currently well-fit parameters to deprecate as we add more possible behaviours to the system. An example of this effect can be seen by comparing the diffusion coefficients fitted by the Brownian and bacteria models in the control oocyte discussed in Figures 2.12 and 2.13. In the Brownian



(a) Control.



(b) FCCP.

Figure 2.14: The motility fraction α fitted at each frequency q for oocyte 7 under both the control (2.14a) and FCCP (2.14b) conditions. Vertical dashed bars correspond to the lower and upper bounds of the fitting region, as in Figures 2.12 and 2.13.

model, we fit $D = 0.99 \times 10^{-3} \mu\text{m}^2/\text{s}$, whilst in the bacteria fitting regime we fit $D = 1.6 \times 10^{-3} \mu\text{m}^2/\text{s}$. These values are of a similar order of magnitude, but as we show later in Section 2.3.4, this difference is significant compared to the variance expected between different eggs using the same model. Choosing a more complex model with (optional) additional terms has led to a different diffusion coefficient being extracted from the same dataset, calling the reliability of our DDM analysis into question in these datasets.

We can provide further evidence that the available ISFs do not account for the observed motion in the cytoplasm, by highlighting anisotropy within the DDM tensor.

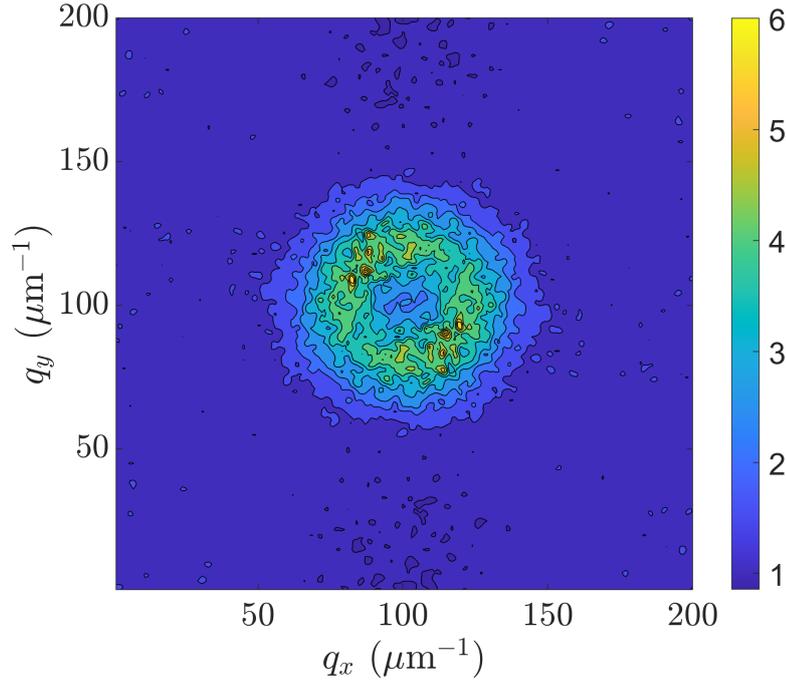
2.3.3 Cytoplasmic movement is anisotropic

In addition to performing fitting of the DDM matrix, we can also generate contour plots of the DDM matrix prior to radial averaging, as seen previously in Figure 2.4, to check for anisotropy. Choosing the smallest time lag available, we plot $\mathcal{D}(\mathbf{q}, \Delta t)$ in Figure 2.15 for the same oocyte (2 in Appendix C) in the control and FCCP datasets. Comparing the positive and negative diagonal slopes, we observe a slight anisotropy in the control oocyte. This anisotropy is also observed in the FCCP images, although to a lesser extent. The source of this anisotropy is unknown; it could be a drift of the entire oocyte over the plate, resulting in a very slow overall velocity emerging. Alternatively, there could be a directional bias in the movement of the cytoplasm which cannot be accounted for by either of the models available. Combined with the previous analysis, that both the Brownian motion and active-diffusion model fails to account for some aspect of the cytoplasmic movement, this leaves us to question what the best model for cytoplasmic movement is for these oocytes.

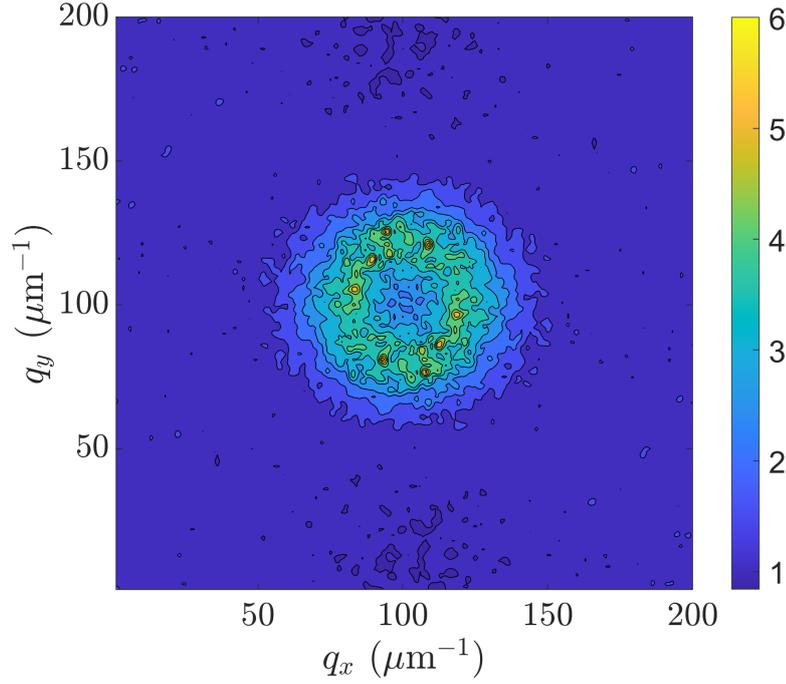
So, we have shown that we lack a good model for the ISF that would allow us to extract accurate and reliable parameters from these oocytes. Can we still extract any meaningful information from the DDM matrix? We now show that although the diffusion coefficients extracted from the oocytes have been shown to not necessarily correlate exactly with the movement of the cytoplasm, they still correlate with the underlying oocyte health.

2.3.4 Diffusion coefficients fall following poisoning, but not significantly

Whilst we have shown the performance of our current fitting approaches to be questionable in the control oocytes, we can still extract out a diffusion coefficient for each oocyte considering only the Brownian model, and compare between the control and FCCP datasets. These diffusion coefficients are given as box plots in Figure 2.16, with connecting lines denoting which points come from the same oocyte. The diffusion coefficient drops following poisoning, which matches the expectation that in the absence of active processes, cytoplasmic movement becomes slower. This suggests that the DDM matrix does contain information that informs us in the change in activity between the FCCP and control datasets. However, since the error bars of the box plots for the control and FCCP groups overlap, the diffusion coefficient cannot be used to identify which group the oocyte originated from. If we cannot predict this rather extreme difference between an alive or dead oocyte based on the parameters extracted from DDM,



(a) Control, Brownian model.



(b) FCCP, Brownian model.

Figure 2.15: Contour plots of the DDM tensor $\mathcal{D}(\mathbf{q}, \Delta t)$ for $t = 1s$, for the same oocyte (2) in both the control (2.15a) and FCCP (2.15a) datasets. A slight anisotropy is observed in the control dataset, with the positive diagonal axis achieving lower values than the negative diagonal. This effect is reduced in the FCCP dataset, but still present.

this would make DDM unviable in the clinic to differentiate between oocytes with more subtle changes in activity levels.

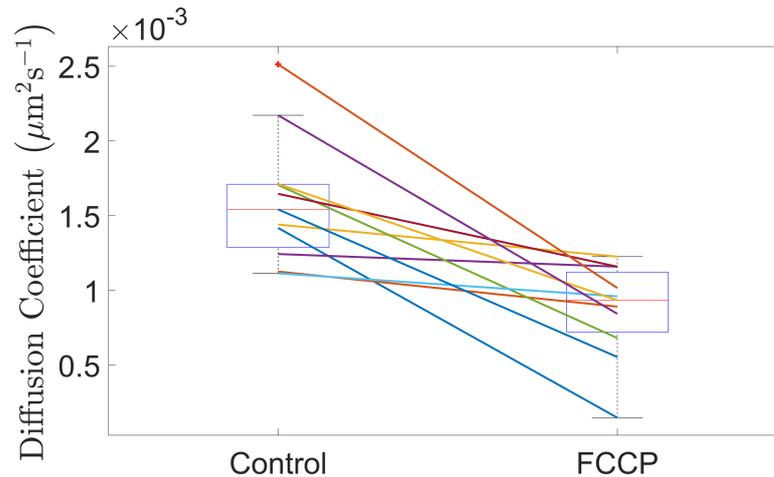


Figure 2.16: Box plots of the fitted diffusion coefficients extracted from the control and FCCP mouse oocyte datasets, using a Brownian model to describe cytoplasmic movement. The diffusion coefficients all drop following poisoning with FCCP.

2.3.5 Identifying areas of improvement for DDM in oocyte data

In summary, it appears that DDM has the potential to inform us about the rate of activity inside of oocytes, but current practices for interpreting the DDM matrix are insufficient. We now outline some of the challenges faced when performing DDM analysis on these oocytes, motivating the work of this chapter.

The first challenge we face is that we need a suitable model for the van Hove function, and therefore ISF, which describes cytoplasmic movement and allows us to interpret the shape of the DDM matrix. This challenge is based on our understanding of the biology of the system. For example, the van Hove function for Brownian motion is defined by the physics that drive it [299], and for bacterial swimmers, the behaviour of an individual swimmer is described from experimental observation [281, 300]. In both of these cases, a single model exists for the behaviour of all particles or cells, allowing us to derive summary statistics for the ensemble. As discussed in Section 1.2.2.1, models of cytoplasmic movement do exist, but these models are spatially heterogeneous and, thus, difficult to reduce to a set of summary statistics. Additionally, these models are defined via computational fluid dynamics, rather than some analytic solution to the Stokes flow, which means they cannot be substituted into the current parameter fitting pipeline. We therefore need to create a method to insert these fluid mechanical models into the fitting process.

As an extension to this challenge, we face difficulty not only in finding the correct van Hove function, but also in cases where the van Hove function changes between two instances of the same dataset. We showed that whilst the Brownian model performed well in the FCCP data (see Figure 2.12b), the bacterial model

did not (Figure 2.13b), despite the potential for us to fit the motility fraction parameter $\alpha = 0$ and reclaim the Brownian fit. It is possible that varying health in the oocytes will be expressed as different van Hove functions distributing the movement of the cytoplasm, but current DDM analysis approaches would fail to identify these changes. In order to classify which behaviour the oocyte follows, we would need to rerun DDM analysis for each different van Hove function in our library, and somehow enumerate the goodness of fit for each method. Each ISF would likely have a different number of parameters, so comparison between models is challenging, e.g. what do we do when some parameters are well fit, but others are not? Additionally, more complex behaviours require more parameters, increasing the likelihood of finding a local error minimiser away from the true estimate of the parameters governing the cytoplasmic movement. It is a challenge to balance choosing the most simple fitting function, with the most reliable numerical outputs, against choosing the most accurate model that best approximates the behaviour of the oocyte cytoplasm.

Another challenge is that current parameter fitting approaches lack mechanisms for validating their performance. Given any DDM matrix and a model for the ISF, we can always undertake parameter fitting and extract some number for each desired parameter, such as the diffusion coefficients in Figure 2.16. However, assuming we have the correct model for the ISF and a set of parameters is returned, how is the accuracy of the fit of these parameters measured? For example, although the Brownian model appeared to well approximate the behaviour of the FCCP oocytes, there was variation in the diffusion coefficients between oocytes, despite all being driven by the same Brownian motion. This variation could reflect a true difference in the speed of cytoplasmic movement between the oocytes, but equally, this variation could be the result of noise in the fitting data, or the fitting function not matching the distribution of the fitting data as a result of choosing the wrong model for the ISF. We therefore need to introduce methods which allow us to quantify the accuracy of the fitting algorithms employed, to give confidence in the values they return.

For some behaviours, validation is possible by comparing the derived movement parameters to a value determined theoretically, or by another technique. In Brownian motion of colloidal dispersions, a theoretical diffusion coefficient is given by the Stokes-Einstein formula [164], whilst in bacteria, experiments can be designed where the density of the population is sufficiently small that tracking is permissible [281], allowing reconstruction of the summary statistics extracted by DDM. However, in oocytes, validation can only be performed in idealised experiments. These experiments would use injected tracer particles, such as gold nanoparticles [303, 304], which can then be imaged and followed using SPT or

PIV, providing a second estimate of movement statistics which can be compared with those determined by DDM. This process is invasive, however, and requires the use of human oocytes, for which the demand vastly outweighs supply. Minimising the volume of data needed for validation relies on designing new mathematical approaches to determine causes and signs of error in DDM analysis, allowing us to identify and, where possible, mitigate such error before real data is even required.

The variation in fitted diffusion coefficients between oocytes under the same condition in Section 2.3 highlights the need for validation in DDM. The Brownian model performed well in the FCCP dataset, but despite this, there is a large variation in the fitted diffusion coefficients amongst the FCCP treated oocytes. We cannot say, however, whether this variation is due to the stochastic nature of the cytoplasmic movement, numerical fitting error, or reflects a variable rate of intra-cytoplasmic movement between the oocytes. By the Stokes-Einstein formula [164], variable kinematic viscosity of the cytoplasm, or a different distribution of particle size in the cytoplasm of different oocytes, could result in variable diffusion coefficients. In turn, these measurements could be indicative of oocyte health. It is therefore a critical challenge in DDM to pair analysis with robust validation methods that guarantee the accuracy of the fitted parameters, so that we can be confident that varying output statistics are accounted for solely by the biology of the studied system.

A final challenge, which is described independently of the previous analysis, is that we must discard large amounts of data from the edges of the oocyte in order to satisfy the boundary assumption in Section 2.2.8.4. The oocytes in Figure 2.1 have a diameter of approximately 320 pixels, but to stay away from the cell wall to satisfy this assumption, we performed DDM over only a square region with length 200 pixels. This means discarding approximately 50% of the cytoplasm image data, which could be critical to understanding the health of the oocyte. Furthermore, limiting the size of the domain reduces the available range of frequencies over which we can perform fitting, since for very large and small frequencies, noise and numerical errors dominate the fitting regime [164, 266]. We therefore hope to adapt the DDM algorithm in such a way that we can include the oocyte cell wall in the construction of the DDM matrix.

These challenges are primarily concerned with the precision of the analysis undertaken in DDM- namely, we expect that there is some movement parameter we wish to approximate as closely as possible. However, one might also consider that precision and usefulness may not be equivalent in the context of DDM applied to oocyte analysis. For example, if we wish to use DDM only to rank-order oocytes within a single cohort, then there may be some acceptable error in our estimation

for which the same rank-ordering is preserved. In this case, further optimisation of the accuracy of DDM provides no benefit to the clinic. Understanding the appropriate level of accuracy to maintain usefulness is a fascinating challenge itself.

Many of these challenges are highly specific to the data in question; to decide what ISF is appropriate for the average oocyte, we must observe a sufficient number of oocytes. Similarly, we cannot quantify the accuracy with which DDM predicts the movement of the cytoplasm or health of the oocyte, without comparing our analysis to invasive experiments that measure the behaviour of tracer particles or biological health markers. This would be expensive and impractical to generate, so tackling many of these challenges head on is beyond the scope of this thesis. For this reason, we also cannot begin to quantify the trade-off between precision and usefulness of DDM analysis in the clinic. However, whilst large volumes of real data are beyond our reach, synthetic data is significantly easier to acquire. In the next chapter, we introduce synthetic data in the context of DDM.

First, we will briefly discuss the current role that synthetic data plays in DDM in Section 3. We also identify that generation of synthetic data introduces several new parameters that, despite being critical to the performance of DDM, currently have no guidelines to motivate their selection. We hence undertake some analysis to determine sensible bounds for these parameters to use when generating synthetic data, supported by simulated experiments. In Section 4, we tackle the challenge of determining the ISF in an unknown system, by introducing a pipeline that enables the numerical approximation of the ISF where an analytic expression is not accessible, such as in oocyte data. A key feature of this pipeline is that it allows us to account for boundary effects, as demonstrated through an application of the pipeline to a Brownian process confined to a circular domain.

Although this pipeline performs well, it still uses the same two-stage fitting process as ‘classical’ DDM. In Section 4.3, we use simulations to tackle the challenge of validating the results of this fitting approach in DDM, by exploring the codependent sensitivity of advection and diffusion behaviours using synthetic data. We identify parameter regions in which the fitting performance suffers, highlighting a weakness of current fitting methods. This motivates the work of Section 4.4, where we introduce a new supervised machine-learning approach to DDM parameter fitting, which can be trained on synthetic data and applied in real data. This brings together the previous results, by numerically generating our fitting function, whilst relying on machine learning to cover the unreliable fitting performance of previous curve-fitting methods. Finally, we discuss our results in Section 4.5.

Chapter 3

Optimal design of synthetic data in DDM

In Section 2.3, we identified a number of challenges we face when performing DDM analysis in oocyte data. Overcoming these challenges will require changes to the current approaches taken in DDM, alongside a large bank of oocyte data to verify that these changes have yielded a significant improvement to the performance of DDM. Acquiring such data is expensive, time consuming, and impractical. We propose, instead, to use synthetic images to drive the creation of new approaches to DDM. Synthetic data, once the initial theory and machinery is introduced, is cheap and fast to produce, allowing us to construct more robust approaches to DDM analysis that can be tested in smaller real datasets. We begin this chapter by first introducing the role synthetic data currently plays in DDM analysis in Section 3.1. We identify several key parameters which must be carefully chosen to ensure the simulated data we generate is suitable for use in DDM analysis. In particular, we optimise the combined selection of the particle brightness α_n and size σ in Section 3.2, and the combined selection of the particle number N_p and number of frame pairs per time lag, M , in Section 3.3.

3.1 Introduction

Synthetic data currently plays several important roles in DDM. Since we know what parameters and statistics were used to build the simulation, we know what to expect from DDM analysis applied to simulated images, which gives us a ‘true value’ against which to compare the performance of DDM. This has led to simulations being used to test new theories in DDM, before being implemented in experiments. For example, simulated images of Brownian particles suggested that DDM remains valid in dark field microscopy, but only where particle displacement is small relative to the gradient of heterogeneous illumination [292].

Simulations have also been used to support statistical analysis when comparing different methods for characterising the noise term in DDM [305]. In these applications, simulations provide supporting evidence for theoretical arguments, and can even be used to explore new applications of DDM before real data is available, saving time and expense on generating real data.

Synthetic data can also be used alongside real data in DDM analysis, as a validation step [163, 281]. For the mouse oocytes analysed Section 2.3, despite having a set of images from an experiment and some model for the ISF, it was difficult to decide whether our model was suitable or not. For some oocytes, the model seemed to be performing well, but for others the performance was poor, and we cannot immediately discern why this happens. The model could be appropriate for some, but not all, oocytes, or there could be numerical problems such as high image noise skewing our results. It is also possible that the good performance in some oocytes could be the result of random chance. Simulations are therefore useful because they can clarify that the results we see in real data match what we see in simulations, where we have prescribed the underlying mechanics of the system. Good performance of DDM analysis in simulated data suggests we have the correct model for the ISF, and good agreement between the performance in simulated and real data suggests that the same ISF holds in both datasets, giving confidence that DDM is performing well in the real dataset too. Examples of simulation validation include characterising bacterial motility [163, 281], as well as confirming the ISF for helical and oscillatory swimmers [306], such as spermatozoa [270].

In this work, we aim to use synthetic data to quantify the performance of current analysis methods, and to drive the development of new alternatives to existing fitting methods. Before we can do so, we must first define how simulations are constructed.

3.1.1 Simulation construction

Here we describe the prevailing simulation method in DDM literature [163, 281, 292, 305, 306]. For some finite number of particles N_p , we generate a set of particle trajectories $\mathbf{r}_n(t)$, where $n \in \{1, \dots, N_p\}$ denotes the particle number and t the time of observation. The initial position of particles can either be determined randomly [306], or placed on a fixed grid [292], with subsequent displacements drawn randomly from the chosen van Hove function. In Brownian motion, for example, each orthogonal axis has displacements drawn from a normal distribution as follows,

$$x_n(t + \Delta t) - x_n(t) \sim N(0, \sqrt{2D\Delta t}). \quad (3.1.1)$$

Next, particle trajectories must be converted into images. The most common approach is to define the PSF such that each particle appears as a Gaussian intensity profile [292, 295, 305, 306]. The contribution to the image brightness from each particle is typically considered as additive, meaning the image intensity at any given pixel is equal to the sum of the contributions from all of the particles. The image stack can thus be expressed as

$$I(\mathbf{r}, t) = \sum_{n=1}^N a_n e^{-\frac{((x-x_n(t))^2 + (y-y_n(t))^2)}{2\sigma_n^2}}, \quad (3.1.2)$$

where a_n describes the brightness of a particle at its centre, and σ_n describes an effective beam spread, or particle size. Comparing to the definition of a Linear space Invariant (LSI) image in Equation (2.2.2), the simulated image defined takes background brightness $i_0 = 0$ and a Point Spread Function (PSF) of a Gaussian beam. In an ergodic system, we want every particle to be representative of the sample, which is made easier when the particles have a homogeneous appearance. Hence, we will simplify our simulations by considering all particles to have the same brightness and size, a and σ , respectively.

An alternative approach to constructing images considers particles to be discrete objects, with brightness ‘spread’ across a local neighbourhood of pixels around each particle centre [163, 281]. A key disadvantage of this method is that particles may have sharp edges, which are problematic under the Fourier transform, so we choose to use Equation (3.1.2) to generate images. Many approaches also account for the position of the particle in the axis orthogonal to the imaging plane, before projecting the particle location back into the imaging plane [163, 281, 306], although in this work we will remain within the 2D imaging plane for simplicity.

An image must have finite size, but we have not discussed any measures which bound the movement of the simulated particles. Given a sufficiently large time, particles can therefore move outside of the image edges. If enough particles leave the image, the assumption that $|\tilde{I}(q, t)|^2$ is constant in time is violated, since later images will have less particles contributing to the overall brightness. To maintain an approximately constant particle number within the frame, simulations can be conducted over a finite domain with periodic boundaries [281, 306], such that particles passing one boundary re-enter the domain on the opposite boundary. This is geometrically equivalent to performing simulations on a torus, a concept we visualise in Figure 3.1. However, if the boundaries are placed at the edge of the image, such that particles moving out of the image immediately re-enter on the opposite side, the positions of these particles will correlate in space when in reality, they should be considered as two different particles. A simple solution

to this problem is to generate images that are smaller than the domain of the simulation, as shown by the dashed red region in Figure 3.1. If the distance between the borders of the image and simulation domain are large enough, spatial decorrelation of the position of particles should occur in the time between leaving the image and rejoining. The number of particles within the image will now no longer be constant, but the combination of uniform distribution of initial particle positions and independent movement should ultimately preserve the homogeneous distribution of particles throughout the frame. As such, we expect only minor variation in particle numbers per frame, and hence only minor effects on DDM analysis..

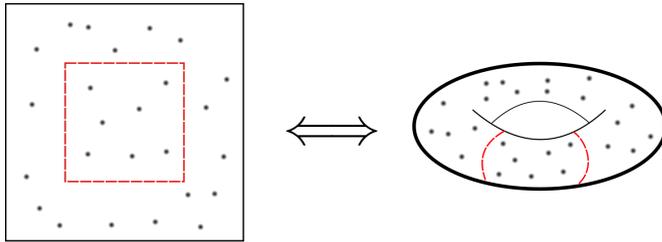


Figure 3.1: Simulations are performed over a larger simulation domain (solid black) with periodic boundaries to maintain a constant image brightness. The images themselves are taken as a central region of the simulation domain (red dashed square), to minimise the correlation of particles which exit the frame and on side and rejoin on the other. The simulation and image domain can be mapped onto a torus.

Given this simulation methodology, we now have a list of variable parameters:

- The image length, L , and simulation domain length, which we define as L_s , in pixels.
- The length of a pixel in real space, ν .
- The brightness, a , and the standard deviation, σ , which controls particle size.
- The number of particles, N_p .
- The movement parameters determining the motion within the movie, such as diffusion coefficients or velocity fields.
- The number of frames within the movie, N_t .
- The imaging frequency, or equivalently, the real time between images, denoted Δt_{\min} .

When emulating an experiment, these parameters can be matched to qualities of the real data. However, there is currently no investigation in literature that determines constraints on these parameters that ensure the validity of DDM applied to these simulated datasets. It is easy to see that such constraints exist, by considering very extreme cases. For example, if there was only one particle in the simulation, $N_p = 1$, the simulation would be too dilute to study bulk movement (although SPT would be appropriate). Equally, if the simulation is too concentrated, the image brightness will become uniform, and we will lose the ability to identify the position of particles within the frame. Since we allow particle overlap and have no maximum brightness, we will reach this over-saturation point at higher particle numbers than in real images where there is typically some maximum exposure, but nonetheless we can still reach state when the particle number is too high. This implies there is an optimal number of particles needed to well characterise the behaviour of the system through DDM, even though the dynamics themselves in the simulated images are independent of particle number. Similar heuristics can be made for other parameters, and hence we will attempt to explore, briefly, some reasonable conditions on the construction of images that maximise the validity of DDM. We will only consider Brownian motion as the source of particle movement here, noting that different parameter choices may be suitable in other types of motion.

Before trying to optimise parameter selection, we can remove some parameters from the system by considering a rescaling of the distance and time variables. Converting from real distance (μm) to pixels (px for brevity), and from time (s) to the number of frames m , we can rescale our movement parameters to achieve the same DDM matrix with new variables. For example, in the case of Brownian motion, we can rescale the diffusion coefficient by

$$\hat{D} = \frac{D\Delta t_{\min}}{\nu^2},$$

where the DDM matrix would be expressed as

$$\mathcal{D}(q\nu, m\Delta t_{\min}) = A(q\nu) \left(1 - e^{-(q\nu^2)\hat{D}m/\Delta t_{\min}} \right).$$

Varying the imaging frequency and pixel size has the same effect as tuning the diffusion coefficient, so under this non-dimensional parameter grouping, we can consider the imaging frequency and pixel size to both be equal to 1. We will drop the hat from \hat{D} for simplicity, and hence, there are only four key parameters remaining which relate specifically to the construction of the simulation, and not to the movement of the particles within the frame: a, σ, N_t and N_p , denoting the brightness, beam spread, number of frames and number of particles respectively.

A natural split exists within these parameters; a and σ control how particles look, whilst N_t and N_p control how many statistical samples of particle trajectories we take. We will hence approach these two parameter groups separately, aiming to identify some parameter region for which DDM performs well.

3.2 Controlling a and σ to optimise the appearance of simulated images

We wish to identify some parameter region for a and σ where the performance of DDM is optimised. Fortunately, a needs little consideration, following the linearity of the Fourier transform. Consider two simulated image stacks, I_1 and I_2 , that depict the same simulated trajectories and differ only by the brightness of the depicted particles, a_1 and a_2 respectively. The relationship between the two stacks can be expressed as

$$I_2 = \frac{a_2}{a_1} I_1. \quad (3.2.1)$$

For each image stack, a DDM tensor can be generated by the definition in Equation (2.2.12). In particular, let $\mathcal{D}_2(\mathbf{q}, \Delta t)$ be the DDM tensor generated from I_2 . Substituting the relationship between I_1 and I_2 from Equation (3.2.1) into the definition of $\mathcal{D}_2(\mathbf{q}, \Delta t)$ from Equation (2.2.12), we show that

$$\begin{aligned} \mathcal{D}_2(\mathbf{q}, \Delta t) &= \left\langle |\Delta \tilde{I}_2^2| \right\rangle_t \\ &= \left\langle \left| \frac{a_2}{a_1} \Delta \tilde{I}_1^2 \right| \right\rangle_t \\ &= \left(\frac{a_2}{a_1} \right)^2 \mathcal{D}_1(\mathbf{q}, \Delta t). \end{aligned}$$

Changing the brightness a therefore results in a rescaling of the entire DDM tensor, without changing its shape. Parameter fitting would thus only be affected by the choice of brightness when a is very large or small, where floating point errors may occur.

To choose a sensible brightness, we may desire images which have similar properties to real biological image data, which are commonly stored in an *unsigned integer* format [307]. For example, in *uint8* images, each pixel has brightness equal to an 8-bit non-negative integer, given by the set $\mathbb{Z} \cap [0, 2^8 - 1]$. Since particle brightness is additive in our simulated images, we want to choose a such that the maximum brightness in the image stack remains within same range as these unsigned integer images; we opt to therefore use the value $a = 50$ to achieve this, although we note that the choice of a is unimportant.

In contrast to the flexibility in the brightness, there are immediately identifiable poor choices for the beam width, σ . If $\sigma \ll 1$, particles are treated as point masses. Taking the limit of Equation (3.1.2) as σ tends to zero, we find

$$\lim_{\sigma \rightarrow 0} \sum_{n=1}^{N_p} a e^{-\frac{((x-x_n(t))^2+(y-y_n(t))^2)}{2\sigma^2}} = \sum_{n=1}^{N_p} a \delta(i - x_n(t)) \delta(j - y_n(t)).$$

Under this problematic expression for the image appearance, a particle can only be seen in a given pixel if its trajectory passes through a small neighbourhood around the centre of the pixel at time t . As particles move across the image, they will spontaneously appear and disappear depending on their location, flickering over time and violating the requirement for constant image brightness. The position of particles which spontaneously appear and disappear cannot be tracked using DDM, so any movement parameters derived from such a movie would have so physical meaning.

There is clearly a need to prescribe some minimum value of σ . To motivate a choice of lower bound, we note that the issues caused by excessively small σ arise due to a loss of conservation of mass of the particles. The appearance of the particle should not depend on its location within the frame, but if σ is too small, particles gain and lose mass depending on their position within a pixel. We demonstrate this effect in Figure 3.2, where we construct an image containing two particles. The first is located at the centre of a pixel, whilst the other is placed at a corner junction between four pixels. On the left, when $\sigma = 3$, the two particles look the same, but when $\sigma = 0.5$, the appearance of the particle changes significantly depending on its location. Although we could artificially weight the contribution of the beams to ensure equal contribution, with dependence on its location within the frame, this would require some additional post-processing which we hope to avoid through careful selection of σ .

We can hence make a heuristic argument to fix σ in such a way that as a pixel moves throughout the frame, it maintains its mass spread across the image. For a very rough approximation, consider a particle initially located at the centre of a pixel, which has the Cartesian coordinate (i, j) . Currently, this pixel has brightness a . Now, let the particle displace by half a pixel length in the horizontal direction, without loss of generality, such that its new position is $(i+1/2, j)$. If the central brightness of the particle was originally a , and the particle now straddles the border between the pixels (i, j) and $(i+1, j)$, we might expect that the sum of the brightness in these two adjacent pixels is equal to a . Hence, we can fix the brightness of the beam to be $a/2$ at a distance of $1/2$ a pixel away from the beam centre. More generally, we can seek σ satisfying

$$e^{-\frac{x^2}{2\sigma^2}} > \epsilon,$$

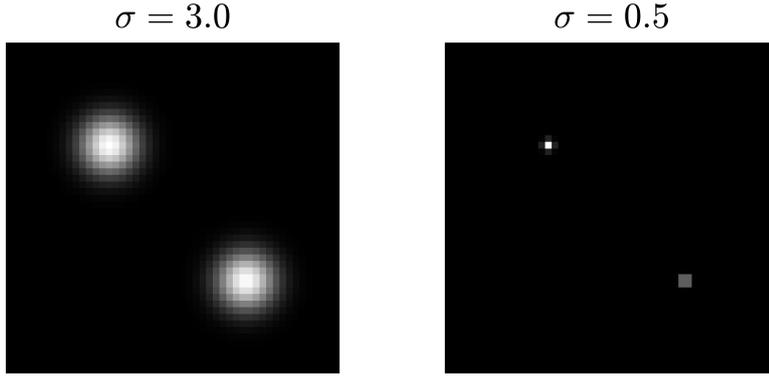


Figure 3.2: Diagram showing loss of conservation of mass in simulated particles when the beam spread parameter, σ , is too small. The left and right images use the same particle brightness, $a = 50$, and the same particle locations. The top left particle is located at the centre of a pixel, whilst the bottom right particle is located at a corner between four pixels. When $\sigma = 3.0$, the particle appearance is invariant of the particle position, but when $\sigma = 0.5$, the particle appearance changes depending on the location within the frame.

where x is the distance in pixels from the beam centre and $0 < \epsilon < 1$ is some tolerance, which denotes the rate of brightness drop experienced moving away from the centre of the particle. Rearranging,

$$\sigma > \sqrt{\frac{-x^2}{2 \ln(\epsilon)}}.$$

Taking $x = 1/2$ and $\epsilon = 1/2$ gives us $\sigma > 0.6$. This value is not necessarily an optimal choice for σ , but is sufficient to guarantee the particle is wide enough to track across the frame.

To numerically validate our exploration of these parameters, we perform a single simulation of Brownian motion, generating a DDM matrix for different pair combinations of a and σ . We use $D = 0.5 \text{ px}^2/\text{frame}$, matching the approach taken in [292], iterating over 2000 frames with 150 particles. For each DDM matrix, we perform fitting over the interval $q \in [0.2, 0.4] \text{ px}^{-1}$, determining some output diffusion coefficient. We plot the error between the fitted and expected diffusion coefficients for each set of image parameters in Figure 3.3. For $\sigma < 0.1$, the error is unpredictable and large, but as σ increases, the error quickly converges to $-0.05 \text{ px}^2/\text{frame}$ for all combinations of a_n and σ . Under-estimation the diffusion coefficient is problematic, but clearly is not a result of the selection of a and σ ; we address the cause of this fitting error later. To guarantee that we are outside of the problematic small σ region, we should pick σ well above the value proposed by our heuristic argument, and within the convergent range of σ proposed in Figure 3.3.

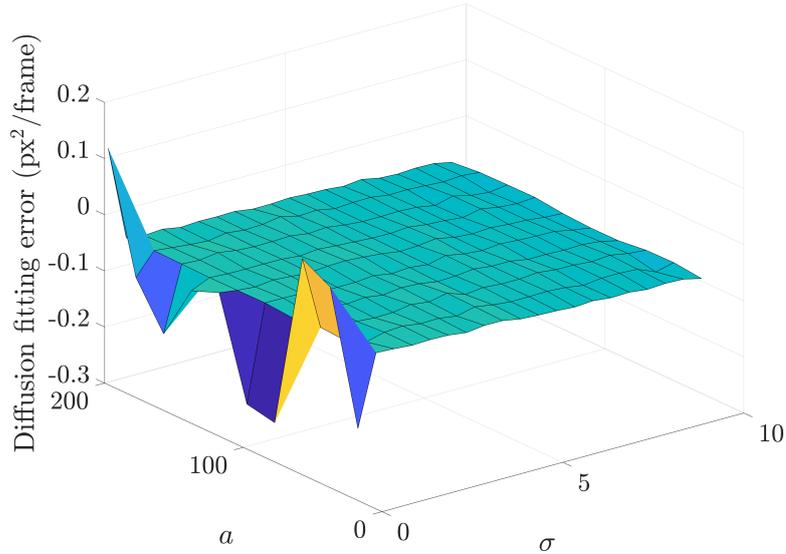


Figure 3.3: The error observed when fitting the diffusion coefficient, D , for different combinations of particle brightness, a , and beam spread, σ , in the same simulated Brownian motion. Error is constant everywhere, except for $\sigma < 0.1$, where the error is large and has high variance.

Having found a suitable lower bound, $\sigma > 0.6$, we turn our attention to determining whether there is a suitable upper bound. Whilst the lower bound was necessary to ensure particle appearance remains invariant of position throughout the frame, the upper bound is needed to ensure that particle movement over small distances can be tracked effectively. Consider the example given in Figure 3.4, which shows the image difference ΔI for a single particle undergoing a displacement of 10 pixels to the right. On the left hand side, the particle beam width is $\sigma = 3$, and on the right, $\sigma = 50$. For the higher value of σ , where the beams are large and relatively flat in the middle, the brightness gradient over small distances in ΔI is smaller than when the beams are sharply defined by a small value of σ . This weakens the signal over small distances, resulting in poor quality fitting over large frequencies. As such, selecting too large a value of σ may cause restriction in the size of the fitting region, which can result in an increased contribution of noise to key fitting parameters.

To demonstrate this detrimental shortening of the fitting region, we perform a new simulation of Brownian motion with $D = 0.5\text{px}^2/\text{frame}$, setting the brightness to $a = 50$, and consider the set of beam spread values $\sigma \in \{0.1, 0.5, 1.0, 3.0, 5.0\}$ px. For each value of σ , we generate the DDM matrix and perform parameter fitting to determine the characteristic timescale τ_D , which we plot in Figure 3.5. Choosing $\sigma = 0.5$ ensures τ_D is fit well, which shows that our lower bound $\sigma > 0.6$ may be too weak. Starting at $\sigma = 0.5$ px, as we increase σ , the lower bound q_l remains constant whilst the upper fitting bound q_u varies, attaining its maximum when $\sigma = 1.0$ px and subsequently decreasing with

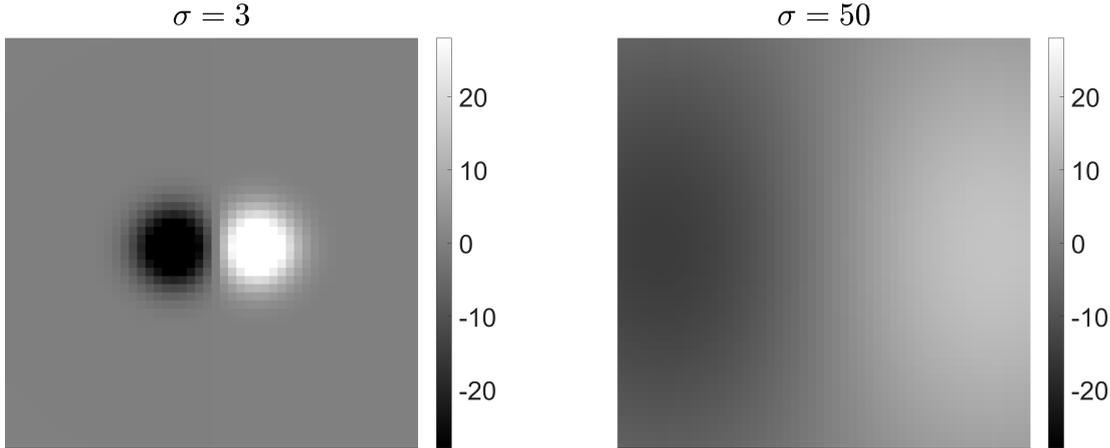


Figure 3.4: Visualisation of the image differences ΔI for a single particle undergoing a displacement of 10 pixels. In the left hand side, $\sigma = 3$, and in the right hand side, $\sigma = 50$. As the beam spread increases, the brightness gradient in the image difference decreases over small distances, resulting in insensitivity of DDM at higher frequency.

increasing σ . High frequencies correspond to smaller distances, so the retreat of the upper limit of the fitting interval is fuelled by the increasing insensitivity to small distances within the simulations.

In summary, the selection of a is not particularly important, but σ is. We choose $\sigma = 3$ for the rest of this work, unless otherwise stated, since this parameter is safely large enough to not violate the conservation of mass requirement, but small enough that the upper bound of the fitting interval remains high.

3.3 Balancing the particle number, N_p , and number of frames, N_t

The remaining parameters to control in our synthetic data are the number of particles contained in the simulation, N_p , and the total number of frames which make up the movie, N_t . We have grouped analysis of these parameters because they are linked by their relationship with a fundamental principle of DDM, that averaging $|\Delta\tilde{I}(\mathbf{q}, t, \Delta t)|^2$ over a large number of image pairs allows us to capture the average behaviour of particles depicted within the movie. In order to describe the ISF from the DDM matrix (and, therefore, the van Hove function), we need to sample a sufficient number of particle displacements at each time lag. The more particles contained within the frame, the more displacements we can sample at each image difference Δt . Similarly, the longer the movie is, the more image difference pairs are available to average over.

In our synthetic data, we can control both particle number, N_p , and the number of frames generated, N_t . Determining what combination of these parameters

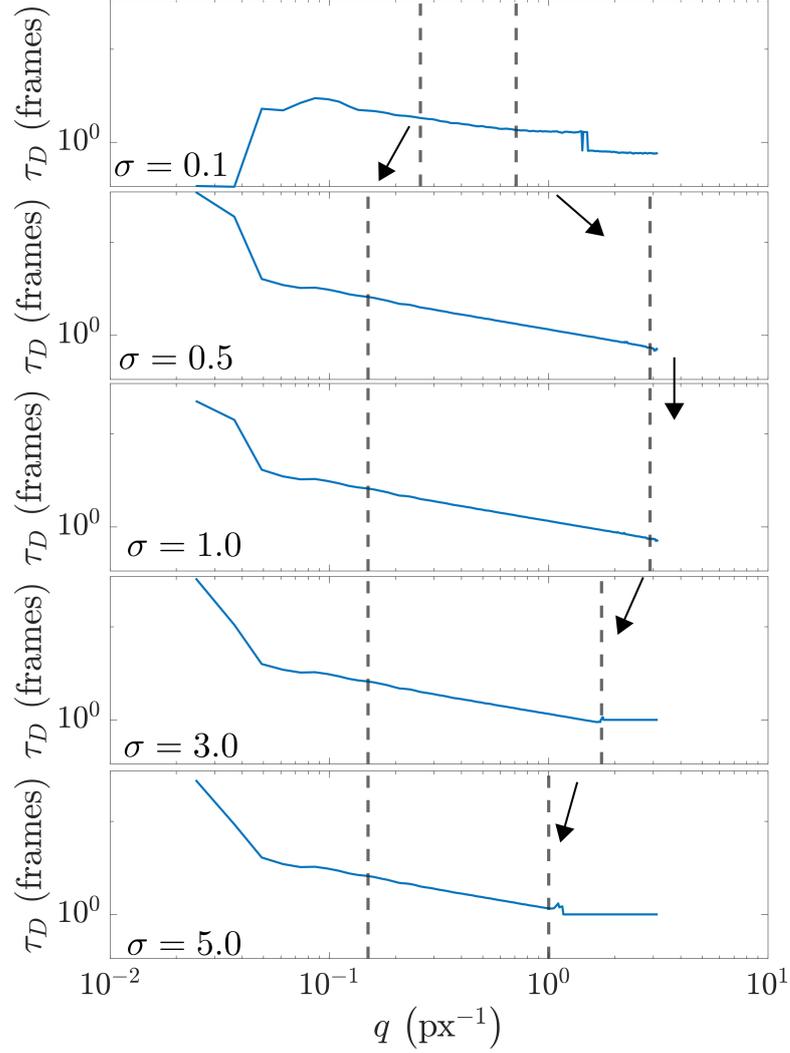


Figure 3.5: Plots of the fitting parameter τ_D for the same simulation, performed using $a = 50$ at different values of $\sigma \in \{0.1, 0.5, 1.0, 3.0, 5.0\}$ px. When $\sigma \geq 0.5$ px, a suitable fitting region can be drawn, but the upper limit of the fitting interval (right-hand grey dashed line) begins to decrease with increasing σ .

is suitable will depend primarily on the van Hove function, since movement with a higher variance will require a larger number of sampled displacements to describe. Consider Brownian motion, where the displacement of a given particle over a time lag Δt is drawn from a Normal distribution with mean 0 and variance $2D\Delta t$ (see Equation (2.2.21)). The variance depends on both D and Δt , so as both of these parameters increase, a larger ensemble is needed to generate a sufficient sample of the expected motion. In particular, the dependence on D makes the task of predetermining what constitutes a sufficient sample more difficult. Designing tests or heuristics which predict the optimal combination of N_t and N_p for every choice of D is beyond the scope of this thesis, but we will instead identify a suitable choice of these parameters which guarantees the validity of the synthetic data used in this work. In particular, this challenge can be simplified

by fixing either N_p or N_t , and identifying a suitable value for the other variable.

In this work, it is more approachable to consider a fixed value of N_t , rather than N_p . This is in contrast to the approach taken in real biological datasets, for which the particle number cannot be controlled (for example, we cannot interfere with the cytoplasm of the oocyte), and neither can the van Hove function, such that we may only control the movie length, N_t . If we attempt to identify the minimum N_t for a given N_p , we first observe that for a movie made of N_t frames, there are only $N_t - m$ possible image difference pairs to sample over for a time lag of m frames. This means that increasing N_t not only increases the maximum number of time lags we may sample over, but also increases the validity of all sampled time lags. The optimal choice of N_t is therefore motivated by first choosing the maximum time lag to be sampled over, and then finding the number of image pairs needed at that time lag.

This problem could be equivalently classed as identifying the optimal value of M , the maximum number of image differences sampled at each stage. This maximum is imposed to optimise DDM performance at small time lags, where the variance is small and the number of samples needed is, therefore, also small. If M is taken to be the smallest number of image pairs needed to generate a sample at the desired Δt , we can create a test to determine the maximum time lag which can be described by DDM in a given image stack, since if $M(m) > N_t - m$ for any given m , then there are insufficient available image pairs to generate an average. Choosing the optimal M is an interesting challenge, but it is as complex to tackle as the original problem of balancing N_p and N_t , so fixing N_p provides no real benefit. Hence, it is preferable to fix N_t , and consider a variable particle number.

Before we undertake any analysis, we consider some extreme choices of N_p to demonstrate the negative effects of poor selection. In DDM analysis, the assumption of ergodicity means that the ISF may be reconstructed by considering a single particle, given a sufficiently long observation time N_t . In practice, however, modelling only a single particle is not representative of the spatially dense arrangement of particles in the cytoplasm. As such, our generated DDM matrices would be ‘too perfect’, containing none of the contributions of the distinct parts of the van Hove function that measure the correlation between different particles in space. New methods and theory that work in our simulated datasets would therefore be likely to fail when employed on real data.

A second, more important consideration is that we would need a large number of frames to study a single particle. The computational complexity of the one dimensional Fast Fourier Transform (FFT) is $(L \log_2(L))$ for a square image L pixels in length [308] (the two dimensional FFT performs the one dimensional algorithm for each row and column, and hence has complexity equal to the square

of the one dimensional case). If the FFT is applied to every frame in the movie, this multiplies the complexity by a factor of N_t , which means that the DDM algorithm becomes very slow for long movies. Additionally, as N_t increases, so too will the computational requirement for storage of larger image stacks. There is, thus, a clear incentive to reduce the required length of the movie.

The solution to bringing down these costs is to consider more particles per frame, such that image difference ΔI contains more information. This reduces the number of Fourier transformed pairs needed, and importantly, does not increase the complexity of the transforms themselves. The drawback of adding more particles per frame is, however, is that the spatial correlation between different particles contributes more strongly to \mathcal{D} . When we introduced the van Hove function in Equation (2.2.4), we presented a partition into the ‘self-correlating’ and ‘distinct’ parts, making the assumption that the self-correlating part dominates the van Hove function as a result of particle trajectories being independent. However, as more particles are added to the system, the number of distinct terms grows. We can see this immediately by substituting the definition of the simulated image (Equation (3.1.2)) into the definition of the DDM tensor (Equation (2.2.12)). First, we note that we can rewrite the definition of the simulated image as the sum of N_p individual images, containing one particle each, as follows:

$$I(\mathbf{r}, t) = \sum_{n=1}^{N_p} I_n(\mathbf{r}, t),$$

where

$$I_n(\mathbf{r}, t) = ae^{\frac{((i-x_n(t))^2 + (j-y_n(t))^2)}{2\sigma^2}}.$$

The DDM tensor is therefore equal to

$$\begin{aligned} \mathcal{D}(\mathbf{q}, \Delta t) &= \left\langle \Delta \tilde{I} \Delta \tilde{I}^* \right\rangle_t \\ &= \left\langle (\tilde{I}(\mathbf{q}, t + \Delta t) - \tilde{I}(\mathbf{q}, t)) (\tilde{I}(\mathbf{q}, t + \Delta t) - \tilde{I}(\mathbf{q}, t))^* \right\rangle_t \\ &= \left\langle \left(\sum_{n=1}^{N_p} \tilde{I}_n(\mathbf{q}, t + \Delta t) - \tilde{I}_n(\mathbf{q}, t) \right) \left(\sum_{n=1}^{N_p} \tilde{I}_n(\mathbf{q}, t + \Delta t) - \tilde{I}_n(\mathbf{q}, t) \right)^* \right\rangle_t \\ &= 2 \left\langle \left| \sum_{n=1}^{N_p} \tilde{I}_n(\mathbf{q}, t + \Delta t) \right|^2 \right\rangle_t \\ &\quad - 2 \left\langle \Re \left(\left(\sum_{n=1}^{N_p} \tilde{I}_n(\mathbf{q}, t + \Delta t) \right) \left(\sum_{n=1}^{N_p} \tilde{I}_n(\mathbf{q}, t) \right)^* \right) \right\rangle_t. \end{aligned}$$

The term inside the real brackets can be split into a self part, where n is the same in both sums, and the remaining distinct part. There are N_p self-correlating

terms, and $N_p^2 - N_p$ distinct terms, so the number of distinct contributions grows faster than the number of self-correlating terms. We assumed in the theory of DDM that the self part of the van Hove function dominates the distinct part, because particles are independent, but for large N_p this assumption becomes less appropriate. This effect can be seen more clearly by taking the limit as N_p tends to infinity; if particles are initially placed with a uniform random probability, for sufficiently large N_p , the brightness in the frame will be uniform at all time points. Hence, the image differences tend to zero, and DDM will fail to characterise the observed motion.

This once again suggests the existence of some ‘Goldilocks’ region where the number of particles, N_p , and the length of the movie, N_t , are chosen to balance the computational requirement for Fourier transforms during DDM, whilst preserving the ability to extract particle movements from the frame. The exact nature of this balance depends on the computational resource available to the individual; use of GPU acceleration or parallel computing can allow for a significant volume of data generation and processing, which may facilitate the use of lower particle numbers. For example, generating and analysing a single DDM matrix with 150 particles, 2000 frames and image length 600 pixels takes 71.1 seconds on a Windows Desktop PC using a Intel(R) Core(TM) i9-9900 CPU @ 3.10GHz (16 CPUs), 3.1GHz processor with 64GB RAM, using parallel processing. For images larger in both space and time, or with less computational resource, the process of generating simulations may become infeasible. For this work, rather than develop this task in detail, we instead simplify this problem by considering a fixed number of image differences, denoted M , and seek to determine the number of particles needed to generate a good approximation of the ISF.

We design a synthetic experiment where we fix $M = 500$ image difference pairs, in a movie of length $N_t = 2000$, and vary the particle number, N_p , to identify the minimum value where good fitting is achieved. We simulate first only 10 particles to generate an image stack, on which we apply DDM, and then iteratively add new particles to the image in multiples of 10, until we reach a total of 300 particles. A diffusion coefficient is generated for each stack. Repeating this experiment 10 times, we plot the mean diffusion coefficients, and their corresponding standard deviations, for each particle number in Figure 3.6. A fitting region of $q \in [0.45, 1.1]\text{px}^{-1}$ is chosen based on the results of Figure 3.5, which highlight this to be an acceptable frequency range for the choice of σ used. The mean diffusion coefficient remains fairly constant at approximately $0.48 \text{ px}^2/\text{frame}$, but as N_p increases, the standard deviation decreases steadily. We suggest that a sufficiently high accuracy is achieved using $N_p \geq 150$.

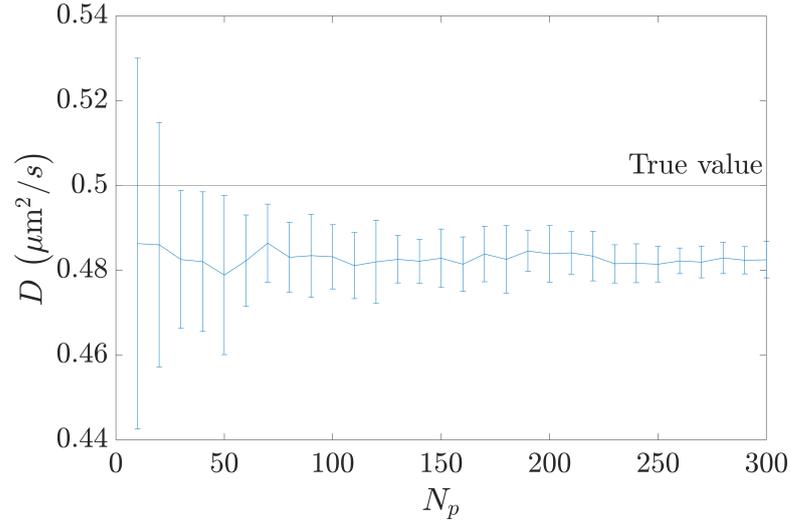


Figure 3.6: Plots of the mean diffusion coefficient, D , fitted to simulated Brownian motion images containing varying numbers of particles, with error bars showing \pm one standard deviation. The mean diffusion coefficient stays constant at around $0.48 \text{ px}^2/\text{frame}$, but with increasing particle number, the standard deviation decreases.

We note that, even for sufficiently large N_p where diffusion coefficients are fitted consistently with low variation, the diffusion coefficient is under fit. For $N_p \geq 150$, the standard deviation of these experiments is small, implying this under-fitting is not a result of random sampling error. Instead, this consistent under-fitting arises due to *localisation* error [292]. Whilst the movement we wish to study is continuous, images are a discrete form of data. This limits how exactly we can identify the location of each particle within the image. Any displacement we measure within the frame is therefore subject to a degree of uncertainty, with a tendency to under-estimate how far particles have travelled. By under-estimating the magnitude of observed displacements, we hence under-estimate the magnitude of the movement statistics sought by DDM, such as the diffusion coefficient. Whilst this error does cause separation between the ‘true’ and observed movement statistics in DDM analysis, we observe from Figure 3.6 that localisation error causes the fitted diffusion coefficient to be fit as a constant fraction of the true value. So long as there exists an injective map between the estimated diffusion coefficient and its true value, localisation error has no negative implications for DDM analysis in oocyte health assessment. Determining whether such an injective map exists is therefore important work, which would require systematic study, although this is beyond the scope of this thesis; we know from literature approaches to the application of DDM in synthetic data that for well conditioned images, the diffusion coefficient is accessible [292], and our work in this chapter is sufficient to verify that the images are well conditioned.

3.4 Summary and conclusion

In Section 2.3, we identified several challenges faced when using DDM as a form of non-invasive oocyte assessment. Overcoming these challenges using current approaches to DDM would require a large amount of oocyte data. Performing experiments to retrieve this data is an expensive and time consuming process, and, when progressing from mouse studies to human, is complicated further by strict ethical guidelines. As such, we seek to design techniques that use synthetic data, which is easy to generate, in order to answer these questions. Simulation methodology is widely available in literature [163, 281, 292, 305, 306], but introduces several new parameters that currently have no associated guidelines to motivate their selection.

In this chapter, we explored optimal parameter selection for generation of synthetic image data to be used in DDM analysis. Using a combination of numerical experiments and heuristic arguments, we have determined a suitable set of simulation parameters, given in Table 3.1, which balance computational complexity with performance in our subsequent work in Chapter 4. These parameters will, unless otherwise stated, be used throughout the rest of our simulations: In par-

Parameter	Description	Chosen value
a	Particle brightness (of beam centre)	50
σ	Particle size (beam spread standard deviation)	3
N_p	The number of particles	150
M	The maximum number of image pairs per time lag	500

Table 3.1: Parameter selection for generation of synthetic image data of the form given by Equation (3.1.2), which has been shown to yield well-conditioned images for DDM application.

particular, we have identified that the most sensitive parameters to control were the width of the particles, σ , and the particle number N_p . We can propose a rough rule of thumb, suggesting that once both N_p and σ surpass their respective minimum viable values, smaller is better; minimising N_p saves on some small amount of computational workload during simulation, and minimises the distinct contribution from different particles correlating with each other, whilst minimising σ maximises the length of the fitting region.

Whilst the parameters we have identified are suitable for our current exploration, there is still work to be done in further optimising the selection of parameters. For example, whilst our proposed minimum value of $\sigma = 0.6$ px is a sufficient choice, it is likely not the optimum, with $\sigma = 0.5$ px achieving good performance in Figure 3.5. More thorough numerical exploration may be able to further refine our bounds on the parameters explored in this work.

Additionally, the choice of N_p was made given a fixed number of image difference pairs, $M = 500$. There is a complex relationship between the image difference pair number, M , the number of particles, N_p , and the van Hove function which describes displacement in the frame. Investigating this relationship in detail, though unnecessary in our application, would make this work more generalisable to other DDM applications.

Further to this point, a more thorough exploration of the relationship between the appearance of a synthetic image stack, $I(\mathbf{r}, \Delta t)$, and the DDM matrix, $\mathcal{D}(q, \Delta t)$, may have important implications for real image data applications. Given an image stack, there are a number of factors which may have a negative impact on the quality of DDM analysis, such as excessive image noise [164, 266], particle size being excessively large (Figure 3.4) or small (Figure 3.2), or the spatial density of the particle number being too high or low (Figure 3.6). Without expert knowledge of how DDM works, however, it may be difficult to decide whether a given dataset is suitable for DDM or not. Exploration of the relationship between image construction and DDM performance in synthetic data could help to define guidelines and automated validation steps, informing users of potential pitfalls. For example, in PIV analysis, automated validation steps can detect anomalous velocity vectors, whose direction and magnitude are incoherent with the surrounding field, helping prevent users from drawing spurious conclusions from their data. Similar mechanisms in DDM would be vital to integration of the technique into regular biological image analysis.

An effective starting point for this exploration could be Approximate Bayesian Computation (ABC), a family of algorithms which sample the posterior distribution of input parameters by finding values for which the simulated data resembles real data [309]. This statistically-enhanced trial and error approach would allow us to identify the parameter space for which our synthetic image data yields DDM matrices that, when analysed, return accurate estimates of the input parameters of the van Hove function. These synthetic data parameters can then be related to properties of real images, hence providing guidance to assess the quality of biological image data with respect to potential DDM analysis.

For now, however, this chapter has achieved the purpose of identifying a synthetic image parameter set, which will allow us to create simulated images suitable for DDM analysis. In Chapter 4, we will use these parameters to generate large volumes of synthetic image data, in order to create new fitting approaches in DDM which overcome the current limitations and challenges faced when tackling oocyte data.

Chapter 4

Simulation-guided approaches to Differential Dynamic Microscopy in IVF

4.1 Introduction

In Section 2.2, we explored the performance of DDM in an example time-lapse oocyte dataset, identifying a number of key challenges faced when using DDM as a form of non-invasive health assessment. First, the library of ISFs available to analyse the DDM matrix did not describe the movement observed in the cytoplasm. Second, in order to justify the assumption that there are no boundary interactions in the time-lapse images, we needed to discard large amounts of image data by taking subsets of the central region of the cytoplasm. Finally, we lack validation to guarantee that the reported statistics from DDM are a true reflection of variable movement rates in the cytoplasm. In this chapter, we present two new approaches to overcoming these challenges using the synthetic data described in Chapter 3.

Given an image stack, one can always generate the DDM matrix, $\mathcal{D}_R(q, \Delta t)$, by following the protocol outlined in Section 2. The matrix must then be interpreted with respect to the underlying motion. The shape of the matrix, particularly along trajectories with fixed spatial frequency q where the scale and noise terms $A(q)$ and $B(q)$ in Equation (2.2.18) can be decoupled, can inform the user of properties of the images; for example, a DDM matrix exhibiting oscillation within fixed q transients is likely to have ballistic or oscillatory properties, or the timescale over which the DDM matrix converges to some constant may be used to interpret the timescale of the underlying behaviour [310]. However, the more common application method by which the DDM matrix is interpreted is through numerical fitting to some model, determined by the type of movement expected to be observed within the images. These models typically take parameters which

can be interpreted through one of three standard measurements used to analyse \mathcal{D}_R [311];

1. The Intermediate Scattering Function (ISF), $f(q, \Delta t)$, which describes the autocorrelation of particles in the system at specific time and frequency combinations. Examples can be found in [142, 280, 312–314]. For Brownian motion, $f(q, \Delta t) = e^{-q^2 D \Delta t}$ in any dimension.
2. The relaxation rate, $\Gamma(q)$, which describes the timescale for which a system perturbed at frequency q will return to its original equilibrium state. Examples can be found in [282, 305, 311]. For Brownian motion, $\Gamma(q) = q^2 D$ in any dimension.
3. The Mean Square Displacement, $\langle \Delta r^2(\Delta t) \rangle$, a measure of the average distance a particle is expected to travel in a given time interval. Examples can be found in [282, 315]. For Brownian motion in N dimensions, $\langle \Delta r^2(\Delta t) \rangle = 2ND\Delta t$ [281].

These measurements are not stand-alone; they are mutually defined as different lenses through which we can study the same displacements. There is no significant advantage to choosing one measurement over another, although as an ensemble, they provide a more complete characterisation. Determining any of these listed quantities from the DDM matrix requires being able to estimate the scale and noise terms defined in Equation (2.2.19), $A(q)$ and $B(q)$, respectively [311].

Additionally, extracting specific parameters from any of these listed quantities requires a model function to interpret the DDM matrix in terms of the chosen measurement. The types of data we can approach using DDM are therefore limited either by our ability to determine a model function, or by the library of model functions available in literature. We are interested in the ISF, for which a considerable number of models exist for very complex behaviours, including diffusion [266], bacteria swim speed distribution [163], rotation [316], stochastically switching two state dynamics [317] and more. The ISF describing motile bacteria has been previously used to characterise cytoplasmic movement in *Drosophila* oocytes [142], but in Section 2.3, we showed that the same approach cannot be applied to the mouse oocytes studied in this work. Since there is no ISF available in literature for this specific application, we must formulate one ourselves in order to use DDM as a predictor of oocyte health.

Although our focus is on characterising cytoplasmic movement in oocytes, the challenge of picking an ISF is not exclusive to this application. The library of currently defined ISFs will never be exhaustive, so it is likely that for a given new application of DDM, it will be necessary to create a new model for the

ISF. This requires some level of mathematical training, which presents a possible barrier to entry for the technique. More importantly, this requires that a suitable analytic form of the ISF exists. In Section 1.2.2.1, we described how the cytoplasm is subject to spatially heterogeneous fluid-mechanical models which are solved numerically, not analytically, and hence are currently inaccessible with DDM.

The second challenge we discussed in the context of DDM applied to oocytes was boundary interaction. Depending on the nature of the boundary, this can be considered a subset of the challenge of determining a new ISF. For example, in [318], DDM was used to quantify the rate of diffusion of particles, where the domain featured a regular lattice of nano-posts, small cylindrical plastic protrusions into the domain with a diameter of 500nm, compared to particles with diameter ranging from 200-500 nm. The diffusive movement of the particles is inhibited by collision with the nano-posts, which becomes more extreme as the post array becomes more densely arranged. It was shown in [318] that this confined diffusion was well described by the Kohlrausch–Williams–Watts model, which is equivalent to the Brownian motion ISF with an additional stretching exponent term, i.e.

$$e^{-q^2 D \Delta t} \rightarrow \left(e^{-q^2 D \Delta t} \right)^{\Upsilon(q)}.$$

The same approach can be taken to describe diffusion in crowded domains [319], where collisions between particles result in sub-diffusive movement. In both of these applications, the newly introduced interaction, either with the boundary or with other particles, is homogeneously spread throughout the domain. In oocytes, however, the boundaries exist on the outside of the domain, which means that displacement of a given particle in the cytoplasm depends on the initial position of the particle, as well as the time lag between observations. We, hence, cannot adopt the same stretching strategy in our application.

A further example of DDM in a bounded domain can be found in [320], in which a variant of DDM is used to probe the relaxation times of the formation of a constrained ferrofluid in the presence of a magnet. This application does not use DDM for parameter fitting, however, instead estimating the time at which the DDM matrix stops changing to gauge the timescale of movement. As such, no alternative strategy for defining the ISF is needed.

With no alternative method of describing the ISF, boundary effects are currently mitigated by cropping the image to remove the boundary from the time-lapse images [142, 280, 302]. This approach removes boundary interaction, resulting in the distribution of displacement throughout the frame being drawn from a single, spatially homogeneous process which DDM analysis can be used to characterise. However, this also results in throwing away important data which could

provide new or complementary information on the health of the oocyte. For example, consider the cytoplasmic streaming model in mouse oocytes we presented in Figure 1.6, where the cytoplasm flows down the the sides of the cell. Missing this behaviour could limit the information we derive from DDM analysis. Furthermore, by cropping the image stack, we limit the size of the imaged domain. The assumption that edge interactions are negligible becomes stronger as the size of the subdomain increases, but smaller images increase the likelihood that we cannot capture a sufficient sample of the observed motion to approximate the ISF, as well as increasing the risk that poor performance at small and large frequencies may distort the fitting of key movement parameters [164].

There is, thus, an incentive to account for boundary effects in the ISF, as opposed to cropping the image data. However, we now demonstrate in Section 4.1.1 that even when boundary effects can be included in the analytic form of the ISF, the resulting form of the ISF may be unsuitable for the parameter fitting algorithms required to extract movement parameters from the DDM matrix. We do so by presenting a relatively simple ‘toy’ problem, of Brownian motion inside a finite square domain with reflective boundaries..

4.1.1 Example: Brownian motion in a square domain with reflective boundary conditions

Before approaching a two-dimensional domain, it is easier to consider the one-dimensional case first. Let the domain be a line of length L with reflective boundaries at 0 and L . The PDE system solved by the van Hove function is the same as in standard Brownian motion, given in Equation (2.2.20), with the addition of a zero-flux Neumann condition on each end of the domain,

$$\begin{aligned}\frac{\partial G(\Delta x, \Delta t, x_0)}{\partial t} &= D\nabla^2 G(\Delta x, \Delta t, x_0), \\ \frac{\partial G(\Delta x, \Delta t, 0)}{\partial n} &= \frac{\partial G(\Delta x, \Delta t, L)}{\partial n} = 0, \\ G(\Delta x, 0, x_0) &= \delta(\Delta x).\end{aligned}$$

An analytic solution to this system can be found by the method of images [321]. This considers the release of ‘ghost’ particles outside of the boundary, at positions $2mL \pm x_0$, where $m \in \mathbb{Z}$. The real and artificial particles each have a respective probability of being found at any point in the domain, represented by a Gaussian, and the van Hove function, given the initial position of the particle in the bounded domain, is therefore equal to the sum of the contributions from each Gaussian over the interval $[0, L]$. This can be written as

$$G(\Delta x, \Delta t, x_0) = \frac{1}{\sqrt{4\pi D\Delta t}} \sum_{m=-\infty}^{\infty} e^{-\frac{(\Delta x - 2mL - x_0)^2}{4D\Delta t}} + e^{-\frac{(\Delta x - 2mL + x_0)^2}{4D\Delta t}}.$$

For large L , with x_0 sufficiently far from boundaries, this rapidly decays to the infinite boundary case. In order to determine the van Hove function, we must average this PDF over all possible positions x_0 , i.e

$$G(\Delta x, \Delta t) = \frac{1}{L} \int_0^L G(\Delta x, \Delta t, x_0) dx_0.$$

For a square domain with side length L , since the x and y displacements are independent, we can trivially generate the van Hove function by taking the product of the one-dimensional van Hove function with its transpose, i.e.

$$G(\Delta x, \Delta y, \Delta t) = \frac{1}{L^2} \int_0^L \int_0^L G(\Delta x, \Delta t, x_0) G(\Delta y, \Delta t, y_0) dx_0 dy_0.$$

However, even for such a simple example, the van Hove function and corresponding ISF are given by an infinite series expansion. Whilst we can take only leading order terms for simplicity, parameter fitting would involve numerically resolving the van Hove function for a given choice of D , before numerically integrating over all potential initial conditions. There are now several steps in which we might introduce numerical error into the fitting process, increasing uncertainty around the validity of the output of DDM. Furthermore, slight changes to the dynamics or boundary could easily result in an ISF which cannot be analytically expressed at all. We therefore require some new approach to characterising the ISF in such a system.

4.2 Simulation-determined fitting functions in DDM analysis.

We introduce a new method for generating an analytic expression of the ISF, by first numerically approximating it using synthetic data, and using polynomial interpolation to reclaim the continuous expression for the ISF. This new method, which requires none of the cumbersome analysis or integral handling of previous analytic approaches, works well in behaviours where the trajectory of a single particle can be well described, even if the average behaviour cannot. Simulations may be constructed using the methodology outlined in Chapter 3.

It is easier to introduce this workflow through a toy problem, for which we choose Brownian motion confined to a circular domain with reflective boundaries. This is a more complex version of the problem introduced in Section 4.1.1, where the boundaries were straight and parallel to each respective axes, allowing the problem to be decoupled into the product of one-dimensional Brownian motions.

The general workflow, also outlined in Figure 4.1, is as follows:

1. First, simulate a large number of particle trajectories. Displacements are drawn from the corresponding van Hove function using a fixed set of parameters (in this case, constant diffusion coefficient D), to create trajectories in the unbounded domain. These are mapped back into the bounded domain, by calculating the effect of boundary interactions. These simulations are repeated for a range of input movement parameters.
2. The simulated trajectories can then be binned into a histogram to numerically approximate the self-correlating van Hove function $G(\Delta\mathbf{r}, \Delta t)|_D$ for the bounded process at every time lag, Δt , where the subscript D denotes the van Hove function is evaluated at a fixed diffusion coefficient.
3. By Equation (2.2.15), which states that the real part of the Fourier transform of the van Hove function is equal to the ISF, we can convert each numerically approximated van Hove function into a numerical approximation of the ISF, by the relation $f(\mathbf{q}, \Delta t)|_D = \Re \left(\tilde{G}(\mathbf{q}, \Delta t)|_D \right)$.
4. We apply the radial average to each ISF, assuming that the observed motion isotropic and the domain is radially symmetric, such that we have $f(q, \Delta t)|_D$.
5. During the DDM fitting process outlined in Section 2.2.7, we aim to fit a continuous diffusion coefficient over a continuous time lag, at each discrete spatial frequency. In order to do the same, we observe that each numerically approximated $f(q, \Delta t)|_D$ forms a matrix, which can be stacked together to form a tensor. This tensor can be sliced into a different set of matrices, which evaluate the ISF at a fixed values of q for variable D and Δt . Each matrix is assigned a label $f(D, \Delta t)|_q$. This process is visualised in Figure 4.2.
6. A continuous spline is fitted to each $f(D, \Delta t)|_q$, resulting in a collection of continuous polynomial fitting functions labelled $\mathcal{P}(D, \Delta t)|_q \approx f(D, \Delta t)|_q$.

Critically, the polynomials $\mathcal{P}(D, \Delta t)|_q$ satisfy the following relation,

$$\mathcal{D}_R(q, \Delta t) = A(q) (1 - \mathcal{P}(D, \Delta t)|_q) + B(q), \quad (4.2.1)$$

implying that they can directly replace the ISF in the fitting stage of DDM analysis.

An important difference between the two-stage fitting approach typical of DDM, and the spline fitting in Equation (4.2.1), is that the diffusion coefficient is output as a function of q . In order to extract a single diffusion coefficient, we

Repeat for range of D :	Output
Simulate trajectories using fixed D	$\mathbf{r}_n(t) _D$
Bin into histograms to approximate van Hove function	$G(\mathbf{r}, \Delta t) _D$
Fourier transform to get ISF at fixed D	$f(\mathbf{q}, \Delta t) _D$
Radial average	$f(q, \Delta t) _D$
Rearrange to get ISF at fixed q	$f(D, \Delta t) _q$
Repeat for all q :	
Fit continuous polynomial to discrete ISF	$\mathcal{P}(D, \Delta t) _q$
Fit polynomial to $\mathcal{D}_R(q, \Delta t)$ to get $D(q)$ at each frequency	$D(q)$
Fit constant diffusion from $D(q)$	D

Figure 4.1: The process flow for the simulation-driven pipeline to fit the diffusion coefficient from Brownian motion in a circular domain with reflective boundaries. The objective is to create a continuous numerical polynomial approximation of the ISF, by first generating it numerically from simulated particle trajectories. These polynomials can be used to extract the diffusion coefficient from the DDM matrix.

Re-ordering

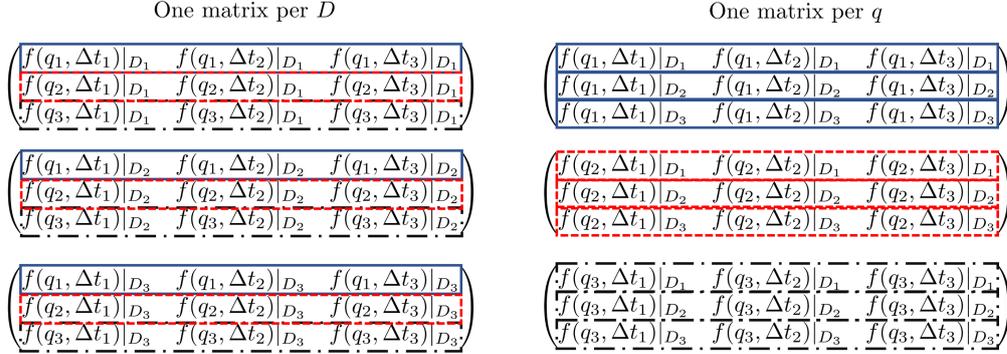


Figure 4.2: Re-ordering allows us to express the ISF at a single frequency q , with varying time and diffusion coefficient, Δt and D respectively.

must consider the average of $D(q)$ over some reasonable fitting interval $[q_l, q_u]$ where $D(q)$ is flat.

To demonstrate the validity of the proposed workflow, we apply it to a simulated, bounded Brownian motion inside a circular domain. Particle trajectories

forming both the synthetic data from which we approximate the ISF, as well as the images used for DDM, are determined using a pre-existing MATLAB function from the Central File Exchange [322]. We consider a circular domain with diameter 500 pixels, inlaid at the centre of a square image with side length 512 pixels. To numerically approximate the ISF, we perform Brownian motion simulations with diffusion coefficients sampled between 0.01 and 3 px²/frame at regular intervals 0.1 pixels²/s apart. We simulate 1000 particles for each diffusion coefficient, over a timespan of 2000 time steps at a non-dimensional frequency of 1. Cubic splines are fit to the numerically approximated ISF values $f(D, \Delta t)|_q$ in MATLAB.

The images used to validate our technique use the same code to generate trajectories, with diffusion coefficients between 0.05 and 3 px²/frame, regularly spaced 0.15 apart. These images consider only 150 particles, consistent with our recommendations from Chapter 3. To benchmark the performance of our new approach to generating the ISF, we also use the current standard fitting approach of taking subdomains of the original image. The purpose of this work is not to optimise the size of the subdomain over which DDM analysis is applied, but having noted that such an optimum may exist, we must consider multiple subsets of the image to ensure we do not negatively bias our comparison against the new fitting pipeline. To define the size of the image subsets considered, if the circular boundary of the system has radius R , the first image subset is defined by the largest square inscribed within the circular domain, which has side length $L_1 = \sqrt{2R^2}$. Subsequent refinements are found by taking the square subset inscribed in the circle whose diameter equals the side length of the previous subset. Hence, the i^{th} image subset has side length

$$L_i = \sqrt{\frac{L_{i-1}^2}{2k}}, \quad k \in \mathbb{N} \setminus 1.$$

We also consider the $i = 0$ refinement, where the subdomain is the whole image. We visualise the refinement scheme in Figure 4.3.

When we apply our new pipeline to fit the diffusion coefficient from the DDM matrix, we also note that we must remove the image windowing described in Section 2.2.8. The image window makes the assumption that dynamics are invariant of position in the image, and in particular, that the dynamics in the centre of the window are equivalent to those at the edge. This is not true for the dynamics we study here, since particles behave differently near the boundary of the domain. Windowing would reduce the contribution of particles near the boundaries, which will skew our ISF towards that of classical Brownian motion. We retain the win-

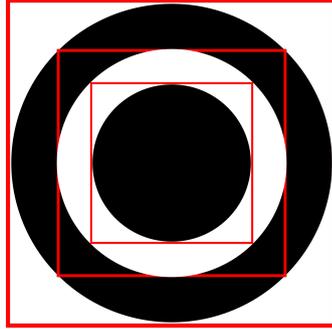


Figure 4.3: Visualisation of the scheme used to refine the image domain length iteratively, defining subdomains (squares) that inscribe circles which, in turn, inscribe the squares defining the previous subdomain.

dow when applying standard DDM approaches, however, since this is consistent with literature recommendation [295].

To each DDM matrix generated, we fit the diffusion coefficient using the new numerically-guided pipeline, as well as using the unaltered DDM code for each consecutively smaller image subdomain, displaying the results in Figure 4.4. All approaches underestimate the diffusion coefficient, but the spline method achieves a consistently closer estimate than standard DDM approaches, regardless of subdomain size. This demonstrates that there is an accuracy benefit achieved by employing this new approach. Underestimation of the diffusion coefficient might be reasonably expected with classical DDM approaches applied to bounded domains; the presence of boundaries limits the displacement any particle may undertake, which means particles appear to move less inside the bounded domain, which is interpreted by DDM as slower movement.

Contrary to expectation, however, the accuracy of the classical DDM approach decreases as we reduce the size of the subdomain, with the optimal choice of subdomain being the entire image. This suggests that the trade-off between information loss from taking sub-images, and the negative effect of including the boundary in the model, does not exist in these simulations. It could be argued that this is because the circular domain diameter (500) is approximately equal to the image length (512), and the Blackman-Harris filter is fairly aggressive; application of the filter may essentially remove the boundary effects by weighting the centre of the domain considerably more than the edges, even when considering the whole image. In this case, there is no benefit to reducing the size of the subdomain, since boundary effects are already accounted for. This theory can be tested by repeating the simulations using a significantly smaller circular domain of diameter 250 pixels, for the same image size (512 pixels long). In this repeated set of simulations, the domain is much farther from the edges of the image, so will not be removed by the Blackman-Harris filter. We use the same subdomain sizes,

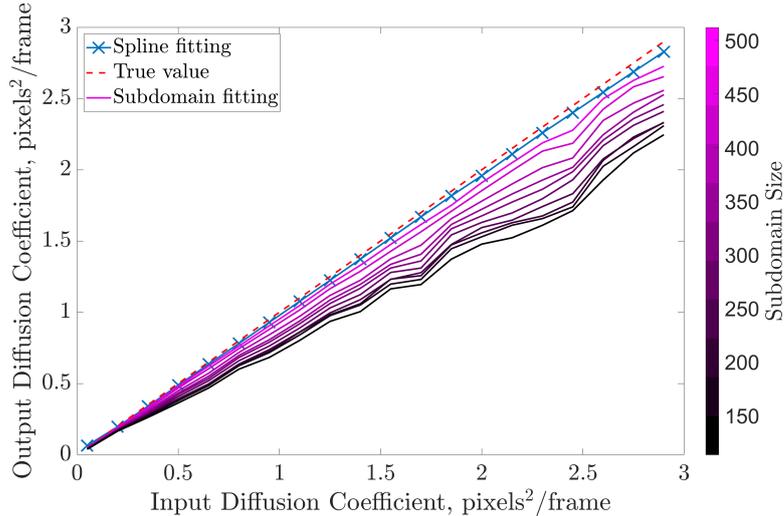


Figure 4.4: Fitted diffusion coefficients from simulated Brownian motion inside of a reflective circular domain of diameter 500 pixels, contained within an image window with length 512 pixels. Fitting is performed both using a classical windowed DDM approach (black-pink gradient denoting varying window size), and using the spline interpolation method proposed (x curve). The ‘true value’ curve is also given (red dash).

as opposed to re-applying the scheme outlined in Figure 4.3 to redraw smaller subdomains. Determining a new set of fitted diffusion coefficients, we show in Figure 4.5 that the same pattern emerges, with the magnitude of under-fitting increasing as the subdomain size decreases.

Making the circular bounded domain smaller should make the boundary effects more significant, but we still achieve the optimal fitting performance by not taking subsets of our image. Therefore, a new explanation is needed to understand why sub-setting results in worse fitting of the diffusion coefficient. The perceived benefit of sub-setting images is that particle interactions with the boundary are unobserved, but we will now show that the probability of boundary interaction is very low in our toy problem, such that we gain little benefit from taking image subsets. In this way, only the negative consequences of considering smaller images are incurred, in that we discard useful image data and restrict the range of frequencies over which we may fit the diffusion coefficient.

To calculate the probability of a particle interacting with a boundary, we observe that if the displacement of a given particle is known, we can determine for which initial locations in the domain the trajectory would intersect with a boundary. Given the distribution of initial locations is uniformly random, this allows us to estimate the probability of a boundary interaction occurring at any given time. Assume that particles are uniformly distributed throughout a circular domain of radius R , centred at $(0,0)$. Furthermore, we assume that after a time period Δt , the displacement a given particle would have made without a boundary

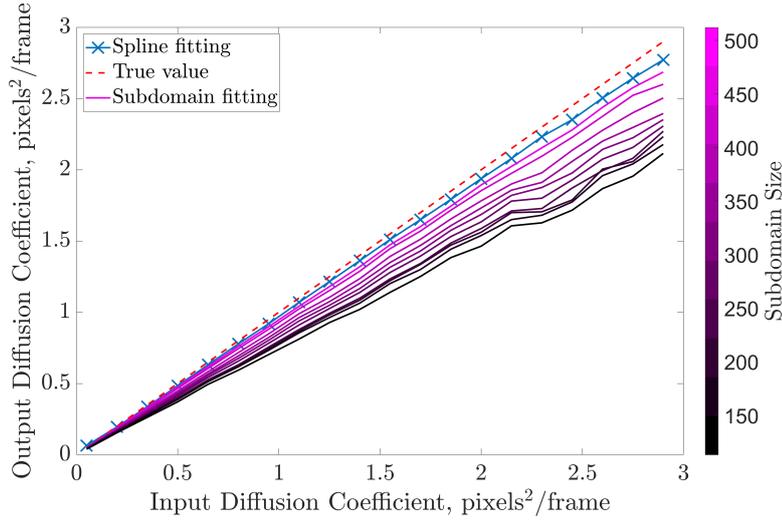


Figure 4.5: Fitted diffusion coefficients from simulated Brownian motion inside of a reflective circular domain of diameter 250 pixels, contained within an image window with length 512 pixels. Comparing to Figure 4.4, almost the exact same distribution of diffusion coefficients is observed despite the diameter of the bounded domain being halved. Additionally, the spline fitting under-estimates the diffusion coefficients more significantly, particularly when $D > 2.25$, compared to Figure 4.4.

is Δx , which, due to the radial symmetry of the domain, can be assumed to be in the horizontal direction without loss of generality, as shown in Figure 4.6.

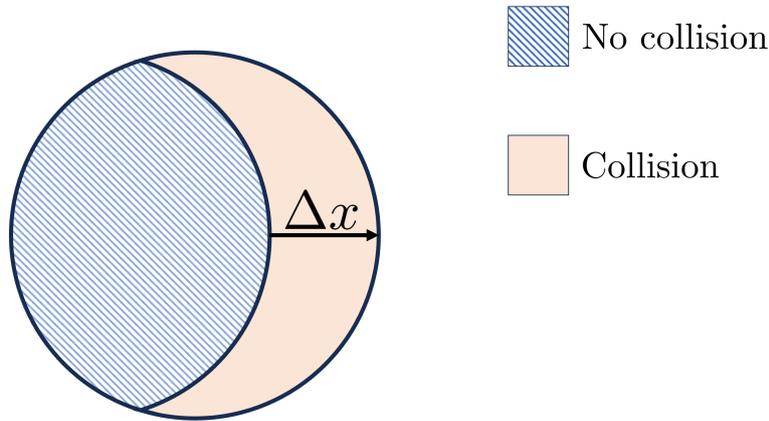


Figure 4.6: For a particle located inside a circular domain which undergoes a horizontal displacement of Δx , collision occurs if the particle was initially located inside of the mono-colour ‘crescent moon’ area.

Boundary interaction occurs when the horizontal distance between the particle and boundary is less than or equal to Δx . The area of the domain for which this condition is met, shown by the peach coloured ‘crescent moon’ shape in Figure 4.6, has area equal to the circular domain, minus the area of the circle with equal

radius centred at $(-\Delta x, 0)$. The area where the particle would not interact with a boundary is equal to

$$A = 2R^2 \cos^{-1} \left(\frac{\Delta x}{2R} \right) - \frac{\Delta x \sqrt{4R^2 - \Delta x^2}}{2}.$$

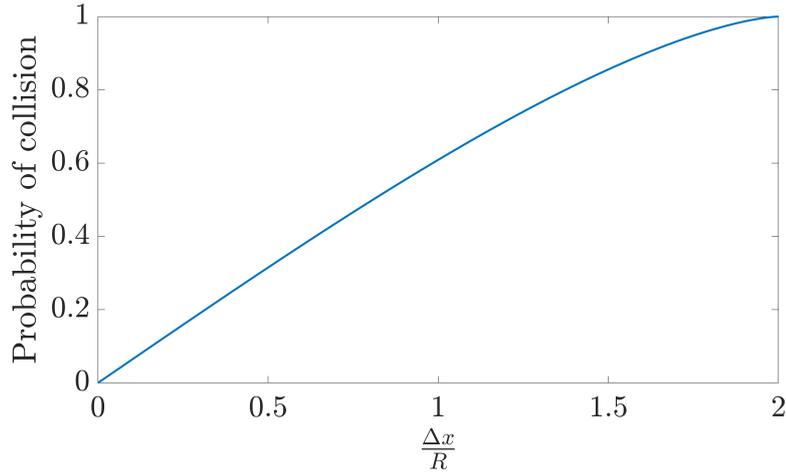
The probability of a boundary interaction given a displacement of Δx , plotted in Figure 4.7a, is the ratio of the non-intersection area over the area of the domain, i.e

$$\frac{\pi R^2 - A}{\pi R^2}.$$

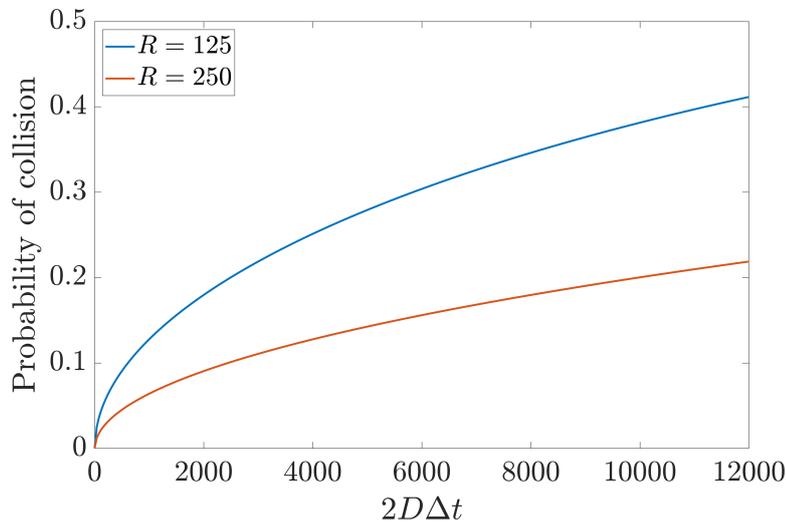
We additionally know that the probability of an orientation-independent displacement Δx occurring is normally distributed, with mean 0 and variance equal to the MSD, $2D\Delta t$ (where orientation-independence reduces the displacement to one dimension). The product of the probability of collision given Δx , and the probability of displacement Δx , integrated over all Δx , yields a general collision probability distribution depending on the variance $2D\Delta t$. We plot the collision probability against the MSD in Figure 4.7b. In our simulations, $\text{MSD} < 12000 \text{ px}^2$ (calculated by considering the largest time lag and diffusion coefficient), but we can see from Figure 4.7b that even for this MSD, the collision probability is less than half for both sizes of circular domain considered.

Even when a collision does occur, we can also consider how significantly the particle trajectory is affected by that collision. For particles located initially very close to a boundary, the displacement before collision is very small, which means only a small amount of displacement is ‘lost’ relative to an unbounded domain. Similarly, particles colliding with a boundary towards the end of their displacement will only lose a small proportion of their displacement. So, even when boundary interactions occur, the size of the interactions can be comparatively small.

To conclude, boundary interactions in these simulations are both unlikely, and may have a minimal impact on the trajectory for short time lags. This means that the boundary does not contribute significantly to the shape of the DDM matrix. As such, including the boundary in the image domain does not have a strong effect on the output statistics from DDM analysis, which means that taking smaller image subdomains has no particular benefit. However, this may not be true if we studied faster-acting Brownian motion, different types of movement, or even different boundary condition (such as a combination of sources and sinks). The uncertainty associated with the choice of subdomain presents a significant benefit of our alternative, simulation-guided approach to fitting bounded



(a) The probability of a boundary interaction occurring for a particle which displaces by a distance $\Delta x/R$, where R is the radius of the bounded domain.



(b) The probability of a boundary interaction occurring as a function of the Mean Square Displacement, $2D\Delta t$, for two different domain radii $R \in \{125, 250\}$.

Figure 4.7: Visualisations of the probability of boundary interactions for particles undergoing Brownian motion inside of a circular domain. These plots demonstrate that, for the small time lags investigated during DDM, collision is unlikely.

domain behaviours, since it can be applied regardless of how significant the impact of boundary interactions is. This mitigates the requirement to estimate the significance of boundary interactions, in order to motivate selection of optimal subdomain size.

Before discussing limitations of our new approach to generating the ISF, we will make a general comment about the performance of DDM in bounded domains, such as our toy problem. When the domain is bounded, the van Hove function will converge to a uniform distribution much faster than in an unbounded domain. When this happens, the DDM becomes constant with increasing

time lags; this shortens the optimal range on Δt for which DDM provides useful information about the spatial autocorrelation of particle positions. For bounded domains, it is therefore recommended to consider a shorter movie length, imaged at a higher frequency. This comment holds true regardless of the technique used to analyse the DDM matrix, and hence affects classical fitting approaches in the same way as our new proposed pipeline. However, we raise this important comment to demonstrate that our new approach does not negate all of the issues we face in bounded domains. We now address some limitations of our analysis approach.

First, simulation and spline generation is slow, computationally expensive and requires the storage of a large simulated ISF bank. In hardware-limited environments, such as the IVF clinic, performing simulations like these may not be viable. If the objective is to automate DDM analysis for one particular kind of problem, spline generation need only be performed once, and simulation work could therefore be performed elsewhere, outside of the lab or clinic, and the splines themselves can be imported. For more variable applications, however, where the domain size or ISF is not fixed, the user would be required to have access to sufficiently powerful hardware.

Another limitation is that when generating splines, we naturally assume there is some upper bound on the input movement parameters. Our proposed pipeline uses polynomial interpolation to fit the diffusion coefficient, which means that if the upper bound on D in the simulated data is smaller than the true value of D in a real dataset, we will be unable to extrapolate and accurately evaluate the diffusion coefficient. In practice, we may not have reasonable bounds for the expected movement parameters before undertaking parameter fitting, so this limits the usefulness of our proposed pipeline when applied to new, poorly understood datasets.

Similarly, we not only had to decide the maximum bound for the diffusion coefficient, but also the grid refinement on D for which we generate simulations. Sampling more diffusion coefficients improves the quality of interpolation, but increases computational costs for generating the splines. Additionally, the grid refinement defines the smallest values of diffusion which we may capture. Knowing what range of movement parameters is appropriate to simulate over *before* undertaking analysis requires physical intuition about the studied system, which may not be available.

Another limitation of our proposed pipeline is that it will become impractical when fitting more than one parameter. Our implementation contained only one movement parameter, the diffusion coefficient, but for more complex systems with more parameters, the spline interpolation method may fail to well characterise the

studied system. The size of the polynomial used to interpolate these parameters grows rapidly with an increasing number of variables, with the appearance of mixed terms of varying order that will make parameter fitting more challenging. It would be a useful direction for future work to explore the limits of our approach, to identify the number of parameters we can accurately account for.

Despite these limitations, the new pipeline introduced in this section has shown a significant improvement in the ability to fit the diffusion coefficient from a bounded domain. However, it is important to note that the fitting algorithms used to generate the diffusion coefficient from the numerically determined ISF are the same as in classical DDM. In Section 4.3, we demonstrate that in when fitting more than one movement parameter from the system, the accuracy of these fitting algorithms can vary strongly depending on the true value of the movement statistics driving motion in the image stack.

4.3 Simulations highlight poor performance of fitting algorithms when some movement sources dominate others

When applying DDM to oocyte data previously in Section 2.3, we identified that both available models for the ISF, Brownian motion and an undirected ‘bacteria’ advection-diffusion, resulted in poor quality parameter fitting. A new model is therefore required, the construction of which must be motivated by the underlying oocyte biology. The strategy proposed in Section 4, for example, could be used to generate this ISF numerically, from simulations which match experimental observations of movement in the cytoplasm of the oocyte. However, even when applying our new ISF generation methodology to synthetic Brownian motion data, we observed a consistent fitting error, showing that even when the ISF is chosen correctly, numerical constraints may play a significant role in the performance of DDM.

If true, the correlation between the output parameters from DDM analysis and the ‘true’ movement statistics of the observed system may be weakened. This would have problematic consequences for the application of DDM in fertility applications, preventing the prediction of oocyte quality. There are many well documented numerical constraints which can cause the output statistics from DDM to become unreliable, but in this section, we show that even when these constraints are controlled for, it is possible to construct ‘well-conditioned’ images that DDM characterises poorly. In order to define a ‘well-conditioned’ image set, we will first briefly review the current suite of additional constraints which must be met in order for DDM to be applicable.

Perhaps the most well known numerical constraint faced in DDM is that it is necessary to decouple the image noise and scaling terms, $B(q)$ and $A(q)$ respectively, from the DDM matrix in order to fit movement parameters. This is only possible when there is a sufficient amount of long and short time lag data available [164, 266]. One method for controlling this numerical error is through experimental design, combining high frequency and long exposure movies [164]. In commercially available bench-top incubators, the imaging frequency is often fixed, which can make controlling for this constraint more challenging. A second strategy is to adopt an altered form of DDM named ‘cross-differential dynamic microscopy’ (c-DDM) [323], where multiple cameras are used to capture copies of the image stack. Image differences are taken to compare frames not only acquired at different times, but from different cameras. This reduces the noise contribution in the DDM matrix by cancelling the different levels of noise from each camera, and also increases the range of viable fitting frequencies.

In this work, we do not add noise to the simulated images generated, so only the scaling term must be decoupled from the DDM matrix. This constraint can therefore be mitigated by sampling at a sufficiently high time lag [164]. Even in the absence of imaging noise, the inclusion of a fitted noise term appears to improve the fitting performance in DDM; this can be attributed to reducing over-fitting [305], by allowing a small error-tolerance. This fitted noise is always several orders of magnitude smaller than $A(q)$, and hence the decoupling of noise and scale need not be considered.

A second numerical constraint arises due to the fact that the DDM matrix is a statistical sample, and will hence always be subject to an element of randomness. This results in fluctuations that can manifest as distortion in the observed pattern of the ISF, outliers and over-fitting, where the model attempts to capture the noise as part of the model [305]. Considerations for robust validation steps to overcome this numerical error have been proposed in the ‘DDM with Uncertainty Quantification’ method (DDM-UQ), enabling the detection of bias from image noise during fitting of the diffusion coefficient in Brownian motion [305]. Following the work in Section 3.3, we have already controlled for this constraint by ensuring our simulated images contain a sufficiently high particle density, which minimises the variation between DDM matrices with the same input movement parameters.

Finally, in Section 3, we discussed localisation error [292], where uncertainty in the estimation of particle position due to the discrete nature of images causes an underestimation of movement parameters. Whilst this is a form of numerical error, we note that this error can be accounted for by an injective, multiplicative mapping of the fitted output of DDM to the true output. Hence, we need not control for this form of error.

We have discussed several sources of numerical error, some of which are controlled for, and others which we justify that no mitigating action is required. We now introduce a new form of error not presently discussed in literature. In order to describe movement of the cytoplasm, we have previously discussed the need to characterise two independent forms of motion, active diffusion and cytoplasmic streaming. When considering more than one source of movement, DDM has to indirectly partition the observed displacement of particles into contributions from each type of motion. This task is theoretically possible, so long as the two forms of motion act on different timescales; for example, in bacterial motion, the decorrelation over short timescales is dominated by diffusive-type motion, but at longer timescales, advection dominates, resulting in two distinct regions in the ISF that guide parameter selection [164]. However, there will exist some central region for which both behaviours act over similar timescales. In this region, similar DDM matrices may be observed for a range of different input movement parameters. When attempting to determine movement parameters from the DDM matrix, this can cause one parameter to be over-estimated, whilst the other is underestimated.

An important challenge is therefore to explore this source of numerical error, identifying when it happens, how significant the error is, and whether it can be mitigated or controlled for. Tackling this challenge is highly difficult in real datasets, as there is often no ‘true value’ against which we may compare fitted parameters. A notable exception is Brownian motion, where we can use the Stokes-Einstein relation to approximate the theoretical diffusion coefficient [164, 266, 270]. To compare other fitted parameters, however, one often requires an additional image analysis technique. For example, fitted velocity fields have been compared between DDM and PIV when studying cytoplasmic streaming in *Drosophila* oocytes, despite observing that the accuracy of PIV is diminished when the displacements due to diffusion and streaming are of a similar order of magnitude [142]. Another example can be found when using DDM to characterise the motility of bovine semen samples, where the mean velocity derived from DDM is compared against SPT experiments [270]. The outcome of the study in [270] was that the predicted movement parameters from DDM and SPT experiments were not identical, but were of the same order of magnitude, implying that error from combining estimates of parameters is bounded.

In the IVF lab, it is not sufficient to evaluate movement parameters up to the correct order of magnitude; DDM may only be used to discriminate between oocytes if the error with which we estimate movement parameters is smaller than the natural variance between two different oocytes. It is therefore crucial to determine how sensitive DDM analysis is to the parameters driving the movement of the cytoplasm. Other sources of numerical error, as described, can be

mitigated by controlling the way images are generated, but if fitting becomes inaccurate for particular combinations of parameters, this type of error cannot be controlled for, and would highlight a limitation of current DDM approaches in oocyte assessment.

Simulations provide a perfect medium for exploring this phenomena, because we already know the ‘true value’ parameters driving movement of particles, against which we may compare the fitting results from DDM. In order to observe this new type of numerical error, we need to define a system in which there are two forms of independent motion. Hence, we consider a simplified advection-diffusion motion, where advection is represented by a homogenous, constant velocity field \mathbf{v} applied to all particles. For a directed, constant velocity field, the van Hove function is given by

$$G(\Delta\mathbf{r}, \Delta t) = \delta(\Delta\mathbf{r} - \mathbf{v}\Delta t),$$

for which the Fourier transform is

$$\tilde{G}(\mathbf{q}, \Delta t) = e^{i\mathbf{q}\cdot\mathbf{v}\Delta t}. \quad (4.3.1)$$

Since the advection and diffusion effects are independent, the theorems outlined in Section 2.2.6 allow us to construct a fitting function by multiplying the individual ISFs for each process. The real part of Equation (4.3.1) is equal to

$$\Re\left(\tilde{G}(\mathbf{q}, \Delta t)\right) = \cos(\mathbf{q}\cdot\mathbf{v}\Delta t).$$

Defining $\mathbf{q} = q(\cos(\theta), \sin(\theta))$ and $\mathbf{v} = v(\cos(\phi), \sin(\phi))$, the radial average of the ISF (see Section 2.2.3) is equal to

$$\begin{aligned} \frac{1}{\|\Gamma\|} \oint_{\Gamma} \cos(\mathbf{q}\cdot\mathbf{v}\Delta t) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \cos(\mathbf{q}\cdot\mathbf{v}\Delta t) d\theta \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \cos(qv\Delta t \cos(\theta - \phi)) d\theta. \end{aligned}$$

The symmetry of the Fourier transform implies the integrand is periodic with period π , so we are integrating over twice the period. We can therefore consider the equivalent integral

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} \cos(qv\Delta t \cos(\theta - \phi)) d\theta = \frac{1}{\pi} \int_0^{\pi} \cos(qv\Delta t \cos(\theta - \phi)) d\theta.$$

This integral, with domain equal to the period of the integrand, is invariant under a constant shift in θ . Since ϕ is constant, for the new parameter $\Theta = \theta - \phi$, the following integral is equivalent to the radial average,

$$\frac{1}{\|\Gamma\|} \oint_{\Gamma} \cos(\mathbf{q}\cdot\mathbf{v}\Delta t) = \frac{1}{\pi} \int_0^{\pi} \cos(qv\Delta t \cos(\Theta)) d\Theta.$$

By definition of the zero-order Bessel function of the first kind [324],

$$\frac{1}{\pi} \int_0^\pi \cos(qv\Delta t \cos(\Theta)) d\Theta = J_0(qv\Delta t). \quad (4.3.2)$$

The ISF for the joint advection-diffusion process is the product of the ISF for each independent process (see Section 2.2.6), yielding the ISF to be equal to

$$f(q, \Delta t) = e^{-q^2 D \Delta t} J_0(qv\Delta t). \quad (4.3.3)$$

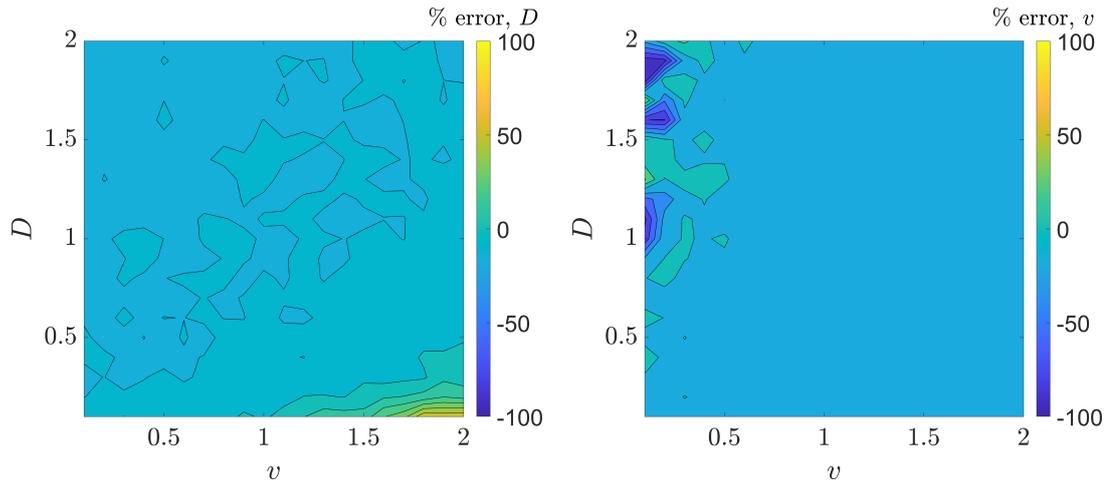
We generate simulations of the described advection-diffusion behaviour for diffusion coefficients sampled at intervals of 0.1 between 0.1 and 2 px²/frame, and sample the velocity over the same interval, units px/frame. Each simulation considers a stack of 2000 images containing 150 particles. These simulations are converted to image stacks, as described in Chapter 3, from which the DDM matrix is generated. To each DDM matrix, standard fitting algorithms are applied, using Equation (4.3.3) as the model of the ISF, to determine the characteristic values $\tau_D = 1/q^2 D$ and $\tau_v = qv$ at each frequency. These fitted values are then plotted against frequency, to determine the velocity field and diffusion coefficients. We expect the absolute approximation error to increase with the size of the displacements observed in the frame; hence, to allow comparison of the fitting error for different input parameter values, we express error as a percentage of the true value. For example, the percentage error of the diffusion coefficient is equal to

$$\frac{D_{\text{fitted}} - D_{\text{true}}}{D_{\text{true}}}.$$

We visualise the fitting error achieved over these simulations in Figure 4.8.

In general, the percentage fitting error for both the diffusion coefficient (Figure 4.8a) and velocity magnitude (Figure 4.8b) is relatively constant for both parameters, with the diffusion coefficient being under-estimated by approximately 8-10% and the velocity being under-estimated by around 5%. Such consistent error, previously attributed to localisation error in Section 3.3, can be adjusted in post-processing, and is therefore not problematic.

Non-constant fitting error does occur in the diffusion coefficient and velocity magnitude, however, when the true value of the parameter in question is approximately zero, whilst the opposite parameter is high. For simulations with approximately zero diffusion and $v > 1$, we overestimate D with increasing severity as the velocity increases. Meanwhile, when the velocity field is approximately zero and $D > 1$, the velocity fit becomes increasingly unstable as D increases, with a tendency to under-estimate the velocity. When one motion becomes dominant, the ISF becomes insensitive to the slower-moving behaviour. For example, when the diffusion is very high, the ISF in Equation (4.3.3) decays to 0 quickly.



(a) Error in fitted diffusion coefficient, D px^2/frame . (b) Error in fitted velocity magnitude, v px/frame .

Figure 4.8: Fitting error for simulations of advection-diffusion behaviour, expressed as a percentage of the true value of each parameter. Input diffusion coefficients and velocity magnitude range between $0.1 \leq D \leq 2 \text{ px}^2/\text{frame}$ and $0.1 \leq v \leq 2 \text{ px}/\text{frame}$, respectively.

As such, small changes in the velocity will have minimal effect on parameter fitting.

We also note that whilst the diffusion coefficient is over-estimated for systems with high v and low D , the velocity is under-estimated when D is high and v is low. This result initially appears counter-intuitive, since we expect that when the ISF is insensitive to a parameter, we should over-estimate and under-estimate with equal probability. However, this clear trend towards prioritising fitting the diffusion coefficient over the velocity magnitude is a result of the ISF being generally more sensitive to diffusion than advection. Referring to Equation (4.3.3), the ISF has an exponential dependence on the diffusion coefficient, which will dominate the rate of decay of the ISF in comparison with the Bessel function depending on v . As such, changes in the diffusion coefficient have a much stronger effect on the shape of $f(q, \Delta t)$ than changes in the velocity. This can cause the fitting algorithm to converge to an approximately correct solution more quickly by controlling the diffusion coefficient, rather than the velocity. This imbalance in the significance of each parameter in the ISF thus causes a detrimental distortion in the approximation of the parameters themselves.

This analysis demonstrates that current parameter fitting methods in DDM struggle to characterise systems where the displacement from one form of motion is dominant over the total displacement. When this occurs, gradient descent methods to fitting movement parameters will prioritise the dominant signal, and hence the weaker parameter is less accurately fit. Returning to the motivating

example of oocytes, this problem can occur when there is an imbalance between the rates of cytoplasmic streaming and active diffusion. A common way to express the balance between these two terms is the Péclet number, defined by

$$\text{Pe} = \frac{vR}{D}, \quad (4.3.4)$$

where R is some characteristic length scale, such as the radius of the tracer particles used, or the length of a bacteria [147, 154]. In our case, we use $R = \sigma$, since this parameter defines the effective width of the particle. A Péclet number of around 1 corresponds to dynamics where velocity and diffusion are matched. Analysis of the particle size inside the cytoplasm of plant cells indicates $\text{Pe} \approx 50$, suggesting dominance of the streaming behaviour over active diffusion [147, 154]. Equivalent analysis remains to be conducted for eukaryotic cells, and in particular, oocytes.

For our simulations, we can observe the effect of variable Péclet number on the fitting error of both velocity and diffusion. Taking $R = \sigma$, we plot the percentage error from our previous simulations against the Péclet number in Figure 4.9.

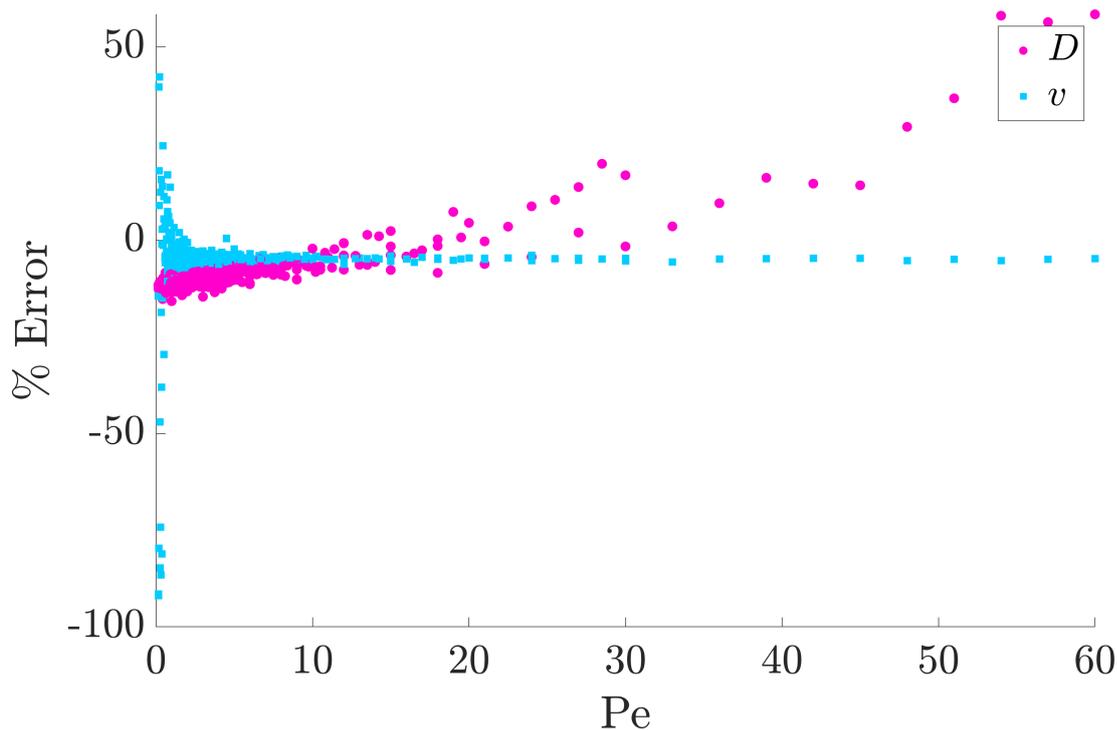


Figure 4.9: Scatter plot of the percentage error for fits of the velocity v and diffusion coefficient D in simulations, as a function of the Péclet number, Pe , defined in Equation (4.3.4). As Pe increases, the fit of the velocity improves, whilst D is fit optimally for $\text{Pe} \approx 10$.

As Pe increases, the accuracy of the velocity fit increases; this matches our expectation that with increasing Pe , the system is dominated by advection, and the

relative effect of the diffusive term decreases. The velocity is still under-estimated, even for the largest Pe , by around 5%. In contrast, the fitting accuracy for the diffusion coefficient is not maximised when diffusion dominates. When $Pe < 10$, the diffusion coefficient is consistently under-estimated, which can be attributed to the localisation error described in Section 3.3. When $Pe \approx 10$, the fit of the diffusion coefficient is optimal. When $Pe > 10$, however, the diffusion coefficient is over-estimated. This is likely related to under-estimation of the velocity observed at high Pe . The movement of each particle can be split into the movements derived from advection and diffusion separately, so when one form of movement is under-estimated, the other form must be over-estimated to generate approximately the same level of displacement. Hence, DDM cannot always separate the two forms of motion accurately, particularly when one dominates the other.

The dependence of the accuracy of DDM analysis on the Péclet number leaves us in a precarious position. The rate of active processes in the oocyte depends on the kinesin activity and cytoskeleton structure [150], which may be indicative of oocyte quality. Therefore, the Péclet number may depend on oocyte quality, such that each oocyte has a different Péclet number. This would imply that the fitting error depends on the qualities we are trying to assess, casting doubts over the usefulness of current fitting methods in DDM. If oocytes are an example of life at high Péclet number, then current fitting methods may be able to account for cytoplasmic streaming, but not active diffusion. Similarly, when active diffusion dominates cytoplasmic movement, we may lose the ability to determine a slow moving cytoplasmic stream.

One direction for future research is, therefore, to more clearly quantify the Péclet number for particles contained in the cytoplasm of oocytes, since this will determine whether we expect DDM to yield accurate, consistent results. This would require experimental evidence, however, so is out of the scope of this thesis. We also note that the simulations performed to generate Figure 4.9 would need to be repeated with real image data, since movement of particles in our simulated images is not dependent on particle size, although Péclet number is. For example, we could double σ to get a wider particle, and obtain the same results from from Figure 4.9 at higher Péclet numbers. This does not invalidate the conclusions drawn from Figure 4.9, since our simulations still allow for exploration of fitting performance as diffusion and advection magnitudes vary. This does, however, mean that the Péclet numbers reported in this work do not correlate exactly with the Péclet numbers of real biological samples.

Another direction for further work is to develop new, more robust fitting methods which are less prone to error in the desired parameter region. In the next section, we introduce a new parameter fitting pipeline, which incorporates

synthetic-data driven machine learning to develop fitting algorithms which can be applied in real data.

4.4 Introducing a simulation-driven machine learning pipeline for parameter fitting in DDM

In Section 4.2, we outlined a new pipeline for parameter fitting in DDM by generating the ISF numerically from particle simulations, enabling parameter fitting where analytic expression of the ISF is difficult or impossible, but individual particle simulation is approachable. This could be paired with fluid dynamic simulations of cytoplasmic streaming and active diffusion, in order to generate the ISF which best describes movement in the cytoplasm of oocytes. However, we also considered that when approaching complex problems such as cytoplasmic movement, which depend on several movement parameters, using splines to approximate the ISF may be ineffective due to the high complexity of the multivariate polynomials required.

Fitting these numerically generated ISFs is made even more difficult in light of the results of Section 4.3, where we demonstrated that the standard parameter fitting algorithms used in DDM suffer from instability and bias that depend on the underlying parameters defining the movement. In particular, when displacement from one type of motion dominates another, the parameter corresponding to the slower movement is fit poorly. This is a particularly challenging feature of current analysis methods in the context of IVF, since we want to use differences in the rate of cytoplasmic movement to predict health. If slower cytoplasmic movement is an indicator of poor quality, but also results in numerical instability, DDM cannot be used as a discriminant of oocyte viability.

Together, these two bodies of work show that whilst we can in theory generate ISFs numerically for analytically challenging behaviours, parameter fitting concerns will still be a limiting factor in the performance and usefulness of DDM in the IVF clinic. This motivates the search for a new approach to parameter fitting. The problem is one of pattern recognition, interpreting the shape of a DDM matrix in terms of the parameters that generated it, which naturally leads us to consider the application of supervised machine learning in the parameter fitting stage.

Machine learning refers to the class of algorithms and approaches in which the algorithm ‘learns’ to identify patterns from raw data [325]. Whilst machine learning is a broad term covering a vast array of algorithms, we are interested in artificial neural networks [325], whose name is derived from their resemblance to models of the brain, having the appearance of many layers of connected ‘neurons’.

An example of a very basic network is presented in Figure 4.10, reproduced from [326]. The network is formed by a series of layers of neurons, where each neuron from one layer feeds into every neuron of the subsequent layer.

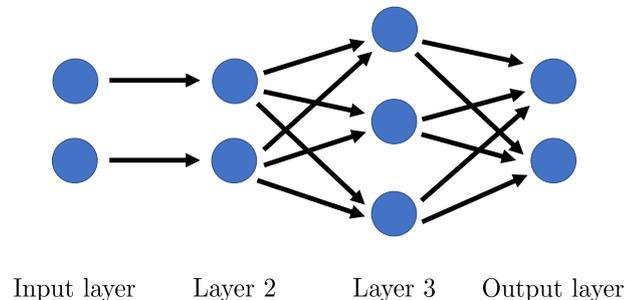


Figure 4.10: A very simple neural network architecture, featuring four layers of neurons. The first layer is determined by the input into the network, and subsequent layers are defined sequentially by summing weighted contributions from previous layers, with an additional bias.

A single layer of the network, containing n nodes, can be expressed as a vector $a \in \mathbb{R}^n$. To define the values of the next layer, an arc connects every neuron in the current layer with every neuron in the next. Each arc is assigned a ‘weight’, such that each layer after the first admits a weight matrix $W \in \mathbb{R}^{n \times m}$, where each of the n rows corresponds to the neuron that we map from, and the m columns relate to the neuron being mapped to. Finally, we may add to each neuron in the current layer some bias, given for each layer as a vector $b \in \mathbb{R}^n$, which allows us to tune the relative importance of a given node independently of the contributions it receives from the nodes in the previous layer.

In the brain, a neuron only fires if the input is large enough to surpass an activation threshold [326]. To mimic this effect in the neural network, we define some activation function ς , which controls how much of the input signal from one layer is transferred to the next. Hence, the l^{th} layer $a^{[l]}$ is defined from the previous layer by the equation

$$a^{[l]} = \varsigma(W^{[l]}a^{[l-1]} + b^{[l]}).$$

The output layer is typically defined by the purpose of the network. In classification tasks, each class has a node, and the value of each node represents the predicted probability of the input being assigned to that class. In regression problems, a single node can be used as the output.

Training the network is equivalent to changing the weights and biases in such a way that the network improves its ability to perform the desired task [326].

This requires training data from which we learn, and some way of measuring performance. In supervised learning, the training data is labelled with some ‘true value’ corresponding to the output we want the network to reproduce. Performance is assessed by defining a cost function that measures the distance between the network output and the ground truth label.

For example, the MNIST dataset consists of images of handwritten digits, each with a corresponding label denoting the digit depicted in the image. A neural network can be trained to identify which digit is depicted in each image. The output layer such a network consists of ten neurons, representing the ten digits, and the value of each neuron corresponds to the probability that the input image contains that digit. The label for each image is a vector with 1 in the correct position, and 0 elsewhere, so a suitable cost function could be the Euclidean distance between the true value and the output data. The cost function is chosen to be continuous, in order for the weights and biases to be tuned by taking partial derivatives of the cost function with respect each weight and bias in the network. Training can be equivalently considered as determining the direction to perturb each weight and bias in order to decrease the cost function.

Whilst we have given only a brief overview of artificial learning here, there is a great wealth of knowledge available around artificial intelligence, machine learning and deep learning [325–327], spanning many varieties of algorithm. However, only one application of machine learning exists in DDM analysis [310]. This work embeds bacteria into a silk hydrogel construct, giving the material new biological properties, such as pH sensitivity. The design of these hydrogels depends on a number of physical parameters, which result in different properties of the final material. The particular desired property for this study was ‘gel time’, the timescale over which the material develops solid-like mechanical properties, at which point the MSD of particles contained in the gel becomes constant. The objective of the study was to train a neural network to identify which combinations of physical parameters would produce hydrogels that had gel time within a given interval. DDM was used to determine the gel time for the training set of hydrogels, by finding the time at which the DDM matrix plateaus. Next, supervised machine learning was used to predict whether a combination of physical parameters relating to the construction of the gel would give a gel time within the desired time interval, using the actual gel time of real experiments as a labelled training set.

This approach does not integrate machine learning into DDM directly, but instead uses parameters determined by DDM as the label for supervised learning. We propose to use machine learning as an alternative method of parameter fitting, by using the DDM matrix itself as the input, and setting the network output to

be the desired movement parameters, such as the diffusion coefficient. To train a neural network to perform parameter fitting, we need a wealth of labelled training data; generating real data to suit this purpose is time-consuming and expensive, and in many cases may simply be practically impossible. Hence, we propose using synthetic data as the training set, since it is cheap and easy to produce. Despite being trained on synthetic data, we demonstrate that the trained network is capable performing parameter fitting in DDM matrices derived from real, unseen time lapse images.

Before describing the approach taken to using machine learning as a parameter fitting tool, we wish to discuss the value of using the DDM matrix as the input of the network, as opposed to the original images. There are some clear disadvantages to doing so; for example, generating the DDM matrix is computationally expensive, which can make development of the training set a slow process. Additionally, since the DDM matrix is subject to both a time and radial average, the DDM matrix loses all information which is heterogeneous in time and space. As such, by using the DDM matrix instead of the original image, we limit the scope of what can be measured from the original image stack. Despite these disadvantages to using the DDM matrix as the network input, there are a number of significant benefits to doing so which outweigh these considerations.

First, using DDM matrices as an input means that we do not need to produce simulations that look like the real data. A network trained to extract parameters from a set of simulated images will not perform well in real data, unless the simulations match the visual appearance of the real data; attempting to produce images which look like real data is a considerable challenge. The DDM matrix is far simpler, and can also be described using only three terms, as given in Equation (2.2.19); the scale term, $A(q)$, the noise term, $B(q)$ and the ISF, $f(q, \Delta t)$. If we simulate data that has the same movement as our real data, then $f(q, \Delta t)$ is equivalent in both the real and simulated DDM matrices, and we only have to consider reconciling the different values of $A(q)$ and $B(q)$ expected between the simulated and real images.

Additionally, a neural network trained on raw image data can potentially learn from features of the images that we do not wish to consider, such as morphokinetic parameters. Whilst these features may actually be indicative of the viability of the oocyte, their consideration may make the algorithm less interpretable, putting distance between the output of the network and our understanding of the biology that caused it. The objective of this thesis was to create tools and metrics that are scientifically motivated, and explainable to patients. Machine learning, however, can often appear to be a ‘black box’, where we do not understand how the output

was determined. Since we know, from Equation (2.2.19) that the DDM matrix can be expressed in the form

$$\mathcal{D}_R(q, \Delta t) = A(q) (1 - f(q, \Delta t)) + B(q),$$

by using the DDM matrix as an input, we are able to discard confounding variables by limiting the available information from which we might learn. This leads us onto an additional benefit of using the DDM matrix; we achieve a significant dimensional reduction when comparing \mathcal{D}_R to the original image stack. Reducing the size of the input layer of the network allows us to potentially reduce the number and size of layers in the network. This means there are fewer parameters to train, which can expedite the training process considerably. Furthermore, more complex networks are at greater risk of over-fitting, since they have more degrees of freedom. Hence, there are considerable advantages to using the DDM matrix as the network input rather than the raw image data.

There is an important difference between the machine learning parameter fitting we propose, and pre-existing fitting methods available for DDM analysis, which we must highlight. In classical methods, even when using a numerically approximated ISF as in Section 4, we could run the same parameter fitting algorithm on any different DDM matrix, regardless of the type of images from which the DDM matrix was generated. In our machine learning approach, this is more difficult. Supervised learning, particularly in biological applications, is known to perform badly when the dataset we apply the algorithm to has different properties than the training data [328]. Hence, in order for the network to learn to decouple the ISF from the scale and noise, some portion of the training data supplied must be similar in appearance to the real image data we wish to process. We therefore require some method to make the DDM matrices in our training set more similar to those generated by the target data.

The first, and most obvious, option available to ensure that the neural network performs well outside of the training data is to train on a sample of real data, although this is an unrealistic option in the IVF clinic, due to the lack of available labelled image data. Another alternative is to train the network in one dataset, and then retrain some, or all, of the network with a small portion of real data. For example, the network could be trained in a mouse oocyte dataset, and then retrained in a small human oocyte sample. This lowers the data requirement, but would still require more human oocyte data than we have available. In the absence of any real data, we only have access to synthetic image data; we therefore propose two alternative, although not necessarily mutually exclusive, strategies to training a network which can tackle real image data using exclusively simulated training images:

1. Transform each simulated DDM matrix under a random perturbation of the scaling and noise terms, so that the neural network can learn to decouple these q -dependent terms from the ISF we wish to classify.
2. Use some reasonable approximation of $A(q)$ and $B(q)$ to transform the DDM matrix into the ISF, so that the input layer takes in $f(q, \Delta t)$ instead of $\mathcal{D}_R(q, \Delta t)$.

The first approach is theoretically more ideal, since it would not require any knowledge of the shapes of $A(q)$ and $B(q)$, but in reality this approach is highly impractical. Both terms are vectors of length $L - 1$ for a square image L pixels in side length, since there is one frequency per pixel distance in the image. This means we have $2(L - 1)$ parameters to randomise for an image of length L . We do not have an upper bound for $A(q)$, or $B(q)$, without knowing what our real data might look like, so each of the new parameters that randomise the scale and noise terms belong to a large interval. As such, we would need a significantly increased training set size to implement this approach.

We therefore choose the second strategy. We can approximate $A(q)$ and $B(q)$ for our simulated datasets either from knowledge of the PSF used to visualise images [292], or by considering high-frequency and long time lag elements of the DDM matrix [164]. For example, when $\Delta t \rightarrow \infty$, the DDM matrix reduces to

$$\lim_{\Delta t \rightarrow \infty} \mathcal{D}_R(q, \Delta t) = A(q) + B(q). \quad (4.4.1)$$

For very small time lags, as $\Delta t \rightarrow 0$,

$$\lim_{\Delta t \rightarrow 0} \mathcal{D}_R(q, \Delta t) = B(q). \quad (4.4.2)$$

Together, these estimates can be used to decouple $A(q)$ and $B(q)$. It may be that the maximum imaging frequency of the camera limits our ability to apply Equation (4.4.2) estimate $B(q)$, in which case, the noise can be approximated by imaging an empty frame. Similarly, if we are limited on the maximum time lag we can image over, or if it is difficult to estimate how large Δt must be for Equation (4.4.1) to hold, the scaling can be roughly approximated by taking the average

$$A(q) = 2\langle |\tilde{I}(q, t)|^2 \rangle_t.$$

The downside to approximating the scale and noise by analysing single frames in this way is that the estimates would be affected by static artefacts in the image, such as scratches or defects in the microscope lens.

Having discussed some important considerations when applying machine learning to parameter fitting in DDM, we now describe the pipeline from which parameter fitting can be achieved in real image data, using a network trained on synthetic data.

1. Randomly generate a large set of movement parameters (diffusion coefficients and velocity magnitudes, for example). Then, generate corresponding simulated particle trajectories by drawing displacements from the van Hove function of the process being observed. For example, Brownian motion displacements are drawn from a normal distribution, as in Equation (3.1.1).
2. Convert trajectories into images, $I(\mathbf{r}, \Delta t)$, using Equation (3.1.2).
3. Apply the DDM algorithm to each image to generate a set of DDM matrices $\mathcal{D}_R(q, \Delta t)$.
4. Transform each DDM matrix to the ISF using the following rearrangement of Equation (2.2.19):

$$f(q, \Delta t) = \frac{\mathcal{D}_R(q, \Delta t) - B(q)}{A(q)},$$

where for our simulated images, which are noiseless by definition,

$$\begin{aligned} B(q) &\approx 0, \\ A(q) &\approx 2\langle |\tilde{I}(q, t)|^2 \rangle_t. \end{aligned}$$

5. Perform training of the neural network, using the input movement parameters for each ISF as a training label.
6. Once trained, the network is ready to be applied to real images. The real image stack should be converted into a DDM matrix, and the same maps defined in step 4 are applied to reduce the DDM matrix to an ISF, where approximations of $A(q)$ and $B(q)$ are available in Equations (4.4.1) and (4.4.2) respectively.

To demonstrate that this approach can be used to tackle real data challenges, we will start by approaching only Brownian motion. We consider images depicting colloidal dispersions, which have been produced and imaged in the Cardiff University School of Physics and Astronomy by Emily Lewis. A colloidal suspension of Polystyrene (PS) beads is used with a nominal radius of 100 nm, with less than 3% coefficient of variance (*cv*) (Alpha Nanotech Colloidal PS Beads NP-PA07CPSX78). These PS beads were dispersed in mixtures containing water

and either 10, 20 or 30% glycerol. The diffusion coefficient of the suspended beads is given by the Stokes-Einstein relation, [164, 266],

$$D = \frac{K_B u}{6\pi\eta r},$$

where K_B J/K is the Boltzmann constant, $u = 291.25$ K is absolute temperature, η mPa s is the dynamic viscosity of a the glycerol solution depending on its percentage concentration, and r the particle radius. Varying the glycerol content of the mixture varies the dynamic viscosity and, therefore, the diffusion coefficient. We can measure the success of our new pipeline by comparing the output diffusion coefficients from the neural network with the ‘true value’ estimates yielded by the Stokes-Einstein relation. The three dispersions imaged, with varying glycerol concentrations of 10, 20 and 30%, have corresponding dynamic viscosities of 1.46, 2.10 and 3.20 mPa s, respectively, resulting in diffusion coefficients of 1.80, 1.23 and 0.83 $\mu\text{m}^2/\text{s}$.

We construct a simple neural network using Tensorflow (Python). The specific choices regarding the number and size of layers, and choice of activation and optimisation methods, formed the basis of a Cardiff Undergraduate Research Opportunities Programme (CUROP), undertaken by Thomas Greatrix, and supervised by Timothy Ostler, Joshua Moore, Thomas Woolley and Katerina Kauri. The input layer takes the DDM matrix, flattened into a vector, whilst there are 5 inner layers, with a monotonically non-increasing number of nodes per layer given by $\{512, 256, 256, 128, 64\}$. The inner layers all use a Rectified Linear Unit (ReLU) activation function, that takes only positive contributions from the previous layer, defined by [329]

$$\varsigma(x) = \max(0, x).$$

This activation function has several benefits, such as allowing sparse representation of the network where specific activations achieve true zero, as well as linear behaviour when activated [329]. The final output layer has only a single node, with the value corresponding to the diffusion coefficient. This layer takes on a linear activation,

$$\varsigma(x) = x.$$

For training, we adopt the Adam optimiser [330], which has been shown to perform well in regression problems and is computationally efficient.

We generate 1500 DDM matrices from simulations of Brownian motion, with diffusion coefficients sampled from a uniform random distribution between 0 and 3. The simulations must match the meta-parameters of the image datasets, so we simulate movies 3310 frames long, at a frequency of 7680/900 frames per second,

for a pixel size of 322.5 nm/pixel. Of this dataset, 80% is assigned to training (1200) and 20% for validation (300). After training, we apply the same network to the DDM matrix of the real images.

To estimate the noise and scaling terms of the real images, we assume that the largest sampled time lag in the DDM matrix is sufficiently large to satisfy Equation (4.4.1). Although this assumption may not be necessarily good in our case, it is easy to make this assumption stronger when designing an experiment, by imaging for a longer period of time. We also note that

$$\lim_{q \rightarrow \infty} A(q) = 0,$$

since larger q relate to smaller distance, and we expect the decorrelation of the particle position to be high at small distances. The noise should be constant in q , so an alternative estimate of $B(q)$ instead of Equation (4.4.2) is to consider the value of the DDM matrix at its greatest frequency and time lag. As such, we estimate $A(q)$ and $B(q)$ by the following,

$$\begin{aligned} B(q) &= \mathcal{D}_R(\max(Q), \max(T)), \\ A(q) &= \mathcal{D}_R(q, \max(T)) - B(q), \end{aligned}$$

where Q and T are the sets of frequencies and time lags, respectively. Using these estimates, we can transform the DDM matrices for the real datasets into ISFs, and apply the neural network to extract a diffusion coefficient.

First, we compare the performance of the neural network against the classical fitting method, outlined in Section 2.2.7, in Figure 4.11. In the centre of the interval of D considered, around $0.5 < D < 2.5 \mu\text{m}^2/\text{s}$, the fitting performance is comparable between the two methods. The machine learning method begins to break down close to the edges of the domain, however; for $D < 0.5 \mu\text{m}^2/\text{s}$, over-estimation occurs, whilst for $D > 2.5 \mu\text{m}^2/\text{s}$, we under-estimate the diffusion coefficient. Poor performance at the end of the intervals is to be expected, since the network cannot extrapolate outside of the training data, and hence maps input data at the extreme ends of the training interval inwards. This results in a sigmoidal relationship between the true and fitted parameter values, which can be mitigated by increasing the sampling rate of the training data at the extreme ends of the interval.

Having verified the performance of the neural network in the validation dataset, we now show that this network, trained exclusively on synthetic data, can be applied to real data. For each of our real colloid datasets, with varying kinematic viscosity due to the variable concentration of glycerol, we can predict the diffusion

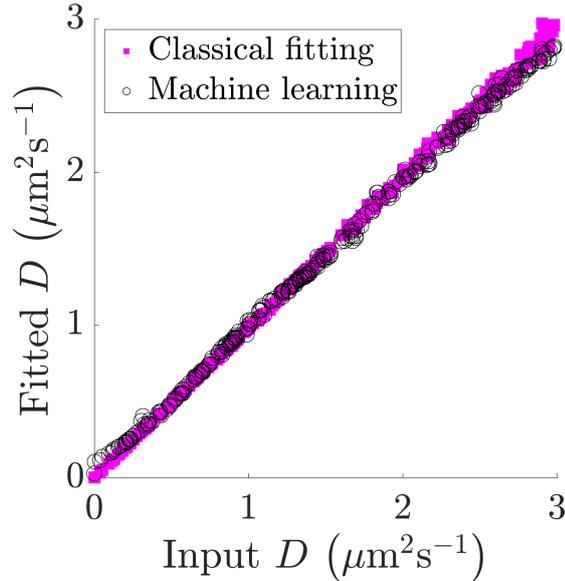


Figure 4.11: Plot of the fitted diffusion coefficients, D , in the test set of simulated Brownian motions, fitted by either both classical fitting approaches outlined in Section 2.2.7 (purple filled square) or machine learning (open circle).

coefficient by the Stokes-Einstein relation. We plot the Stokes-Einstein relationship in Figure 4.12 as a function of ν , where the other parameters remain constant. We also plot the diffusion coefficients fitted using classical approaches, and our machine learning approach. Our new method of parameter fitting achieves the correct decreasing relationship between ν and D , and is relatively close to approximating the Stokes-Einstein relationship. However, particularly for the colloidal dispersion at a 20% concentration with $\nu = 2.10$ mPa s, the machine learning approach is less accurate than the classical fitting methodology.

Although the neural network does not outperform classical fitting methods, the results in Figure 4.12 are positive, showing a proof of concept that a network, purely trained on synthetic data, can be applied to real data and achieve relatively low error rates. The actual performance of this network could easily be improved by increasing the size of the training set, using a small amount of real data to ‘retrain’ the network, or pruning the frequencies used in the input of DDM to reduce the amount of ‘bad’ data considered. Additionally, we could optimise the architecture used in this network, either by improving on the number of, and size of, layers used, or by using an altogether different architecture. Despite the relatively simple approach we have taken, using only a few layers with relatively few nodes and a minimal training set, we have still been able to approach a real dataset using only synthetic data to inform our parameter selection, and obtain respectable results, implying this fitting framework has potential.

In conclusion, we have created a new framework to fit parameters from DDM matrices using supervised machine learning, where training is performed on syn-

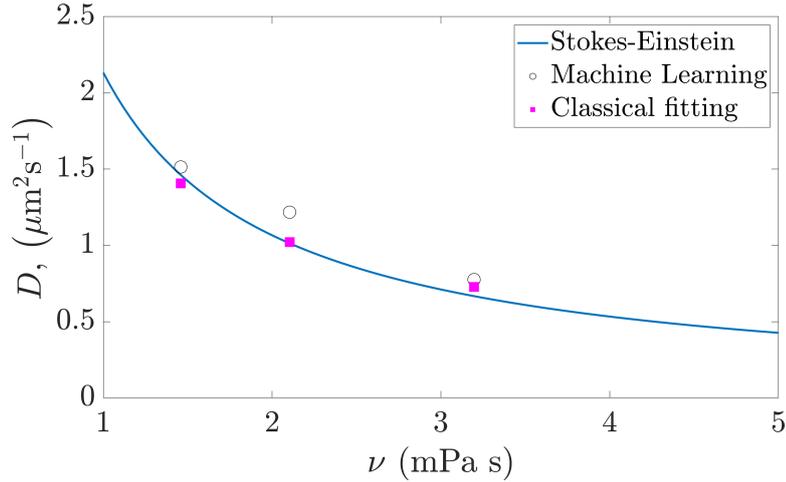


Figure 4.12: Plot of the diffusion coefficients fitted to the three colloidal image datasets with varying kinematic viscosity ($\nu \in \{1.46, 2.10, 3.20\}$ mPa s), using the classical fitting approach (purple filled square) and machine learning (open circle). The Stokes-Einstein relation, showing the expected diffusion coefficient for each dataset, is given by the blue line.

thetic data. The power of this approach lies in its generality, since it can be implemented in any dataset in which the behaviour of an individual tracer particle can be simulated. For some behaviours, such as cytoplasmic movement in oocytes, the ISF might be inaccessible or prone to numerical instability, but this new methodology has the potential to tackle a wide variety of problems. A more general pipeline can be defined, allowing us to employ this machine learning approach tailored to a specific problem.

The first step towards this goal is to design experiments in which tracer particles, injected into the cytoplasm, can be used to derive movement statistics using SPT or PIV approaches. This data can then be used to motivate the design of fluid mechanical models, such as those discussed in Section 1.2.2.1, which will generate the synthetic data required to train the neural network. We visualise this workflow in Figure 4.13. This pipeline is flexible, and allows for effective analysis of complex systems using only a very small amount of real image data.

However, a more desirable pipeline is one in which no pre-requisite knowledge of the system we wish to study is needed at all. In this theoretical pipeline, we could train the network on a vast synthetic dataset, containing combinations of a large pool of different behaviours. In this way, the network can learn not only to fit parameter values for a set combination of movement types, but also to identify the type of motion observed within a supplied image stack. Such a neural network would be expected to output non-zero values corresponding to the parameters which govern the motion in that data, and 0 for absent effects, thus allowing DDM to be used as a diagnostic tool, as well as for parameter

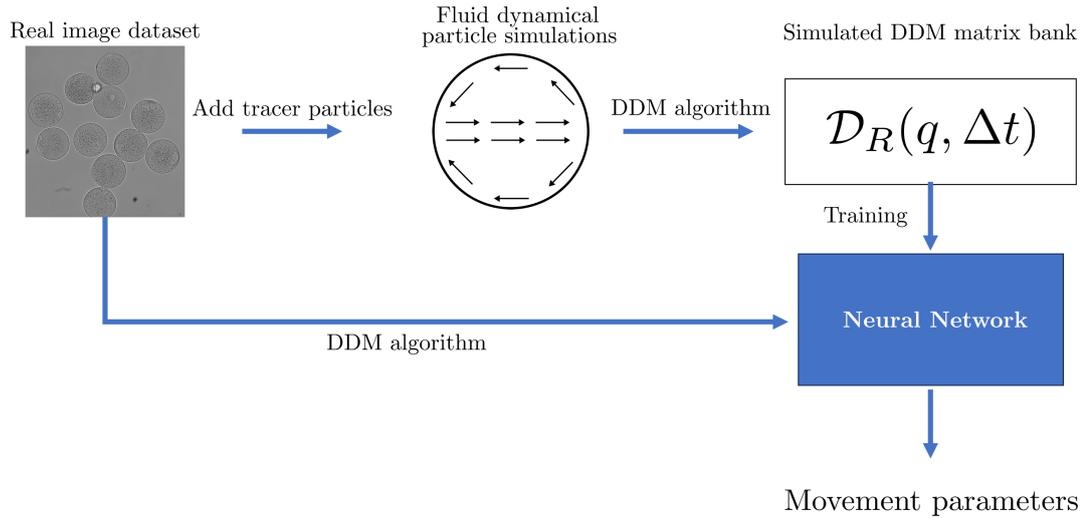


Figure 4.13: A proposed future workflow for designing new DDM fitting methods for use in the IVF clinic. A small amount of real data is studied using invasive imaging techniques with injected tracer particles, to motivate fluid mechanical models from which synthetic particle trajectories may be generated. The synthetic data may be used to train a neural network, which may then be applied to the real data for non-invasive extraction movement statistics.

fitting. If such a network could be trained, the positive implications would be significant, since any biologist wishing to analyse a dataset with DDM could use the network to immediately summarise their data. However, this is not an immediately achievable goal; there are several limitations, even of our much more simple approach, which must be overcome first, taking us step by step towards a more general and powerful clinical assessment.

Many limitations of this machine learning approach are shared with the new ISF generation method we proposed in Section 4. Training the network requires a large amount of data and computational power; training the network may not be viable in resource-constrained environments such as an IVF lab. Choosing the correct size for the training dataset is itself a challenge, as we must use enough data to effectively train the model, without over-fitting the training data. The amount of training data needed also depends on the parameter range over which we simulate. As shown in Figure 4.11, the network performs poorly around the extreme values of the range of simulated input parameters, so it may be necessary to simulate movement parameters over a wider interval than we expect to see in reality. However, the wider the range of the simulated parameters, the larger the training set we need to guarantee sufficient performance of the model.

An additional concern with parameter fitting through machine learning is that

when fitting fails, it is much harder to diagnose than in the existing curve-fitting methods. In Section 2.3, we were able to see when the quality of curve fitting was low, which gave an indication that the results from DDM analysis may be inaccurate. However, when using machine learning, it is far more difficult to identify when or why the model performs poorly. This is a concern in the clinic, where it is crucial that we are confident in the fitted output of the model. It can be difficult to justify the application of a neural network to patients, as well as industry regulators, when we do not understand how the output parameter is achieved. This limitation can prevent even highly effective neural networks from seeing practical application.

Finally, in its current state, the chosen architecture for the neural network demands that the input takes the same shape as real data, which makes it inflexible when applied to different problems. Different clinics using different incubators, with variable image stack sizes and imaging frequencies, would require a bespoke trained network, limiting the ability to create a single ‘master’ network which can be exported to all clinics. This limits the generality of the model in its current state, since for every new size of problem, we would need to repeat the training strategy.

4.5 Summary and Discussion

Oocyte health is critical in determining the viability of embryos in IVF [119–122]. Methods of assessing oocyte health are in high demand, but are limited by the requirement that they must be non-invasive [119]. Attempts to infer health from embryo and oocyte morphology have had mixed results [128–132], but with the emergence of commercially available time-lapse incubators [54, 57], attempts to incorporate timings of morphological changes during embryo development, given the term ‘morphokinetics’, have been a prevalent direction of research in this field [57, 59–62]. Despite some promising results, these new morphokinetic considerations have failed to yield significant clinical improvement in success rates, attributed to the lack of biological explanation for how morphokinetic features relate to embryo quality [55, 56, 60, 203–205]. Furthermore, these methods are embryo-centric, resulting in a lack of more advanced time-lapse approaches to classifying oocyte quality.

Motivated by [135], we propose that the rate of movement in the cytoplasm could have the potential to be a non-invasive indicator of oocyte quality. In order for this metric to be useable in the clinic, we need tools capable of accurately extracting movement statistics from image data of oocytes. In particular, we set out to explore the use of Differential Dynamic Microscopy (DDM) [266], since

this method is has been shown to perform well at extracting movement statistics from the mixture of active diffusion [140, 141] and cytoplasmic streaming [146, 147] expected in the cytoplasm.

In Chapter 2.3, we identified some of the challenges faced when using DDM for oocyte quality assessment. We demonstrated in Section 2.3.1 that DDM can capture Brownian motion in poisoned oocytes, but neither of the available advection-diffusion or Brownian models were a good fit for healthy oocytes, so the first challenge we identified is to find a suitable model for the ISF that describes movement of the cytoplasm in oocytes.

The second challenge we identified in Section 2.3.4 is that, without additional validation steps to evaluate the accuracy of DDM, we cannot be confident that the output parameters from DDM reflect variable health between oocytes. By making the imperfect assumption that active diffusion dominates the movement of the cytoplasm, we showed in Figure 2.16 that poisoning is reflected by a drop in the diffusion coefficient of each oocyte, but high variation in D between oocytes meant that we could not discriminate between healthy and poisoned oocytes from diffusion coefficient alone. Identifying whether intra-egg variance indicates differences between the rate of movement of the cytoplasm in each egg, or is attributed to numerical error, is critical to ensuring the validity of the results from DDM analysis.

Finally, we identified in Section 2.3.5 that in order to respect the assumption in DDM that boundary effects are negligible, we had to remove approximately 50% of the oocyte from the analysed image by taking a subset of the original image. Spatial heterogeneity of the cytoplasm means that the information determined from movement at the cell wall may be important in determining oocyte quality, so we need to adapt the fitting functions used in DDM to enable the potential description of boundary effects.

Using the synthetic data described in Chapter 3, we tackled the challenges outlined by our exploration of the mouse oocyte dataset. First, in Section 4.2, we outlined a new methodology for numerically approximating the ISF in behaviours where the trajectory of individual particles can be described more easily than the ensemble. This new strategy for generating the ISF serves two purposes. First, it can be used to generate the ISF for complex, spatially heterogeneous behaviour as seen in the cytoplasm of oocytes. Second, this new approach is capable of dealing with boundary effects, so that we may image the entire domain and are no longer required to discard significant portions of the available images in order to apply DDM analysis.

This new method for generation of the ISF uses the same curve fitting approaches as before, however. In Section 4.3, we showed that there exists particular

parameter regions for which these curve fitting methods are prone to error. In particular, we were able to show that the accuracy when fitting Brownian motion with a constant, directed advection was dependent on the non-dimensional Péclet number, Pe , which balances the relative rates of advection and diffusion. For $Pe > 50$, a velocity-dominated regime, the velocity field is well fit but the diffusion is poorly characterised, whilst for $Pe < 1$, a diffusion-dominated regime, the velocity is poorly characterised. This fitting insensitivity means that if the Péclet number is a function of oocyte quality, then so to will be the error of DDM analysis, resulting in diminished ability to assess oocytes using DDM.

To mitigate this concern, we introduced a new fitting approach in Section 4.4 that uses synthetic data to generate DDM matrices, which are used to train a machine learning model to perform parameter fitting. We demonstrated that this model can then be applied to real image datasets, despite being trained purely with synthetic data. Compared to previous curve-fitting, which required us to perform fitting iteratively over slices of the DDM matrix, machine learning can use information at all times and frequencies to make decisions, resulting in a more powerful fitting approach. We demonstrated the success of this method when applied to colloidal images, achieving within an order of magnitude of the correct diffusion coefficients.

Collectively, the results from this work set out a basic framework for how we might explore the use of DDM in real datasets, detailed in Figure 4.13. Using synthetic images depicting simulated trajectories, we can either construct a numerical approximation of the ISF as in Section 4.2, and use curve-fitting techniques to extract movement parameters, or instead train a neural network as in Section 4.4 which removes the requirement for a model altogether. Our long-term vision for DDM analysis is driven by machine learning, but there are many hurdles to clear before such a technique can be clinically viable.

Even when only describing Brownian motion, the accuracy of the parameter fitting through a neural network was less than the accuracy achieved through curve-fitting approaches. We expect this can be improved by increasing the size of the training set, or by the use of a different architecture, both of which would require further research to optimise. When moving away from Brownian motion, we would then need to train a new network, repeating the optimisation of the training set and architecture for each new application. This is cumbersome and, in some ways, less flexible than the current DDM fitting approach, in which we need only change the model of the ISF, assuming it is known.

Additionally, our current work makes the strong assumption that parameter fitting is even possible from the DDM matrix presented. Behaviour which evolves too quickly, or slowly, within a given image stack cannot be assessed accurately

using DDM, as well as behaviours which are non-stationary and change over time. There are likely also limits on the complexity of the behaviour we can consider, where the parameter space can become large and the ISF becomes dominated by non-linear interactions. The combination of these factors means that when applying DDM analysis to an image stack, the output statistics may have no correlation with the underlying behaviour of the system, but as shown in Section 2.3, diagnosing these issues can be difficult. An important avenue for future work is to develop guidelines which can inform prospective users of the limitations of DDM, identifying features of their data which could cause problems in DDM.

Assessing the power of DDM in the IVF clinic is difficult without a sufficiently large dataset on which to experiment, however. In the absence of real image data, we propose a realistic direction for future work would be to benchmark DDM against other similar image analysis techniques when applied to synthetic image data. Such analysis would provide insight into the suitability of each technique in different problems, such that when real image data becomes available, we may suggest which technique to apply. These other techniques, namely SPT and PIV, access movement directly from the image stack, and rely on post-processing to reconstruct movement distributions. SPT is the ideal approach to use when individual trajectories can be traced, since it allows for the reconstruction of any parameters derived from DDM or PIV. SPT is also the most restrictive of the three techniques, however, since we require a low particle density and individually discernable trajectories. Hence, we only really need to compare PIV and DDM, which occupy a similar role in image analysis.

DDM is essentially a more general form of PIV. In PIV, we can access the velocity field evaluated over a grid on the domain, allowing us to characterise flow which is spatially heterogeneous. Since PIV does not use any time averaging, we can also resolve a velocity field which changes over time, such as characterising the spasms of a fertilised oocyte [135]. However, PIV requires localised bulk flow, which means objects in the same location need to move in the same direction with the same speed. DDM is the opposite in all three of these descriptions; information is accessed over the whole image, the van Hove function must be stationary, particles need not be subject to localised bulk flows. Hence, although DDM and PIV access similar information, they do so in different ways.

PIV and DDM are therefore complementary techniques, covering each others weaknesses. Whilst PIV is well adapted to characterise cytoplasmic streaming, and fails when in diffusion-dominated environments (where Péclet number is low) [142], we have shown DDM to characterise diffusion well in low Péclet number environments. As such, it is likely that some combination of the two approaches will

allow us to extract the most accurate movement parameter from our system, validating the output from one method using the other. Determining measurements which identify the technique most suitable to analyse a given dataset would be a useful contribution to image analysis in biological datasets. Alternatively, the combination of the two techniques could provide a robust and effective method for classifying both ballistic and diffusive properties of image data. In this potential multi-scale approach, applying PIV over longer time periods where ballistic motion is dominant could provide accurate estimates of the local velocity field within the system. Over the same interrogation windows, DDM may then be applied, where the pre-determined local velocity field is supplied to the fitting model in order to increase the accuracy of the estimation of the diffusion coefficient. Pairing these techniques presents an exciting opportunity to expand on the strengths of each method and approach more complex systems which combine sources of directed and undirected motion.

Chapter 5

Differential Dynamic Microscopy can be applied to Differential Interference Contrast images despite shadowing effects

5.1 Introduction

In Chapter 4, we discussed several complex interactions between experimental design factors and the accuracy of DDM, but we did not discuss the choice of microscopy technique used to generate the time lapse image stack. DDM has been successfully applied in a wide variety of different microscopy settings, including bright field [266], fluorescence [331], confocal fluorescence [332], dark-field [292], and phase-contrast [163, 281], a particular form of which is Differential Interference Contrast (DIC) [142].

Phase-contrast imaging is of interest in this work, because LWC uses the Embryoscope[®] bench-top incubator, which generates time-lapse images using a form of phase-contrast imaging known as Hoffman Modulation Contrast (HMC) [333]. Images generated by HMC are considered functionally equivalent to those generated by DIC [334], despite the mechanism for image generation being different between the two techniques. Due to this equivalence in the resulting image, we will limit ourselves to discussing DIC for brevity, noting that our analysis will likewise apply in HMC setting too. A thorough explanation for the theory of DIC imaging, provided by Wolfgang Langbein of the School of Physics and Astronomy, Cardiff University, can be found in the related work [335], in which we presented some of the results of this chapter. The explanation of the theory of DDM is also reproduced in Appendix C.1 for convenience. This thesis does not aim to advance the theory on this topic, so we will summarise some of the basic concepts behind the formation of a DIC image here.

A light wave is a periodic signal, where the phase is defined as the relative position of the wave in its cycle at any given time [268], equivalent to the angle of the oscillation as a fraction of 2π . When light is passed through an object, the emergent wavefront experiences a phase shift ϕ , defined by

$$\phi = 2\pi \frac{(n_2 - n_1)t}{\mathcal{T}},$$

where \mathcal{T} is the period of the wave, t is the time over which the wavefront passes through the sample and n_2 , n_1 are the refractive indexes of the object and background, respectively [268]. In DIC, a single beam is split into closely spaced polarised beams before being passed through the sample, separated in space by the so-called shear vector, \mathbf{s} , which is typically similar to the optical resolution in the focal plane [268]. Each of the separated beams passes through a different part of the imaged object and, therefore, experiences a different amount of phase shift. The difference in phase shift is denoted $\delta(\mathbf{r})$, and is converted into image intensity via the following formula [336],

$$I_t(\mathbf{r}, \psi) = \frac{I_e}{2} [1 - \cos(\psi - \delta(\mathbf{r}))], \quad (5.1.1)$$

where I_e is the excitation intensity, \mathbf{r} is the position in the imaging plane and $\psi \in [0, \pi/2]$ is the phase offset, a parameter derived from the experimental setup that controls the contrast. When $\psi = 0$, a dark-field type image is observed, whereas for $\psi = \pi/2$, a bright-field type image is observed [336].

A simple way to understand DIC is that it visualises gradients in the Optical Path Length (OPL), defined by [337]

$$\text{OPL} = d \frac{c}{v},$$

where d is the distance travelled by the beam through the sample, c is the speed of light in a vacuum and v is the speed of light through a particular homogeneous medium. The brightness in the image at each location is approximately determined by the local difference in OPL, or $\nabla(\text{OPL})$, between the beams passing through that element of the sample. This approach allows imaging with very high resolution and contrast [338], even at very large magnification [336]. However, the directional nature of $\nabla(\text{OPL})$ means that particles often appear dark on one side, and light on the other, resulting in DIC images appearing ‘shadow-cast’ [268], as if a directed light source is applied. In Figure 5.1, we sketch the rough idea behind the formation of the shadow in an idealised semi-spherical object.

We also illustrate the shadowing effect observed in DIC in real data in Figure 5.2, which depicts a mouse oocyte obtained from experiments in the Swann lab in the School of Biosciences at Cardiff University. Imaging is performed using

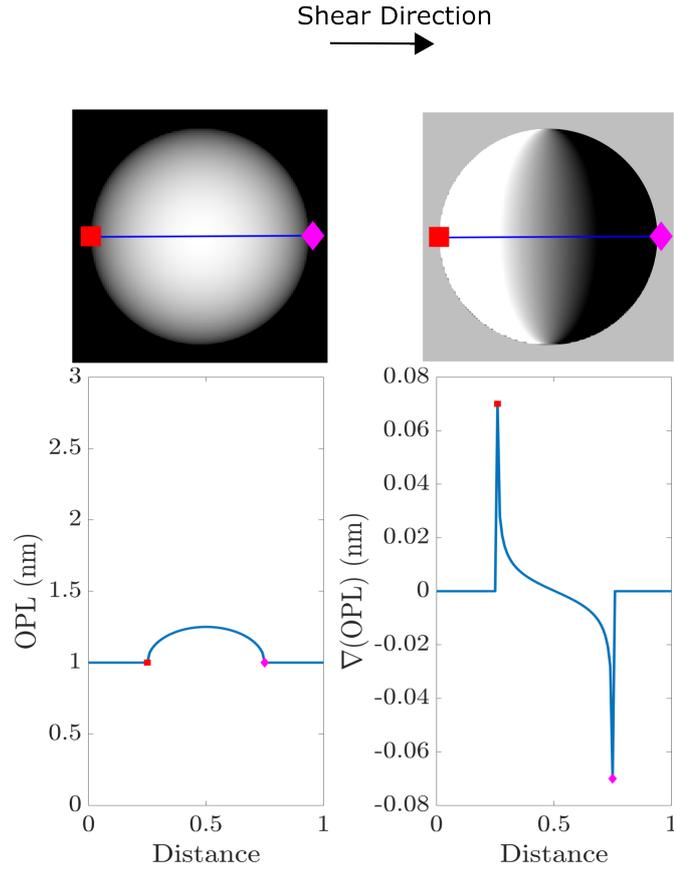


Figure 5.1: An example of the relationship between OPL and its spatial gradient $\nabla(\text{OPL})$ in a homogeneous semi-spherical particle. The OPL and $\nabla(\text{OPL})$ are plotted along a single line oriented parallel to the DIC shear direction, with the square and diamond on the plots corresponding to the location on the image of the particle. Where $\nabla(\text{OPL})$ is positive on the left of the particle, the particle is light, and gets darker as it approaches the right hand edge. Based on a similar diagram from [268].

a custom-built DIC system in the School of Physics and Astronomy at Cardiff University, further details for which can be found in [336], and experimental information regarding mice is given in Appendix A.

In the mouse oocyte in Figure 5.2, the cytoplasm of the oocyte appears to have a rough texture, interspersed with a few larger objects. This roughness corresponds to changes in the phase gradient across the cytoplasm, meaning we observe movement not only of the larger particles in the frame, but of the cytoplasm itself. The level of detail these images can provide make them ideal for studying the relationship between DDM-determined parameters and oocyte viability, since we can track even the smallest movements within cytoplasm to a highly accurate degree. However, shadowing as a result of DIC gives the illusion of heterogeneous illumination along the direction of the shear within the images. This is a source of anisotropy in the image, which violates a key assumption of DDM [164, 266].

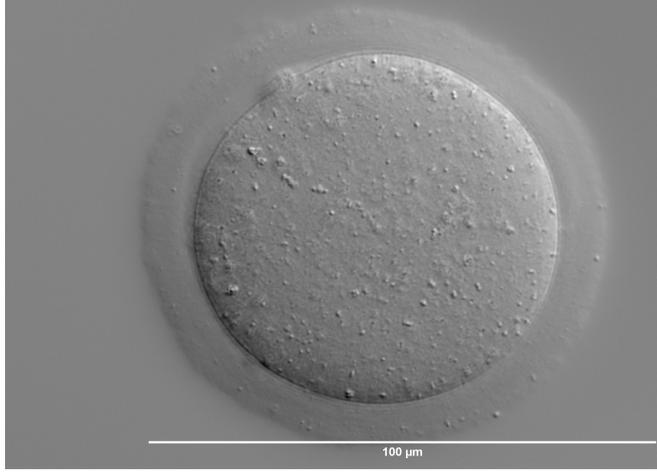


Figure 5.2: Image of a single mouse oocyte taken with a DIC microscope setup.

To demonstrate this anisotropy, we perform DDM analysis on two image stacks of colloidal dispersions, and plot a slice of the DDM tensor $\mathcal{D}(\mathbf{q}, \Delta t)$ at the smallest time lag available in Figure 5.3. Figure 5.3(a) shows the DDM tensor derived from bright field images [164], with concentric rings demonstrating that the underlying process is isotropic [164]. However, Figure 5.3(b) is derived from DIC images of a similar colloidal dispersion, as used previously in Section 4.4, and displays an asymmetric ‘kidney bean’ shape, where the highest peaks align along the orientation of the shear \mathbf{s} in the DIC image, and the lowest peaks are orthogonal to it.

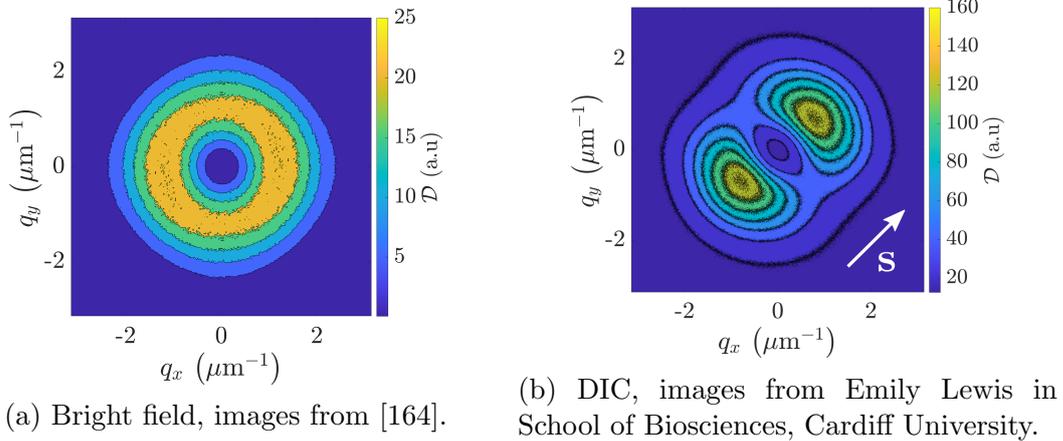


Figure 5.3: Isocontours of $\mathcal{D}(\mathbf{q}, \Delta t)$ for (5.3a) bright field [164] and (5.3b) DIC images of colloidal dynamics. Both sets of images are displayed for the smallest available time separation: $\Delta t = 2.5$ ms for bright field and 0.12s for DIC. The bright field images yield concentric circular isocontours, but in DIC images, there is a clear asymmetry aligned with the direction of the DIC microscope shear, denoted by the white arrow \mathbf{s} .

Since we know that colloidal dispersions involve only Brownian motion, which is isotropic [339], the anisotropy in Figure 5.3 is an artefact of the DIC imaging

process. This anisotropy raises the question of how to interpret output statistics from DDM in terms of the underlying object motion. Although DDM has been used for phase-contrast microscopy in various studies [142, 163, 280, 281], at the time of writing we are not aware of investigations into how shadowing affects the output statistics from DDM.

There are a number of approaches to dealing with similar forms of spatial anisotropy. In [292], simulated dark-field images were used to show that DDM is robust to non-uniform background brightness, an inherent form of anisotropy, under the condition that gradients in the background illumination are small relative to the length scale of the imaged motion. This condition is equivalent to saying that when particles remain in areas of relatively constant local illumination, heterogeneous illumination does not skew the DDM analysis. In DIC, anisotropy in the particle appearance is not dependent on particle's location within the frame, however, so we cannot apply the same logic to our problem.

Another anisotropic setting in which DDM has been successfully applied is found in [340], in which 'Janus' particles were studied. These particles are spherical, with different chemical coatings applied to each half in such a way that the particles themselves can be made to be visually anisotropic (see Figure 5.4), which draws a parallel with the asymmetric particle appearance expected in DIC microscopy. Although each individual Janus particle may be visually heterogeneous, the particles are free to rotate in space; the study in [340] makes use of this freedom to rotate, by coating the particles in such a way that rotational dynamics are induced. DDM is then used to characterise the speed of this rotation. In [340], no consideration is needed to account for the anisotropic appearance of the particles because they may rotate independently of one another, such that the ensemble average of all particles within the frame will remain isotropic. Hence, this application is not directly comparable with the DIC case, where the anisotropic appearance of particles is oriented across the whole image.

The most relevant application of DDM to the challenge of applying DDM in DIC images is [280], which explored the movement of ovoid silica-coated spindle type haematite particles. These particles diffuse at varying speeds along and orthogonal to their long axis. The particles were aligned through the use of a magnetic field, which resulted in anisotropy in the appearance of the image, as well as the displacements observed. The proposed solution in [280] was to apply the radial averaging step over small directed bands coinciding with the major and minor axes of the particles, rather than the entire circular contour Γ . This approach could be applied to phase-contrast images; whilst the orthogonal axes were chosen to align with the major and minor axis of the magnetically oriented particles in [280], the axis in a DIC application could be chosen parallel

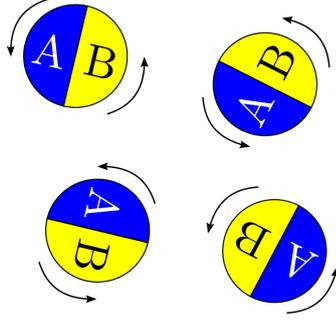


Figure 5.4: A schematic for a series of rotating Janus particles, each half coated with two chemicals ‘A’ and ‘B’. It is possible to tether these particles in such a way that rotation occurs only in the axis parallel to the imaging plane, ensuring that in the resulting images, each particle has exactly half its visible surface coated with each type of chemical [340].

to, and perpendicular with, the shear direction. In particular, perpendicular to the direction of the shear, the shadowing effect should disappear. However, this solution is imperfect, since the accuracy of the radial average is diminished by restricting the orientations over which the average is applied. This is important in cases where we expect the motion to be isotropic; we should be able to use data for all orientations of the spatial frequency \mathbf{q} , since the underlying movement itself has no preferred direction. Furthermore, this solution brings us no closer to understanding what effect, if any, phase-contrast shadowing has on DDM analysis.

The above approaches in [280, 292, 340] show that there are several approaches to dealing with forms of anisotropy in DDM, but none directly solve the problem we would like to tackle in DIC microscopy. We aim to extract motility parameters from time-lapse images of visually heterogeneous particles, whose displacement distributions are independent of the asymmetry observed in the images of the particles. Despite the heterogeneity induced by DIC images, which should violate the isotropic assumption required in DDM analysis, previous applications of DDM to DIC images have been shown to be highly successful, yielding parameters validated by Particle Image Velocimetry (PIV) [142]. We will therefore explore the effect of DIC on DDM analysis, to understand any conditions required in order to justify good use of DDM in this setting.

5.2 An analytic expression for the DDM matrix of a DIC image stack for isotropic motion

A single DIC image is generated by the interference between two beams separated by the shear vector \mathbf{s} , where a change in the optical path of each beam results in a phase difference that is converted into image intensity [341]. Since the shear distance is small relative to the optical limit of the microscope, we can make a

rough approximation that the difference in optical path between the beams is linear, which means the DIC image can be expressed as the difference between overlapping identical offset images [341–343], where each polarised beam in the pair generates a different image. For a DIC image stack $I(\mathbf{r}, t)$, we define image stacks generated by each beam to be $X(\mathbf{r}, t)$ and $Y(\mathbf{r}, t)$, respectively, such that

$$I(\mathbf{r}, t) = Y(\mathbf{r}, t) - X(\mathbf{r}, t), \quad (5.2.1)$$

where

$$Y(\mathbf{r}, t) = X(\mathbf{r} + \mathbf{s}, t), \quad (5.2.2)$$

for a shear $\mathbf{s} = (x_0, y_0)$. This approximation does not take into account refraction of the beam, or birefringence [341], and, further, it does not hold when changes in the phase gradient occur on a scale smaller than the optical resolution of the microscope. However, for now, we restrict ourselves to an idealised set of particles meeting this assumption.

For some initial time, t , and some time displacement Δt , let $I_2 = I(\mathbf{r}, t + \Delta t)$ and $I_1 = I(\mathbf{r}, t)$, with the relation $I_j = Y_j - X_j$, $j \in \{1, 2\}$. We may then write the image difference $\Delta I(\mathbf{r}, t, \Delta t)$ as

$$\begin{aligned} \Delta I &= I_2 - I_1, \\ &= \Delta Y - \Delta X. \end{aligned}$$

Given this decomposition of ΔI , we can now construct the DDM matrix for the DIC image, which we denote by \mathcal{D}_I , in terms of only one of the constituent images, X . By the linearity of the Fourier Transform,

$$\Delta \tilde{I} = \Delta \tilde{Y} - \Delta \tilde{X}.$$

So the definition of the DDM tensor from the image stack I , reproduced below from Equation (2.2.12), can be expressed as

$$\begin{aligned} \mathcal{D}_I(\mathbf{q}, \Delta t) &= \left\langle |\Delta \tilde{I}|^2 \right\rangle_t = \left\langle (\Delta \tilde{I})(\Delta \tilde{I}^*) \right\rangle_t, \\ &= \left\langle (\Delta \tilde{Y} - \Delta \tilde{X})(\Delta \tilde{Y}^* - \Delta \tilde{X}^*) \right\rangle_t \\ &= \left\langle \Delta \tilde{Y} \Delta \tilde{Y}^* + \Delta \tilde{X} \Delta \tilde{X}^* - \Delta \tilde{Y} \Delta \tilde{X}^* - \Delta \tilde{X} \Delta \tilde{Y}^* \right\rangle_t \\ &= \left\langle |\Delta \tilde{Y}|^2 + |\Delta \tilde{X}|^2 - \Delta \tilde{Y} \Delta \tilde{X}^* - \Delta \tilde{X} \Delta \tilde{Y}^* \right\rangle_t \\ &= \left\langle |\Delta \tilde{Y}|^2 \right\rangle_t + \left\langle |\Delta \tilde{X}|^2 \right\rangle_t - \left\langle \Delta \tilde{Y} \Delta \tilde{X}^* \right\rangle_t - \left\langle \Delta \tilde{X} \Delta \tilde{Y}^* \right\rangle_t. \end{aligned} \quad (5.2.3)$$

It is reasonable to assume that $\left\langle |\Delta \tilde{Y}|^2 \right\rangle_t \approx \left\langle |\Delta \tilde{X}|^2 \right\rangle_t$, since these are both statistically averaged quantities of two image stacks depicting the same type of

movement. In the case of X and Y arising from DIC imaging, this assumption is even more reasonable, since shear distances are typically around 0.2-0.3 μm [268, 343, 344], meaning the shear is small relative to the size of the image. Hence, X and Y do not just depict the same type of motion; they depict the same particles exactly.

Next, we expand the final two terms of Equation (5.2.3), dropping the time average brackets for simplicity, as

$$\begin{aligned}
& \Delta\tilde{Y}\Delta\tilde{X}^* + \Delta\tilde{X}\Delta\tilde{Y}^* \\
&= (\tilde{Y}_2 - \tilde{Y}_1)(\tilde{X}_2^* - \tilde{X}_1^*) + (\tilde{X}_2 - \tilde{X}_1)(\tilde{Y}_2^* - \tilde{Y}_1^*) \\
&= (\tilde{Y}_2\tilde{X}_2^* + \tilde{X}_2\tilde{Y}_2^*) + (\tilde{Y}_1\tilde{X}_1^* + \tilde{X}_1\tilde{Y}_1^*) - (\tilde{Y}_2\tilde{X}_1^* + \tilde{X}_1\tilde{Y}_2^*) - (\tilde{Y}_1\tilde{X}_2^* + \tilde{X}_2\tilde{Y}_1^*).
\end{aligned} \tag{5.2.4}$$

Both $\tilde{Y}_2\tilde{X}_2^*$ and $\tilde{Y}_1\tilde{X}_1^*$ describe the correlation between an image taken at the same time, before and after it is shifted by the DIC shear, whilst $\tilde{Y}_2\tilde{X}_1^*$ and $\tilde{Y}_1\tilde{X}_2^*$ describe the correlation between an image and its shifted partner at different times. By the definition of Y in Equation (5.2.2), the translation property of the Fourier transform states that for some known shear \mathbf{s} ,

$$\tilde{Y}_j(\mathbf{q}) = \tilde{X}_j(\mathbf{q})e^{-i(\mathbf{s}\cdot\mathbf{q})}, \tag{5.2.5}$$

or for the analogous discrete Fourier transform case [345],

$$\tilde{Y}_j(q_x, q_y) = \tilde{X}_j(q_x, q_y)e^{-2\pi i\left(\frac{lx_0 + my_0}{L}\right)},$$

where l, m are references to the index of the pixels in the horizontal and vertical axis respectively. Thus, Equation (5.2.5) gives

$$\begin{aligned}
\tilde{Y}_j\tilde{X}_j^* + \tilde{X}_j\tilde{Y}_j^* &= \tilde{X}_j\tilde{X}_j^*e^{-i(\mathbf{s}\cdot\mathbf{q})} + \tilde{X}_j\tilde{X}_j^*e^{i(\mathbf{s}\cdot\mathbf{q})} \\
&= |\tilde{X}_j|^2 (e^{-i(\mathbf{s}\cdot\mathbf{q})} + e^{i(\mathbf{s}\cdot\mathbf{q})}) \\
&= 2|\tilde{X}_j|^2 \cos(\mathbf{q} \cdot \mathbf{s}).
\end{aligned} \tag{5.2.6}$$

The assumption that the process is stationary implies that $|\tilde{X}_j|^2$ is constant with respect to time, so Equation (5.2.6) holds for all initial times, and

$$\tilde{Y}_1\tilde{X}_1^* + \tilde{X}_2\tilde{Y}_2^* = \tilde{Y}_1\tilde{X}_1^* + \tilde{X}_1\tilde{Y}_1^*.$$

An alternative derivation of Equation (5.2.6) can be made by finding the auto-correlation of particles moving with a constant velocity $\mathbf{v} = \mathbf{s}/\Delta t$ [267], which can be seen by solving the Fourier transform of the advection equation,

$$\begin{aligned}
\frac{\partial u}{\partial(\Delta t)} &= \mathbf{v}\nabla u, \\
u(\mathbf{r}, 0) &= \delta(\mathbf{r}),
\end{aligned} \tag{5.2.7}$$

solved by

$$\begin{aligned} \implies \frac{\partial \tilde{u}(\mathbf{q}, \Delta t)}{\partial(\Delta t)} &= (-i\mathbf{q}\mathbf{v}) \tilde{u}(\mathbf{q}, \Delta t), \\ \implies \tilde{u}(\mathbf{q}, \Delta t) &= e^{-i\mathbf{q}\cdot\mathbf{v}\Delta t}. \end{aligned} \quad (5.2.8)$$

Applying Equation (5.2.5) to the final two terms in Equation (5.2.4), we get

$$\begin{aligned} \tilde{Y}_1\tilde{X}_2^* + \tilde{Y}_2\tilde{X}_1^* + \tilde{X}_1\tilde{Y}_2^* + \tilde{X}_2\tilde{Y}_1^* &= \tilde{X}_1\tilde{X}_2^*e^{-is\cdot\mathbf{q}} + \tilde{X}_2\tilde{X}_1^*e^{is\cdot\mathbf{q}} + \tilde{X}_1\tilde{X}_2^*e^{is\cdot\mathbf{q}} + \tilde{X}_2\tilde{X}_1^*e^{-is\cdot\mathbf{q}} \\ &= \tilde{X}_1\tilde{X}_2^* (e^{is\cdot\mathbf{q}} + e^{-is\cdot\mathbf{q}}) + \tilde{X}_2\tilde{X}_1^* (e^{is\cdot\mathbf{q}} + e^{-is\cdot\mathbf{q}}) \\ &= \left(\tilde{X}_1\tilde{X}_2^* + \tilde{X}_2\tilde{X}_1^* \right) (2\cos(\mathbf{s}\cdot\mathbf{q})) \end{aligned} \quad (5.2.9)$$

We note that $\tilde{X}_2\tilde{X}_1^*$ is the Fourier transform of the autocorrelation of the images, which we can express in terms of the Fourier transform of the underlying van Hove function [290], describing the probability of displacement within the image stack, $G(\mathbf{q}, \Delta t)$ (see Equation (2.2.11)),

$$\tilde{X}_2^*\tilde{X}_1 = |\tilde{X}_1|^2 \tilde{G}(\mathbf{q}, \Delta t),$$

Hence, Equation (5.2.9) becomes

$$\begin{aligned} \tilde{Y}_1\tilde{X}_2^* + \tilde{Y}_2\tilde{X}_1^* + \tilde{X}_1\tilde{Y}_2^* + \tilde{X}_2\tilde{Y}_1^* &= |\tilde{X}_1|^2 \tilde{G}^*(q, \Delta t) (e^{is\cdot\mathbf{q}} + e^{-is\cdot\mathbf{q}}) + |\tilde{X}_1|^2 \tilde{G}(q, \Delta t) (e^{is\cdot\mathbf{q}} + e^{-is\cdot\mathbf{q}}) \\ &= |\tilde{X}_1|^2 (e^{is\cdot\mathbf{q}} + e^{-is\cdot\mathbf{q}}) \left(\tilde{G}(\mathbf{q}, \Delta t) + \tilde{G}^*(\mathbf{q}, \Delta t) \right) \\ &= 4|\tilde{X}_1|^2 \cos(\mathbf{s}\cdot\mathbf{q}) f(\mathbf{q}, \Delta t). \end{aligned} \quad (5.2.10)$$

Substituting (5.2.6) and (5.2.10) into (5.2.3),

$$\begin{aligned} \langle |\Delta \tilde{I}|^2 \rangle_t &= 2|\Delta \tilde{X}|^2 - 4|\tilde{X}|^2 \cos(\mathbf{q}\cdot\mathbf{s}) + 4|\tilde{X}|^2 f(\mathbf{q}, \Delta t) \cos(\mathbf{q}\cdot\mathbf{r}), \\ &= 2|\Delta \tilde{X}|^2 - 4|\tilde{X}|^2 \cos(\mathbf{q}\cdot\mathbf{s}) (1 - f(\mathbf{q}, \Delta t)). \end{aligned} \quad (5.2.11)$$

Recalling Equation (2.2.13),

$$\begin{aligned} \mathcal{D}_X(\mathbf{q}, \Delta t) &= |\Delta \tilde{X}|^2 \\ &= 2|\tilde{X}|^2 (1 - f(\mathbf{q}, \Delta t)), \end{aligned}$$

so we can express Equation (5.2.11) as

$$\mathcal{D}_I(\mathbf{q}, \Delta t) = 4|\tilde{X}|^2 (1 - f(\mathbf{q}, \Delta t)) (1 - \cos(\mathbf{q}\cdot\mathbf{s})) \quad (5.2.12)$$

$$= 2\mathcal{D}_X(\mathbf{q}, \Delta t) (1 - \cos(\mathbf{q}\cdot\mathbf{s})). \quad (5.2.13)$$

We now make the assumption that we are dealing with isotropic motion, and take a radial average on \mathcal{D}_I to find $\mathcal{D}_{I,R}$. Let us assume that the underlying particle

motion in the images we want to analyse, and therefore the ISF, is isotropic i.e constant on a circular contour centred on $\mathbf{q} = (0, 0)$. Taking the radial average of \mathcal{D}_I in Equation (5.2.12), as described in Section 2.2.3, we find

$$\mathcal{D}_{I,R}(q, \Delta t) = \frac{1}{|\Gamma|} \oint_{\Gamma} \mathcal{D}_I(\mathbf{q}, \Delta t) d\Gamma \quad (5.2.14)$$

$$\begin{aligned} &= \frac{1}{|\Gamma|} \oint_{\Gamma} 4|\tilde{X}|^2(1 - f(\mathbf{q}, \Delta t))(1 - \cos(\mathbf{q} \cdot \mathbf{s})) d\Gamma \\ &= \frac{(1 - f(q, \Delta t))}{|\Gamma|} \oint_{\Gamma} 4|\tilde{X}|^2(1 - \cos(\mathbf{q} \cdot \mathbf{s})) d\Gamma \end{aligned} \quad (5.2.15)$$

$$= A_I(q)(1 - f(q, \Delta t)), \quad (5.2.16)$$

where

$$A_I(q) = \frac{1}{|\Gamma|} \oint_{\Gamma} 4|\tilde{X}|^2(1 - \cos(\mathbf{q} \cdot \mathbf{s})) d\Gamma, \quad (5.2.17)$$

is a time independent scaling term. Hence, the form of the DDM matrix $\mathcal{D}_{I,R}(q, \Delta t)$ given in Equation (5.2.16) matches that of Equation (2.2.18), which expresses that the DDM matrix as

$$\mathcal{D}_R(q, \Delta t) = A(q) (1 - f(q, \Delta t)).$$

Despite the shadow in DIC images, when the analysed motion is isotropic, the only change from shadowing occurs within the scaling parameter $A(q)$, and the ISF remains unchanged. This means that parameter fitting applied to either D_X or D_I will yield equivalent motility parameters. A more specific form of $A_I(q)$ arises when $|\tilde{X}|^2$ has no dependence on the orientation of \mathbf{q} . This is a realistic assumption when the underlying particle motion is ergodic, since we expect a large enough number of particles to be uniformly distributed throughout the sample, such that there is no directional bias in the image appearance. Evaluating the contour integral in Equation (5.2.15) is simplified by observing that, as a consequence of the images being real valued,

$$\Delta I(\mathbf{r}, \Delta t) \in \mathbb{R} \implies \Delta \tilde{I}(-\mathbf{q}, \Delta t) = \Delta \tilde{I}(\mathbf{q}, \Delta t)^*.$$

Hence,

$$\mathcal{D}(-\mathbf{q}, \Delta t) = \mathcal{D}(\mathbf{q}, \Delta t),$$

showing that the circular contour along Γ in the integral of Equation (5.2.14) along $\mathcal{D}_I(\mathbf{q}, \Delta t)$, centred on $\mathbf{q} = (0, 0)$, is periodic with period π . This can be

easily verified from Equation (5.2.13). Applying this π -periodicity to the radial average allows us to express the contour integral in Equation (5.2.14) as

$$\begin{aligned}
\frac{1}{\|\Gamma\|} \oint_{\Gamma} \mathcal{D}_I(\mathbf{q}, \Delta t) d\Gamma &= \frac{1}{\|\Gamma\|} \oint_{\Gamma} 4|\tilde{X}|^2(1 - f(\mathbf{q}, \Delta t))(1 - \cos(\mathbf{q} \cdot \mathbf{s})) d\Gamma, \\
&= 4|\tilde{X}|^2(1 - f(q, \Delta t)) \frac{1}{\|\Gamma\|} \oint_{\Gamma} (1 - \cos(\mathbf{q} \cdot \mathbf{s})) d\Gamma \\
&= 4|\tilde{X}|^2(1 - f(q, \Delta t)) \left(1 - \frac{1}{\|\Gamma\|} \oint_{\Gamma} \cos(\mathbf{q} \cdot \mathbf{s}) d\Gamma \right) \\
&= 4|\tilde{X}|^2(1 - f(q, \Delta t)) \left(1 - \frac{1}{2\pi q} \int_{-\pi}^{\pi} q \cos(qs \cos(\theta - \phi)) d\theta \right) \\
&= 4|\tilde{X}|^2(1 - f(q, \Delta t)) \left(1 - \frac{1}{2\pi} \int_{-\pi-\phi}^{\pi-\phi} \cos(qs \cos(\Theta)) d\Theta \right) \\
&= 4|\tilde{X}|^2(1 - f(q, \Delta t)) \left(1 - \frac{1}{\pi} \int_0^{\pi} \cos(qs \cos(\Theta)) d\Theta \right), \tag{5.2.18}
\end{aligned}$$

where $\Theta = \theta - \phi$ and $\mathbf{s} = s(\cos(\phi), \sin(\phi))$. By definition of the zero-order Bessel function of the first kind [324],

$$J_0(x) = \frac{1}{\pi} \int_0^{\pi} \cos(x \cos(z)) dz, \tag{5.2.19}$$

so the specific form of the scaling term $A_I(q)$ in Equation (5.2.17) reduces to

$$A_I(q) = 4|\tilde{X}|^2(1 - J_0(qs)) = 2A_X(q) (1 - J_0(qs)). \tag{5.2.20}$$

Comparing with the definition of \mathcal{D}_X in Equation (2.2.12), we determine that

$$\mathcal{D}_I(q, \Delta t) = 2(1 - J_0(qx_0))\mathcal{D}_X. \tag{5.2.21}$$

We therefore conclude that for isotropic motion, there exists a time-independent map between the DDM matrix of the phase images, $\mathcal{D}_{X,R}$, and the DDM matrix of the DIC images, $\mathcal{D}_{I,R}$. Critically, the anisotropy due to the DIC shadow can be accounted for exclusively by the q -dependent scaling parameter $2(1 - J_0(qs))$, leaving the ISF unchanged. This means DDM applied to DIC microscopy of isotropic particle motions can be used without needing to account for the artificial anisotropy introduced by DIC, since the same motility statistics will be derived from the DIC image stack as would be determined in another form of imaging without shadowing. This justifies the previous application of DDM to DIC images of *Drosophila* oocytes, which did not adjust analysis to account for shadowing from the choice of microscopy [142]. In order to validate our findings, we now present both simulated and real data, which allow us to confirm the relationship (5.2.21).

5.2.1 Verifying the use of DDM in DIC through simulations and colloid images

5.2.1.1 Validation with simulations

We generate simulations of 150 particles following Brownian motion inside a 490×490 window, with diffusion coefficient $D = 0.5 \mu\text{m}^2/\text{s}$. One pixel represents $1 \mu\text{m}$, for a duration of 4000 frames, at a frequency of 8 frames per second. From these simulations, we generate the phase image stack $X(\mathbf{r}, t)$, and the corresponding DIC image stack $I(\mathbf{r}, t)$ using a DIC shift of 1 pixel to the right, i.e $\mathbf{s} = (1, 0)$, which is small compared to the size of the image. The first frame of each image stack is visualised in Figure 5.5. From each image stack, we generate the corresponding DDM matrices $D_{X,R}$ and $D_{I,R}$, respectively. The code is available at https://github.com/OstlerT/DIC_DDM_Ostler.git.

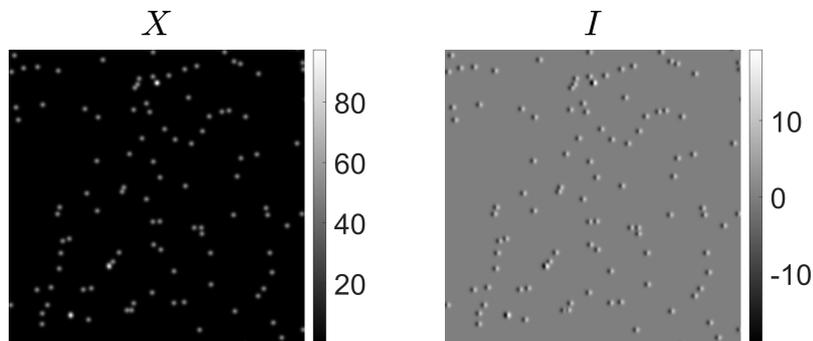


Figure 5.5: The first frame of the simulated phase image, X , and the corresponding DIC image, I , which is generated by $X(\mathbf{r} + \mathbf{s}, t) - X(\mathbf{r}, t)$.

In Figure 5.6, we demonstrate that the decomposition in Equation (5.2.21) holds for all frequencies $q < 2 \mu\text{m}^{-1}$, by evaluating the left and right hand sides of (5.2.21) and comparing them. For different sets of simulation parameters, the agreement can deteriorate at very high and very low frequencies, but, as discussed in Section 2, it is standard practice in DDM to restrict fitting to some q -interval away from the low frequencies, where statistical sampling of the radial average is poor, and away from large frequencies, where the signal to noise ratio is poor [164]. This interval, called the range of accessible frequencies, is where τ_D can be fit. In this simulation, we choose conservative estimates for these limits, choosing the lower limit on the fitting region to be $q_l = 0.5 \mu\text{m}^{-1}$, and the upper limit to be $q_u = 1.2 \mu\text{m}^{-1}$, as shown in Figure 5.7. As long as the relationship holds over accessible frequency range, we can conclude that the analysis of $\mathcal{D}_{X,R}$ and $\mathcal{D}_{I,R}$ will yield equivalent results.

We now show that the diffusion coefficients generated by fitting \mathcal{D}_X and \mathcal{D}_I are equivalent. We generate 150 movies of diffusing particles, choosing the diffusion coefficient to be $D = 0.5 \mu\text{m}^2/\text{s}$ following [292], and perform fitting over the

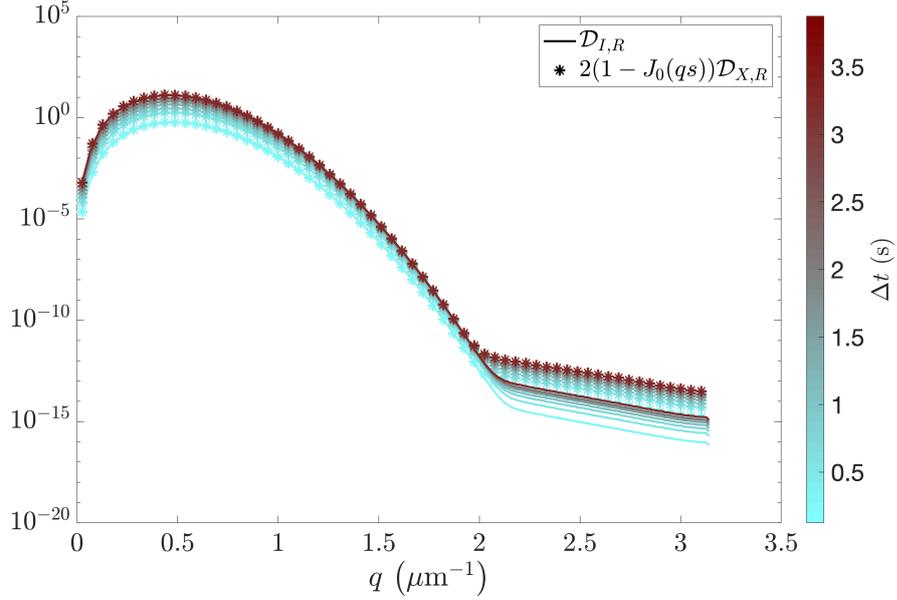


Figure 5.6: Snapshots of $\mathcal{D}_{I,R}(q, \Delta t)$ and the rescaled form $2(1 - J_0(qs))\mathcal{D}_{X,R}(q, \Delta t)$, plotted at fixed times over all spatial frequencies. There is excellent agreement for all frequencies $q < 2\mu\text{m}^{-1}$ hence the relationship (5.2.21) is valid over this interval.

frequency interval $0.5\mu\text{m}^{-1} < q < 1.2\mu\text{m}^{-1}$, where the bounds are chosen to be sufficiently far from noisy regions as seen in Figure 5.7. The average diffusion coefficients generated by the simulations of X and I are $\bar{D}_X = 0.481 \pm 0.009\mu\text{m}^2/\text{s}$, and $\bar{D}_I = 0.482 \pm 0.010\mu\text{m}^2/\text{s}$, respectively, with the distribution of the fitted diffusion coefficients plotted in Figure 5.8a. Although a wide spread of fitted diffusion coefficients occur, plotting the values $D_X - D_I$ as a histogram in Figure 5.8b shows the fitting difference is small between \mathcal{D}_X and \mathcal{D}_I . Hence, DIC shadowing has a minimal effect on the fitted diffusion coefficient.

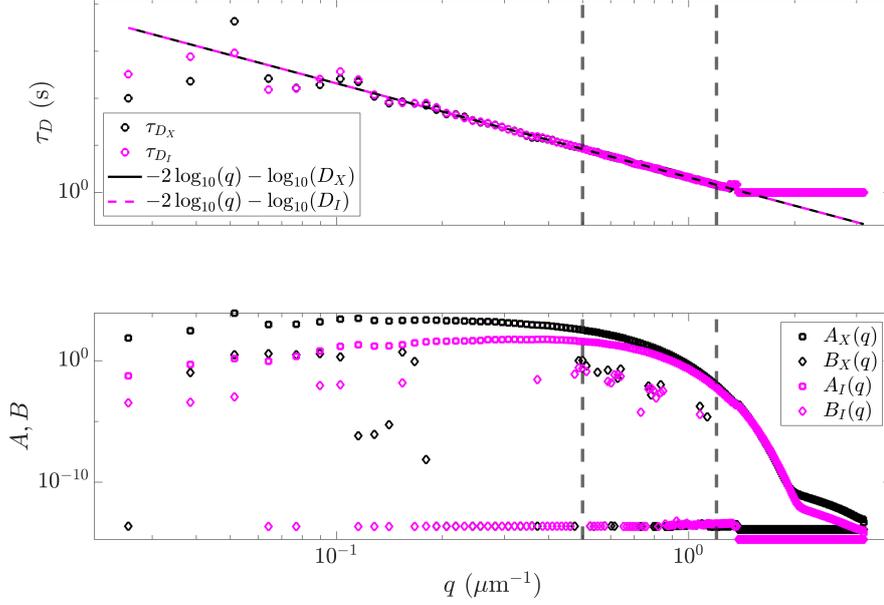
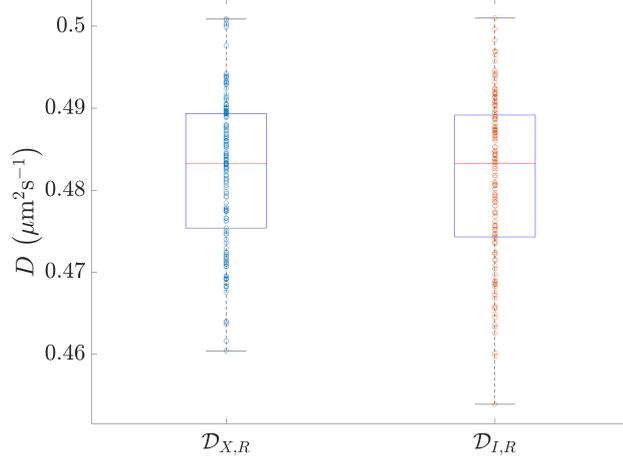


Figure 5.7: The upper plot depicts the fitted values of τ_{D_X} and τ_{D_I} for the image stacks X and I , respectively, in a single simulation, on a logarithmic scale. The linear regression (2.2.27) is also plotted, and agreement of the intercepts demonstrates equivalent diffusion coefficients are determined. The lower plot visualises the fitting parameters $A(q)$ and $B(q)$ for each stack.

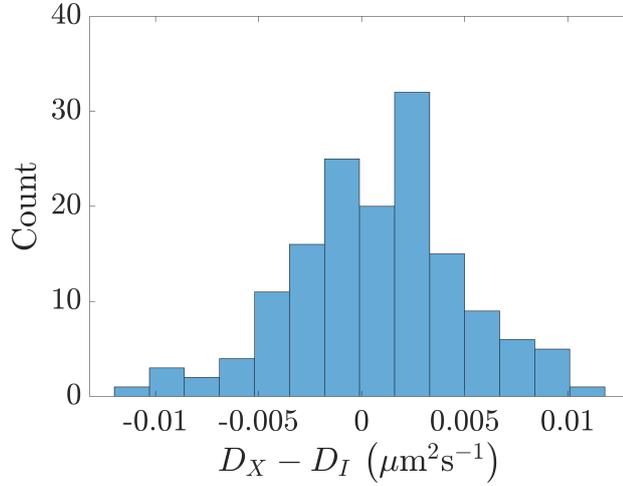
5.2.1.2 Validation with experimental data

Additionally, we show that DDM is accurate in real DIC images. We consider a single image stack formed of 7680 DIC images of a colloidal dispersion, taken by Emily Lewis in the School of Physics and Astronomy, Cardiff University. This suspension uses Polystyrene (PS) beads with a nominal radius of 100 nm, with less than 3% coefficient of variance (cv) (Alpha Nanotech Colloidal PS Beads NP-PA07CPSX78). These PS beads were dispersed in 30% v/v glycerol/water mixtures to a particle concentration of 0.1 mg/mL, and 13 μl were pipetted into a 0.12 mm high and 13 mm diameter chamber made of a Grace Bio-Labs (Bend, US) SecureSeal imaging spacer on a $(76 \times 26 \times 1) \text{ mm}^3$ microscope slide (Menzel Gläser). The chamber was then sealed with a $(24 \times 24) \text{ mm}^2$ #1.5 coverslip (Menzel Gläser), and stored in a 100% humidity environment at 7 until use.

The corresponding images of the colloidal dispersion are 1344×1024 pixels representing a $433.44 \times 330.24 \mu\text{m}$ real region (pixel length 322.5 nm), from which we take a square subset with a side length of 1024 pixels. The shearing distance is measured to be $238 \pm 10 \text{ nm}$ [344], with a shear angle of $\phi = \pi/4$ and a phase offset of 90° . We note that whilst the spheres themselves are sub-pixel in size, and therefore violate the linearity assumption required for separating the DIC image into separate components X and Y , the high volume of spheres means we can study the bulk movement of the sample rather than individual sphere locations.



(a) Distributions of the diffusion coefficients fitted from $\mathcal{D}_{X,R}$. and $\mathcal{D}_{I,R}$.



(b) Distribution of the difference $D_X - D_I$ for each simulation, with mean 6.9×10^{-4} .

Figure 5.8: For 150 simulations of Brownian motion, we generate phase stacks $X(\mathbf{r}, t)$ and corresponding DIC images $I(\mathbf{r}, t)$. DDM analysis on each stack yields a diffusion coefficient, D_X and D_I , respectively. (5.8a) Although the fitted diffusion coefficients have some spread around the true value $D = 0.5 \mu\text{m}^2/\text{s}$, (5.8b) the difference $D_X - D_I$ is small, so the DIC shear has minimal effect on the fitting accuracy.

From the corresponding DDM matrix, \mathcal{D}_I , we fit a diffusion coefficient of $0.676 \mu\text{m}^2/\text{s}$. Comparing this to the Stokes-Einstein relation [164, 266],

$$D = \frac{K_B u}{6\pi\eta r},$$

the expected diffusion coefficient is $0.667 \mu\text{m}^2/\text{s}$, where K_B J/K is the Boltzmann constant, $u = 291.25\text{K}$ is absolute temperature, $\eta = 3.20$ mPa s is the dynamic viscosity of a 30% glycerol solution, and r the particle radius. Hence, D is comparable to the fitted value. We visualise $\tau_{\mathcal{D}_I}$, $A(q)$, $B(q)$ and the fitting regime in Figure 5.9.

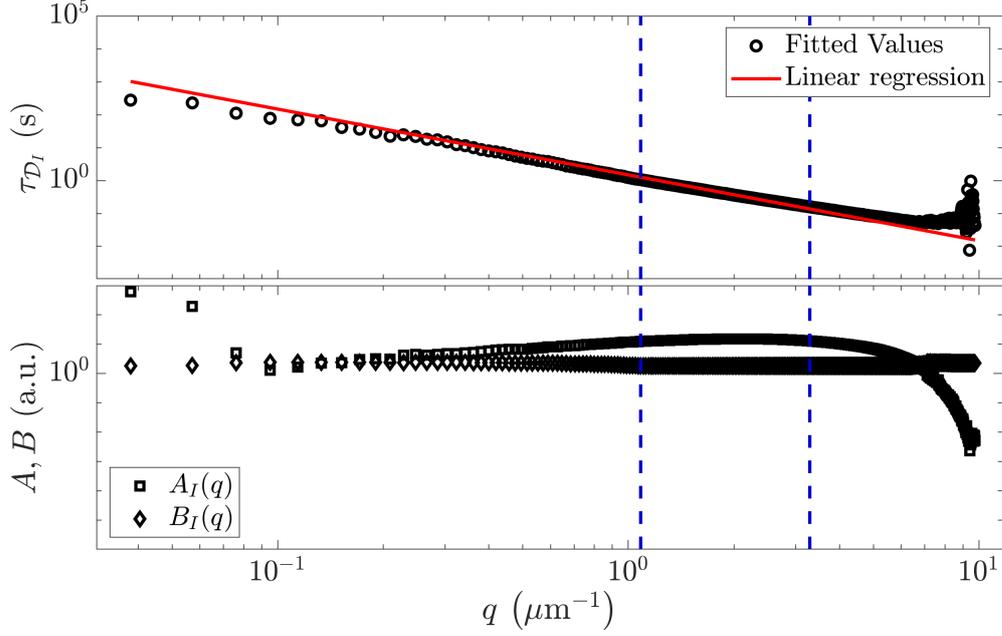


Figure 5.9: DDM analysis of the diffusion coefficients for fitted DIC images, from colloidal suspensions, yields a fitted diffusion coefficient of $0.676\mu\text{m}^2/\text{s}$. The fitting region $1.1\mu\text{m}^{-1} < q < 3.3\mu\text{m}^{-1}$, marked between the vertical blue dashed lines, is chosen so that the signal ($A_I(q)$) to noise ($B_I(q)$) ratio is sufficiently high.

We note that at no point have we used prior knowledge on the DIC shear length or direction in analysing this colloidal data. In practice, the shear orientation is determined by the setup of the microscope, and the shear distance may be computed based on information about the magnification and focal length of the lenses used [342, 344]. We consider now the interesting question of whether DDM can provide an alternate derivation of this information. In the next section, we show that this is theoretically possible, but limited by numerical constraints.

5.2.2 DDM in isotropic motion can extract the direction, but not the magnitude, of the shear

Consider a DIC image stack I , for which no corresponding image stacks X or Y exist, with an associated shear, $\mathbf{s} = s(\cos(\phi), \sin(\phi))$. From Equations (5.2.20) and (2.2.19), we have the relation

$$\mathcal{D}_{I,R}(q, \Delta t) = 2A_X(q)(1 - J_0(qs))(1 - f(q, \Delta t)) + B(q).$$

Since both $A_X(q)$ and $(1 - J_0(qs))$ are functions of q only, we can fit their product, but cannot determine which parts of the scaling term are attributed to the underlying DIC shear and which arise from the brightness of the image stack $I(\mathbf{r}, \Delta t)$. Thus, when dealing with the radially averaged DDM matrix, $\mathcal{D}_{I,R}$, we cannot say anything about \mathbf{s} , so we must instead approach the DDM tensor, \mathcal{D}_I .

Consider a circular contour inscribed into $\mathcal{D}(\mathbf{q}, \Delta t)$. Assuming that $|\tilde{X}|^2$ is isotropic, along this circular contour we may write

$$\mathcal{D}_I(q, \theta, \Delta t) = 4|\tilde{X}|^2(1 - f(q, \Delta t))(1 - \cos(qs \cos(\theta - \phi))) + B_I(q), \quad (5.2.22)$$

where θ and q are the direction and magnitude of the spatial frequency vector, respectively. DDM analysis of the radially averaged DDM matrix allows us to generate the noise term $B_I(q)$, and to describe $f(q, \Delta t)$. Therefore, if we rewrite Equation (5.2.22) in the form

$$\frac{\mathcal{D}_I(q, \theta, \Delta t) - B_I(q)}{1 - f(q, \Delta t)} = 4|\tilde{X}|^2(1 - \cos(qs \cos(\theta - \phi))),$$

the left side of the equation is made up of parameters which are known or may be derived, and the right hand side contains three unknown parameters, $4|\tilde{X}|^2$, qs and ϕ . Hence, we can consider the following model function to be fit, depending on some vector of unknowns \mathbf{b} as follows:

$$\frac{\mathcal{D}_I(q, \theta, \Delta t) - B_I(q)}{1 - f(q, \Delta t)} = b_1(q)(1 - \cos(b_2(q) \cos(\theta - b_3(q)))) \quad (5.2.23)$$

where

$$\mathbf{b}(q) = (b_1, b_2, b_3) = (4|\tilde{X}|^2, qs, \phi).$$

Since the right hand side of (5.2.23) does not depend on Δt , and so neither should the left hand side, we may average its value over all time lags. Consequently, for each q , Equation (5.2.23) generates a single curve, on which we may apply fitting algorithms to determine the three parameters in $\mathbf{b}(q)$.

In practice, however, it is necessary to add a q dependent noise term to the right hand side of Equation (5.2.23), because the fit of $B_I(q)$ is not exact, and for small frequencies, $1 - f(q, \Delta t)$ is close to 0. Any error created by incorrect fitting of $\mathbf{b}(q)$ is therefore amplified when dividing through by $1 - f(q, \Delta t)$, blowing up the left side of (5.2.23). Therefore, we should instead fit

$$\frac{\mathcal{D}_I(q, \theta, \Delta t) - B_I(q)}{1 - f(q, \Delta t)} = b_1(q)(1 - \cos(b_2(q) \cos(\theta - b_3(q)))) + b_4(q) \quad (5.2.24)$$

$$\mathbf{b}(q) = (b_1, b_2, b_3, b_4) = (4|\tilde{X}|^2, qx_0, \phi, b_4).$$

Using the DIC colloid image data, we attempt to determine the unknowns in $\mathbf{b}(q)$ by iterating over the frequency q . In Figure 5.10, we plot some of the curves generated by the right side of (5.2.24), with numerically determined $\mathbf{b}(q)$, which are shown to describe the average behaviour of the left side of (5.2.24) well. However, there is a large spread of data about these curves, due to noise

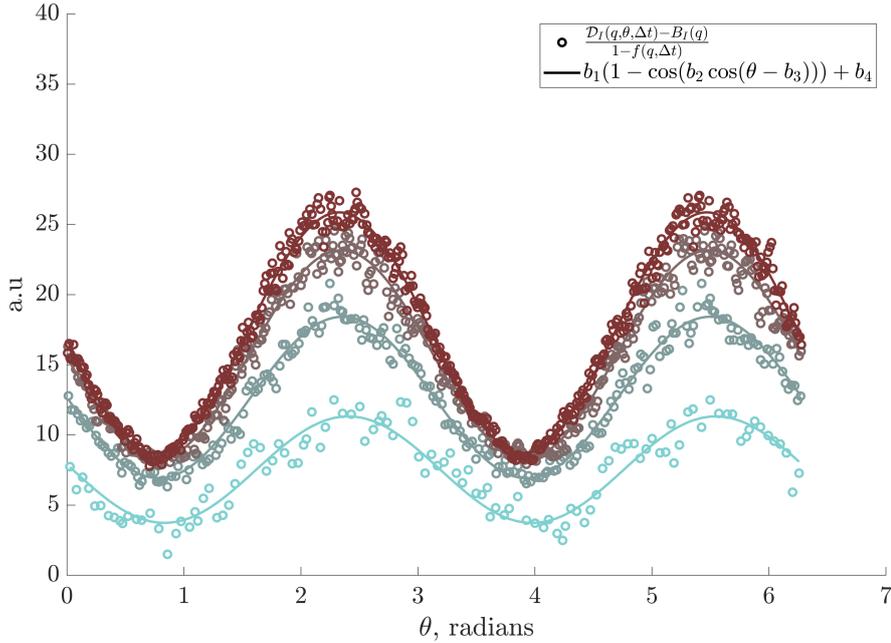
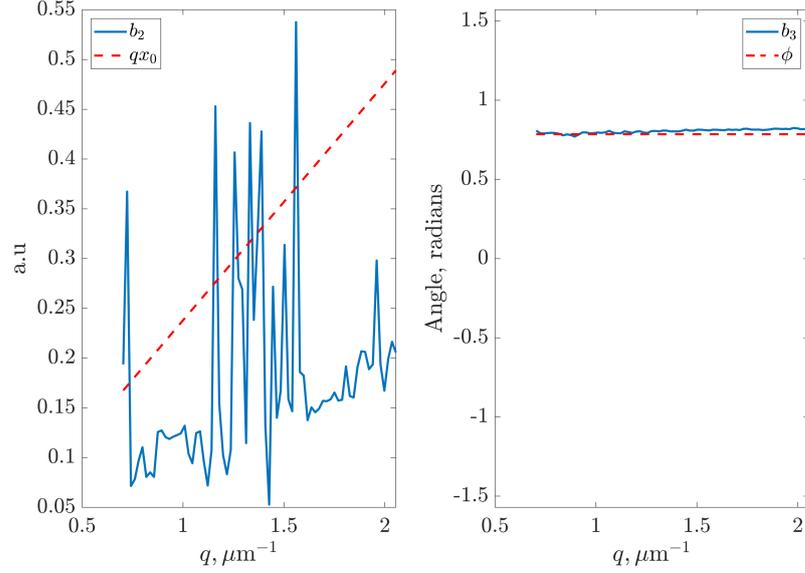


Figure 5.10: Scatter plots of the left hand side of (5.2.24), with fitted curves as described in the right hand side of Equation (5.2.24), applied to the DDM tensor generated by DIC colloidal data. Curves are plotted for $q \in \{0.40, 0.78, 1.16, 1.54\}$.

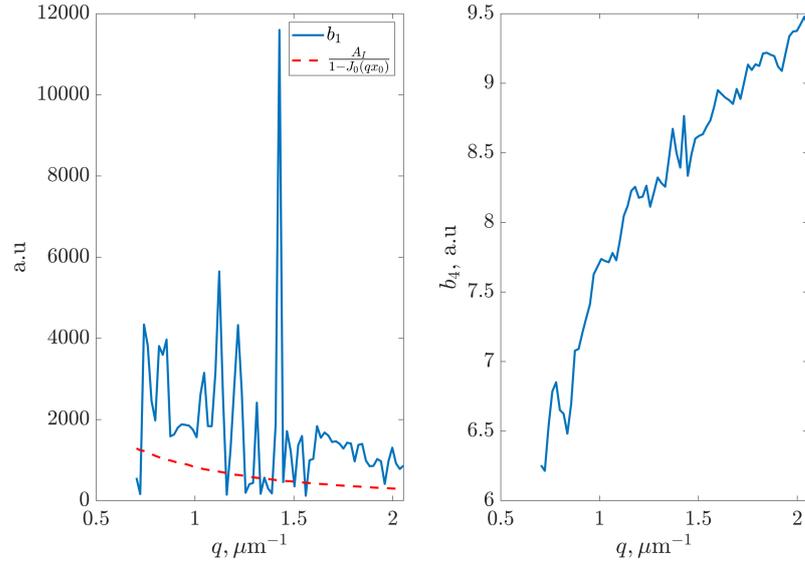
that would normally be mitigated by a radial average. This error is likely carried through in the fit of the parameters in $\mathbf{b}(q)$, which are plotted in Figure 5.11.

In Figure 5.11a, we plot b_2 and b_3 , which are intended to approximate qs and ϕ respectively. Whilst it is clear that the shear angle, ϕ , is very well fit for all q , b_2 does not approximate s well; it under fits the expected value, and is subject to a large amount of noise. This suggests that although ϕ is a very sensitive parameter in this system, s is much less sensitive, so may be more challenging to fit.

We now turn our attention to the parameters b_1 and b_4 , which approximate the scale of the plotted contour of the DDM matrix, and the level of noise in the fit, respectively, plotted in Figure 5.11b. The fitted noise term b_4 is shown to be non zero, and depends on q . This agrees with our previous assertion, that any error in fitting the frequency-dependent noise term $B(q)$ will become larger when dividing by $1 - f(q, \Delta t)$ in (5.2.24), and as such, the additional noise term is needed to improve the quality of fit. However, comparing the magnitude of b_4 to the scale term b_1 , the signal amplitude is orders of magnitude greater than the noise, such that the signal-to-noise ratio is sufficiently high to assume noise is not important in this system. The fit of the scaling term b_1 is also poor, however, when compared with the expected value plotted by the red dashed line. In summary, DDM is too insensitive to the shear distance to be of any use in classifying this parameter.



(a) Fits of the parameters b_2 and b_3 in Equation (5.2.24), compared to their known true values, qx_0 and ϕ , respectively. In particular, the fit of ϕ is excellent, but fitting x_0 by a linear regression on the left hand figure would be unreliable due to the large amount of noise present.



(b) Fits of the parameters b_1 and b_4 in (5.2.24). In the left plot, we also plot the expected value of b_1 (dashed red line) as given by rearranging (5.2.20) to solve for $4|\tilde{X}|^2$, which shows b_1 is not fit accurately. In the right plot, for a noiseless system, b_4 should be equal to zero everywhere, but is clearly non-zero here. The signal b_1 is orders of magnitude greater than the noise b_4 , however, so there is a good signal to noise ratio.

Figure 5.11: Plots of the fitted parameters $\underline{b}(q)$ in (5.2.24), compared to the known parameter values given by the magnitude and directions of the DIC shear, x_0 and ϕ respectively.

A corollary to this conclusion is that using known information about the DIC shear during analysis will not improve fitting performance. One might consider that when fitting the DDM matrix of a DIC image stack, supplying known inform-

ation about the DIC shear to the fitting process could help improve the quality of parameter fitting. However, we may only use known information about the shear direction, ϕ , when fitting parameters from the DDM tensor, since $\mathcal{D}_{I,R}(q, \Delta t)$ is invariant under the shear angle. When trying to fit any parameters from the DDM tensor, as shown in Figure 5.10, we lose the noise-reducing benefit of the radial average. Thus, knowing the angle of the DIC shear provides no benefit to parameter fitting. Furthermore, supplying the known shear distance s when fitting the DDM matrix has no particular benefit, since we still have to fit the coefficient $A(q)$. In fact, it should be easier to fit this coefficient if we do not supply s , since $(1 - J_0(qs)) < 1$ when qs is sufficiently small, a condition easily met due to the small DIC shear size. This means that $A_I(q)$ is expected to be smaller than $A(q)$, meaning the domain over which we must search for $A_I(q)$ is smaller and, thus, the numerical stability of the parameter fitting should be improved when fitting $A_I(q)$ compared to $A(q)$ in DIC images.

In conclusion, we have shown that using DIC images adds artificial anisotropy to DDM analysis. However, we have shown that in isotropic motion, DDM applied to the DIC images yields equivalent results to DDM applied to the phase images constructing the DIC image stack. We have additionally shown that in cases where we do not have any prior knowledge of the shear, we can use DDM to extract the shear direction, but not the shear magnitude. Although we have only verified our results for Brownian motion, relationship (5.2.20) holds for all forms of isotropic motion. Does this conclusion extend beyond isotropy, however? It is possible that some new relationship replaces Equation (5.2.20) when motion is anisotropic, or perhaps no such relationship exists at all. In the next section, we relax the assumption of isotropy, focussing on a directed advection-diffusion process to explore the problem.

5.3 Anisotropy

For anisotropic processes, the derivation of $\mathcal{D}_{I,R}(\mathbf{q}, \Delta t)$ follows the same argument as in Section 5.2, up until taking the radial average in Equation (5.2.12). For this section, we will assume that the underlying images have no spatial heterogeneity, such that the scaling term $|\tilde{X}|^2$ has no dependence on the orientation of q , for simplicity. Applying the azimuthal average to Equation (5.2.12) yields the DDM

matrix as

$$\begin{aligned}
\mathcal{D}_{I,R}(q, \Delta t) &= \frac{4|\tilde{X}^2|}{\|\Gamma\|} \oint_{\Gamma} (1 - f(\mathbf{q}, \Delta t)) (1 - \cos(\mathbf{q} \cdot \mathbf{s})) d\Gamma \\
&= \frac{4|\tilde{X}^2|}{\|\Gamma\|} \oint_{\Gamma} 1 - f(\mathbf{q}, \Delta t) - \cos(\mathbf{q} \cdot \mathbf{s}) + f(\mathbf{q}, \Delta t) \cos(\mathbf{q} \cdot \mathbf{s}) d\Gamma \\
&= 4|\tilde{X}^2| \left(1 - f(q, \Delta t) - J_0(qs) + \frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) \cos(\mathbf{q} \cdot \mathbf{s}) d\Gamma \right).
\end{aligned} \tag{5.3.1}$$

The only difference between the isotropic case in Equation (5.2.18), and the anisotropic case in Equation (5.3.1), is how we deal with the final term,

$$\frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) \cos(\mathbf{q} \cdot \mathbf{s}) d\Gamma.$$

When the ISF is isotropic, such that $f(\mathbf{q}, \Delta t) = f(q, \Delta t)$, we can take the ISF out of the radial average, evaluate the contour integral of the $\cos(\mathbf{q} \cdot \mathbf{s})$ term, and factorise out a q -dependent scaling factor to reclaim the standard DDM format in Equation (2.2.19). When the ISF depends on the orientation of \mathbf{q} , however, this final term must be evaluated as a product integral. This means that we cannot immediately recover the generic DDM fitting form of Equation (2.2.19), and must evaluate some new ISF which defines the DDM matrix. This new expression of the DDM matrix will depend on the specific form of $f(\mathbf{q}, \Delta t)$, making it challenging to explore a general anisotropic motion. Instead, we will study a specific advection-diffusion behaviour below, motivated by cytoplasmic streaming as presented in [142] and [154]. In this simple example, we consider a constant, directed advection applied to all particles.

Deriving the ISF for a directed advection-diffusion process follows the same argument as Section 2.2.5, where we discussed the ISF for bacterial advection. The van Hove function $G(\mathbf{r}, \Delta t)$ that describes the distribution of displacements of particles as they both diffuse and move with velocity \mathbf{v} , satisfies the advection-diffusion equation given in Equation (1.2.1),

$$\frac{\partial G(\mathbf{r}, \Delta t)}{\partial(\Delta t)} = \nabla (D \nabla G(\mathbf{r}, \Delta t)) - \mathbf{v} \cdot \nabla G(\mathbf{r}, \Delta t),$$

where D is the diffusion coefficient, as in Section 2, and \mathbf{v} (μm^{-1}) is assumed to be a constant velocity vector. The solution to the Fourier transformed advection-diffusion equation with a constant diffusion coefficient is

$$\tilde{G}(\mathbf{q}, \Delta t) = e^{-q^2 D \Delta t} e^{-i\mathbf{q} \cdot \mathbf{v} \Delta t},$$

and taking the real part gives the ISF to be [267]

$$f(\mathbf{q}, \Delta t) = e^{-q^2 D \Delta t} \cos(\mathbf{q} \cdot \mathbf{v} \Delta t).$$

Substituting the new ISF into the definition of $\mathcal{D}_I(\mathbf{q}, \Delta t)$ yields

$$\mathcal{D}_I(\mathbf{q}, \Delta t) = 4|\tilde{X}|^2(1 - e^{-q^2 D \Delta t} \cos(\mathbf{q} \cdot \mathbf{v} \Delta t) - \cos(\mathbf{q} \cdot \mathbf{s}) + e^{-q^2 D \Delta t} \cos(\mathbf{q} \cdot \mathbf{v} \Delta t) \cos(\mathbf{q} \cdot \mathbf{s})).$$

We now consider the azimuthal average on \mathcal{D}_I , for which it is convenient to convert to polar coordinates, using

$$\mathbf{v} = v(\cos \phi_v, \sin \phi_v), \quad \mathbf{s} = s(\cos \phi_s, \sin \phi_s), \quad \mathbf{q} = q(\cos \theta, \sin \theta),$$

and

$$\frac{1}{\|\Gamma\|} \oint_{\Gamma} \mathcal{D}(\mathbf{q}, \Delta t) d\Gamma = \frac{1}{2\pi q} \int_{-\pi}^{\pi} \mathcal{D}(q, \theta, \Delta t) q d\theta = \frac{1}{2\pi} \int_{-\pi}^{\pi} \mathcal{D}(q, \theta, \Delta t) d\theta.$$

Since the contour on which we take the azimuthal average of $\mathcal{D}(q, \theta, \Delta t)$ is π -periodic, we restrict $\phi_v, \phi_s \in [-\pi/2, \pi/2)$ and take the azimuthal average as

$$\begin{aligned} \mathcal{D}_I(q, \Delta t) = \frac{4|\tilde{X}|^2}{\pi} \int_{-\pi/2}^{\pi/2} & (1 - e^{-\kappa} \cos(\lambda \cos(\theta - \phi_v)) - \cos(\xi \cos(\theta - \phi_s)) \\ & + e^{-\kappa} \cos(\lambda \cos(\theta - \phi_v)) \cos(\xi \cos(\theta - \phi_s))) d\theta, \end{aligned} \quad (5.3.2)$$

where we have introduced the non-dimensional parameters

$$\lambda = qv\Delta t, \quad \xi = qs, \quad \kappa = q^2 D \Delta t.$$

We consider $|\tilde{X}|^2$ constant in θ since we assume no visual asymmetry or spatial heterogeneity exists in the phase images X . We will deal with this integral term-wise, defining the following terms.

$$T_1 = \frac{4|\tilde{X}|^2}{\pi} \int_{-\pi/2}^{\pi/2} 1 d\theta = 4|\tilde{X}|^2. \quad (5.3.3)$$

$$T_2 = \frac{4|\tilde{X}|^2}{\pi} \int_{-\pi/2}^{\pi/2} e^{-\kappa} \cos(\lambda \cos(\theta - \phi_v)) d\theta, \quad (5.3.4)$$

$$T_3 = \frac{4|\tilde{X}|^2}{\pi} \int_{-\pi/2}^{\pi/2} \cos(\xi \cos(\theta - \phi_s)) d\theta, \quad (5.3.5)$$

$$T_4 = \frac{4|\tilde{X}|^2}{\pi} \int_{-\pi/2}^{\pi/2} e^{-\kappa} \cos(\lambda \cos(\theta - \phi_v)) \cos(\xi \cos(\theta - \phi_s)) d\theta. \quad (5.3.6)$$

Each of these terms has a physical interpretation, which we explain through Figure 5.12. The image difference ΔI , formed by a pair of DIC images of a single particle, will show 4 ‘spots’. These correspond to the position of the particle before and after shearing, at times t_1 and t_2 . The DDM tensor is the sum of all spatial correlations in ΔI , with each term T_1 to T_4 corresponding to the spatial correlation between different pairs of spots. The term T_1 (Equation (5.3.3))

corresponds to the self correlation of each spot, which has no spatial or time dependence, and is therefore constant. The spatial correlation between spots at different times both before shearing, or both after shearing, gives rise to T_2 (Equation (5.3.4)), which would be the ISF in bright-field imaging (see Equation (4.3.3)). The spatial correlation between DIC-shifted spot pairs at the same time give rise to T_3 (Equation (5.3.5)). Finally, T_4 is generated by the spatial correlation of pairs of spots formed by particles at different times, where one is sheared and the other is not.

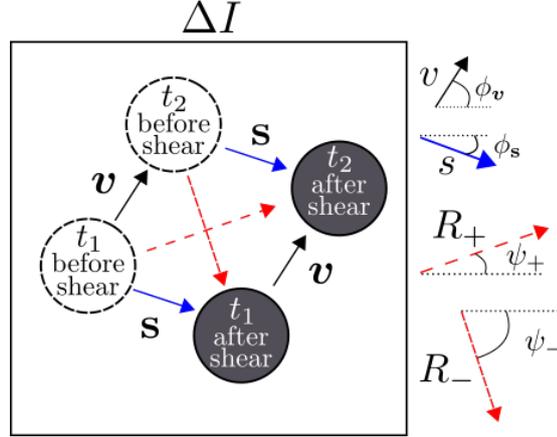


Figure 5.12: The image difference ΔI of one DIC imaged particle will contain four spots, corresponding to the position of the particle at times t_1 and t_2 , both before and after shearing. The DDM tensor is formed by the spatial correlations between spots, which can be broken down into terms T_1 - T_4 in Equations (5.3.3)-(5.3.6). The vectors between spot pairs at different times and shears are described by R_{\pm} .

An important observation from Figure 5.12 is that the magnitude and direction of the vectors between spots at different times and shears, denoted R_+ and R_- and shown by red dashed arrows, will be dependent on both \mathbf{s} and \mathbf{v} . In isotropic motion, where the direction of particle movement is uniformly randomly distributed, the ensemble average of these red dashed vectors will also be uniformly randomly distributed, whereas in directed movement, their orientations will always be the same. It is this distinction which causes the product integral to form in Equation (5.3.1).

The terms T_2 and T_3 in Equations (5.3.4) and (5.3.5), by the definition of the zero-order Bessel function of the first kind in Equation (5.2.19), evaluate to

$$T_2 = \frac{1}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} e^{-\kappa} \cos(\lambda \cos(\theta - \phi_v)) d\theta = e^{-\kappa} J_0(\lambda), \quad (5.3.7)$$

$$T_3 = \frac{1}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\xi \cos(\theta - \phi_s)) d\theta = J_0(\xi). \quad (5.3.8)$$

To evaluate T_4 in Equation (5.3.6), we use the trigonometric identity

$$\begin{aligned}\cos(\lambda \cos(\theta - \phi_{\mathbf{v}})) \cos(\xi \cos(\theta - \phi_{\mathbf{s}})) &= \frac{1}{2} \cos(\lambda \cos(\theta - \phi_{\mathbf{v}}) + \xi \cos(\theta - \phi_{\mathbf{s}})) \\ &\quad + \frac{1}{2} \cos(\lambda \cos(\theta - \phi_{\mathbf{v}}) - \xi \cos(\theta - \phi_{\mathbf{s}})).\end{aligned}$$

Hence, the final term in (5.3.2) can be written as

$$\begin{aligned}&\frac{e^{-\kappa}}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\lambda \cos(\theta - \phi_{\mathbf{v}})) \cos(\xi \cos(\theta - \phi_{\mathbf{s}})) d\theta \\ &= \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\lambda \cos(\theta - \phi_{\mathbf{v}}) + \xi \cos(\theta - \phi_{\mathbf{s}})) d\theta \\ &\quad + \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\lambda \cos(\theta - \phi_{\mathbf{v}}) - \xi \cos(\theta - \phi_{\mathbf{s}})) d\theta \\ &= \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\lambda(\cos\theta \cos\phi_{\mathbf{v}} + \sin\theta \sin\phi_{\mathbf{v}}) + \xi(\cos\theta \cos\phi_{\mathbf{s}} + \sin\theta \sin\phi_{\mathbf{s}})) d\theta \\ &\quad + \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\lambda(\cos\theta \cos\phi_{\mathbf{v}} + \sin\theta \sin\phi_{\mathbf{v}}) - \xi(\cos\theta \cos\phi_{\mathbf{s}} + \sin\theta \sin\phi_{\mathbf{s}})) d\theta \\ &= \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\cos\theta(\lambda \cos\phi_{\mathbf{v}} + \xi \cos\phi_{\mathbf{s}}) + \sin\theta(\lambda \sin\phi_{\mathbf{v}} + \xi \sin\phi_{\mathbf{s}})) d\theta \\ &\quad + \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\cos\theta(\lambda \cos\phi_{\mathbf{v}} - \xi \cos\phi_{\mathbf{s}}) + \sin\theta(\lambda \sin\phi_{\mathbf{v}} - \xi \sin\phi_{\mathbf{s}})) d\theta. \quad (5.3.9)\end{aligned}$$

Using the identity,

$$R_{\pm} \cos(\theta - \psi_{\pm}) = R_{\pm} \cos(\theta) \cos(\psi_{\pm}) + R_{\pm} \sin(\theta) \sin(\psi_{\pm}),$$

and comparing to Equation (5.3.9), we derive

$$\begin{aligned}R_{\pm} \cos(\psi_{\pm}) &= \lambda \cos(\phi_{\mathbf{v}}) \pm \xi \cos(\phi_{\mathbf{s}}), \\ R_{\pm} \sin(\psi_{\pm}) &= \lambda \sin(\phi_{\mathbf{v}}) \pm \xi \sin(\phi_{\mathbf{s}}),\end{aligned}$$

where

$$R_{\pm} = \sqrt{\lambda^2 + \xi^2 \pm 2\lambda\xi \cos(\phi_{\mathbf{v}} - \phi_{\mathbf{s}})}, \quad (5.3.10)$$

and

$$\tan \psi_{\pm} = \frac{\lambda \sin(\phi_{\mathbf{v}}) \pm \xi \sin(\phi_{\mathbf{s}})}{\lambda \cos(\phi_{\mathbf{v}}) \pm \xi \cos(\phi_{\mathbf{s}})}. \quad (5.3.11)$$

Hence, we can express Equation (5.3.9) as

$$\begin{aligned}
T_4 &= \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\cos\theta(\lambda\cos\phi_v + \xi\cos\phi_s) + \sin\theta(\lambda\sin\phi_v + \xi\sin\phi_s)) d\theta \\
&+ \frac{e^{-\kappa}}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\cos\theta(\lambda\cos\phi_v - \xi\cos\phi_s) + \sin\theta(\lambda\sin\phi_v - \xi\sin\phi_s)) d\theta \\
&= \frac{e^{-\kappa}}{2\pi} \left(\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(R_+ \cos(\theta - \psi_{\pm})) + \cos(R_- \cos(\theta - \psi_{\pm})) d\theta \right). \quad (5.3.12)
\end{aligned}$$

The representation of T_4 in Equation (5.3.12) is useful, because it has split T_4 into two terms, corresponding to each of the red dashed arrows in Figure 5.12. We only needed one term to describe T_2 , because the vectors which contribute to it (black vectors \mathbf{v} in Figure 5.12) have equal magnitude and are parallel; the same is true of the terms forming T_3 (blue vectors \mathbf{s}). However, the T_4 terms will only be parallel when \mathbf{s} and \mathbf{v} are, or equivalently, $\phi_v = \phi_s$. Similarly, these terms only have the same magnitude when $v = s$. When these conditions are not met, each red dashed vector requires its own term in T_4 . The values of R_{\pm} denote the length of these vectors, whilst ψ_{\pm} denotes their orientations.

To simplify Equation (5.3.12), since neither R_{\pm} or ψ_{\pm} depend on θ , both are constant on the circular contour Γ . Also, the integral in the definition of the Bessel function of the first kind in Equation (5.2.19) is invariant under a constant shift in θ , since the domain of integration is equal to the period of the integrand. Therefore, Equation (5.3.12) can be written as

$$T_4 = \frac{e^{-\kappa}}{2} (J_0(R_+) + J_0(R_-)). \quad (5.3.13)$$

Combining (5.3.3)-(5.3.8) and (5.3.13), we can express the radial average of \mathcal{D}_I for advection-diffusion problems as

$$\mathcal{D}_I = 4|\tilde{X}|^2 \left(1 - e^{-\kappa} J_0(\lambda) - J_0(\xi) + \frac{e^{-\kappa}}{2} (J_0(R_+) + J_0(R_-)) \right), \quad (5.3.14)$$

which, by Equation (2.2.18), gives us the ISF for the DDM matrix \mathcal{D}_I as

$$\begin{aligned}
f_I(q, \Delta t) &= J_0(\xi) + e^{-\kappa} J_0(\lambda) \\
&- \frac{e^{-\kappa}}{2} \left(J_0 \left(\sqrt{\lambda^2 + \xi^2 + 2\lambda\xi \cos(\Phi)} \right) + J_0 \left(\sqrt{\lambda^2 + \xi^2 - 2\lambda\xi \cos(\Phi)} \right) \right), \quad (5.3.15)
\end{aligned}$$

where $\Phi = \phi_v - \phi_s$. Unlike the case of isotropic motion, the ISF now depends on the DIC shear parameters ξ and ϕ_s , and we cannot immediately conclude that the DDM analysis on DIC images will yield accurate fitting parameters for an advection-diffusion process. In Section 5.3.1, we consider what happens when we attempt to include these parameters in the fitting stage of DDM.

5.3.1 Fitting the DDM matrix to a simulated advection-diffusion process

The new ISF stated in Equation (5.3.15) is more complex than the ISF for the same process in bright-field images, so we aim to evaluate the practicality of its use in parameter fitting. We simulate image data in the same way as Section 5.2.1, with the additional feature that each particle has velocity $v = 1\mu\text{m s}^{-1}$ with angle $\phi_v = 0$. For this simulation, we generate three DIC image stacks, with the intention of testing how DDM performance varies with the angle difference Φ . Each image stack uses the same shear distance $s = 0.238\mu\text{m}$, reflecting the realistic shear condition of the DIC microscope used to image the colloidal data in Section 5.2.1, but has a different shear angle $\phi_s \in \{0, \pi/4, \pi/2\}$, such that $\Phi \in \{0, \pi/4, \pi/2\}$.

With the addition of the velocity term, and the dependence of the ISF on the DIC shear, we must adapt the DDM fitting approach if we hope to determine the unknown movement parameters in our images, namely the diffusion coefficient, D , the velocity magnitude, v , and the velocity direction ϕ_v , where we assume that the shear \mathbf{s} is known through experimental design. Although we would like to fit ϕ_v directly, this is not possible, since Equation (5.3.15) is even with respect to the angle difference $\Phi = \phi_v - \phi_s$. That means for a fixed shear orientation ϕ_s , the velocity orientations ϕ_v and $2\phi_s - \phi_v$ will yield same equivalent ISF. We may therefore only fit the angle difference Φ , where we restrict the domain for Φ to non-negative angles.

The first stage of fitting will therefore aim to determine the following three parameter groups,

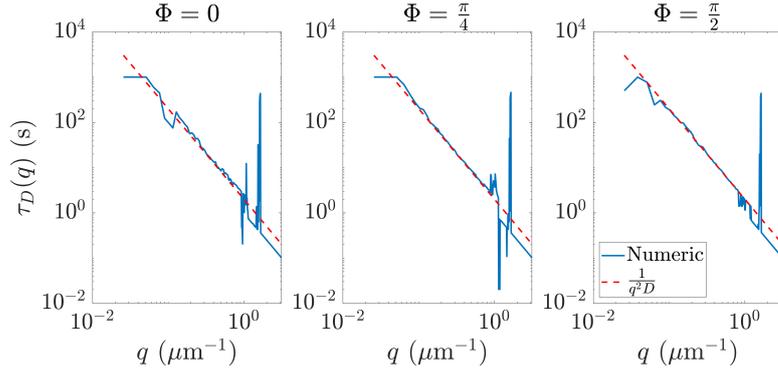
$$\tau_D(q) = \frac{1}{q^2 D}, \quad (5.3.16)$$

$$\tau_v(q) = 1/qv, \quad (5.3.17)$$

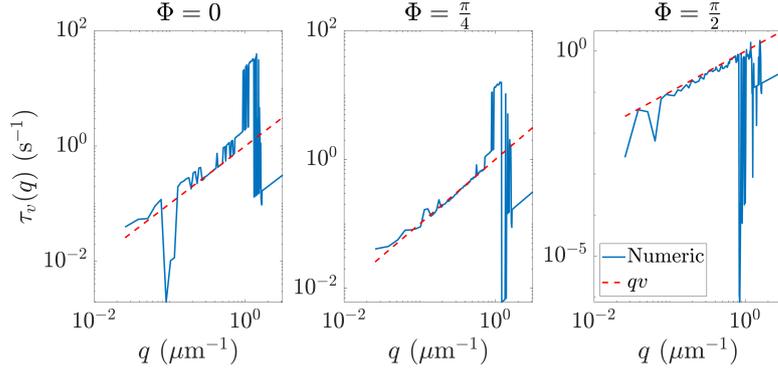
$$\Phi(q) = \phi_v - \phi_s, \quad (5.3.18)$$

An additional adaption we implement is normalisation of the DDM matrix, by dividing though by its largest value. This reduces the size of the parameter search domain for the scaling term $A(q)$, which improves the efficiency of the fitting algorithm. Normalisation was not implemented in Section 5.2.1, because the parameter space was much smaller, and good quality fitting was achievable without rescaling. We plot the fitted parameters in Figure 5.13.

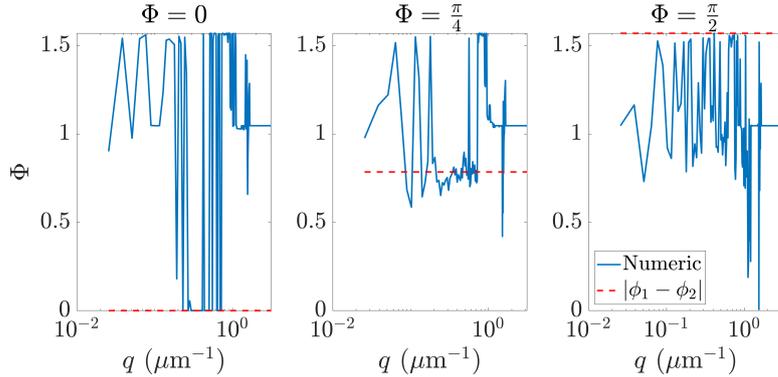
Based on Figure 5.13, regardless of the angle between the DIC shear and advection, the critical movement parameters τ_D and τ_v are generally well fit, although they suffer from small noisy regions. The approximation of Φ , however, is very poor. To understand why, we visualise the ISF for a range of parameters



(a) Fitting diffusion from $\tau_D(q)$.



(b) Fitting velocity from $\tau_v(q)$.



(c) Fitting the absolute angle between the DIC shear and the velocity, Φ .

Figure 5.13: Fits of the three key parameters in Equation (5.3.14); (5.13a) τ_q (5.3.16), (5.13b) τ_v (5.3.17) and (5.13c) Φ (5.3.18), each at different true values of the angle difference Φ . The fit for τ_D and τ_v is universally good, but Φ is always poorly fit.

in Figure 5.14. When the shear is small, the ISF appears to be only weakly dependent on Φ , so fitting is insensitive to the Φ . However, as the shear distance increases, for example $s = 5$, the ISF demonstrates a non-linear relationship on Φ , which could result in the existence of several values of Φ which locally minimise fitting residuals.

In summary, we have shown that for a directed advection-diffusion problem, the anisotropy of the observed movement interacts with the artificial anisotropy

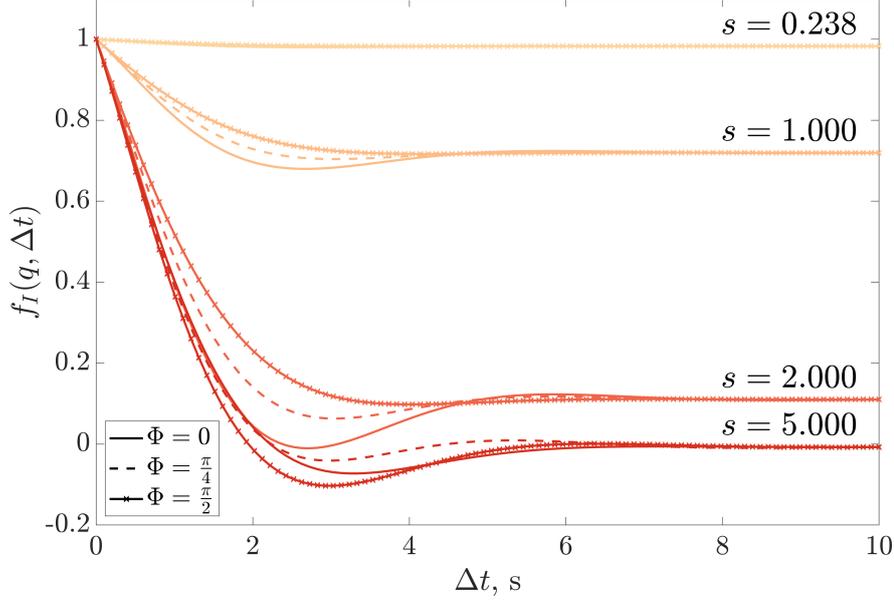


Figure 5.14: The expected ISF $f_I(q, \Delta t)$ (5.3.15) in our simulated data for DIC shear values of $s \in \{0.238, 1.000, 2.000, 5.000\}$. For small shears, the functions are tightly constrained and indiscernible from one another, but both the function range and the significance of Φ increase as the shear increases.

induced by the DIC shadow. This interaction, described entirely within the term T_4 in Equation (5.3.6), results in the need to adapt the ISF to account for shadowing. This DIC-adjusted ISF, $f_I(q, \Delta t)$ in Equation (5.3.15), can be used in parameter fitting in order to accurately approximate the velocity and diffusion coefficient. However, we have also shown that the non-linearities and large parameter space in the DIC-adjusted ISF cause parameter fitting to become unstable, so caution is advised when undertaking DDM analysis of DIC image data to ensure error is avoided.

Whilst our approach to generating the DIC-adjusted ISF can theoretically be adopted in other problems than advection-diffusion, this is likely to be challenging. In more complex problems, an analytic expression DIC-adjusted ISF may be difficult to determine, or may not exist if the term T_4 has no analytic solution. Even when such a function does exist, it may not be useable; although we have been able to well approximate the diffusion coefficient and advection velocity in our example cases in Figure 5.13, we cannot assume this will be true for all combinations of D , v , Φ and s due to the non-linearity of the ISF as shown in Figure 5.14. These problems with DIC-adjustment would be naturally mitigated by the machine learning pipeline proposed in section 4.4, however, since the DIC shadow can be simulated in the training data.

Since we are the first to observe interactions with DDM as a result of artificial anisotropy induced by phase-contrast shadowing, it is likely that DDM analysis,

without adjustment for the phase-contrast shadow, has already been applied to phase-contrast images of anisotropic movement. If true, this could raise concerns that the inference drawn from this analysis is flawed. However, in Section 5.3.2, we will demonstrate that it is possible to determine a good approximation of the ‘correct’ diffusion coefficient and advection velocity by fitting the ‘incorrect’ non-adjusted ISF to $\mathcal{D}_{I,R}$, under the realistic assumption that the shear size is small.

5.3.2 Approximating advection and DIC shears as separable effects

In this section, we will describe a condition for which parameter fitting, without adjusting the ISF for phase-contrast shadowing, will still result in a reasonable approximation of the diffusion coefficient and advection velocity. This condition is true when the shear distance is small compared to the distance over which particles move under advection.

Briefly recapping our description of the DIC microscope in Section 5.1, images are formed by finding the difference in phase shift between two closely separated beams [268]. The size of this separation must be small in order for the difference in phase shift to correspond to local changes in optical path length, and hence, generate high resolution images. The DIC shift is typically 0.2-0.3 μm [268, 343, 344], approximately equal to the length of a pixel (optical resolution) [344]. We compare the size of the shear to the size of the displacements resulting from advection, equal to $v\Delta t$ at time lag Δt . If the displacement resulting from velocity is smaller than one pixel length at all time lags, then the velocity can scarcely be said to exist; in this case, we can ignore advection, and movement is reduced to the isotropic diffusion behaviour studied in Section 5.2.1.

We therefore consider the case where the displacement from advection is greater than the DIC shear. It is reasonable to assume, more strongly, that $v\Delta t \gg s$, since we can choose to record images for a longer period of time to enforce this assumption. Under this assumption, we can use Taylor approximations to simplify the ISF in Equation (5.3.15). In Appendix C.2, we use such Taylor approximations to show that for DIC images, small shear distances mean the ISF can be bounded above and below by its behaviour when $\Phi = 0$ and $\Phi = \pi/2$, which enhances our understanding of the ISF behaviour in DIC images and demonstrates that the ISF is only weakly dependent on the angle difference Φ . More importantly, we now present an approximation of the ISF, from which we conclude that DDM analysis of advection-diffusion behaviour in DIC images

may be undertaken without consideration of shadowing effects, under the condition that displacements from advection are much greater than the DIC shear distance.

To motivate the following approximation, we refer back to the radial average applied to \mathcal{D}_I in Equation (5.3.1), where the key difference in the isotropic and anisotropic case is the integral of the product

$$\oint_{\Gamma} f(\mathbf{q}, \Delta t) \cos(\mathbf{q} \cdot \mathbf{s}).$$

In the isotropic case, f does not depend on the orientation of \mathbf{q} and can be taken outside of the integral, whilst in the anisotropic case, the product must be dealt with in more detail. Before this work, however, there was no indication that the shear from DIC might even affect the DDM tensor. A misguided user, unaware of this issue, might assume that the DDM matrix can still be defined by Equation (2.2.18), that is,

$$\mathcal{D}_R(q, \Delta t) = A(q) (1 - f(q, \Delta t)),$$

for some scaling term $A(q)$. By fitting Equation (2.2.18) to \mathcal{D}_I , the user is unknowingly making the assumption that

$$\frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) \cos(\mathbf{q} \cdot \mathbf{s}) d\Gamma \approx \left(\frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) d\Gamma \right) \left(\frac{1}{\|\Gamma\|} \oint_{\Gamma} \cos(\mathbf{q} \cdot \mathbf{s}) d\Gamma \right). \quad (5.3.19)$$

In general, Equation (5.3.19) is clearly not true when f depends on the orientation of the spatial frequency vector q . However, we can use the small-shear properties of DIC to argue that when $\mathbf{q} \cdot \mathbf{s} \approx 0$ for all \mathbf{q} , $\cos(\mathbf{q} \cdot \mathbf{s})$ is approximately constant on the whole contour. In this case, the product integral would once again be treated as separable, and we could rewrite the DDM matrix as given in Equation (5.3.1) as

$$\begin{aligned} \mathcal{D}_I(q, \Delta t) &= 4|\tilde{X}^2| \left(1 - f(q, \Delta t) - J_0(qs) + \frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) \cos(\mathbf{q} \cdot \mathbf{s}) \right) \\ &\approx 4|\tilde{X}^2| (1 - f(q, \Delta t) - J_0(qs) + f(q, \Delta t) J_0(qs)) \\ &= 4|\tilde{X}^2| (1 - J_0(qs)) (1 - f(q, \Delta t)), \end{aligned}$$

and hence, we can account for the DIC shear through only a scaling term once more. This argument is not as strong as the isotropic case, where f had no dependence on orientation at all. However, we will show that for small DIC shears, the approximation in Equation (5.3.19) holds in the case of advection-diffusion behaviours, for which we can undertake DDM analysis using the original,

unaltered fitting approach as stated in Chapter 2.2. We do so taking the Taylor series of the left and right hand sides of Equation (5.3.19) about $\mathbf{s} = 0$, and matching up leading order terms.

The left-hand side of Equation (5.3.19) is equal to T_4 , i.e.

$$T_4 = \frac{e^{-\kappa}}{2} \left(J_0(\sqrt{\lambda^2 + \xi^2 + 2\lambda\xi \cos(\Phi)}) + J_0(\sqrt{\lambda^2 + \xi^2 - 2\lambda\xi \cos(\Phi)}) \right).$$

The Taylor expansion of $T_4(\lambda, \xi, \Phi)$ around $\xi = 0$ is equal to,

$$T_4 = e^{-\kappa} \left(J_0(\lambda) - \frac{2 \sin^2(\Phi) J_1(\lambda) + \lambda \cos^2(\Phi) (J_0(\lambda) - J_2(\lambda))}{4\lambda} \xi^2 + \mathcal{O}(\xi^4) \right),$$

which by the recurrence relation on the Bessel Functions [324],

$$\frac{2a}{x} J_a(x) = J_{a-1}(x) + J_{a+1}(x), \quad (5.3.20)$$

can be expressed as

$$T_4(\lambda, \xi, \Phi) = e^{-\kappa} \left(J_0(\lambda) + \frac{J_2(\lambda) (\cos^2(\Phi) - \sin^2(\Phi)) - J_0(\lambda)}{4} \xi^2 + \mathcal{O}(\xi^4) \right). \quad (5.3.21)$$

meanwhile, the right hand side of Equation (5.3.15) is equal to

$$\left(\frac{1}{\|\Gamma\|} \oint_{\Gamma} f(\mathbf{q}, \Delta t) \right) \left(\frac{1}{\|\Gamma\|} \oint_{\Gamma} \cos(\mathbf{q} \cdot \mathbf{s}) \right) = e^{-\kappa} J_0(\lambda) J_0(\xi),$$

with the Taylor expansion about $\xi = 0$ equal to

$$J_0(\lambda) e^{-\kappa} \left(1 - \frac{1}{4} \xi^2 + \mathcal{O}(\xi^4) \right). \quad (5.3.22)$$

The approximation error is found as the difference between Equations (5.3.21) and 5.3.22, equal to

$$\frac{e^{-\kappa} \xi^2}{4} J_2(\lambda) (2 \cos^2(\Phi) - 1) + \mathcal{O}(\xi^4), \quad (5.3.23)$$

which is maximised when $\Phi = 0$ or $\Phi = \pi/2$ and minimised when the $\mathcal{O}(\xi^2)$ terms vanish at $\Phi = \pi/4$. The error is small when the quadratic term is small, i.e

$$\frac{\xi^2}{4} J_2(\lambda) (2 \cos^2(\Phi) - 1) \ll 1 \quad (5.3.24)$$

Since $|J_2(\lambda)| < 0.5$ [346], and $|(2 \cos^2(\Phi) - 1)| < 1$, $\xi \ll 8$ gives us a weak bound on the size of ξ .

When ξ is sufficiently small for this approximation to hold, $\mathcal{D}_I(q, \Delta t)$ as stated in Equation (5.2.13) can be fitted without consideration of the DIC shear, i.e

$$\mathcal{D}_I = A(q) (1 - e^{-\Delta t / \tau_D} J_0(\tau_v \Delta t)) + B(q), \quad (5.3.25)$$

We fit Equation (5.3.25) to the same simulation from Figure 5.13, where $s = 0.238 \mu\text{m}$ and $q < \pi$ means $\xi = 0.75$, which means the condition on the size of the approximation error in Equation (5.3.24) is satisfied. We plot the results of our parameter fitting in Figure 5.15, which demonstrates excellent agreement between the expected and fitted diffusion coefficients. The fitted velocity curves derived in Figure 5.15b are smoother than when fitting the full, correct form of the ISF as in Figure 5.13. This is likely because the simplified ISF in (5.3.25) contains fewer terms, and is thus easier to fit. The velocity fits are generally good, although when $\Phi = \pi/2$, we slightly under-estimate the velocity, and slightly over-estimate the velocity parameter when $\Phi = 0$.

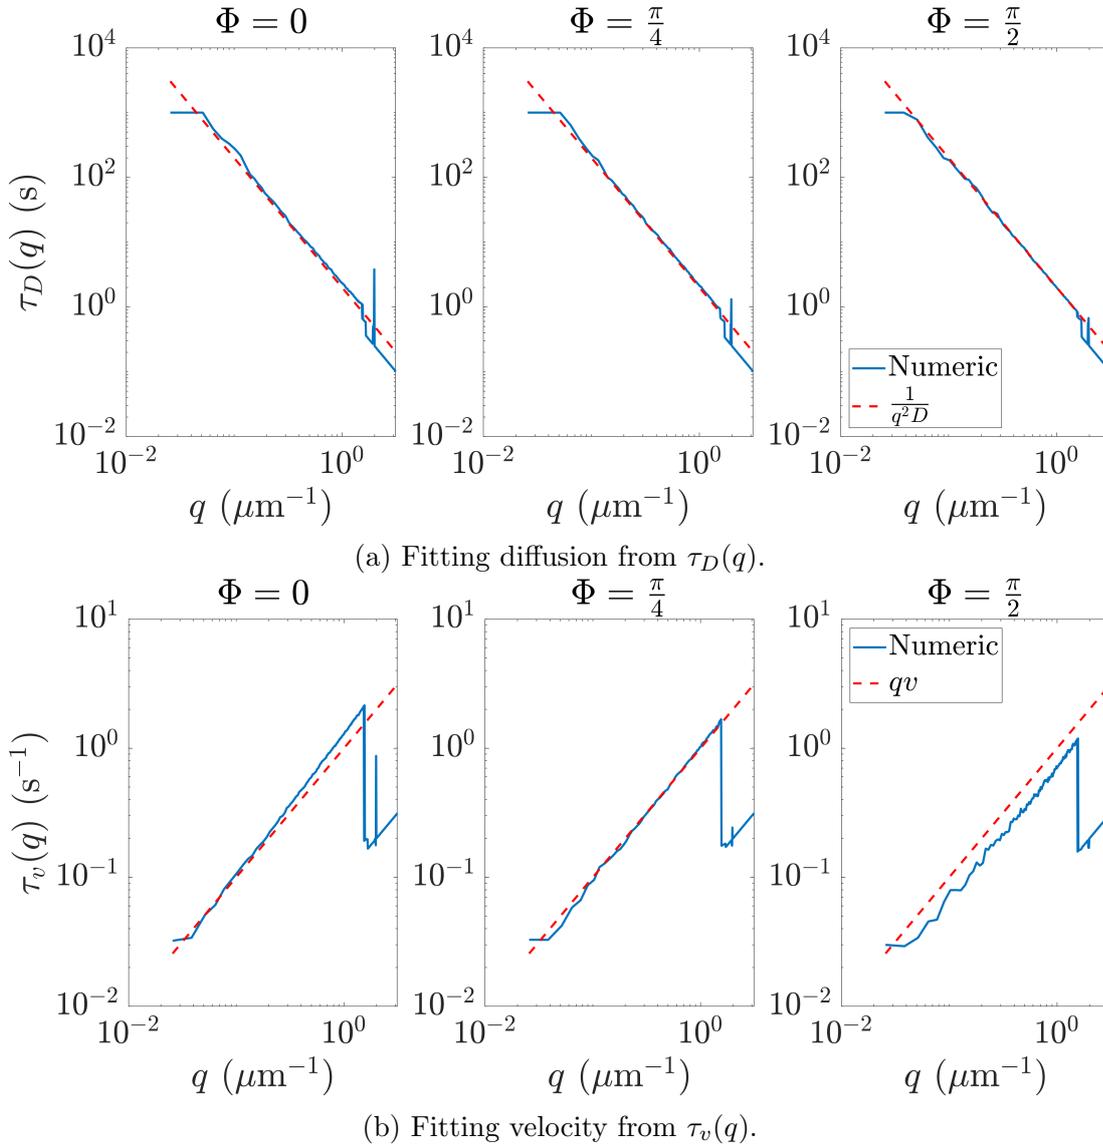


Figure 5.15: Fits of the movement parameters τ_D and τ_v as in Figure 5.13 using the simplified Equation (5.3.25) to fit the DDM matrix as if no DIC shear existed. Diffusion is well fit in all cases, but although the velocity curves are fitted more smoothly than using other fitting functions, the gradients are less accurately fit when $\Phi = 0$ or $\pi/2$.

In conclusion, we have shown that it is possible to approximate both the diffusion coefficient and velocity magnitude from DDM analysis of DIC images of an advection-diffusion behaviour, without adjusting the ISF for phase-contrast shadowing as presented in Equation (5.3.15). This is a useful result that considerably simplifies DDM analysis in real advection-diffusion systems, since we can use a much more simple form of the ISF with fewer parameters, which is easier to fit.

There are, however, some challenges associated with the approximation in Equation (5.3.19). First, we observe in Figure 5.15 that the accuracy with which we fit the velocity term varies with respect to Φ . If you can choose $\Phi = \pi/4$, the approximation becomes more accurate, but in the context of time-lapse oocyte assessment, we cannot control Φ ; the DIC shear direction ϕ_s is fixed, whilst the velocity direction ϕ_v is unknown until the oocyte is placed into the incubator. Rotating the oocyte would require opening the incubator, which is invasive. Therefore, different oocytes will have different associated values of Φ , which may compromise comparison of fitted velocities between oocytes. If the effect of Φ on velocity fitting is consistent and predictable, however, we may be able to adjust our fitted velocities accordingly to mitigate the Φ -induced error. It is therefore an important challenge to describe how parameter fitting accuracy varies in DDM when we ignore the DIC shear.

A second challenge is that we have yet to quantify for what range of shear sizes s the approximation in Equation (5.3.19) holds. Although small ξ justifies our assumption, as $\xi \rightarrow 0$, we note that

$$\lim_{\xi \rightarrow 0} D_{I,R}(q, \Delta t) = 0,$$

because with no shear, there is no image. Hence, if the shear (and therefore ξ) is too small, parameter fitting will be compromised by numerical instability, as we try to fit curves which are approximately zero everywhere. There is, hence, a lower bound on the range of viable ξ for which we may undertake fitting. There will also be an upper bound on ξ for which the approximation in Equation (5.3.19) holds, above which we cannot justify ignoring phase-contrast shadowing on based on approximating T_4 (Equation (5.3.6)) as the product of two contour integrals. Hence, we wish to explore the range of ξ for which this approximation is viable.

Furthermore, we observe that the approximation in Equation (5.3.19) is not the only condition which justifies ignoring the shadowing effect during fitting. For very large ξ ,

$$\lim_{\xi \rightarrow \infty} J_0(\xi) = 0,$$

and the term R_{\pm} in Equation (5.3.10) is dominated by ξ , which means

$$\lim_{\xi \rightarrow \infty} \frac{J_0(R_+) + J_0(R_-)}{2} = 0.$$

Taking the limit of Equation (5.3.14) as ξ tends to infinity,

$$\lim_{\xi \rightarrow \infty} \mathcal{D}_{I,R}(q, \Delta t) = 2\mathcal{D}_R(q, \Delta t).$$

That means for sufficiently large ξ , we can once more ignore the shadow during fitting. This demonstrates that there are other conditions on the input parameters of the studied data which can justify parameter fitting without considering the DIC shadow, particularly for larger ξ . Although we are not currently aware of any data which is equivalent to DIC images with a large shear, such data may one day be available, which justifies exploring the behaviour of DDM analysis in larger shear datasets.

These challenges, collectively, can be summarised as the requirement to determine when we can fit the non-DIC ISF in DDM analysis of a DIC dataset, and to quantify the fitting error we expect when doing so. We will briefly explore these challenges further in Section 5.4, once again through the simple example of a directed advection-diffusion problem.

5.4 Exploring the use of the ‘bright-field’ ISF in DDM analysis of DIC images

In this section, we aim to describe when we can accurately fit the diffusion coefficient D and the velocity magnitude v from phase-contrast images of an advection-diffusion process, when ignoring shadowing effects. Furthermore, when parameter fitting becomes inaccurate, we aim to explore whether we can predict and quantify this inaccuracy, and therefore still achieve accurate fitted movement statistics through post-processing. We could explore this problem through simulation, by generating DDM matrices using randomly chosen values for D , \mathbf{v} and \mathbf{s} , from which we may fit the diffusion coefficient and velocity using the non-DIC ISF. The random variation in particle trajectories, however, would be a confounding variable in our analysis. Additionally, generating and analysing DDM matrices is slow. We therefore present a formal analytical description of parameter fitting which allows us to simplify our exploration.

The objective of the first stage of DDM parameter fitting, as introduced in Section 2.2.7, is to determine the values of a set of variables which minimise the distance between some model curve, given by the ISF, and the data available, given by the DDM matrix $\mathcal{D}_R(q, \Delta t)$. These curves are plotted for fixed spatial frequency q , over variable time lag Δt . For a directed advection-diffusion

behaviour, we have previously stated the ISF in Equation (4.3.3), without DIC adjustment, to be

$$f(q, \Delta t) = e^{-q^2 D \Delta t} J_0(qv \Delta t),$$

Considering the ‘noiseless’ definitions of the DDM matrix and ISF in Equation (2.2.18), parameter fitting in images not generated by phase-contrast is equivalent to finding the values of D , v and $A(q)$ which minimise the frequency-dependent residual function

$$E(q) = \int_0^\infty \left| \mathcal{D}_R(\mathbf{q}, \Delta t) - A(q) \left(1 - e^{-q^2 D \Delta t} J_0(qv \Delta t) \right) \right|^2 d\Delta t. \quad (5.4.1)$$

When Equation (2.2.18) describes the DDM matrix, then there exists a ‘correct’ choice of D , v and $A(q)$ for which $E \equiv 0$, and the residual achieves its minimum. When considering DIC images of an advection-diffusion behaviour, we observed the DDM matrix $\mathcal{D}_I(q, \Delta t)$ was described by Equation (5.3.14) instead of Equation (2.2.18). Despite this, in Section 5.3.2 we showed that when the DIC shear is small, we can approximate the diffusion coefficient and velocity magnitude by fitting $\mathcal{D}_{I,R}$ to the ‘incorrect’ ISF model in Equation (2.2.18). This means that the same values of D and v that minimise the residual in Equation (5.4.1) would additionally minimise the residual

$$E_I = \int_0^\infty \left| \mathcal{D}_{I,R}(\mathbf{q}, \Delta t) - A(q) \left(1 - e^{-q^2 D \Delta t} J_0(qv \Delta t) \right) \right|^2 d\Delta t. \quad (5.4.2)$$

where $A(q)$ in Equation (5.4.2) may not be equal to $A(q)$ in Equation (5.4.1). It is important to note that E_I will not be equal to zero for the ‘correct’ choice of D and v , but will be minimised at some non-zero value. Critically, we have found two possible fitting functions which, when the small DIC shear condition is met, minimise their distance to the DDM matrix at the same values of D and v . This gives an alternative interpretation of the validity of ignoring phase-contrast shadowing effects in DDM analysis; we can fit the DDM matrix using the non-DIC ISF if both E and E_I are minimised by the same values of D and v .

This alternative expression of parameter fitting allows us to numerically explore the validity of ignoring the phase-contrast shear, without needing to generate DDM matrices. Substituting the definition of $\mathcal{D}_{I,R}$ from Equations (5.3.14) and (5.3.15) into Equation (5.4.2),

$$E_I = \int_0^\infty \left| 4|\tilde{X}|^2 (1 - f_I(q, \Delta t)) - A(q) (1 - f(q, \Delta t)) \right|^2 d\Delta t.$$

Taking the constant $4|\tilde{X}|^2$ outside of the integral,

$$\begin{aligned} E_I &= 16|\tilde{X}|^4 \int_0^\infty \left| (1 - f_I(q, \Delta t)) - \frac{A(q)}{4|\tilde{X}|^2} (1 - f(q, \Delta t)) \right|^2 d\Delta t. \\ &= 16|\tilde{X}|^4 \int_0^\infty \left| (1 - f_I(q, \Delta t)) - \hat{A} (1 - f(q, \Delta t)) \right|^2 d\Delta t. \end{aligned} \quad (5.4.3)$$

where for simplicity we define

$$\hat{A} = \frac{A(q)}{|\tilde{X}|^2}.$$

The constant outside of the integral does not affect the minimisation problem and may therefore be discarded. Our ‘data’ to be fit, therefore, does not have any scaling term, whilst our model function is scaled by the constant \hat{A} . We can therefore find the parameters D and v which minimise E_I by instead minimising Equation (5.4.3), which does not require the DDM matrix D_I . The parameters we need to assign in order to plot $f_I(q, \Delta t)$ are D , v , s , q and Φ , as well as choosing an interval for the time lag Δt . We can remove the frequency and time lag consideration, and hence further simplify our exploration, by using the non-dimensional parameter groups introduced in Section 5.3, namely

$$\kappa = q^2 D \Delta t, \quad \lambda = qv \Delta t, \quad \xi = qs.$$

If the maximum time-lag simulated is Δt_{\max} , we can define the following parameters,

$$\kappa_{\text{in}} = q^2 D \Delta t_{\max}, \tag{5.4.4}$$

$$\lambda_{\text{in}} = qv \Delta t_{\max}, \tag{5.4.5}$$

such that we may consider time lag to belong to the unit interval, $\Delta t \in [0, 1]$. Hence, we only need to randomly generate κ_{in} , λ_{in} , ξ and Φ . We simulate 10,000 simulated ISFs $f_I(q, \Delta t)$, where parameters are drawn from the following uniform distributions

$$\kappa_{\text{in}} = q^2 D \sim U(0, 10), \tag{5.4.6}$$

$$\lambda_{\text{in}} = qv \sim U(0, 10), \tag{5.4.7}$$

$$\xi = qs \sim U(0, 30), \tag{5.4.8}$$

$$\Phi \sim U(0, \pi/2). \tag{5.4.9}$$

We then perform non-linear least squared curve fitting to identify the parameters κ_{out} , λ_{out} and \hat{A} which minimise the distance between $1 - f_I(q, \Delta t)$ and the model function

$$\hat{A} (1 - e^{\kappa_{\text{out}} \Delta t} J_0(\lambda_{\text{out}} \Delta t)).$$

We plot the pairwise distributions of our four input parameters (Equations (5.4.7)-(5.4.9)) to our two output parameters κ_{out} and λ_{out} in Figure 5.16. We also plot the distribution of the scaling parameter \hat{A} against ξ in Figure 5.17. The scaling term \hat{A} very strongly approximates the curve $1 - J_0(\xi)$, which would be the

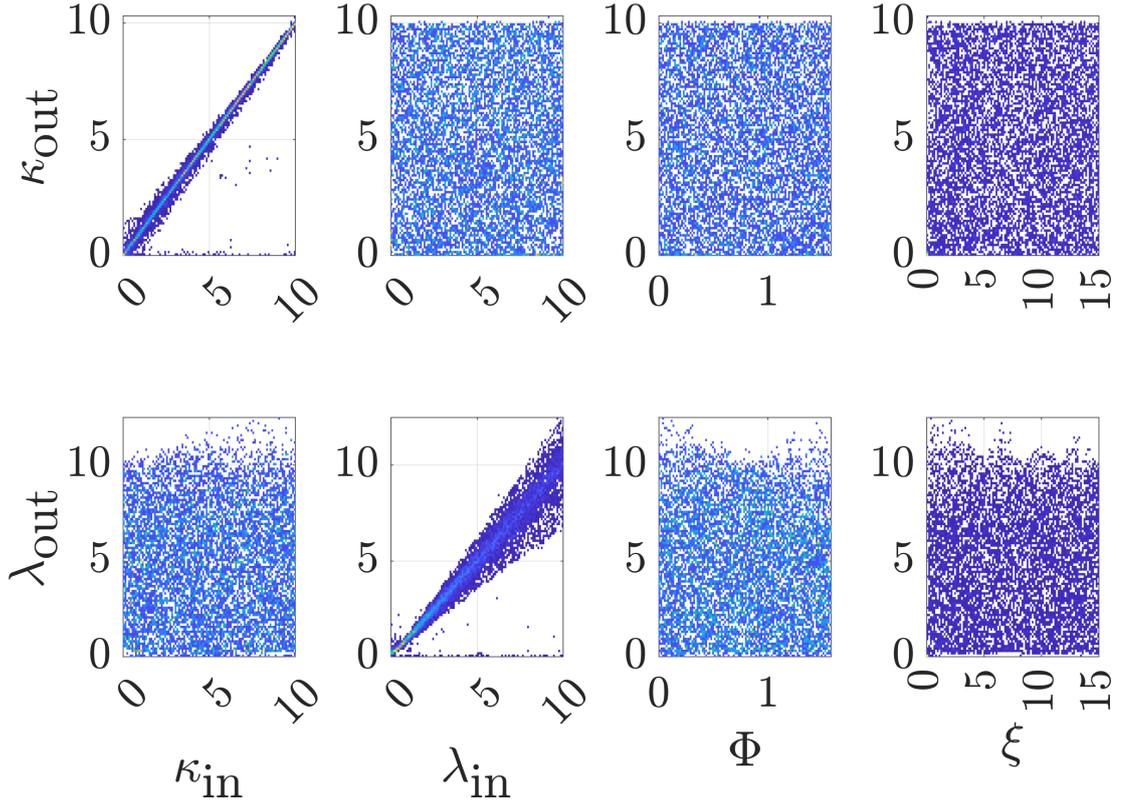


Figure 5.16: For 10,000 simulated instances of the DIC ISF, $f_I(q, \Delta t)$, with randomly generated input parameters (Equations (5.4.7)-(5.4.9)), we fit the output parameters κ_{out} and λ_{out} by minimising the residual function E_I in Equation (5.4.3).

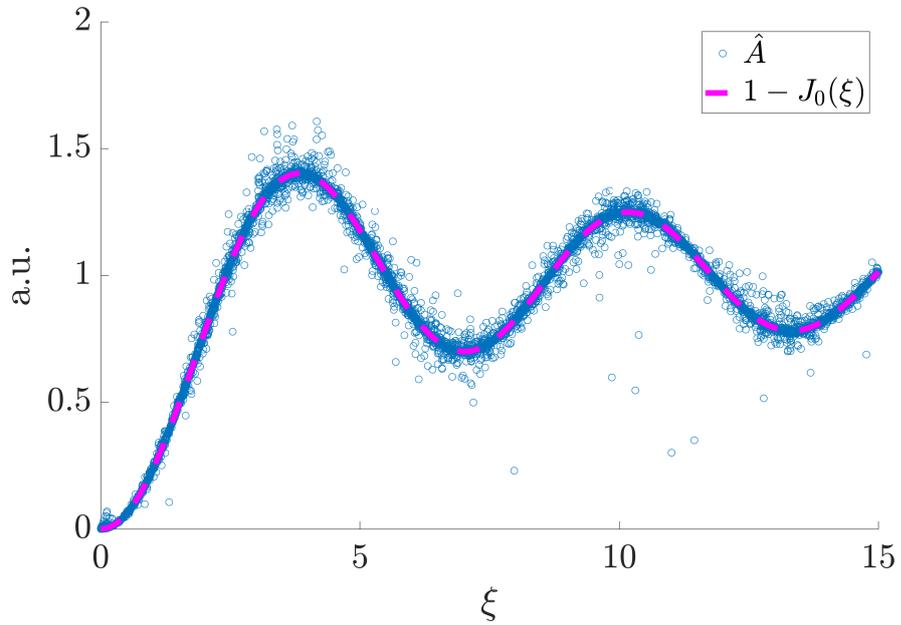


Figure 5.17: Scatter plot values of the fitted scaling term \hat{A} , alongside the curve $1 - J_0(\xi)$.

expected value \hat{A} takes when Equation (5.3.19) holds under the small-shear approximation.

There is a linear correlation between κ_{out} and the input parameter κ_{in} , as well as between λ_{out} and λ_{in} . Furthermore, κ_{out} and λ_{out} depend only on their corresponding input parameter, being uniformly distributed with respect to each other parameter. In particular, we note that the spread about the linear $\kappa_{\text{out}} = \kappa_{\text{in}}$ curve is greatest for small values of κ_{in} , whilst the opposite is true for the relationship between λ_{out} and λ_{in} , suggesting that this fitting regime is optimal in low Péclet number regimes (see Section 4.3).

It is also illuminating to study the error in our parameter estimation, i.e. $\kappa_{\text{out}} - \kappa_{\text{in}}$ and $\lambda_{\text{out}} - \lambda_{\text{in}}$, shown in Figure 5.18.

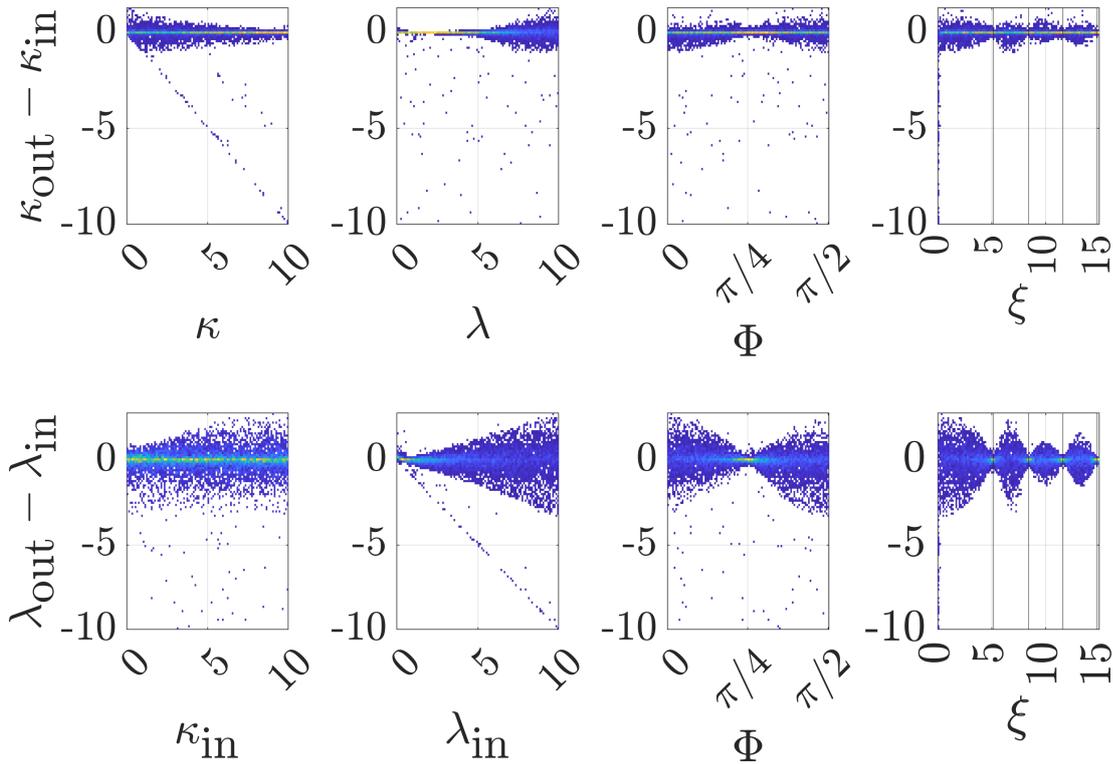


Figure 5.18: Paired distribution histograms of estimation errors $\kappa_{\text{out}} - \kappa_{\text{in}}$ and $\lambda_{\text{out}} - \lambda_{\text{in}}$, against the input parameters κ_{in} , λ_{in} and Φ . On the right-most plots where ξ is on the horizontal axis, vertical bars denote the values of ξ for which $J_2(\xi) = 0$, coinciding with the minimisation of both estimation errors.

Again, we see that the optimal regime for fitting both κ_{out} and λ_{out} coincides with small λ_{in} and larger κ_{in} , but there are several new features we will highlight. The first is that there is a clear subset of simulations for which κ_{out} and λ_{out} scale linearly with the negative of their respective input parameter, representing simulations where either κ_{out} or λ_{out} are fitted as zero incorrectly. Comparing with the right-hand graphs where error is plotted against the value of ξ , we see that these points occur when $\xi \approx 0$. This validates our claim that for excessively small ξ , numerical instability prevents us from achieving a good approximation

of the diffusion coefficient and velocity magnitude. These points appear to occur only when $\xi < 0.3$.

For $\xi > 0.3$, the diffusion error $\kappa_{\text{out}} - \kappa_{\text{in}}$ is low for all ξ , distributed mostly between ± 1 . The spread of the velocity error term, $\lambda_{\text{out}} - \lambda_{\text{in}}$, is initially very high at low ξ , however. Comparing to the error distribution with respect to λ_{in} and Φ , we can see that the combination of high λ_{in} , low ξ and Φ being close to 0 or $\pi/2$ result in poor quality fitting. The wide spread of error for high λ_{in} is likely a product of reduced sensitivity, but the high error variation even at small ξ has interesting implications for our small-shear approximation. This does not contradict our conclusion from Section 5.3, that a small-shear approximation justifies ignoring the DIC shear; instead, high spread of fitting error implies the range of ξ for which the approximation does hold may be highly restricted.

We additionally note that the error in both λ_{out} and κ_{out} decreases to almost 0 when $J_2(\xi) = 0$, shown by black vertical bars in right-most plots in Figure 5.18. This unexpected observation, which is significantly more pronounced in the fit of λ_{out} than κ_{out} , is not predicted by the small-shear approximation argument in Section 5.3. This feature will be explained in Section 5.4.2.

Away from the regions where $J_2(\xi) = 0$, we finally observe that as ξ increases, the spread of the error in λ_{out} decreases. We previously discussed that as $\xi \rightarrow \infty$, $\mathcal{D}_{I,R} \rightarrow 2\mathcal{D}_{X,R}$. Our numerical exploration appears to suggest that this relationship may hold for ξ as low as 10.

In conclusion, our numerical exploration has allowed us quantify the parameter regions for which we can accurately approximate the diffusion coefficient and velocity magnitude when ignoring the DIC shadow. This region is characterised by the following conditions:

- High κ , particularly $\kappa > 3$.
- Low λ , particularly $\lambda < 5$.
- $\Phi \approx \pi/4$.
- ξ chosen such that $J_2(\xi) \approx 0$, or $\xi > 10$.

We additionally may conclude that the spread of the fitting error for the velocity is maximised when $\xi < 5$, Φ is close to 0 or $\pi/2$, and λ is large. This may imply that the small-shear approximation that justifies Equation (5.3.19) may only be valid in very small parameter region, outside of which, fitting the non-DIC ISF to \mathcal{D}_I induces error. Alternatively, the poor accuracy with which we fit λ in this region may be a result of λ being very large, resulting in rapid decay of the ISF to zero and a resulting numerical instability in the fitting algorithm. Understanding exactly how the choice of λ and Φ affects the small shear

approximation is an important consideration if the simplified form of the ISF is to be used in practice.

We now show that the fitting error is predictable for some combinations of parameters, which may allow us to reclaim the correct velocity, even when the small-shear approximation does not hold.

5.4.1 Finding parameter regions where error in the fitted velocity is predictable

In order to prove that fitting error is predictable for some combinations of input parameters, we will first need to simplify the residual E_I in Equation (5.4.3). We observe that κ_{out} and \hat{A} are much easier to fit than λ_{out} . We propose that this is because these parameters both control a unique aspect of the shape of the ISF, such that they may almost always be approximated regardless of the input parameter. The scaling term, \hat{A} , controls the long term behaviour of the ISF, and therefore can be fit so long as the velocity and diffusion parameters are sufficiently high that the ISF decays to zero.

Meanwhile, κ_{out} uniquely determines the rate of convergence to the long-term behaviour of the ISF. For sufficiently large λ , the Bessel function of the first kind of order 0 is approximated by [324].

$$J_0(\lambda) \approx \sqrt{\frac{2}{\pi\lambda}} \left(\cos \left(\lambda - \frac{\pi}{4} \right) \right).$$

Hence, the rate at which $J_0(\lambda)$ converges to 0 with increasing Δt is approximately $1/\sqrt{\lambda}$, which is slower than the rate of decay of $e^{-\kappa}$ unless $v \gg D$. In general, the diffusion coefficient strongly determines the rate of convergence of the DDM matrix, which helps us to fit κ_{out} well regardless of input parameters.

If the fit of κ_{out} and \hat{A} is invariant of our input parameters, we can make the assumption that these parameters are known and constant, in order to further simplify the residual F_I to explore how the final parameter λ_{out} will be fit. Assuming $\kappa_{\text{out}} = \kappa_{\text{in}}$ and $\hat{A} = 1 - J_0(\xi)$, Equation (5.4.3) may be expanded as follows,

$$E_I = 16|\tilde{X}|^4 \int_0^\infty \left| \left(1 - e^{-\kappa_{\text{in}}\Delta t} J_0(\lambda_{\text{in}}\Delta t) - J_0(\xi) + \frac{e^{-\kappa_{\text{in}}\Delta t}}{2} (J_0(R_+) + J_0(R_-)) \right) - (1 - J_0(\xi))(1 - e^{-\kappa_{\text{in}}\Delta t} J_0(\lambda_{\text{out}}\Delta t)) \right|^2 d\Delta t,$$

which may be simplified to

$$E_I = 16|\tilde{X}|^4 \int_0^\infty e^{-2\kappa_{\text{in}}\Delta t} \left| (J_0(\lambda_{\text{out}}\Delta t) - J_0(\lambda_{\text{in}}\Delta t)) + \left(\frac{J_0(R_+) + J_0(R_-)}{2} - J_0(\lambda_{\text{out}}\Delta t)J_0(\xi) \right) \right|^2 d\Delta t, \quad (5.4.10)$$

Only λ_{out} remains to be fit, but determining the value of λ_{out} that minimises this residual is not a trivial task, because the terms inside the absolute value are all oscillatory. Pairing the terms as we have in the brackets of Equation (5.4.10) feels natural; the first bracket is equal to zero when $\lambda_{\text{in}} = \lambda_{\text{out}}$, and the remaining residual is equal to the error of the approximation in Equation (5.3.19). In practice, however, we cannot treat these brackets separately because positive contributions from one bracket can balance negative contributions from the other, leading to the possible existence of multiple local minimum points of the residual.

Despite the complexity of Equation (5.4.10), we will demonstrate that for some input parameter values, we can predict the error with which we fit λ_{out} . In particular, we assume the error can be described by some multiplicative constant that relates λ_{in} and λ_{out} , namely

$$\lambda_{\text{out}} = \gamma\lambda_{\text{in}},$$

where γ is the ‘adjustment factor’ which predicts the error with which we fit the velocity, such that $\gamma = 1$ corresponds to an accurate parameter fit. Substituting this definition into Equation (5.4.10), the residual can be expressed as

$$E_I = 16|\tilde{X}|^4 \int_0^\infty e^{-2\kappa_{\text{in}}\Delta t} \left| (J_0(\gamma\lambda_{\text{in}}\Delta t) - J_0(\lambda_{\text{in}}\Delta t)) + \left(\frac{J_0(R_+) + J_0(R_-)}{2} - J_0(\gamma\lambda_{\text{in}}\Delta t)J_0(\xi) \right) \right|^2 d\Delta t, \quad (5.4.11)$$

We study an asymptotic expansion of Equation (5.4.11), that allows us to identify particular cases where γ is guaranteed to be not equal to 1. We first consider an asymptotic expansion of $J_0(R_\pm)$, given by [347],

$$J_0(\sqrt{\lambda^2 + \xi^2 - 2\lambda\xi \cos(\Phi)}) = J_0(\lambda)J_0(\xi) + \sum_{n=1}^{\infty} 2J_n(\lambda)J_n(\xi) \cos(n\Phi). \quad (5.4.12)$$

Noting that Bessel functions of the first kind of odd order are themselves odd,

$$\frac{1}{2}(J_0(R_+) + J_0(R_-)) = J_0(\lambda)J_0(\xi) + \sum_{n=1}^{\infty} 2J_{2n}(\lambda)J_{2n}(\xi) \cos(2n\Phi). \quad (5.4.13)$$

We also consider an asymptotic expansion of $J_0(\gamma\lambda)$ [324],

$$J_0(\gamma\lambda) = \sum_{n=0}^{\infty} \frac{1}{n!} \left(\frac{(1-\gamma^2)\lambda}{2} \right)^n J_n(\lambda), \quad (5.4.14)$$

valid when $|\gamma^2 - 1| < 1$. The $n = 0$ term of (5.4.14) is equal to $J_0(\lambda)$. Substituting Equations (5.4.14) and (5.4.13) into the residual in Equation (5.4.11), we define an asymptotic expansion of the residual as

$$E_I = \int_0^{\infty} e^{-2\kappa_{\text{in}}\Delta t} \left| \sum_{n=1}^{\infty} 2J_{2n}(\lambda_{\text{in}}\Delta t) J_{2n}(\xi) \cos(2n\Phi) + \frac{1}{n!} \left(\frac{(1-\gamma^2)\lambda_{\text{in}}\Delta t}{2} \right)^n J_n(\lambda_{\text{in}}\Delta t) (1 - J_0(\xi)) \right|^2. \quad (5.4.15)$$

We therefore have an expansion of the residual in terms of Bessel functions of order greater than 0, where changing the value of γ can increase or decrease the residual. Using this asymptotic expansion, we may prove that specific combinations of parameters will result in fitting γ incorrectly. In particular, we consider two cases where $\Phi = 0$ and $\Phi = \pi/2$

5.4.1.1 Case 1: $\Phi = 0$

Let the k^{th} positive zero of the Bessel function of order n be denoted by $j_{n,k}$. Let $\xi < j_{2,1}$ and $\lambda < j_{0,1}$. Since the Bessel function zeros are monotone increasing for increasing order n , we can say that

$$J_{2n}(\lambda) J_{2n}(\xi) > 0.$$

Hence, the signs of the terms in the expansion in Equation (5.4.13) are defined by $\cos(2n\Phi)$. If we take $\Phi = 0$, every term in Equation (5.4.13) is positive over the region of λ observed. Meanwhile, considering the sum,

$$\sum_{n=1}^{\infty} \frac{1}{n!} \left(\frac{(1-\gamma^2)\lambda}{2} \right)^n J_n(\lambda) (1 - J_0(\xi)).$$

When $\gamma = 1$, these terms are identically equal to zero. We note that if $\gamma > 1$, then $1 - \gamma^2$ is negative, whereas for $\gamma < 1$, $1 - \gamma^2$ is positive. Therefore, we may use γ to tune not only the size of each term in the sum, but also the sign. If we can make the terms of the sum negative, these would reduce the positive terms from the first sum and therefore result in a smaller residual. Unfortunately, this is not immediately possible; although $J_n(\lambda)(1 - J_0(\xi)) > 0$ when $\lambda < j_{1,1}$, such

that for $\gamma > 1$ and odd n the terms of this sum are negative, all even n terms are positive. However, we can instead show that the sum of each pair of consecutive terms is negative, by showing that the absolute size of each term is monotone decreasing. The absolute ratio of terms in the sum can be written as

$$\left| \frac{\frac{1}{(n+1)!} \left(\frac{(1-\gamma^2)\lambda}{2} \right)^{n+1} J_{n+1}(\lambda) (1 - J_0(\xi))}{\frac{1}{n!} \left(\frac{(1-\gamma^2)\lambda}{2} \right)^n J_n(\lambda) (1 - J_0(\xi))} \right| = \left| \frac{1}{n+1} \left(\frac{(1-\gamma^2)\lambda J_{n+1}(\lambda)}{2J_n(\lambda)} \right) \right|.$$

By the recurrence relation in Equation (5.3.20),

$$\frac{2J_n(\lambda)}{\lambda} = \frac{J_{n-1}(\lambda) + J_{n+1}(\lambda)}{n},$$

which is permissible since neither n nor λ is equal to zero. Hence,

$$\left| \frac{1}{n+1} \left(\frac{(1-\gamma^2)\lambda J_{n+1}(\lambda)}{2J_n(\lambda)} \right) \right| = \left| \frac{n}{n+1} \right| |1-\gamma^2| \left| \frac{J_{n+1}(\lambda)}{J_{n-1}(\lambda) + J_{n+1}(\lambda)} \right|.$$

The expansion in Equation (5.4.14) required that $|1-\gamma^2| < 1$. Since $\lambda < j_{0,1}$, then $J_{n-1}(\lambda)$ and $J_{n+1}(\lambda)$ are both positive, and hence

$$\frac{J_{n+1}(\lambda)}{J_{n-1}(\lambda) + J_{n+1}(\lambda)} < 1$$

The ratio of terms is therefore strictly monotone decreasing. Hence, each consecutive pair of terms in the sum is negative, such that the whole sum is negative. If we therefore take $\gamma > 1$, the contribution of these negative terms will cancel some of the positive terms from the first sum. As such, the residual is minimised for some value of γ strictly less than 1, in the parameter region denoted by

$$\lambda < j_{0,1}, \quad \xi < j_{2,1}, \quad \Phi = 0,$$

5.4.1.2 Case 2: $\Phi = \pi/2$

When $\Phi = \pi/2$, the expansion in Equation (5.4.13) becomes an alternating series,

$$\sum_{n=1}^{\infty} J_{2n}(\lambda) J_{2n}(\xi) \cos \left(2n \left(\frac{\pi}{2} \right) \right) = \sum_{n=1}^{\infty} (-1)^n J_{2n}(\lambda) J_{2n}(\xi).$$

By a similar argument as the $\Phi = 0$ case, we can show that pairs of terms in the series expansion of h are negative. When $\lambda < j_{0,1}$ and $n \geq 1$,

$$\left| \frac{J_{n+1}(\lambda)}{J_n(\lambda)} \right| < 1.$$

Recursively, this implies

$$\left| \frac{J_{n+2}(\lambda)}{J_n(\lambda)} \right| < 1,$$

and thus

$$\left| \frac{J_{2(n+2)}(\lambda)}{J_{2n}(\lambda)} \right| < 1.$$

This result holds for $\lambda < j_{1,1}$. As such, when $\lambda, \xi < j_{1,1}$,

$$\left| \frac{J_{2(n+1)}(\lambda) J_{2(n+1)}(\xi)}{J_{2n}(\lambda) J_{2n}(\xi)} \right| < 1,$$

which means successive pairs of terms in the expansion of Equation (5.4.13) from $n = 1$ sum to a negative value, since the first term in the sum is negative. When $\lambda < j_{1,1}$, choosing $\gamma < 1$ ensures that every term in the expansion in Equation (5.4.14) yields only positive terms in the sum, allowing us to reduce the residual. Hence, it is beneficial to under-estimate the velocity (take $\gamma < 1$) when

$$\lambda < j_{1,1}, \quad \xi < j_{1,1}, \quad \Phi = \frac{\pi}{2}.$$

5.4.2 Estimating the error factor γ

We have shown that when $\Phi = 0$ or $\Phi = \pi/2$, we can supply conditions on the size of λ and ξ which cause the velocity term to be over or under-estimated, respectively. This shows that for some input parameter combinations, the existence of error in the fit of the velocity is predictable. We now pose the question of whether the size of this error is also predictable, which would allow for the extraction of accurate velocity through post-processing.

We can approximate the value of γ which minimises Equation (5.4.15) by considering only leading order terms. Immediately, this explains why the velocity term is fit perfectly when $J_2(\xi) \approx 0$, as observed in Figure 5.18. When $J_2(\xi) \approx 0$, the smallest order term in the left hand sum in Equation (5.4.15) is $J_4(\lambda_{\text{in}} \Delta t)$. Hence, the leading order terms in Equation (5.4.15) all belong to the second sum. Setting $\gamma = 1$ removes all leading order terms up to $J_4(\lambda_{\text{in}} \Delta t)$, and therefore minimises the residual.

When $J_2(\xi) \neq 0$, we propose to approximate γ by identifying the value it must take so that the $J_1(\xi)$ terms are equal to zero. Using the recurrence relation in Equation (5.3.20), we can rewrite

$$2J_2(\lambda)J_2(\xi) \cos(2\Phi) = \frac{\lambda}{2} (J_1(\lambda) + J_3(\lambda)) J_2(\xi) \cos(2\Phi).$$

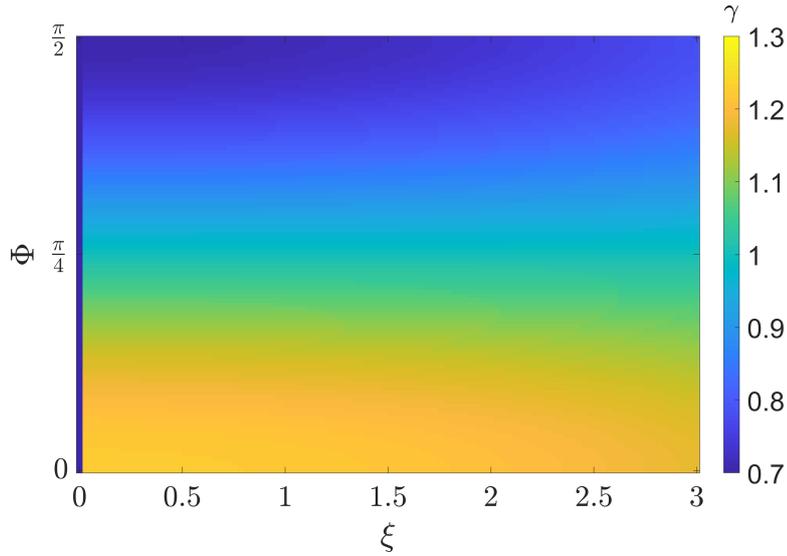
Hence, we can set all the $J_1(\xi)$ terms to be equal to zero in the residual to determine

$$\gamma = \sqrt{\frac{1 - J_0(\xi) + J_2(\xi) \cos(2\Phi)}{1 - J_0(\xi)}}. \quad (5.4.16)$$

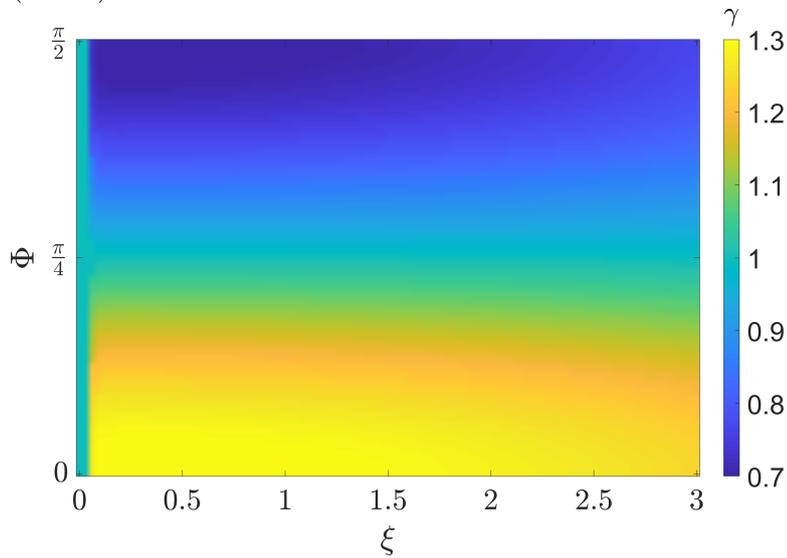
In Figure 5.19a, we plot this predicted value of γ for $\xi \in [0, j_{1,1}]$, such that the conditions for both cases in Sections 5.4.1.1 and 5.4.1.2 should hold. Additionally, we perform numerical fitting for these combinations of ξ and ϕ to determine the value of γ which minimises the residual E_I in Equation (5.4.11) before taking the asymptotic expansion, plotted in Figure 5.19b, for $\kappa = 0$. There is excellent agreement between the predicted and numerically determined values of γ in Figures 5.19a and 5.19b, respectively. This means that for some of the parameter combinations where the velocity is fit poorly when using the non-DIC fitting model, we can estimate the fitting error. Following post-processing, we can still therefore determine the velocity magnitude v for these parameter spaces.

To conclude, we investigated the validity of fitting advection-diffusion parameters from the DDM matrix $\mathcal{D}_{I,R}$ when using the non-DIC ISF, rather than the DIC-adjusted form in Equation (5.3.15). We have highlighted in Figures 5.16, 5.17 and 5.18 that small λ , $\Phi \approx \pi/4$ and $J_2(\xi) \approx 0$ are conditions for which the diffusion coefficient and velocity magnitude are fit with high accuracy. However, when these conditions are not met, the velocity magnitude can be poorly fit. We have shown that for a limited parameter space, namely $\lambda < j_{0,1}$ and $\xi < j_{1,1}$, the velocity fitting error can be predicted, which means that post-processing could restore the ‘correct’ velocity magnitude.

Previously in Section 5.3, we proposed that we can fit $\mathcal{D}_{I,R}$ using the non-DIC ISF if the approximation in Equation (5.3.19) holds. However, Figures 5.16 and 5.18 show that conditions on λ and Φ may mean that Equation (5.3.19) may only provide an approximation of the desired fitting parameters, even when ξ is small, with some parameter regions yielding a constant and predictable error in the fit of the velocity term. However, if this error can be predicted for all λ and ξ , then Equation (5.3.19) is only a sufficient, not necessary, condition to justify using the non-DIC ISF during parameter fitting. Instead, the justification for using the simplified ISF is derived from the knowledge that whilst the residual between the DDM matrix and fitting function is minimised at some non-zero value, the output from parameter fitting can be used to identify the true velocity and diffusion coefficients, even if the output parameters are, themselves, not an accurate estimate of their true counterparts. Although we have shown that it is possible to predict the true velocity from the fitted velocity in some parameter regions, we cannot claim that this is always true; predicting the error is a difficult task, because the residuals in Equation (5.4.11) are non-linear, but this remains an interesting direction for future work.



(a) Predicted scaling values γ calculated by Equation (5.4.16).



(b) Numerically determined values of γ which minimise the residual in Equation (5.4.11).

Figure 5.19: For the limited parameter region $\lambda < j_{0,1}$ and $\xi < j_{1,1}$, the predicted value of the velocity scaling term γ (Figure 5.19a) aligns with the numerically determined scaling term (Figure 5.19a).

5.5 Discussion

When applying image analysis techniques to time-series images, we make the unspoken assumption that the results of our analysis depend exclusively on the underlying movement within the movie. This is not always true; for example, we know that when the scale of the displacement in the image stack is greater than the gradient of the heterogeneous background illumination in dark field imaging, different fitting functions are potentially required for accurate analysis in DDM [292]. It is important to understand how features of the technique used to generate

images could affect subsequent analysis, either to improve our confidence that our analysis is accurate, or to identify potential problems and mitigate them.

In this chapter, we considered Differential Interference Contrast (DIC) microscopy (and thus, due to their similarity, Hoffman Modulation Contrast [334]), due to its application in oocyte imaging and use in IVF clinics. We identify that the characteristic shadowing in these microscopy settings introduces anisotropy into the DDM tensor (Figure 5.3), potentially violating assumptions in DDM that enable the use of the azimuthal average [164, 266]. We thus set out to determine exactly what effect shadowing has on DDM analysis.

By presenting an approximation of DIC images as the difference between spatially shifted pairs of identical phase images, X and Y , we determined a new expression for the DDM tensor of a DIC image stack, $\mathcal{D}_I(\mathbf{q}, \Delta t)$ in (5.2.13). For isotropic motion, Equation (5.2.21) gives a time-independent map between the DDM matrix on the DIC image, $\mathcal{D}_{I,R}(q, \Delta t)$, and the DDM matrix on the phase image, $\mathcal{D}_{X,R}(q, \Delta t)$, following an azimuthal average. This map, validated in simulated and real data in Section 5.2.1, implies that the shadowing effect in DIC imaging has no effect on fitting results when undertaking DDM analysis, since it can be accounted for by a q -dependent scaling term.

This powerful result ensures that existing DDM implementations are immediately useable in DIC image stacks without adaption, so long as the underlying motion is known to be isotropic, allowing confidence that the output statistics from DDM analysis match the underlying dynamics observed. Furthermore, an extension of this conclusion in Section 5.2.2 demonstrates attempting to adapt the technique by accounting for the anisotropy in the DDM tensor will likely reduce the accuracy of subsequent analysis, due to the weak dependence of $\mathcal{D}_{I,R}$ on the DIC shear and the increased effect of noise when foregoing the use of a radial average.

When we no longer assume isotropy, however, the picture becomes much less clear. Section 5.3 demonstrates that the ISF for anisotropic processes is dependent on the DIC shear, and not exclusively on the underlying motion. We derive a new specific form the ISF for advection-diffusion which can be used for parameter fitting in Equation (5.3.15), verifying in Figure 5.13 that its use in simulations yields approximately correct solutions. However, this new ISF has additional non-linear dependencies on the velocity term compared to the ISF expected in non-DIC images, resulting in diminished numerical stability during fitting.

Users of DDM who are unfamiliar with the effect of phase-contrast shadowing on DDM would perform fitting using the non-DIC ISF, rather than the new ISF model we propose in this work. Whilst such an approach might initially appear misguided, we identified a condition in Section 5.3 for which approximations of

the correct diffusion coefficient and velocity magnitude can be determined when using this overly simple approach. Namely, we showed that the product integral in Equation (5.3.6) can be approximated as the product of two contour integrals (i.e. Equation (5.3.19)) when the shear distance s is much smaller than the distance over which particles travel by advection. Under this assumption, which is realistic in DIC imaging, we have shown in Figure 5.15 that it is possible to approximate D and v in this way.

Whilst Figure 5.15 is a working example, we showed in Figure 5.18 that although ignoring the shadowing effect during parameter fitting allows us to estimate the diffusion coefficient well, the error in the velocity fit depends not only on the input velocity, but also on Φ and ξ , as predicted by analysis of the Taylor approximation error in Equation (5.3.24). We have therefore made some efforts to understand the exact nature of this dependence in Section 5.4. By considering an asymptotic expansion of the residual error during this fitting process, we identified that for a restricted parameter space, we could not only predict whether we would under or over-estimate the velocity magnitude (Cases 5.4.1.1 and 5.4.1.2), but that we could also predict how large this error would be. This means that, for this restricted parameter region, post-processing would allow us to extract the exact velocity magnitude, despite the error we encounter by using the ‘incorrect’ ISF.

The phase-contrast shadow in anisotropic motion is a complex consideration. There are some datasets where, due to the combination of input parameters which define the studied images, fitting without consideration of the shadowing effect will yield the correct movement statistics. However, there are also datasets where such fitting will yield incorrect estimates of the velocity v . Although the velocity fit may develop some additional error when using the non-DIC ISF, doing so reduces the complexity of the model function and improves the stability of parameter fitting algorithms, and may therefore be beneficial. Neither approach is ideal, since one must choose to have either a numerically unstable approximation of the true velocity, or a stable fitting algorithm which can return the incorrect movement parameters. Until we have more tangible guidelines which can inform us of where this ‘simplified’ approach is suitable, we would recommend using the full, DIC adjusted form of the ISF in fitting to avoid error.

We are the first to explore DDM parameter fitting in the context of DIC, highlighting issues which were previously unexpected. However, much work remains to be done to fully understand how parameter fitting performs in DIC images. In the context of advection-diffusion, we have primarily focussed on determining whether the ‘incorrect’ non-DIC ISF can be used to extract the true velocity and diffusion coefficient from DIC image data. Based on our recommendation that

the DIC adjusted ISF should be used preferentially, we have yet to explore the effect of different input parameters on the performance of the DIC-adjusted ISF. It is important to quantify how the interaction between λ , ξ and Φ might affect the accuracy of parameter fitting, so that we can ensure the validity of DDM analysis in real data. For example, if we discover that convergence to the true velocity is slower in the DIC-adjusted ISF, we can propose guidelines to increase the number of iterations used by curve-fitting algorithms, or to reduce the error tolerance to improve fitting quality.

A key limitation of this work is that we only address one specific form of anisotropic movement, namely diffusion paired with a constant, directed velocity field. This particular type of motion is even more simple than the typical advection-diffusion behaviour studied in DDM literature, which assumes that velocity magnitudes are drawn from an underlying Schulz distribution [142, 163, 164]. Expanding on this analysis for more complicated behaviours may identify even more unpredictable interactions between the imaging mode, and the underlying motion within the movie.

The complex and non-linear interaction between DIC imaging and DDM analysis provides further incentive to pursue machine-learning fitting approaches in DDM. In Chapter 4, we showed that even in behaviour as simple as directed advection-diffusion, parameter fitting already suffers from error and instability, problems which will only become worse with adaption of the ISF to account for DIC shadowing. Furthermore, for other behaviours, adapting the ISF to account for shadowing may not even be analytically possible. All of these problems would be mitigated using the pipeline we proposed in Section 4.4, however, since the shadowing in DIC would be built into the images forming the training dataset. Whilst further exploring the interaction between DIC and DDM would form an interesting direction for future work, this may be unnecessary if a machine-learning approach to DDM analysis is developed successfully.

Chapter 6

Modelling the effect of vitrification protocols on oocyte or embryo cooling rates

Whilst Chapters 2-5 were concerned with the application of DDM in oocyte health assessment, Chapters 6 and 7 instead focus on aspects of cryopreservation. In this chapter, we explore optimal vitrification protocols, concerning the number and arrangement of embryos or oocytes simultaneously placed onto the cryopreservation device. The work discussed in this chapter has been published in the *Journal of Cryobiology* [348], with accompanying code made available at <https://github.com/OstlerT/MultipleEmbryoModels>.

6.1 Cryopreservation and vitrification protocols

We will briefly recapitulate some of the key concepts relayed in the Introduction which relate to this work here. Cryopreservation is used to maintain oocytes or embryos for use in future cycles, by cooling them to the point at which cessation of intracellular metabolic processes occurs. Freezing is typically lethal to most living systems for many reasons [217–219]. The formation of ice crystals is the predominant cause of cell damage during freezing [220]. Vitrification avoids this by combining high cooling rates [222, 224, 227] with small $< 1\mu\text{l}$ volumes [224, 230, 231] of concentrated high concentration permeating cryoprotectant agents (CPAs) [223, 229]. Together, these cause biological material to enter a glass-like state of thermodynamic non-equilibrium [31, 222, 223]. There are several commercially available cryopreservation devices onto which the oocyte or embryo can be placed [233–235]. The device we focus on in this Chapter is the Open Cryotop[®] [224], which is reported by the manufacturer, Kitazato[®], to be used in over 3000 clinics and egg banks, making it the most widely used vitrification protocol in the world [349]. In particular, this device is also used by LWC, the IVF

clinic we collaborate with. We will briefly outline the protocol for vitrification of oocytes and embryos using the Open Cryotop[®] device, which is shown in Figure 6.1. We note that although LWC uses the Open Cryotop[®] device, they use a proprietary variant of the standard manufactures protocol. Hence, the following description represents a generic approach to vitrification, and may not reflect the exact procedure used in the clinic.

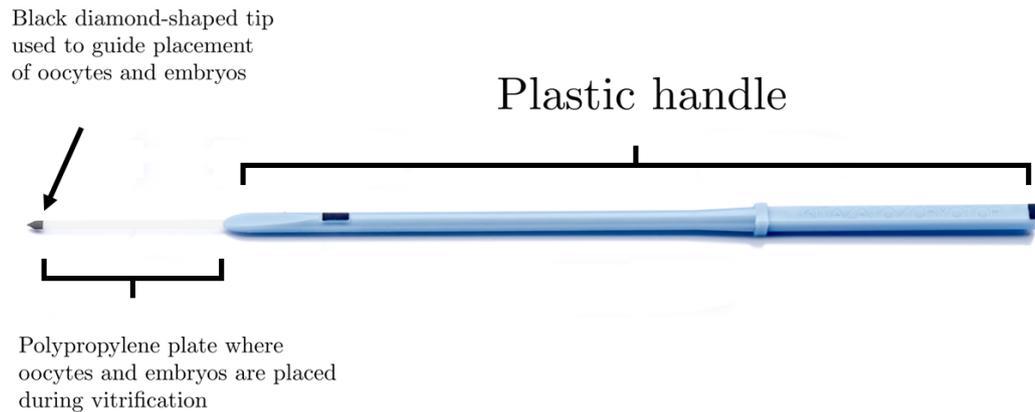


Figure 6.1: A labelled image of the Open Cryotop[®] device used to vitrify oocytes and embryos, modified from the Kitazato website [350]. The blue plastic part on the right hand is the handle of the device. The embryo is mounted onto the thin transparent plastic (polypropylene) plate on the left, just below the black diamond-shaped tip of the device.

In order to perform vitrification, the oocyte or embryo is suspended in a droplet of a concentrated CPA. However, CPAs are toxic, and changing the CPA concentration too quickly results in rapid loss of water from cells, causing damage [351]. A continuous and smooth increase in CPA concentration would be the safest protocol, but this is not feasible in the clinic [351]; hence a two-stage protocol is most commonly followed [224, 226]. Alongside the Cryotop[®] device, Kitazato[®] additionally provides 3 different mixtures, named in the protocol as the ‘basic solution’ (BS), ‘equilibration solution’ (ES) and ‘vitrification solution’ (VS) [248]. Both ES and VS contain CPAs, whilst BS is used to prepare oocytes for vitrification.

The first stage, equilibration [248], is used to bring the concentration of CPA up slowly in the oocyte or embryo. This stage is different for oocytes and embryos; oocytes are transferred to 20 μ l of BS, after which an additional 20 μ l is added immediately. After three minutes, 20 μ l of ES is added (the first introduction

of CPA), followed by $240\mu\text{l}$ of ES at the six minute mark. The oocyte is then left until the 12-15 minute mark, depending on when equilibration is complete; this is when the perivitelline space width becomes equal to the width before immersion in ES. The equilibration stage is more simple for embryos, which are placed directly into $300\mu\text{l}$ of equilibrating solution and left for 10-15 minutes. The protocol is subsequently the same for oocytes and embryos, which from now on, we will refer to as ‘samples’ for simplicity.

The second stage is vitrification. Two wells in a well plate are filled with $300\mu\text{l}$ of VS, denoted VS1 and VS2 respectively. The sample is placed into VS1, and undergoes a repeated process of stirring, aspiration and blowing out to wash out any remaining ES from the previous stage. This process should last no more than 30 seconds, after which the sample is transferred to VS2 and the previous process is repeated for a further 30 seconds. Having reached the required CPA concentration, the sample is loaded onto the Open Cryotop[®] device, pictured in Figure 6.1. The Open Cryotop[®] device consists of a plastic (polypropylene) plate with a depth of 0.1mm and a width of 0.7mm [248], attached to a blue plastic handle. The samples are transferred onto the plastic plate, just to the right of the black diamond-shaped tip. Particular care must be taken with loading to ensure a minimal amount of CPA is transferred alongside the sample, as shown in Figure 6.2. Once excess media is aspirated, the plate and mounted embryo are plunged directly into a bath of liquid nitrogen, to achieve rapid cooling. During equilibration and treatment in the VS solution, embryologists work on a bench top which is warmed to maintain a temperature of 310.15 K (37°C) [352], whilst the liquid nitrogen at standard atmospheric pressure has temperature 77.15 K (-196°C) [242].

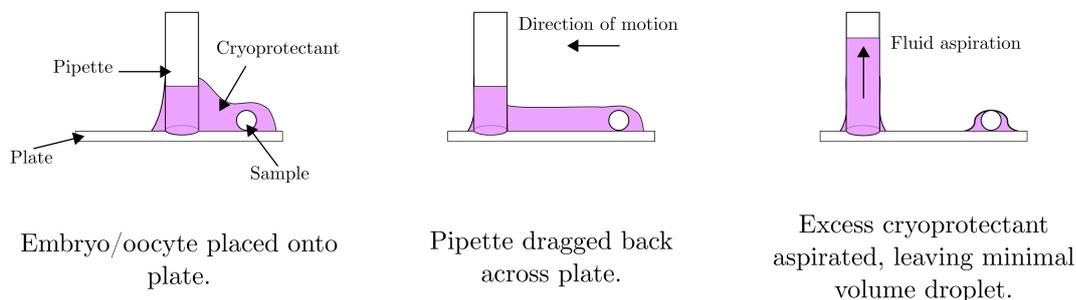


Figure 6.2: Diagram depicting the process of transferring a sample from a dish onto the plate of the Open Cryotop[®], and removing excess cryoprotectant through aspiration to leave a ‘minimal droplet’. Adapted from the Open Cryotop[®] user manual [248].

Whilst the Open Cryotop[®] protocol is very detailed and specific about the timing of each stage, there is flexibility about the number of samples cryopreserved on a single device; up to four samples can be vitrified simultaneously [237]. Since clinical practice is flexible, questions arise about the existence of an optimal protocol. A retrospective study reports that post-thaw survival rates of embryos are affected by the number of embryos vitrified simultaneously, but also states that further work is still required to determine the optimum number, without offering further insight [238].

Another question arises about the optimal arrangement of samples vitrified simultaneously. With more than one sample on the plate, the relative positioning of the samples becomes important; samples positioned closer together are likely to affect the cooling rate of their neighbours more than those spaced far apart. Hence, it is likely that even within a constant number of samples, there exists an optimal arrangement that embryologists should aim for.

It is important to determine whether the variability in vitrification protocols between different clinics, embryologists and even different procedures by the same embryologist affects survival rates for the samples. However, answering this question through mathematical modelling is a complex challenge, due to the significant number of confounding variables and uncertainty relating to survival during vitrification. We can, however, model the effect of different sample arrangements on cooling rates, which, as previously discussed, are strong indicators of survival. We focus on modelling the vitrification protocol for the Cryotop[®] device, as described in the user manual [248].

Some mathematical models of vitrification in the Cryotop[®], as well as other cryopreservation devices, have already been presented in the literature. A series of papers by Santos, Sansinena, Zaritzky and Chirife present a number of mathematical models of cryopreservation, used to estimate model parameters such as heat transfer coefficients [242, 244] and predict cooling rates [243]. In their model, they reproduce the 3D geometry of the cryopreservation device using COMSOL Multiphysics. The CPA is modelled as a stationary fluid. The heat equation is solved numerically, with Robin boundary conditions modelling the relative cooling of the liquid nitrogen. The models predict the temperature at any given point within the device over time, which allows estimating of the cooling rate of the sample. The model is validated by experimental data and the parameter choices are well evidenced [224, 226, 229, 233, 246, 353]. However, none of these models deals with several oocytes or embryos; therefore, our work here is novel.

We also highlight work on modelling the warming process after vitrification [245]. These works use similar techniques to the vitrification process, but switch

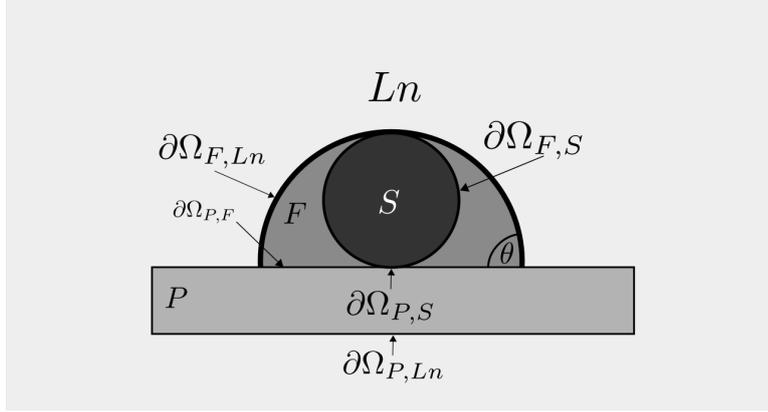


Figure 6.3: A cross section of the model geometry.

the initial and external temperatures. Whilst important in the context of understanding current thermal modelling of cryopreservation, these models are not directly relevant to the cooling process, because warming and cooling are fundamentally different processes. Where relevant, we will highlight the differences in these processes when justifying our assumptions.

6.2 Mathematical model of the vitrification process

We model the Open Cryotop[®] device, which is made of a polypropylene plate [237]. A schematic of the model geometry is displayed in Figure 6.3 with labelled domains, interfaces and boundaries. A sample placed onto the plate, contained within a droplet of CPA, and the system is assumed to be submerged into a liquid nitrogen bath. We, thus, define the following domains (see Figure 6.3):

- The liquid nitrogen, denoted by Ln ,
- The CPA droplet, denoted by F , assumed to be a spherical cap with a contact angle, θ ,
- The sample. We have a disjoint set S_i for $i \in \{1, 2, 3, 4\}$, assuming up to four samples modelled as per [248],
- The polypropylene plate, denoted by P .

The shared interface between two domains A and B is denoted $\partial\Omega_{A,B}$.

The spatiotemporal evolution of temperature, $u = u(x, y, z, t)$, for each model component determined is determined by solving the heat equation [354],

$$\frac{\partial u}{\partial t} = D_u(u) \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2} \right), \quad (6.2.1)$$

with initial temperature of $u(x, y, z, 0) = 310.15\text{K}$ (37°C), to match the temperature of the warmed bench-top [352]. The thermal diffusion coefficient, $D_u(u)$ ($\mu\text{m}^2/\text{s}$), is defined by [354]

$$D_u(u) = \frac{k(u)}{\rho(u)c(u)}, \quad (6.2.2)$$

where:

- $c(u)$ is the specific heat capacity of the component, at temperature u , ($\text{J}/(\text{kg K})$),
- $k(u)$ is the thermal conductivity of the component, at temperature u , (W/mK),
- $\rho(u)$ is the density of the component, at temperature u , (kg/m^3).

We note that the same equation is used to describe Brownian motion (2.2.20) and heat transfer (6.2.1), although the diffusion coefficient in each context has different interpretation and scale. This explains our choice to differentiate the diffusion coefficient in each context by applying the subscript u to the thermal case, denoting temperature. A general comment can be made that the diffusion coefficients in each context are related; one can compare the Stokes-Einstein relation [164, 266] with the definition of the thermal diffusion coefficient in Equation (6.2.2) for more insight into this relation, although such detail is not relevant to this work.

A second comparison to be drawn between this chapter and Chapters 2.2-5 is the omission of cytoplasmic streaming, which would couple with the thermal physics of the system. We do not consider this behaviour due to vitrification being a much faster process than cytoplasmic streaming; the average rate of cooling in vitrification is reported to be typically over $10,000^\circ\text{C}/\text{min}$ (see Table 6.1), whilst the velocity of cytoplasmic streaming in mice is reported to be less than $1\mu/\text{min}$. Hence, complete cooling will occur well before any significant movement of the cytoplasm is seen. The local cytoplasmic velocity will further decrease as cooling progresses, strengthening the assumption that cytoplasmic movement plays almost no role in intracytoplasmic heat transfer.

Whilst $D_u(u)$ is a function of temperature [245], other models assume k , ρ and c (and thus D) are constant [240]. We further examine this assumption in Section 6.2.1.3. Additionally, the rate of temperature change at the surfaces in contact with liquid nitrogen is described by Robin heat flux boundary condition [354],

$$-k \frac{\partial u}{\partial n} = h(u - u_{\text{ext}}), \quad (6.2.3)$$

where u_{ext} is the temperature of the liquid nitrogen, $\partial u/\partial n$ is shorthand for the directional derivative $\nabla(u) \cdot n$ with n the normal vector to the boundary, and h ($\text{W}/\text{m}^2\text{K}$) denotes the surface heat transfer coefficient, which defines the rate at which the CPA, or the plate, transfers heat to the surrounding liquid nitrogen. The thermal conductivity, k , features in both the definition of the thermal diffusion coefficient D_u (Equation (6.2.2)) and the boundary condition (Equation (6.2.3)), whilst the specific heat capacity is only present in the definition of D_u . Whilst the thermal conductivity represents the ability of the material to receive thermal energy, the specific heat capacity represents the materials ability to hold the same energy; this explains why the thermal diffusion coefficient scales as a ratio of the two parameters, whilst the boundary condition is only concerned with how readily the material gains thermal energy from some external source. In practice, whilst the two parameters are indistinguishable from one another throughout the bulk of the domain, they are not indistinguishable at the boundary, implying that we need to know both values in order to solve the heat equation over this system.

It can be assumed with good accuracy that u_{ext} is constant; we assume the same here, and examine the validity of the assumptions related to the boundary conditions in Section 6.2.1.1.

Let $u_P = u_P(x, y, z, t)$, $u_F = u_F(x, y, z, t)$ and $u_{S_i} = u_{S_i}(x, y, z, t)$ be the temperatures in the plate, cryoprotectant fluid and the i^{th} sample, respectively, at location (x, y, z) and time t . The temperature in the system is described by the following system of equations:

$$\frac{\partial u_P}{\partial t} = D_{u,P} \left(\frac{\partial^2 u_P}{\partial x^2} + \frac{\partial^2 u_P}{\partial y^2} + \frac{\partial^2 u_P}{\partial z^2} \right) \quad \text{in } P, \quad (6.2.4)$$

$$\frac{\partial u_F}{\partial t} = D_{u,F} \left(\frac{\partial^2 u_F}{\partial x^2} + \frac{\partial^2 u_F}{\partial y^2} + \frac{\partial^2 u_F}{\partial z^2} \right) \quad \text{in } F, \quad (6.2.5)$$

$$\frac{\partial u_{S_i}}{\partial t} = D_{u,S_i} \left(\frac{\partial^2 u_{S_i}}{\partial x^2} + \frac{\partial^2 u_{S_i}}{\partial y^2} + \frac{\partial^2 u_{S_i}}{\partial z^2} \right) \quad \text{in } S_i, \quad (6.2.6)$$

$$k_P \frac{\partial u_P}{\partial n} = h(u_P - u_{\text{ext}}) \quad \text{on } \partial\Omega_{P,Ln}, \quad (6.2.7)$$

$$k_F \frac{\partial u_F}{\partial n} = h(u_F - u_{\text{ext}}) \quad \text{on } \partial\Omega_{F,Ln}, \quad (6.2.8)$$

$$k_{S_i} \frac{\partial u_{S_i}}{\partial n} = k_F \frac{\partial u_F}{\partial n} \quad \text{on } \partial\Omega_{F,S_i}, \quad (6.2.9)$$

$$k_{S_i} \frac{\partial u_{S_i}}{\partial n} = k_P \frac{\partial u_P}{\partial n} \quad \text{on } \partial\Omega_{P,S_i}, \quad (6.2.10)$$

$$k_P \frac{\partial u_P}{\partial n} = k_F \frac{\partial u_F}{\partial n} \quad \text{on } \partial\Omega_{P,F}, \quad (6.2.11)$$

$$u_P(x, y, z, 0) = u_F(x, y, z, 0) = u_{S_i}(x, y, z, 0) = 310.15\text{K}, \quad (6.2.12)$$

$$u_{\text{ext}} = 77.15\text{K}. \quad (6.2.13)$$

Equations (6.2.4)-(6.2.6) govern the temperature in the three domains. Equations (6.2.7) and (6.2.8) describe heat conduction through the boundaries of the plate and cryoprotectant into the liquid nitrogen. Equations (6.2.9)-(6.2.11) describe conduction between the samples and cryoprotectant, samples and plate, and plate and cryoprotectant, respectively (where sample/sample conduction is not included as the contact surface would be much smaller than the interface between sample and droplet, resulting in minimal effect on heat transfer). Equations (6.2.12) and (6.2.13) describe the initial temperature (assumed constant throughout the system) and the external liquid nitrogen temperature, respectively.

Our objective is to study how the rate of temperature change in the samples during vitrification is affected by the arrangement and number of samples on the Open Cryotop[®] device. To compare arrangements, we need to choose how sample temperature is measured. We can either use the maximum temperature within the sample, or the average temperature. The maximum temperature informs us about the hottest part of the sample, and hence gives us a lower bound on the cooling rate, whilst the average temperature gives a more general picture of the rate of cooling in the sample. We choose the average temperature, noting that the high cooling rates and small sample volume prevent these metrics from differing significantly. The average temperature within the samples is calculated as a volume average, by integrating the temperature within the sample domain, and dividing by the sample volume [355]. The average temperature is denoted by [354]

$$\bar{u}(t) = \frac{\int u(x, y, z, t) dV}{\int dV}, \quad (6.2.14)$$

where V is the sample volume. Regardless of the choice of measurement of sample temperature, comparison of models in literature often focusses on ‘cooling rate’, defined by the difference in temperature over the time elapsed i.e

$$\frac{\bar{u}(t_2) - \bar{u}(t_1)}{t_2 - t_1} \quad (6.2.15)$$

with units °/min (or equivalently, K/min). Whilst the upper temperature, $\bar{u}(t_2)$, is usually fixed at 20°C [229, 240, 243, 246], there is no consensus for the choice of the final temperature $\bar{u}(t_2)$, with values of −120°C (153.15K) [229, 246], −130°C (143.15K)[240] and −150°C (123.15K)[242] reported. This variable definition is problematic, because heat transfer is a non-linear process, which can be seen by considering a separable variables approach to solving the heat equation in Equation (6.2.1). We propose a solution of the form

$$u(\mathbf{r}, t) = R(\mathbf{r})T(t),$$

Reference	Temperature interval (°C/min)	Cooling rate (°C/min)
[243]	Not reported	37,500
[242]	[-150,20]	41250
[224](used by the manufacturer, Kitazato [®] [350])	[-100, -20]	22,800
[229]	[-120,20]	69,250 ± 4285
[246]	[-120,20]	69,250 ± 4285
[240]	[-130,20]	10,465-60,000

Table 6.1: Reported cooling rates of the Open Cryotop[®] device. In particular, [224] is used by the manufacturer, Kitazato[®] as the official reported device cooling rate [350]. The interval over which cooling rate is defined varies, which by the exponential temperature decay predicted in Equation (6.2.16), may affect reported cooling rates. The variable cooling rates reported in [240] correspond to different definitions of the heat transfer coefficient, see Section 6.2.1.3. Cooling rates are equivalent in K/min.

where $\mathbf{r} = (x, y, z)$ denotes the Cartesian spatial coordinates, whilst $R(\mathbf{r})$ and $T(t)$ denote some functions to be determined, depending only on \mathbf{r} and t respectively. Substituting this solution form into Equation (6.2.1), we derive

$$\begin{aligned}
 R(\mathbf{r})T'(t) &= DR''(\mathbf{r})T(t), \\
 \implies \frac{R''(\mathbf{r})}{R(\mathbf{r})} &= \frac{T'(t)}{DT(t)}.
 \end{aligned}$$

Since the left and right hand sides depend on different variables, they must both be equal to some constant, which we call ϱ . Ignoring the spatial component $R(\mathbf{r})$, we consider the time-dependent solution, from which we derive a first order ordinary differential equation,

$$T'(t) = D\varrho T(t),$$

which has the solution

$$T(t) = c_1 e^{D\varrho t}. \quad (6.2.16)$$

where c_1 is another constant, which like ϱ , depends on the domain geometry, initial condition and boundary conditions. Hence, temperature decays *exponentially* with time, up to an equilibrium temperature [354]. Choosing a larger temperature interval over which the cooling rate is evaluated will therefore decrease the cooling rate significantly. We present some of these reported cooling rates from literature in Table 6.1, alongside the interval over which they are defined, demonstrating a general lack of agreement between reports. An additional concern with reporting cooling rate is that the timescale over which cooling occurs is very small. An analytical estimate of this timescale requires that we first define the parameters in Equations (6.2.4)-(6.2.13), and is hence later provided in Equation (6.2.18) after we discuss the assumptions made in our model. For now, we plot experimental

measurements of the temperature of the Open Cryotop[®] once plunged into liquid nitrogen in Figure 6.4, taken from [356]. During this experiment, the device is initially at room temperature (23°C, or 296.15K), but reaches the temperature of the liquid nitrogen (-196°C, or 77.15K) within less than half a second. The change in temperature during vitrification is on the order of hundreds of degrees centigrade, whilst the timescale for this change is thousandths of minutes; combined, this makes the defined cooling rate, in C/Min, highly sensitive to small changes in time, which can make comparison between models with very similar cooling times misleading.

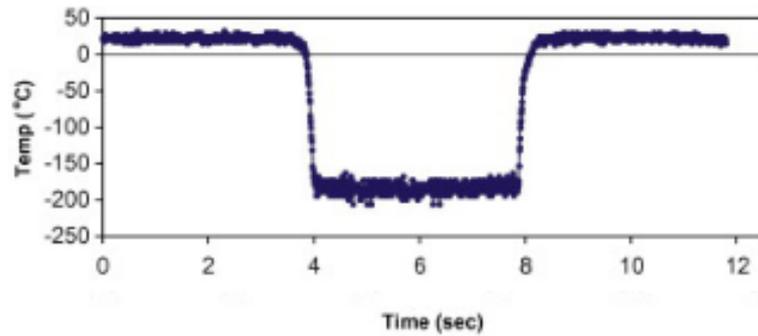


Figure 6.4: Experimental temperature reading of the Open Cryotop[®] device after being plunged into liquid nitrogen, taken from [356]. Initially at ambient room temperature, the device is first submerged in liquid nitrogen at around 4 seconds, cooling completely within less than half a second.

Despite the issues we report with the use of cooling rate as a comparison between our model and others, it is nonetheless the current standard measurement. We therefore also report cooling rates, although our temperature interval is again different to those reported in Table 6.1, because we take the upper temperature to be 310.15K (37°C), in accordance with our initial condition, and a lower limit of 143.15K (-132°C), which is consistent with previous literature [240, 246]. The lower limit coincides with the glass transition temperature (T_g) of the material, which is the temperature at which vitrification is considered to occur, as reported for similar cryoprotectant solutions [241].

Before proceeding to analyse the model, we will first examine some key assumptions made in the model.

6.2.1 Assumptions

6.2.1.1 Liquid nitrogen is a stationary, isothermal liquid

Liquid nitrogen is assumed here to be a stationary and isothermal liquid with a fixed temperature of 77.15K, as assumed recently in a vitrification model of the Open Cryotop[®] in [240]. Assuming the liquid nitrogen to be isothermal is

reasonable, because the liquid nitrogen is actively boiling, which buffers against temperature increase. The volume of liquid nitrogen in the bath is large enough that the bath will not ‘boil dry’ over the course of vitrification; this means that the Cryotop[®] device remains submerged in the isothermal domain for the duration of the protocol.

However, the assumption that the liquid nitrogen is stationary is less accurate, in light of the active boiling. Boiling itself can be characterised as being in one of two different regimes: Nucleate boiling is where bubbles form at specific sites on the heated object, growing in size until they detach once buoyancy exceeds cohesive surface tension forces [357]. Film boiling, meanwhile, is a more extreme case where bubble formation is so rapid on the surface that bubbles do not have time to detach, instead coalescing into a vapour film which insulates the surface [357]. These two regimes are fundamentally opposed, since nucleate boiling helps to cool the surface, whilst film boiling prevents cooling. Therefore, it is important to know which regime we are in. Plunging the device directly into liquid nitrogen has been shown to favour film boiling in oocyte vitrification [242], but in this work, it is also shown that one can account for this film boiling by assuming a stationary isothermal domain, and tuning the heat transfer coefficient h to match experimental cooling data. Our assumption is therefore justified, although more complex models of fluid dynamics and boiling could be developed to better represent the state of liquid nitrogen. In such a model, explicit modelling of the film boundary layer whose width and thermal properties vary with the temperature gradient could be undertaken.

6.2.1.2 Samples are spheres

Human oocytes can vary in size and shape, but are typically spherical, with a diameter of approximately 0.12mm [358]. Including the Zona Pellucida and perivitelline space, the total oocyte diameter is approximately 0.15mm [359]. Cumulus cells are not included, as these are removed according to the Open Cryotop[®] user manual [248]. Only spherical samples are modelled here, as deviation from a spherical shape is considered a form of dysmorphia [360], and hence poor oocyte quality. We choose 0.1mm as a conservative size estimate; this accounts for loss of water volume [361] and artificial collapse [362] during the vitrification process. Assuming this size estimate allows more samples to fit together in dense arrangements that are more likely to have slower cooling rates, so this assumption offers conservative estimates of the time taken for vitrification.

For simplicity, we assume the same model may be used to describe oocytes and embryos. This makes the assumption that oocytes and embryos have the same size, and thermal properties. In comparison to oocytes, embryo size remains

constant throughout the early development stages, with the mean diameter of early blastocysts being reported to be $140 \pm 9 \mu\text{m}$ in [363]. During the latter stages of the development of the blastocyst, a small expansion occurs, with expanded blastocysts having a diameter of $190 \pm 9 \mu\text{m}$ [363]. Hence, the validity of our assumption varies depending on the stage of development of the embryo at the point of vitrification.

6.2.1.3 Thermophysical properties and parameters

To develop a good mathematical model, parameter values must be chosen for c, k and ρ , for the plate and for the CPA. These values are ‘characteristic of, and measured by, different experimental situations’ [354], and can only be accurate if measured for the specific case examined. The vitrification solution contains 15% ethylene glycol and dimethyl sulphoxide (DMSO) v/v each, as well as a 0.5M sugar solution [237]. The concentration of solutions of ethylene glycol in water has a documented effect on the thermal properties of the system [364]. Although experimental derivations of temperature-dependent thermophysical properties exist for various compounds and tissues [365], parameters for the VS solution supplied by Kitazato[®] are not well documented, so it is necessary that some simplifying assumptions are made [240, 245]. First, we assume that the pressure within the liquid components of the system remains constant, such that the thermophysical properties of the modelled materials depend only on temperature. This assumption is consistent with similar models [240, 243], although discussion of the use of increased pressure to expedite the vitrification process can be found here [366].

Where temperature-dependent data is available for similar compounds, it is possible to assume that the CPA is made up of this compound, and interpolate between recorded parameter values, an approach which has been taken for the separate ‘reverse’ problem of sample warming [245]. Experimental measurements of chemical and tissue thermophysical properties are often incomplete, however, with only a few experimentally derived temperatures over a smaller temperature range than the cooling observed during vitrification [365]. The most complete dataset available is for the temperature-dependent thermophysical properties of water [365, 367, 368]. Therefore, we can undertake our simulations here assuming that the CPA is water. This assumption is not imperfect, but as this is the best data available, this assumption allows us to explore the effect of variable thermophysical parameters on the vitrification model.

In general, $k(u)$ increases as temperature decreases, whilst $c(u)$ and $\rho(u)$ decrease [365]. Hence, Equation (6.2.2) implies that $D(u)$ increases as temperature decreases. This increase in $D(u)$ leads to overestimation of cooling rates if

parameters are assumed constant [365]. In the absence of suitable temperature-dependent data for CPA mixtures, we could alternatively assume that thermophysical parameters are constants [240]. Hence, in this work, we use three sets of thermophysical parameters. First, we assume that the CPA is comprised of ethylene glycol only, as it is reasonable to assume that the vitrification solution behaves more similarly to the CPA it contains than to water (otherwise, the CPA would serve no purpose in the vitrification process). The solution is still water-based, however, so we consider a second case in which the CPA consists of vitreous water, as in [240]. Thirdly, to compare, we consider the case where the CPA is purely water, with variable thermophysical parameters.

Values for c , ρ and k are reported in Table 6.2, for ethylene glycol and vitreous water, whilst the thermophysical properties of water are shown in Table 6.3. During simulations, linear interpolations are used to generate parameters between these recorded values. Outside the reported temperature range, thermophysical parameters are assumed to be extrapolated as constant. Based on the reported parameter values, we can plot the diffusion coefficient for all three cases, shown in Figure 6.5.

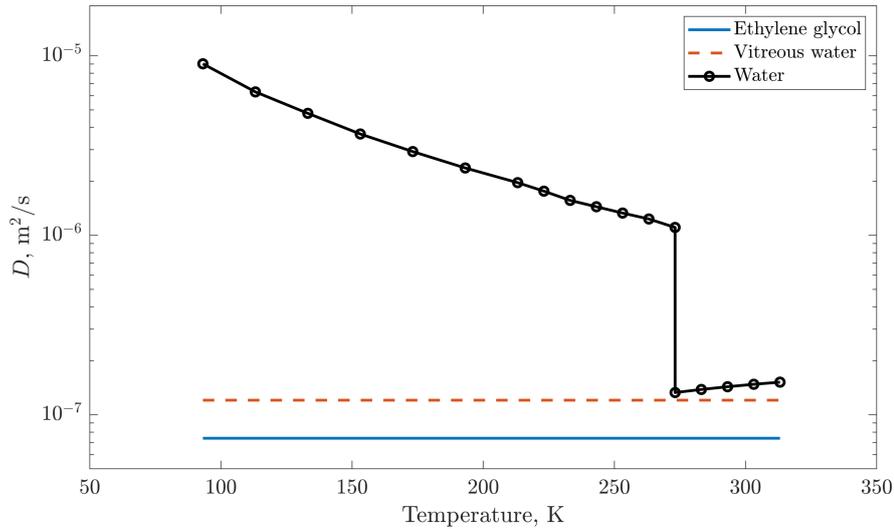


Figure 6.5: The diffusion coefficient for each the different thermophysical parameter choices in Section 6.2.1.3: ethylene glycol, vitreous water and water (with variable thermophysical parameters), under the assumption of constant pressure. The vertical axis is on a log scale. For water, the sharp transition occurs at the freezing point, when water will solidify into ice during cooling. Since we model the cooling problem, temperature goes from right to left.

Another parameter of interest is the surface heat transfer coefficient, h , which along with thermal conductivity, affects the rate of heat transfer at a boundary with a fluid. The boundary conditions in equations (6.2.7)-(6.2.11) arise from Newton’s Law of Cooling, for objects being cooled by forced convection [354],

so h is dependent on the fluid dynamic behaviour at the interface between the polypropylene plate/CPA droplet and the liquid nitrogen. As previously discussed, the interface between the plate/CPA and liquid nitrogen is, first, subject to insulating film boiling [357], and subsequently, as the system cools, nucleate boiling re-establishes itself [242]. In computational thermal models, these different boiling regimes are often represented by changing the value of h [244, 369]. Choosing appropriate values of h that are determined experimentally can therefore justify the assumption 6.2.1.1, since we can, in this way, replicate realistic results without needing to model the fluid mechanics of liquid nitrogen. It is critical that we choose the correct value of h , however, since this determines the rate at which the boundary controls the cooling rate. At very extreme cases, $h \approx 0$, which corresponds to an insulated system with no cooling, whilst $h \rightarrow \infty$ corresponds to a Dirichlet condition in which the boundary is set to a fixed temperature with cooling rates then maximised. We will show that the results of the model depend strongly on the choice of h .

Choosing h accurately is difficult given the wide variety of reported values. These values reflect specific details of the experiment modelled, such as the choice of device, or whether the device is being warmed or cooled. For example, a different type of cryopreservation device, known as an open pulled straw, has been shown to be subject to nucleate boiling, with the heat transfer transfer coefficient reported to be $h = 1,355\text{W/m}^2\text{K}$ [245], yet for warming of the same device in a sucrose solution, the range for h is given as $1800 - 2200\text{W/m}^2\text{K}$ depending on the loading volume. For the Open Cryotop[®], numerical solutions taking $h \approx 2000\text{W/m}^2\text{K}$ match experimentally determined cooling rates [240, 246], although separate experimental measurements suggest a vastly different range of $9000 < h < 10000\text{W/m}^2\text{K}$ [227]. In this work, we choose $h = 2000\text{ W/m}^2\text{K}$ as in previous similar models of cooling of the Open Cryotop[®] [240], as well as to match experimental cooling rates [246].

One explanation for the variability in h arises from the variable definition of the cooling rate, as shown in Table 6.1. The choice of h is often made in order to make the model in question achieve target cooling rates [240, 246], but this means that different definitions of cooling rate will hence result in different heat transfer coefficients. Hence, choosing h to match cooling rates in literature results is a selection that depends not only on the thermophysical properties of the system, but also on an arbitrary definition within the system; this is a limitation of our model, and previous models. Nonetheless, our model constitutes a framework, in which any choice of thermophysical parameters can be easily and quickly implemented and simulated. As more accurate and applicable thermo-

physical parameter data become available, our model could be quickly rerun to yield more accurate temperature values.

Material	Parameter	Value	Source
Ethylene glycol	ρ	$1.11 \times 10^3 \text{ kg/m}^3$	COMSOL Multiphysics [355]
	c	3140 J/kg K	
	k	$2.58 \times 10^{-1} \text{ W/mK}$	
	D	$7.40 \times 10^{-8} \text{ m}^2/\text{s}$	
Vitreous water	ρ	$0.983 \times 10^3 \text{ kg/m}^3$	[240]
	c	4218 J/kg K	
	k	$5.0 \times 10^{-1} \text{ W/mK}$	
	D	$1.21 \times 10^{-7} \text{ m}^2/\text{s}$	
Polypropylene	ρ	$9.46 \times 10^2 \text{ kg/m}^3$	[370]
	c	1900 J/kg K	
	k	$1.50 \times 10^{-1} \text{ W/mK}$	
	D	$8.35 \times 10^{-8} \text{ m}^2/\text{s}$	
Liquid nitrogen boundary	h	1850 W/m ² K	[245]

Table 6.2: The thermophysical properties for the materials involved in the Open Cryotop[®] system.

Temperature	$k(\text{W/mK})$	$\rho \text{ (kg/m}^3\text{)}, \times 10^2$	$c \text{ (J/kg K)}, \times 10^2$	$D \text{ (m}^2/\text{s)}, \times 10^{-5}$
93.15	7.00	9.34	8.30	0.90
113.15	5.70	9.33	9.70	0.63
133.15	4.90	9.31	1.10	0.48
153.15	4.20	9.31	1.23	0.37
173.15	3.70	9.29	1.36	0.29
193.15	3.30	9.27	1.5	0.24
213.15	3.00	9.25	1.65	0.20
223.15	2.80	9.24	1.72	0.18
233.15	2.60	9.23	1.80	0.16
243.15	2.50	9.22	1.88	0.14
253.15	2.40	9.20	1.96	0.13
263.15	2.30	9.19	2.03	0.12
273.14 (ice)	2.14	9.17	2.11	0.11
273.16 (water)	0.56	10.0	4.22	0.01
283.15	0.58	10.0	4.19	0.01
293.15	0.60	9.98	4.18	0.01
303.15	0.62	9.96	4.18	0.01
313.15	0.63	9.92	4.18	0.02

Table 6.3: Thermophysical properties of water [368]. Two values are reported at 273.15K for ice or water; with linear interpolation, we determine the values of ice at 273.14K and the value for water at 273.16K.

6.2.1.4 Cryoprotectant droplets are hemispheres

In the Open Cryotop[®] user manual, a configuration shown in Figure 6.2 depicts the CPA droplets as concave after the aspiration of excess media. Additionally, the motion of sliding the pipette back may have an effect on the shape of the droplet [371]. In our work, aspiration and pipette motion are ignored, and it is assumed that droplets are spherical caps, as spheres minimise the surface tension in a droplet (as a result of the isoperimetric inequality in three dimensions [372]). Observations from the clinic suggest spherical caps are a good assumption [371].

If the droplet is a spherical cap, the contact angle of the droplet must be described. The contact angle, θ , is the angle formed between the plate and the tangent plane to the sphere surface at the point of contact, such that $\theta \in (0, \pi)$. This contact point is the meeting place of three domains in Figure 6.3: the fluid droplet, F , the solid plate, P , and the liquid nitrogen, Ln . A tension force $\varkappa_{A,B}$ exists at the interface of any two domains, A and B , and these tension forces may be used to define the contact angle through the Young's Equation [373],

$$\varkappa_{P,Ln} - \varkappa_{P,F} - \varkappa_{F,Ln} \cos \theta = 0.$$

The tension forces depend on the chemical properties of each domain, and environmental factors such as temperature; this makes it difficult to model them accurately. However, since the droplet acts structurally as a solid during vitrification [223], the contact angle should not change as a function of temperature and we, thus, assume it to be constant. The droplet must have a height greater than 0.2mm above the plate in order to cover the samples, and for a fixed droplet height, the droplet volume decreases with increasing θ . The Open Cryotop[®] protocol desires a minimal droplet, and so we take the highest realistic value of θ . Clinical observation suggests that droplets do not actually appear concave at the base, so $\theta = \pi/2$. We assume a radius of 0.21mm to ensure the sample is fully covered in CPA. A droplet with this particular contact angle and radius has volume $0.019\mu\text{l}$, which satisfies the condition given in the Open Cryotop manual [248] that droplet volume is less than $0.1\mu\text{l}$. Similar modelling approaches use a droplet whose volume is equal to $0.1\mu\text{l}$ [245], but choosing a smaller volume allows us to increase the number of samples used without exceeding this upper volume limit.

Now that we have defined the dimensions and parameter values for each component of the domain, we present a schematic diagram of the system in Figure 6.6. The Open Cryotop[®] plate has a depth of 0.1mm and a width of 0.7mm [237]. We do not model the handle of the device, which is held above the surface of the liquid nitrogen, since it is sufficiently far away from the vitrified sample

to have no effect on cooling. We additionally simplify the shape of the plate to be rectangular, rather than ‘diamond-shaped’ as shown in Figure 6.1. This is justified since the shortest diffusion pathway for the sample is through the top of the CPA droplet, or the bottom of the plate, so the effect of the diamond-shaped tip on cooling is minimal.

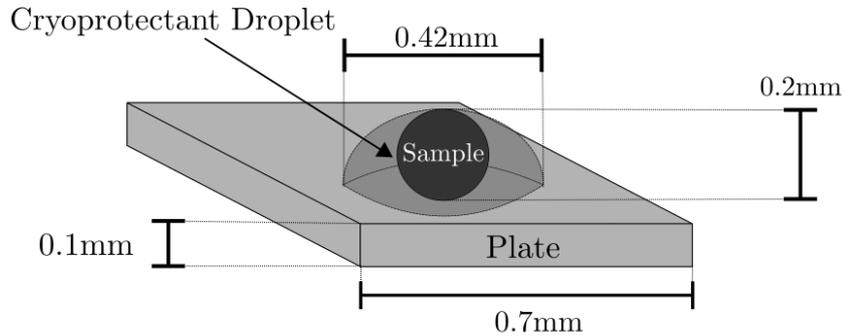


Figure 6.6: A schematic diagram of the Cryotop[®] system, including dimensions of the three key domains: the sample, the plate and the cryoprotectant droplet. The entire system is submerged in liquid nitrogen.

6.2.1.5 Samples have the same chemical properties as the cryoprotectant solution

During vitrification, much of the water volume is removed from the sample by osmosis, falling to as low as 5% within two minutes [229]. Additionally, in blastocysts, artificial collapse of the blastocoel may be used to reduce the water volume [362]. During the water loss process, the sample replaces some of its volume with the surrounding CPA solution, and its chemical composition becomes similar to that of the CPA [374]. We, therefore, assume that the sample has the same thermal properties as the CPA solution. Since samples are not insulators, or conductors, and have a very small volume compared to the total volume of liquid nitrogen that surrounds the Open Cryotop[®], this assumption is justified. Note that this assumption has been made indirectly in previous literature, since only a droplet without a suspended sample was modelled [240, 245].

6.2.1.6 The transition from liquid to glass-like state occurs instantaneously, no moving boundary effects

Vitrification is a complex problem to model, because in reality, the phase change from a liquid to a glass-like solid does not happen instantaneously when the temperature of the medium reaches the vitrification point. During a typical freezing process, the particles in the fluid lose entropy as they enter their new crystalline structure, resulting in a release of latent heat energy [375]; this energy counteracts

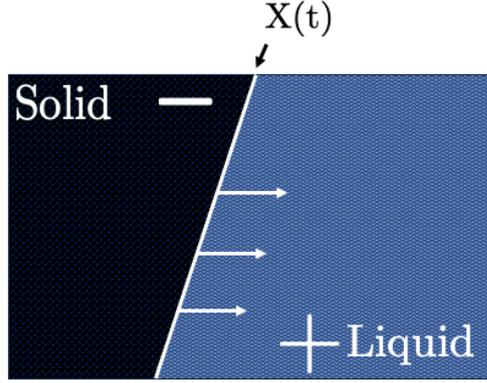


Figure 6.7: The moving boundary problem encountered during modelling of a cooling problem where phase change occurs in the domain. As the temperature falls to the freezing point u_m of the liquid, parts of the domain solidify, resulting in a solid-liquid interface, with position $X(t)$ at time t and fixed temperature $u = u_m$. The interface moves through the domain as more of the liquid solidifies.

the temperature drop and slows the movement of the solid-liquid boundary [375]. This results in a the formation of a moving solid-liquid interphase, positioned at $X(t)$ at time t . We visualise the moving boundary problem in Figure 6.7.

It is possible to model the position of the solid-liquid interface, known as a Stefan front [375], by adding to the system of PDEs in equations (6.2.4)-(6.2.13) a new set of equations referring to the interface position $X(t)$ at time t . First, the temperature at the interface is equal to the freezing temperature u_m ,

$$u(X(t), t) = u_m.$$

Additionally, the Stefan condition [376] must be satisfied, given in one dimension as

$$\rho\ell \frac{\partial X(t)}{\partial t} = -k_{\text{liquid}} \frac{\partial u(X(t)^+, t)}{\partial x} + k_{\text{solid}} \frac{\partial u(X(t)^-, t)}{\partial x} \quad (6.2.17)$$

where k_{liquid} and k_{solid} are the thermal conductivities of the material in its liquid and solid phases respectively, ℓ is the latent heat of freezing and the positive/negative superscript refers to the liquid/solid phase respectively. The Stefan condition represents energy conservation across the boundary of the interface, and gives an expression for the velocity of the Stefan front.

In this work, we make the simplifying assumption that there is no moving boundary between the liquid CPA and the vitrified part of the droplet, equivalent to assuming there is no delay between the material reaching freezing point and the interface moving into this region. The strength of this assumption can be tested if parameter values for k_{solid} , k_{liquid} , ρ and ℓ can be determined [376, 377], by

comparing the timescales of diffusion and front propagation. The characteristic timescale of diffusion is given by

$$t_D = \frac{L^2}{D},$$

where L is the characteristic length scale, for example the CPA droplet radius. The characteristic timescale of the Stefan front propagation, t_S , can be determined from the Stefan condition in (6.2.17), which determines the speed at which the boundary moves. Our assumption that we do not need to model the moving boundary problem is valid when $t_D \gg t_S$, since the position of the solid-liquid interface follows the changing temperature predicted by the heat equation. If $t_D \approx t_S$ or $t_D \ll t_S$, however, then the interface moves more slowly, and our assumption will not be valid. Since the thermophysical and chemical properties of the VS CPA provided by Kitazato[®] are not publicly available, we cannot undertake the outlined analysis to test our assumption. However, we observe that the moving boundary problem is not included in previous modelling approaches [240, 242, 244, 245], which match experimental results [229, 246], suggesting this is a good assumption to make in this problem.

6.2.2 Computational modelling

A Finite Element Method (FEM) is used to simulate the temperature throughout the models, solving Equations (6.2.4)-(6.2.13). Employing COMSOL Multiphysics 5.5, adaptive meshes are constructed for each given geometry to balance accuracy and computational cost. A mixture of triangle and tetrahedral elements are used, as determined by the solver. A ‘fine’ element discretisation is employed, which uses a minimum of 2182 tetrahedral elements and 1320 triangular elements across each of the models developed. Before solving the model, we first must determine the time interval over which we run simulations. This is aided by determining the characteristic timescale,

$$\tilde{t} = \frac{L^2}{D}, \quad (6.2.18)$$

where L is the ‘characteristic length’ of a system and D is the diffusion coefficient. The characteristic timescale gives an approximation of the timescale for the system to reach equilibrium, which in this case is 77.15K. Taking the characteristic length to be the droplet radius, $L = 0.21\text{mm}$, and using the D values for ethylene glycol and vitreous water in Table 6.2, a timescale of approximately $\tilde{t} = 0.59\text{s}$ is reported for ethylene glycol, and 0.36s for vitreous water. The diffusion coefficient is even higher in the variable water case (see Figure 6.5), implying the timescale

will be even smaller in if these parameters are used. This approximation informs us equilibrium is likely reached on the order of seconds, demanding that we simulate temperatures at time steps much smaller than one second to achieve a smooth approximation of the temperature. This approximation also suggests that the timescale for equilibrium to be reached is on the order of a second, allowing us to limit the length of time over which we simulate. As such, we choose to simulate only over 2 seconds, which is shown later in Figures 6.12 and 6.13 to be more than long enough for the temperature everywhere in the domain to reach equilibrium. COMSOL automatically determines the time-stepping involved in the system, which we will discuss now, but we interpolate the temperature at 0.02 second intervals to ensure we record a smooth temperature change over time at each point.

In COMSOL, we use a backwards differentiation formula (BDF) solver, which is a linear multistep method suitable for solving stiff initial value problems [378], which have the general form [379]

$$u' = f(t, u), \tag{6.2.19}$$

$$u(t_0) = u_0, \tag{6.2.20}$$

with the implicit BDF scheme given by [379]

$$\sum_{k=0}^s a_k u_{n+k} = h\beta f(t_{n+s}, u_{n+s}), \tag{6.2.21}$$

where h denotes the ‘step size’ and $t_n = t_0 + nh$. The parameter s denotes the order of the scheme, with $s = 1$ reducing the BDF scheme to a backward Euler method, up to $s = 6$, after which the scheme becomes unstable [380]. In COMSOL, the order is determined automatically at each iteration, with only order one or two being used. COMSOL documentation highlights that at low orders, solvers are known to have dampening effects which smooth out sharp gradients [355], but this does not present a problem in this work, since our initial condition in Equation (6.2.12) is uniform and the heat equation generates smooth solutions.

Numerical performance of the model can be partially measured using the convergence of the automatically-determined step size. Adaptive stepping helps to balance computational workload and numerical error [381]. Whilst smaller step sizes can provide more accurate approximation of the solution, making larger steps improves the speed of the solver [381]. Hence, the adaptive solver only increases step size when doing so is not detrimental to accuracy [381]. Since heat transfer is exponential in time (see Equation (6.2.16)), we expect that step size will increase

over throughout the duration of the simulation, as the temperature becomes more uniform throughout the domain and, hence, changes more slowly. Since step size is often very small initially, it can be more helpful to plot the reciprocal, and look for convergence of to a small reciprocal step size. In Figure 6.8, we show an example COMSOL-generated geometry, which depicts the ‘Triangle’ embryo configuration described in Section 6.2.3, and show the convergence plot in Figure 6.9 which demonstrates the solver is stable. This stability is reflected in the rest of the models.

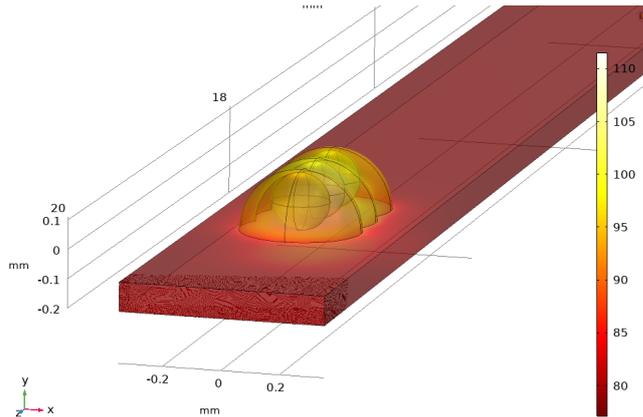


Figure 6.8: A snapshot of the model rendered in COMSOL for the three sample ‘Line’ arrangement described in Section 6.2.3. The snapshot is taken at 0.3 seconds into the simulation, with temperature in Kelvin, corresponding to colour, determined by solving Equations (6.2.4)-(6.2.13). The sample in centre is lighter than those on either side, implying it is warmer, which demonstrates that sample arrangement does affect cooling.

We can also discuss numerical performance by estimating the error. Our numerical solution is an approximation of the solution to Equation (6.2.21), but will in practice generate some small numerical error. COMSOL uses iterative refinement to minimise this error, which we briefly describe here [382]. Since the BDF scheme is linear, the derived temperatures satisfy systems of linear equations throughout the domain. The true temperature, u_* , satisfies the linear system of equations given by

$$Au_* = b,$$

whilst our numerical solution, u , approximates u_* . Iterative refinement defines a set $\{u_1, u_2, u_3, \dots\}$ which converges to u_* , with $u_1 = u$. Each refinement in the sequence is determined by finding the residual error,

$$r_m = b - Au_m,$$

and then solving the system for some correction term c_m that would remove the residual error,

$$Ac_m = r_m.$$

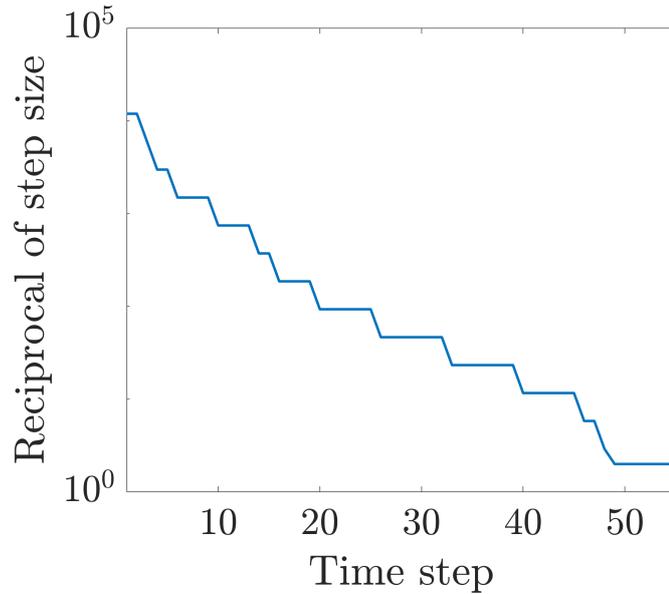


Figure 6.9: The reciprocal of the automatically-determined step size for the three sample ‘Line’ arrangement shown in Figure 6.8. The step size converges, implying stability. Similar convergence is observed for all arrangements, which suggests the model is stable.

Adding the correction term to the current solution generates the next solution,

$$u_{m+1} = u_m + c_m.$$

During iterative refinement, COMSOL reports a relative error estimate, defined by

$$\text{LinErr} = \text{Tol} \frac{\|c_m\|_2}{\|u_m\|_2},$$

where Tol is some constant factor used to weight the relative error (in this case, Tol= 0.1) and $\|\cdot\|_2$ denotes the Euclidean norm. The relative error estimate is therefore used to confirm that the solution is converging during iterative refinement. In all arrangements modelled, LinErr is on the order of 1×10^{-15} or less, suggesting the model converges to a suitable approximation.

6.2.3 Spatial arrangements of oocytes/embryos

During vitrification, the samples are placed by the embryologist on the plate such that they just touch each other, or have only a very small separation. Once on the plate, the samples may move, resulting in a number of different arrangements which we model. After samples are placed, the covering CPA droplet is constructed in the model by taking the union of the individual droplets that would have otherwise surrounded the individual samples, as shown. We use COMSOL Multiphysics 5.5 to construct the model, with an example shown in Figure 6.10 demonstrating an arrangement of two samples. Thus, the volume of the combined

CPA droplet grows with the number of samples vitrified, but this growth depends on the amount of overlap between each individual droplet. The overlap volume between the droplets depends on how close the samples are placed to one another, and on which pattern they are arranged (called ‘arrangement’ from now on). In this way, the model represents clinical practice well by modelling the aspiration of excess CPA, rather than modelling each sample in a fixed volume [248].

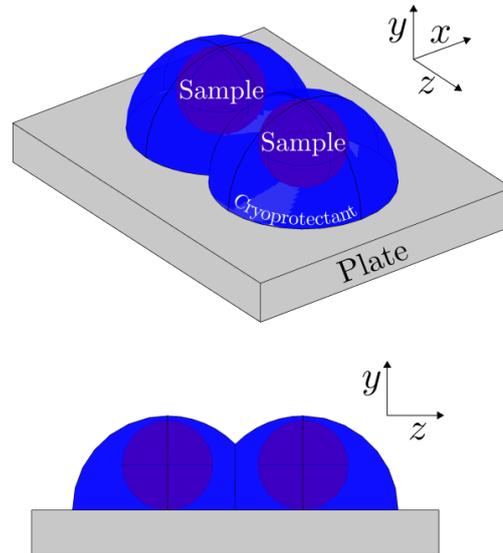


Figure 6.10: Two merged droplets rendered in COMSOL Multiphysics 5.5. The same set-up is displayed from two different angles. The droplet is formed from the union of the two separate droplets accompanying each sample.

The first arrangement studied in this work is a single sample, which provides a reference case for arrangements with more than one sample. Up to three more samples are then added in a line running down the length of the plate, to study the effect of increasing the number of samples (see Figure 6.11a); we call this the ‘Line’ arrangement. Considering longer lines of samples is unnecessary, since the cooling rate of samples located in the centre of the line will be dictated by the shortest distance to the liquid nitrogen, and will hence be unaffected by the length of the sample line. Additionally, the upper limit allowed by the Open Cryotop[®] protocol is 4 [248].

The next arrangement we consider is a variation of the ‘Line’ arrangement. We consider the possibility that by adding more than one sample to the plate, the aspiration process may be more difficult to perfect, resulting in more CPA remaining on the plate. Instead of forming a CPA droplet from the union of the droplets around it, as for the ‘Line’ arrangement, we now add to the CPA volume assuming it forms a cylinder with rounded quarter sphere ends. This geometry,

which we refer to as the ‘Cylinder’ arrangement (Figure 6.11b), is considered to test how perturbations to the droplet geometry compared to the ‘Line’ arrangement might alter cooling rates. We also account for potential aspiration errors which, as shown in the Open Cryotop[®] manual, are considered poor protocol.

Finally, we consider cases where the samples are not aligned along the centre of the plate, but group into more complex arrangements (see Figure 1.8 for practical examples). Choosing close aggregation of samples allows us to study various types of grouping behaviour. Three different arrangements are considered, a ‘Triangle’ (Figure 6.11c), a ‘Square’ (Figure 6.11d) and a ‘Cross’ (Figure 6.11e). In reality, the locations of the samples on the plate may vary, but modelling every possible arrangement is unnecessary. The sample arrangements we consider give a good representation of the expected behaviour.

In the ‘Triangle’ case, samples are placed such that their centres lie on the corners of an equilateral triangle with side length 0.2mm. In the computational model, such an arrangement may cause issues where COMSOL Multiphysics 5.5 cannot generate a mesh around the contact points, which appear to overlap. As such, in practice the leftmost sample is shifted slightly to prevent mesh generation failures (see Figure 6.11c). In the Square case, sample centres form the corners of a square whose side lengths are 0.2mm. In the Cross case, an arrangement similar to that of the square is desired, but with a sample in the centre. This arrangement would cause the droplet to hang off the plate, however, so samples are instead placed on the corners of a rectangle with side lengths 0.28mm wide by 0.3mm long, leaving enough space to place a sample in the centre of the rectangle. The measurements shown in Figure 6.11 are calculated based off of these rules.

Manufacturer’s guidelines allow a maximum of four samples per device [248], which means that the Cross case is not considered a valid arrangement within an IVF clinic. This case is tested nonetheless, because it is an extreme case that allows us to quantify the extreme scenario in which a sample is surrounded by other samples, so that the validity of the guidelines may be confirmed. If the cross case does not preclude negative outcomes for embryo cooling rates, then no other case we can feasibly consider would cool slower.

The sixth case considered is the ‘Worst Case’ benchmark, in which the droplet has radius equal to the width of the plate (Figure 6.11f), and contains only a single sample at its centre. The volume of this droplet is $0.09\mu\text{l}$, which is just less than the recommended maximal volume of $0.1\mu\text{l}$, and as such, this represents a case fits within current guidance and should be safe [248]. As the droplet volume is maximised, this will represent the case with the slowest cooling rate

that is still considered to be viable for sample vitrification. Whilst the model cannot necessarily be used to predict a minimum safe, or effective, cooling rate, any case in which cooling rates are greater than those observed in the Worst Case simulations can be considered to have sufficiently fast vitrification.

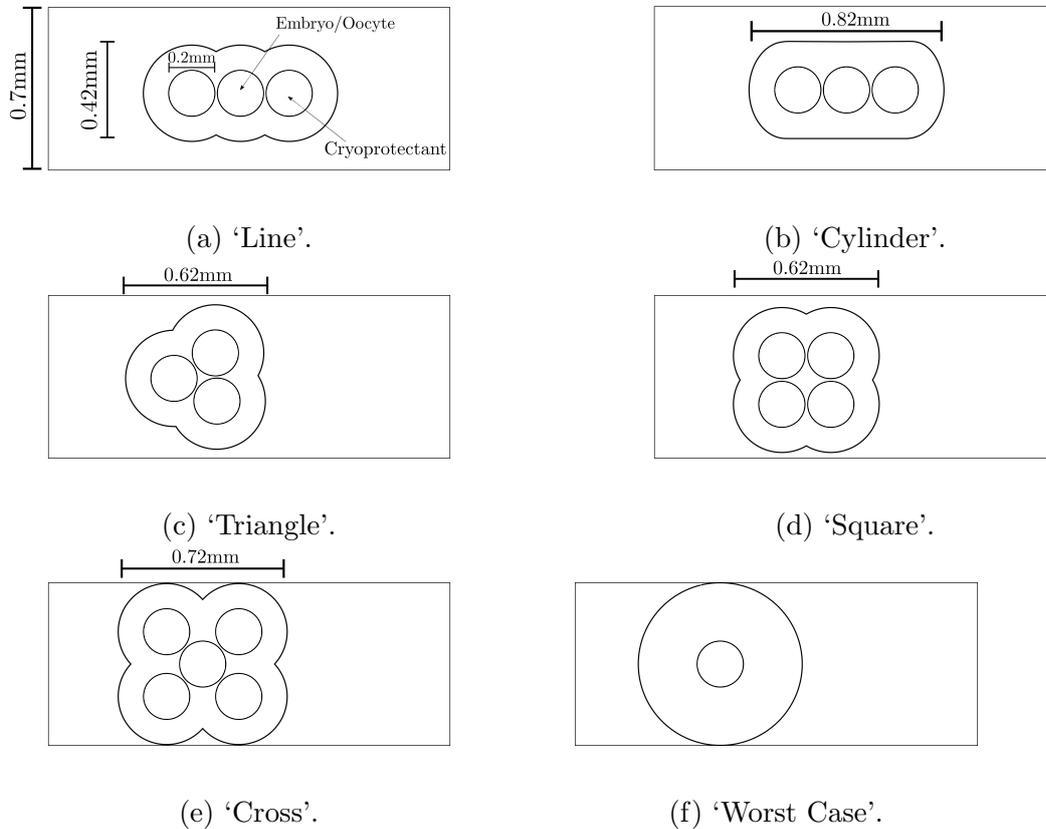
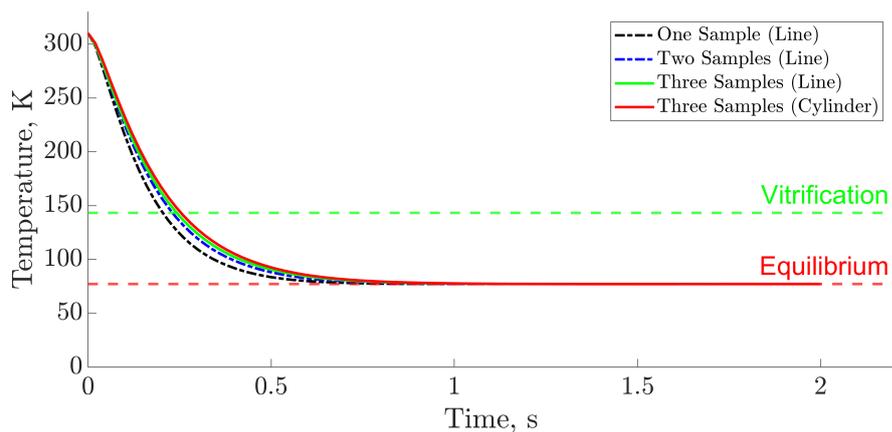


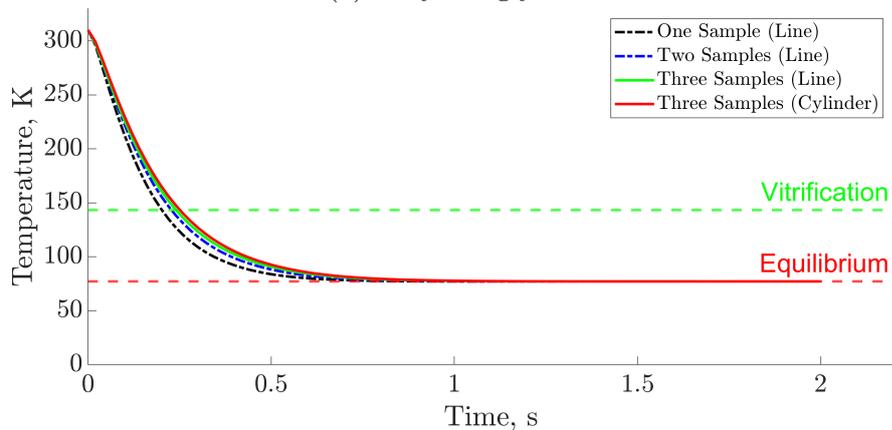
Figure 6.11: Top-down visualisations of the six arrangements our model. The ‘Line’ arrangement is considered in three permutations with variable sample number, named ‘Line-1’, ‘Line-2’ and ‘Line-3’ according to the number of samples considered.

6.3 Results

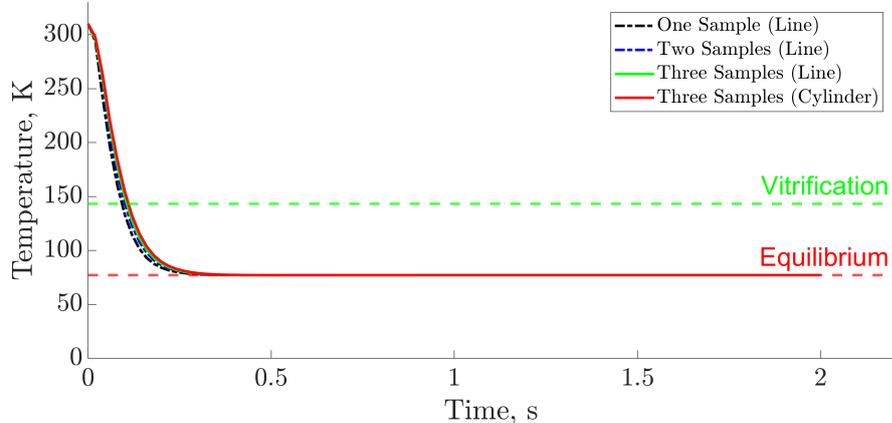
Figure 6.12 shows the average temperature over time for arrangements in which between one and three samples are placed along the centre of the plate (see Figure 6.11a and Figure 6.11a), with separate plots for each set of the three cases assumed for the thermophysical parameters. Figure 6.13 shows the temperature in the ‘Triangle’, ‘Square’, ‘Cross’ and ‘Worst Case’ arrangements, with a minimum of three samples included. The time taken to reach the target temperature of 143.15K, for all cases, is shown in Table 6.4, and the associated cooling rates are displayed in Table 6.5. We will now discuss some key features of these figures.



(a) Ethylene glycol.

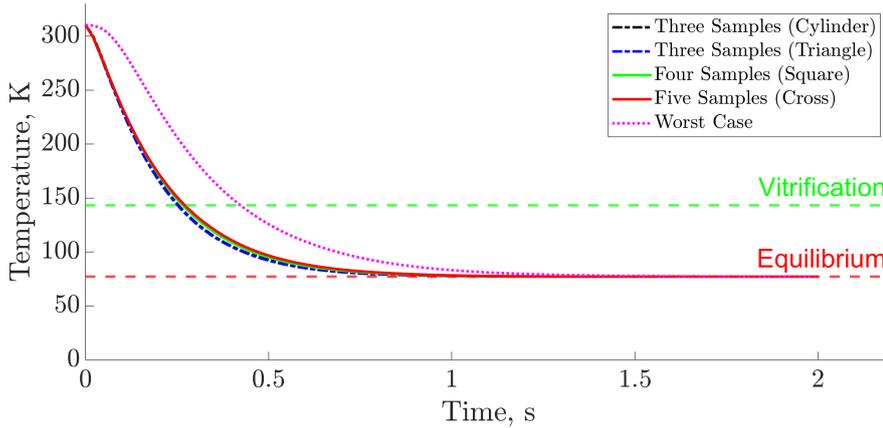


(b) Vitreous water.

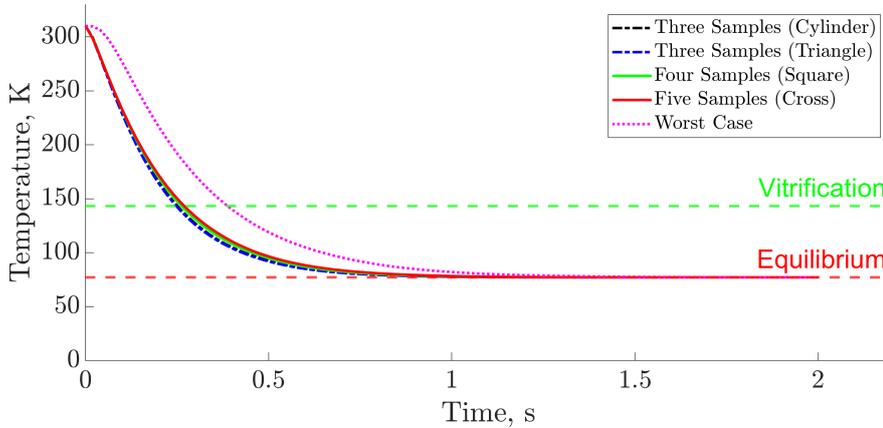


(c) Water with variable thermophysical parameters.

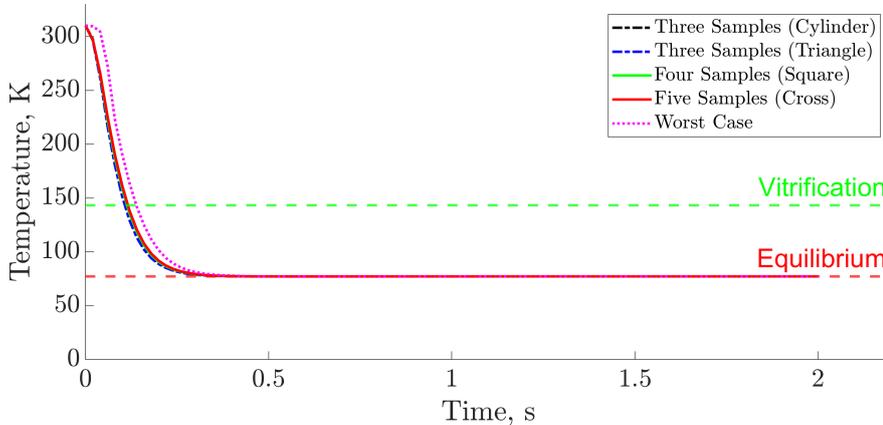
Figure 6.12: The average temperature (Equation (6.2.14)) in samples arranged in a straight line on the Open Cryotop[®] device, namely the ‘Line’ arrangement (Figure 6.11a) with one to three samples, and the ‘Cylinder’ arrangement (Figure 6.11b), which is equal to the ‘3-Line’ arrangement, with a semicircular droplet cross-section. Temperature is determined by numerically resolving the PDE system given by Equations (6.2.4)-(6.2.12) in COMSOL Multiphysics 5.5. Each sub-figure uses a different choice of thermophysical parameters for the CPA solution, with the constant parameters for ethylene glycol and vitreous water given in Table 6.2, and the temperature-dependent thermophysical parameters for water given in Table 6.3.



(a) Ethylene glycol.



(b) Vitreous water.



(c) Water with variable thermophysical parameters.

Figure 6.13: The average temperature (Equation (6.2.14)) in samples in the ‘Cylinder’, ‘Triangle’, ‘Square’, and ‘Cross’ arrangements (Figures 6.11b, 6.11c, 6.11d and 6.11e respectively), as well as the ‘Worst case’ arrangement (Figure 6.11f) that considers a much larger droplet. Temperature is determined by numerically resolving the PDE system given by Equations (6.2.4)-(6.2.12) in COMSOL Multiphysics 5.5. Each of the figures uses a different choice of thermophysical parameters for the CPA solution, with the constant parameters for ethylene glycol and vitreous water given in Table 6.2, and the temperature-dependent thermophysical parameters for water given in Table 6.3.

Model	Number of samples	Time, s		
		EG	vitW	W
Line-1	1	0.26	0.26	0.14
Line-2	2	0.28	0.28	0.16
Line-3	3	0.30	0.30	0.16
Cylinder	3	0.30	0.30	0.16
Triangle	3	0.30	0.30	0.16
Square	4	0.32	0.32	0.16
Cross	5	0.32	0.32	0.16
Worst Case	1	0.48	0.44	0.18

Table 6.4: Time taken (seconds) for the average temperatures inside the samples, plotted in Figures 6.12 and 6.13, to reach the target vitrification temperature, 143.15K. Each column denotes a separate case of the choice of thermophysical parameters, with constant parameters for ethylene glycol (EG) and vitreous water (vitW) available in Table 6.2, as well variable thermophysical parameters for water (W) in Table 6.3.

Model	Cooling rate, K/min		
	EG	vitW	W
Line-1	38541	38541	71576
Line-2	35788	35788	62629
Line-3	33402	33402	62629
Cylinder	33402	33402	62629
Triangle	33402	33402	62629
Square	31314	31314	62629
Cross	31314	31314	62629
Worst Case	20876	22774	55670

Table 6.5: Cooling rates from 310.15K to 143.15K in the simulations, calculated by dividing the temperature difference by the cooling times from Table 6.4. Results shown for the three choices of thermophysical parameters; ethylene glycol (EG), vitreous water (vitW) or water with variable thermophysical parameters (W).

Figure 6.12 indicates that increasing the number of samples leads to a decreased cooling rate. However, we also observe from Table 6.4 that the difference in time taken to reach 143.15K is very small. The difference in cooling times between one and three samples is at most 0.04 seconds in the ethylene glycol and vitreous water cases, and only 0.02 seconds in the case of water with variable thermophysical parameters. Hence, whilst the number of samples does affect cooling rates, the difference is insignificant. Furthermore, we observe that the distance between the two/three sample curves is smaller than between the one/two sample curves in all three variable thermophysical parameter cases. We can hence predict the cooling curve for a four-sample ‘Line’ arrangement will be almost identical to the three-sample arrangement.

Additionally, in Figures 6.12a and 6.12b, we can see that the cooling rate is slightly slower in the ‘Cylinder’ case compared to the three-sample ‘Line’. However, comparison of the cooling times in Table 6.4 show that in all three parameter cases, the actual difference in cooling time is smaller than 0.01 second. Therefore, it can be concluded that these the ‘Cylinder’ and three-sample ‘Line’ arrangements yield almost identical cooling rates, which supports assumption 6.2.1.4, that minor perturbation in droplet shape is unimportant.

Having shown that the cooling rate is dependent on the sample number, up to a point, we next query whether the reduction in cooling rate is large enough to have negative implications for the vitrified samples, by comparing these curves to the Worst Case benchmark. Since the ‘Cylinder’ arrangement is a lower bound for the cooling rate of all ‘Line’ arrangements, it suffices to compare only this arrangement, and not the ‘Line’ arrangements, to the benchmark. Figure 6.13 shows the ‘Cylinder’, ‘Worst Case’ and all other non-linear sample arrangements. The temperature profile for the ‘Cylinder’ and ‘Triangle’ arrangements are almost identical, with Table 6.4 showing that all arrangements with three samples, namely ‘3 Line’, ‘Cylinder’ and ‘Triangle’, all have the same cooling times, regardless of their arrangement. Arrangements with more than three samples, namely ‘Square’ and ‘Cross’ arrangements, have a higher cooling time than 3-sample arrangements, but have no difference between themselves.

The ‘Worst Case’ in Figure 6.13 displays has the smallest cooling rate out of all cases considered. This should be the arrangement with the lowest cooling rate which is still safe, so it can be concluded from Figures 6.12 and 6.13 that all of the other arrangements presented in this work, whose temperature profiles sit well below the ‘Worst case’, are considered safe with respect to cooling rates.

The comparisons we have drawn hold for all three choices of thermophysical parameters (ethylene glycol, vitreous water and temperature-dependent water). By comparing Figures 6.12 and 6.13, we can see that the choice of thermophysical parameters does change the expected cooling rates, as expected by Equation (6.2.16) which shows the temperature has an exponential dependence on the diffusion coefficient. However, the relative shape and distribution of the simulated thermal profiles is constant regardless of the choice of diffusion coefficient, which controls how quickly the curve evolves. For example, cooling is fastest when using the temperature-dependent thermophysical parameters of water, shown in Figures 6.12c and 6.13c, which minimises the effects of the variable geometry and sample number, whilst also minimising the difference between all simulated cases and the Worst Case benchmark. We have already demonstrated in Figure 6.5 that for variable thermophysical parameters, the diffusion coefficients are of order $10^{-5}\text{m}^2/\text{s}$, which is much greater than those reported for ethylene glycol

or vitreous water of the order 10^{-7} . This implies that the samples will cool much faster, despite the increasing thermal conductivity, which suggests that the rate at which the boundary warms up should decrease (see Equation (6.2.3)). The simulations in Figures 6.12c and 6.13 show that whilst actual cooling rates vary depending on thermophysical parameter selection, the relative hierarchy of cooling rates and comparative spatio-temporal thermal profiles remain similar.

The cooling rates in Table 6.5 are of similar order of magnitude to the $23,000^\circ/\text{min}$ (equivalent to K/min) figure reported by Kitazato[®] [350], but we highlight a discrepancy between the cooling rates in our work and [240], despite using the same thermophysical parameters for vitreous water. The droplets modelled in [240] are also spherical caps, with a height of 0.1mm and a width of 0.35mm yielding a droplet volume of approximately $0.02\mu\text{l}$. In our work, the arrangement closest to this is the ‘Line-1’ arrangement, where the droplet volume is also $0.02\mu\text{l}$ [240], yet we report approximately half the cooling rate.

Several factors contribute to this discrepancy, including the different geometry of the two models. In [240], axisymmetric geometry is assumed, which means their model has ‘shaved off’ the rectangular corners of the plate and is, instead, a cylinder. This will increase the surface area to volume ratio, and therefore increases the cooling rate of the plate, compared to the model we use here. Furthermore, the plate material is assumed to be polyethylene in [240], with a diffusion coefficient (thermal diffusivity) of $1.4 \times 10^{-7}\text{m}^2/\text{s}$, which is around double that of polypropylene, that is assumed here (see Table 6.2); this further contributes to an increased cooling through the plate. This effect is demonstrated when tracking the warmest temperature point in each model; in our work, the warmest point is always at the contact between the sample and the plate, whilst in [240] the warmest point is 0.04mm above the plate, almost halfway into the droplet itself. This implies heat transfer through the plate is larger than in our model. Additionally, the contact angle in [240], although not reported, is visibly less than $\pi/2$. The droplet in [240] has height 0.1mm and diameter 0.35mm , relative to our droplet which has height 0.21mm and diameter 0.42mm . Based on these measurements, droplet in [240] has approximately 1.5 times the surface area of the droplet in our work, and as such has a greater exposure to the liquid nitrogen. All these differences contribute to discrepancy in cooling rate between our model and the model in [240].

6.4 Discussion

Determining cooling rates during vitrification of multiple oocytes or embryos is important, as they have direct clinical implications. Comparing the cooling rate

of different sample arrangements cannot be undertaken experimentally, since inserting a thermocouple into an embryo is impossible without the embryos destruction. In this work, we have developed a new mathematical model that predicts the temperature of oocytes and embryos during vitrification. We have shown that differences in cooling rates for different sample arrangements are not large enough to affect survival rates, so long as droplet volumes are constant and within the manufacturer’s limits.

All simulated arrangements have higher cooling rates than those that would be observed in the laboratory, represented by our Worst Case simulation, and are therefore considered to be valid configurations. The small differences in the times taken to reach the target temperature of 143.15K that are observed between different arrangements would be within the error of experimental readings. Although our results are not quantitatively exact, *they do offer a robust “rule of thumb” time scale of around half a second for time to vitrification.* This means that there is no need for embryologists to spend time arranging loaded samples on the plate; instead they should focus on aspirating the medium efficiently. This can save precious time for embryologists, and justifies operating according to their personal preference within the confines of standard operating procedures.

The differences in the cooling rate obtained from our model and other models can be attributed to variations in the geometry and assumptions made in the model construction. Despite the small difference of around 9.2% between our Worst Case cooling rates (22774 K/min in vitreous water) and the 23000 K/min value reported by Kitazato [224], our model yields suitable lower and upper bounds for the cooling rates of samples vitrified in varying numbers and spatial arrangements, both of which are within the safe and appropriate bounds reported for the Open Cryotop[®] device.

Although cooling rates have been used metric for comparing between our model and previous works, we highlight some issues associated with this metric. First, cooling rates have a non-linear dependence on the predetermined temperature interval used (Equation (6.2.16)), so when different authors use different temperature intervals, cooling rates cease to be comparable. Additionally, temperature evolution in the system occurs within the order of half a second, which means that very small absolute differences in the time taken for cooling can result in very large variation in the cooling rate, since it is proportional to the inverse of the time taken for cooling (Equation (6.2.15)). As such, vastly differing cooling rates may appear to justify selecting one technique, or arrangement, over another, when in reality vitrification always occurs within less than half a second. As such, we recommend caution when relying solely on cooling rates as a metric of the quality of vitrification. In order for cooling rates to be truly comparable

between two experiments or models, the same start and end temperature must be agreed upon; in this case, it would be better to report the time elapsed rather than the reciprocal cooling rate, to prevent insignificant differences in vitrification time exploding into seemingly significant differences in cooling rates.

The results of our work are dependent on the validity of the assumptions made. The assumption in Section 6.2.1.4 (droplets are hemispheres) may not be accurate, as different-shaped droplets such as those depicted in [245], which is a rounded cuboid, or the significantly more flat droplet in [240] (although we note that a droplet of height 0.1mm will not actually contain the average early blastocyst, which has mean diameter 0.15mm [363]). The arrangement of samples in different droplet shapes may have different characteristic lengths, and therefore different cooling rates. Such perturbations would lead to small differences in our predicted results.

The question of how many samples can be mounted on a single device is becoming increasingly important. With vitrification, survival rates are always a primary concern [246]. A recent study has shown that survival rates vary depending on the number of embryos placed on a single Open Cryotop[®][238], but were unable to explain why. We therefore questioned whether thermal effects could explain variable survival rates. Our model shows that the only limitation from a thermal perspective is the skill of the embryologist. Specifically, they may mount any number of samples that can be properly stored within 0.1 μ l of vitrification solution within a reasonable time frame. Our work does not assess other factors associated with loading of multiple samples, such as the time taken to process samples, or cost efficiency; we instead only validate that loading of multiple samples per device is thermally justifiable. As a result, we cannot comment on the optimality of any specific arrangement, or number of samples. Instead, it can only be concluded that thermal differences arising from variable sample number and arrangement are an unlikely explanation for variable survival rates.

Further work to refine the model and achieve greater accuracy is possible by revisiting the assumptions in Section 6.2.1. Specifically, the assumptions relating to the size and shape of the samples and droplets, 6.2.1.2 and 6.2.1.4, can be easily altered to account for more realistic droplet shapes, larger droplet volumes, and different sample locations and sizes. Another avenue for further work is to test whether these conclusions also hold true in other devices; the computational model in COMSOL Multiphysics can be easily adapted to achieve this.

Furthermore, one might wish to relax the assumption in Section 6.2.1.1 that liquid nitrogen is a stationary, isothermal liquid. In practice, the liquid is boiling, which involves additional mechanisms of temperature change, such as boiling-induced convection, film boiling and other boiling specific effects [357]. The effect

of these additional mechanisms may differ between arrangements, and therefore may affect cooling rates. This would involve more complex numerical models, which would account for fluid motion at the interface of the droplet and the liquid nitrogen [383].

Additionally, we could relax assumption 6.2.1.6, that there is no Stefan front. This assumption essentially assumes that the solid-liquid interface is unaffected by latent heat effects, and moves instantaneously with the change in temperature of the CPA. A more detailed model would involve a moving boundary, which would reduce the cooling rate. We note that this change would affect all arrangements, including the ‘Worst Case’ benchmark, causing all of them to have a decreased cooling rate. Hence, we would expect the comparison between the ‘Worst Case’ and all other arrangements presented in this work to remain valid, but a more complex moving boundary model is needed to verify this claim.

The choice of parameters in our model could be refined in further work. As discussed, parameter choices for the CPA, polypropylene and liquid nitrogen are taken from literature and COMSOL Multiphysics 5.5, but may not reflect those observed in the clinic. Therefore, there is scope for more experimentally validated parameter choices for the specific materials and environment considered. That said, it is likely the real thermophysical parameters of the CPA will be similar to the value we assumed for ethylene glycol or vitreous water, as stated in Table 6.2. Vitreous water and ethylene glycol diffusion coefficients are of the same order of magnitude (10^{-8}), which means that if the true thermophysical parameters of the CPA are interpolated between these values depending on the relative concentration of each constituent chemical, replacing the parameters used in our model will have little effect on the reported temperature profiles. Additionally, we assumed that samples have the same thermophysical properties of the CPA. Although we justified that samples replace some of their water content with CPA, we cannot quantify the difference between reality and our assumption. We have also not accounted for the effect of the physical structure of the oocyte or embryo on cooling.

In particular, the surface heat transfer coefficient in particular is of interest, since the heat flux at the boundary strongly affects the cooling rate. Making a good choice for the value of h is crucial for making the model as accurate as possible. In our model, it is very easy to replace the values of the parameters, so as more accurate parameter values become available, the model can be quickly updated to generate more accurate approximations of cooling rates.

Despite the emphasis of our work on maximising cooling rates, there is evidence to suggest that this is not the optimal approach to improve survival rates [384], and that there is instead an optimal range for cooling that reduces both the

risk of damage to cells from the CPA for high cooling rates and the risk of intracellular ice formation for low cooling rates. As such, it is appropriate to first look to define this range, and then find arrangements that give cooling rates within this range. Additionally, there is evidence that cooling rates are less significant for the survival of oocytes than warming rates [246]. Indeed, it may be that some balance between cooling rates and warming rates should be sought. Future work could therefore model the effect of spatial arrangement on the warming process.

With some alterations, our model could also be used to simulate the warming process. This would involve setting the initial condition (Equation (6.2.12)) to be a uniform temperature of 77.15K, and to the external temperature u_{ext} to be the temperature of the solution used to warm vitrified samples (in [245], a sucrose solution is used). In accordance with [245], the heat transfer coefficient h would also have to change, since the warming solution does not exhibit the boiling properties of liquid nitrogen. The absence of this boiling phenomena, however, could require a more material alteration of the model if a lower value of h being taken in the warming model results in significantly slower temperature change; this may require that the moving boundary problem be re-introduced to tackle the ‘reverse’ challenge of sample warming.

Despite the shortcomings of our model, the simulations provide the correct order of magnitude for the cooling rates, in line with experimental findings. Specifically, as long as the volume of the CPA droplet does not exceed the maximum volume given by the manufacturer’s guidelines, the cooling process should be on the order of half a second, or less, regardless of the spatial arrangement or number of mounted samples. This means that all arrangements considered in this work exhibit high enough cooling rates to facilitate successful vitrification.

Chapter 7

Statistical analysis of time-lapse images of post-thaw embryos

7.1 Introduction

An embryo which has been safely cryopreserved, as discussed in Chapter 6, may then be thawed and used in a future treatment cycle. We recall that the process of using an embryo which has been previously cryopreserved, and then thawed, is known as a Frozen Embryo Transfer (FET) [66, 385], whilst the term ‘fresh embryo transfer’ refers to those which do not undergo cryopreservation. During vitrification, fluid is removed from the blastocoel of the embryo (see Figure 1.7), either from loss due to osmosis or through intervention in the clinic. This is desirable, since it minimises the volume of fluid inside the blastocyst which must be replaced during vitrification [36]. As a result of this fluid loss, the embryo undergoes collapse, and is physically reduced in size. During thawing, re-expansion occurs, such that the embryo regains fluid and is restored to its original size. The objective of this chapter is to determine whether measurements of the re-expansion of the embryo can be used to predict viability and potential to form a clinical pregnancy.

An embryo at the blastocyst stage is a multicellular structure, and study of the intra-cytoplasmic behaviour is considerably more challenging than in oocytes; we cannot apply DDM or PIV as we did for oocytes in Chapter 2.2. Hence, we are therefore required to study the blastocyst through morphokinetics. Specifically, here we study the re-expansion of embryos following warming from cryopreservation. In Figure 7.1, which is a reproduction of Figure 1.2 in Chapter 1, we show the re-expansion process for a single blastocyst. The re-expansion process begins with the cells of the blastocyst being shrunken and compacted into a tight ball inside the zona pellucida, and is expected to end with the blastocyst either returning to its fully expanded pre-vitrified state (Figure 1.7), or hatching from the zona pellucida. Some embryos may not fully expand, however.

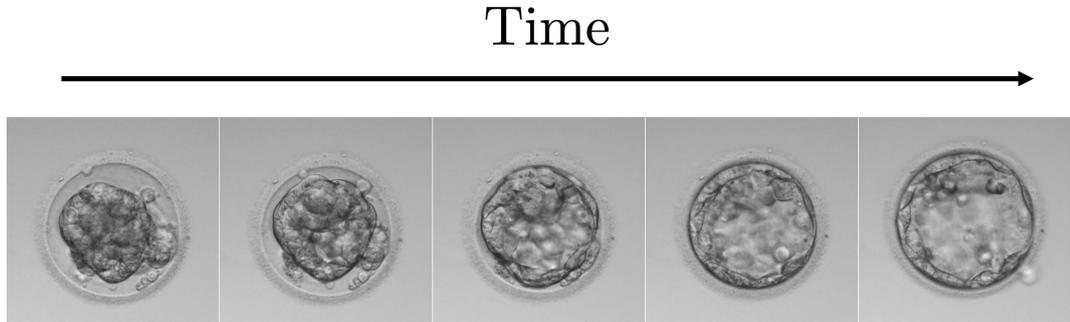


Figure 7.1: A reproduction of Figure 1.2 in Chapter 1, the re-expansion of a single blastocyst is shown. The images are approximately 25 minutes apart, showing the initially compacted and shrunken blastocyst re-expanding to its pre-vitrified state.

In this chapter we wish to perform analysis of time-lapse images of cryopreserved embryos throughout this re-expansion process, in order to identify parameters which can predict clinical pregnancy (recall that clinical pregnancy is defined by the established presence of a fetal heartbeat via ultrasound [65]). First, in Section 7.2, we introduce new time-lapse images of expanding blastocysts, made available by LWC. Then, we apply a neural network approach to automatically extract the blastocyst area from the images in this dataset. Using this machine-learning method to process the images, in Section 7.3 we undertake statistical analysis to identify a metric that allows us to partition embryos into sets with different clinical pregnancy rates, suggesting a new criterion for post-thaw embryo selection.

7.2 Machine-learning image segmentation

Images and other data were gathered from IVF treatments conducted in LWC from 2018-2020. Consent to use outcome data was obtained for each patient, and records are anonymised. For each of the 147 FET recorded cycles, in addition to time-lapse imaging, the following were collected: the presence of a positive pregnancy test, registered clinical pregnancy and live birth. Time-lapse movies acquired in a GERI[®] incubator capture images every 5 minutes, showing the re-expansion of each embryo immediately following thawing. Once re-expanded, the

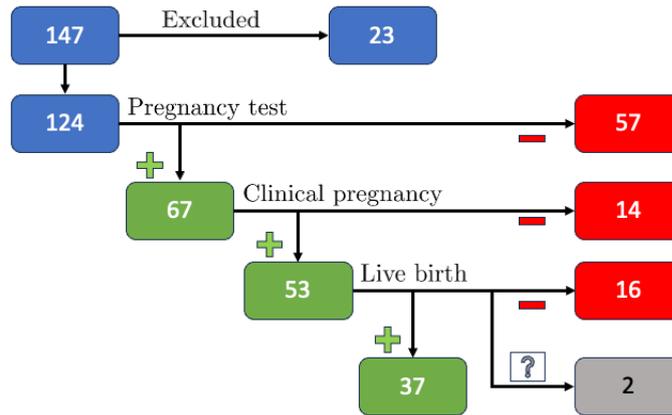


Figure 7.2: Visualisation of the outcomes of the 147 embryos from LWC data. We excluded 23 embryos from the study; 22 were missing corresponding movies, and one did not survive cryopreservation and was not transferred. Of the remaining 124 embryos, 67 resulted in a positive pregnancy test (54.0%), 53 in a clinical pregnancy (41.7%), and 37 in live birth (29.1%). These rates are in line with expected European live birth rates [85]. The two embryos listed as '?' had no recorded live birth outcome data.

embryo remains in culture until the time of transfer. This time period often varies within the clinic, with some embryos remaining in culture far longer than their re-expansion time [36]. The length of the movies is therefore highly variable, with some lasting longer than an hour, whilst others cover as little as 15 minutes.

Of the 147 embryos for which we were given outcome data, 22 embryos were discounted from the study because corresponding movie data was missing, whilst another embryo was not considered because it did not survive cryopreservation and was, therefore, never transferred. Therefore, our dataset consists of 124 embryos. We visualise the outcomes for these embryos in Figure 7.2. The two embryos listed as '?' had no recorded live birth outcome data.

We note that the choice of what we consider a 'positive clinical outcome' is highly important to any statistical analysis we undertake. As we descend through each level of the outcomes (Figure 7.2), the number of positive outcomes decreases. In this work, we choose positive clinical pregnancy as the success outcome, primarily because this outcomes is used in other literature [66, 247], which makes it easier to compare between studies. The overall objective of IVF should obviously be a live birth, but there are a number of confounding variables which may cause a clinical pregnancy to not result in live birth which do not relate to the health of the embryo prior to transfer, such as maternal age and lifestyle [386]. Hence, we focus on potential of the embryo to implant successfully, and focus only on positive clinical pregnancy.

Additional data provided by LWC for each embryo include:

- The date of thaw.
- The number of days between fertilisation and vitrification, either 5 or 6.
- The number of embryos transferred per patient, always one.
- The number of surviving embryos, equal to one for all patients.
- The number of embryos thawed per FET, only greater than one if previous thawed embryos did not survive. This was the case for a single patient, who has two embryos thawed.

Additional comments were also recorded for a selection of embryos with important features to highlight, such as partial hatching, absence of the zona pellucida (ZP) or the presence of twins.

We must now decide what measurements to take from these images. We aim to further the scope of current morphokinetic analysis of post-thaw embryos, by including information from all available time-points, rather than only a select few. It makes sense to therefore measure the blastocyst area as in previous studies [69, 247, 249, 250, 253], i.e. the combined area of the Inner Cell Mass (ICM), trophoctoderm and blastocoel, the fluid filled cavity inside the trophoctoderm (see Figure 1.7). The blastocyst area is, in other works, often estimated by considering an ellipsoidal approximation, and only the semi-major and semi-minor axis are measured. However, this approximation is clearly imperfect, as can be seen in Figure 7.3, where the outer edge of the the blastocyst boundary is not smooth, and clearly asymmetric. Additionally, under such measurement schemes, the area estimates can depend on the orientation of the axes of the ellipse. Here, we introduce a more reliable metric, measuring the blastocyst area instead by the number of pixels it occupies in the image.

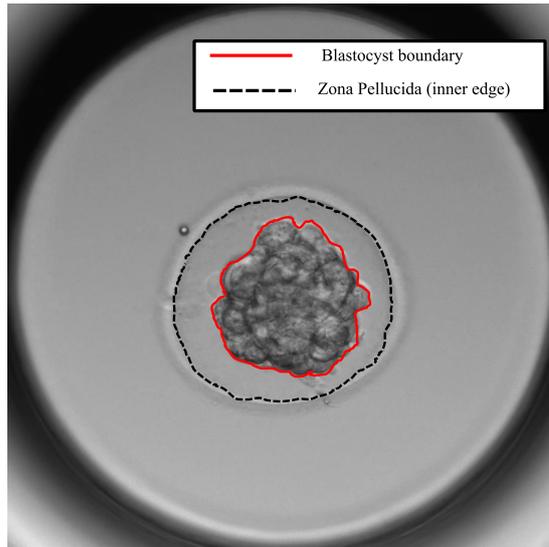


Figure 7.3: A single frame from one time lapse video, at the start of the re-expansion process when the blastocyst is compacted. The two major features of interest are labelled: the border of the blastocyst (red solid line) and the inner edge of the Zona Pellucida (ZP, black dashed line), labelled by hand. The blastocyst area is contained within the red contour, and the perivitelline space is the area contained between the red and black dashed contours.

Automating the measurement of the blastocyst area is also desirable, since in the clinic, manual labelling is an impractical and very time-consuming approach. Evaluating the blastocyst is equivalent to enumerating which pixels within the image correspond to the blastocyst area. This is therefore a *segmentation* task, for which a collection of computer-vision techniques already exist [387]. However, the complexity of biological image data, and in particular that of embryo image data, often prevents the implementation of these techniques. For example, edge detection to identify the boundaries of the zona pellucida (ZP) can frequently fail, since the illumination gradient between the ZP boundary and the background is often locally very small. Additionally, cell fragments and overlapping cells can create challenges in determining the edge of the blastocyst; this challenge becomes increasingly likely as the blastocyst expands and the blastocoel emerges [387].

Machine learning is an optimal tool for tackling this challenge, since the blastocyst is a visually discernable feature in the images, and hence there is a strong pattern to detect. This particular type of machine learning is referred to as semantic segmentation, and has been recently implemented successfully in [124] to segment images of oocytes into regions denoting the cytoplasm, first polar body, zona pellucida, perivitelline space and cumulus/corona cells. In [124], 71 different architectures were tested in order to determine the optimal performing structure. A variant of DeepLab-v3-ResNet-18 was identified to be the strongest performer [124], where the number 18 references the number of layers in the

network. We therefore also adopt the ResNet-18 architecture for segmentation here.

First introduced in 2015, Residual Networks (ResNet) were developed to overcome the problem of ‘degradation’, where adding more layers to a neural network causes first saturation and then a drop in the accuracy of the network [388]. The ResNet architecture overcomes this issue by adding shortcuts between layers in a standard neural network, as shown in Figure 7.4. These shortcuts do not have weights and biases to train, so do not add to the complexity of the network, but have been shown to boost training accuracy [388]. ResNet18 is a popular choice of architecture for biological image segmentation tasks [389–391], and is therefore considered an appropriate choice in this work.

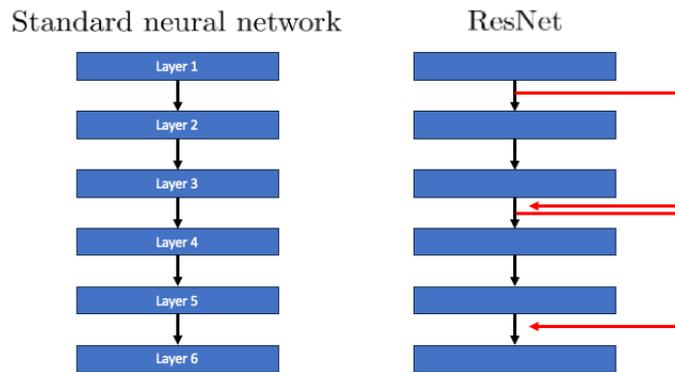


Figure 7.4: ResNet architectures in machine-learning models differ from standard neural networks due to the inclusion of ‘shortcuts’ that can skip layers. Adapted from [388].

The purpose of the segmentation is to assign a label to each pixel in the image, where the labels are elements of a set of ‘classes’. We define three classes,

- l_1 = Background.
- l_2 = Blastocyst area.
- l_3 = Perivitelline space.

The neural network we implement is a form of *supervised* machine learning, which means learning is driven by encouraging the network to reproduce segmentation in a *labelled* training dataset. We use the MATLAB 2022b image labelling app to assign pixel labels to each frame of 40 embryos which we decide to use as a training set. This results in a total of 986 labelled frames; an example is shown in Figure 7.5. By considering whole movies in the training set, we ensure that we capture a roughly equal sample of different stages of expansion for each of the 40 embryos.

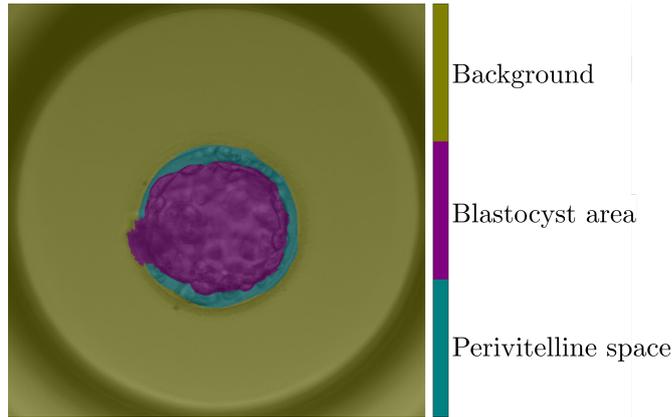


Figure 7.5: A single frame from an embryo expansion movie provided by LWC, manually labelled using the MATLAB 2022b image labelling app.

The set of labelled images is split into training, testing and validation subsets using a 60/20/20% split, with random allocation to each set by performing a random permutation on the ordered list of images, then taking the first 60% for training, the next 20% for testing, and the remaining images for validation. The cost function for a semantic segmentation needs careful consideration due to a disparity between the relative number of labels assigned to each class; the majority of pixels are labelled as Background (l_1), such that there are much fewer labels assigned to Blastocyst area (l_2) or the Perivitelline space (l_3). If I pick a random pixel from the image, the probability it will belong to the background is much higher than any other class, which can bias the network towards over-estimating the probability that any given pixel is part of the background. This issue is known as class imbalance [392], which we control for by weighting the cost function using the prevalence of that label in the training set. Based on some trial and error, we choose to train the network over a maximum of 5 epochs (the number of times the training data is passed through completely), in batches of 16 (the number of images passed through before updating weights and biases). These values optimised fitting performance in the validation set, ensuring we train sufficiently without over-fitting. We adopt the pre-built ResNet-18 architecture available through MATLAB 2022b, achieving an accuracy of 98.14% in the validation set. With the trained network, we can now process the image data available in order to identify morphokinetic parameters which may correlate with success.

7.3 Non-linear expansion rates predict negative clinical pregnancy

We use the trained machine learning algorithm to segment all 124 time-lapse movies. This reduces each movie to a set of pixel labels, which we enumerate

to generate dataset for the blastocyst area at each time point. Since the frame size is constant, the blastocyst area can be represented as a percentage of the frame which it covers. We can plot this area percentage as a function of time for a single embryo, as shown in Figure 7.6a. The objective of our analysis is to infer embryo health from these plots, but comparing the plots of two embryos is made difficult by the variable embryo size, and variable length of imaging. The plot for each embryo will have an entirely different range and domain, making it difficult to spot patterns which discriminate between positive and negative clinical pregnancy. Instead, we choose to compare these plots for different embryos by plotting the blastocyst area in frame number t on the horizontal axis, against the area in the $(t + 1)^{\text{th}}$ frame on the vertical, as in Figure 7.6b. Although the initial and final size of the blastocyst area will now define the range and domain of these plots, the length of imaging is no longer an issue when comparing two embryos. The benefit of this representation is that we can more easily spot patterns relating to the *uniformity* of the embryo expansion; embryos which re-expand at a slow, uniform rate will remain close to the $y = x$ curve, whilst embryos with sudden and large changes in area will produce points far from the positive diagonal.

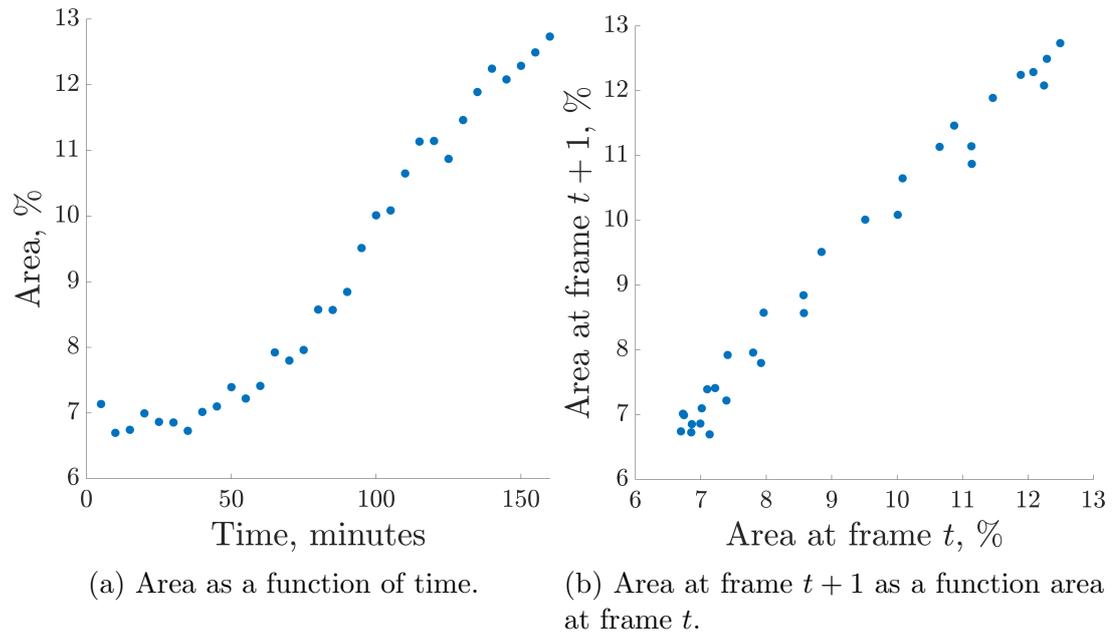


Figure 7.6: The blastocyst area in a single embryo, plotted over time, in two different representations. In Figure 7.6a, the area is plotted as a function of time, in frames (where one frame is taken every 5 minutes), whilst in 7.6b, we plot the blastocyst area in the frame $t + 1$ against the blastocyst at frame t .

Using this new representation, we draw a scatter plot of all measured blastocyst areas from all embryos in Figure 7.7, such that each point represents a single frame from one embryo. We denote the embryos which resulted in a clinical pregnancy

with open circles, and those which did not with a cross. It appears to be a trend that the further away the data points are from the $y = x$ curve, the more likely that the embryo did not result in a clinical pregnancy.

Different colours are used to denote which embryo a point belongs to, so we note that many of the points farthest from the $y = x$ curve originate from the same blastocyst. Hence, one should interpret these points as saying that for some embryos, which did not result in a clinical pregnancy, the expansion or contraction jump was large for many time points. These embryos may be considered to behave more erratically.

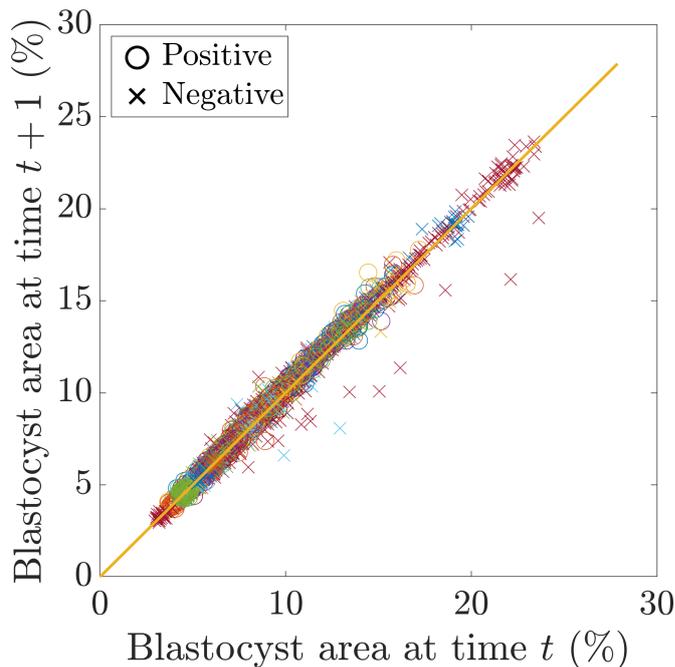


Figure 7.7: Scatter plot of the blastocyst area, at time $t + 1$ against the area at time t , for the 124 embryos, provided by LWC. Each point represents a single frame of an embryo, with crosses marking embryos which did not result in a clinical pregnancy, and open circles marking those which did. In general, the points farthest from the $y = x$ curve tend to be originate from embryos which did not result in a clinical pregnancy. Colour denotes which point originated from which embryo.

The correlation between large jumps in expansion, or contraction, and negative clinical outcomes suggests that the size of each change in blastocyst area might be an indicator of poor health of the embryo; this motivates a new metric which enumerates the magnitude of the change of the area of the blastocyst, as follows. For an embryo whose blastocyst area expands or contracts within certain bounds, denoted $\pm\Delta\mathcal{A}$, the graph in Figure 7.6b would be of the form

$$y = x \pm \Delta\mathcal{A}. \tag{7.3.1}$$

whilst embryos with area remaining constant will generate plots whose points all sit in the same location on the line $y = x$, since the area never changes. An embryo whose blastocyst area changes by a value greater than $|\Delta\mathcal{A}|$ at any given time will result in a curve which moves outside of these bounds. Therefore, we can think of $\Delta\mathcal{A}$ as a threshold value, against which we can compare the magnitude of the change in blastocyst area. We can plot the blastocyst area for an embryo as in Figure 7.8, draw lines a distance $\Delta\mathcal{A}$ above and below the line $y = x$, and classify the embryos as either:

- In-Bounds (IB), if all area points sit inside the bounding lines.
- Out-of-Bounds (OB), if at least one point is outside of the bounding lines.

We give an example of the above classification method in Figure 7.8; the classification of embryos changes from IB to OB as the size of $\Delta\mathcal{A}$ decreases. It is our hypothesis that more erratic re-expansion is an indicator of lower probability of clinical pregnancy, so we propose that there exists a particular choice of $|\Delta\mathcal{A}|$ for which embryos classified as IB have a higher probability of resulting in clinical pregnancy than those assigned the class of OB. It remains now to identify if such a value of $|\Delta\mathcal{A}|$ exists.

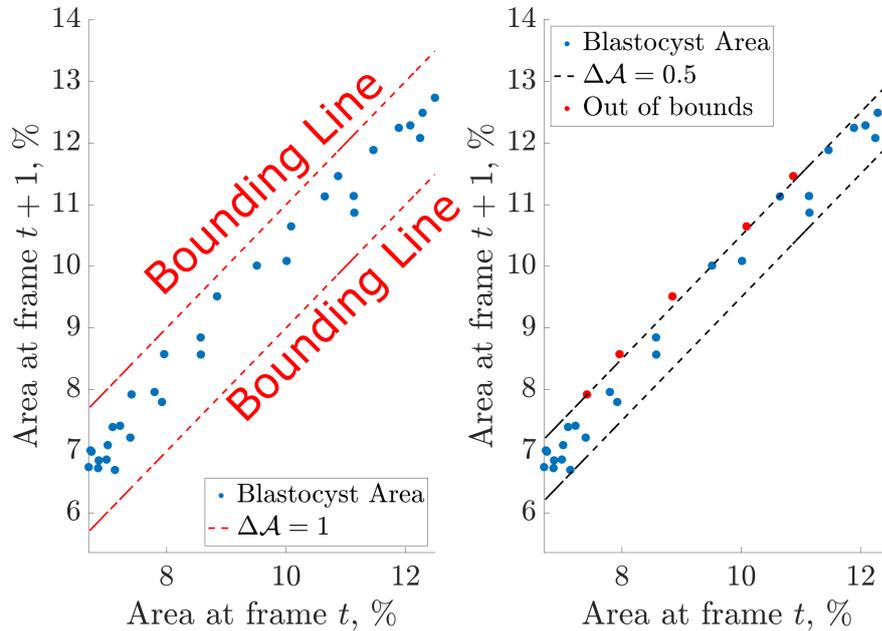


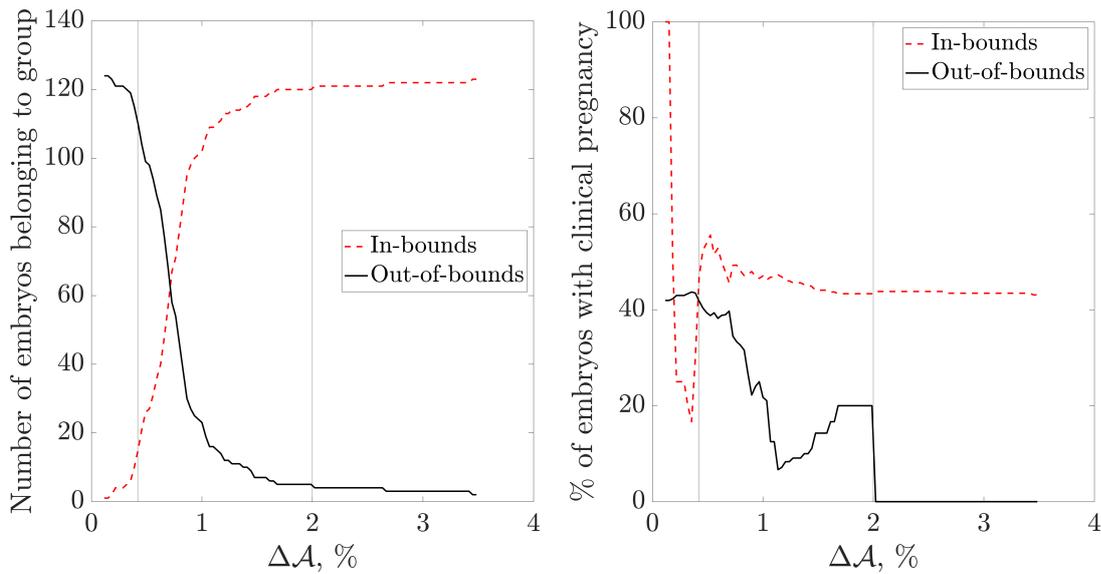
Figure 7.8: The same embryo data plotted in Figure 7.6, with bounding lines given by Equation (7.3.1). When $\Delta\mathcal{A} = 1\%$, the embryo remains inside these bounds since the biggest change in the blastocyst area is smaller than this value, but for $\Delta\mathcal{A} = 0.5\%$, a number of points (filled red circles) sit outside of the boundary. These points would cause the embryo to be classified as ‘OB’

This metric moves beyond current morphokinetic approaches, by considering information from every frame in the movie, rather than only comparing the start

and end points. In order for this metric to be useful, we would need to find some value of $\Delta\mathcal{A}$ which maximises the partition between embryos resulting in positive clinical pregnancy, and those which do not. If ‘cp+’ denotes the outcome of positive clinical pregnancy, and IB the outcome that an embryo is classified as in bounds, the conditional probability P of clinic pregnancy given IB is defined by Bayes theorem to be

$$P(\text{cp} + | \text{IB}) = \frac{P(\text{cp} + \cap \text{IB})}{P(\text{IB})},$$

which can be evaluated as the number of IB positive clinical pregnancy embryos divided by the total number of IB embryos. When there are no embryos classified as IB, we assign the probability of clinical pregnancy to be zero, since otherwise the conditional probability is undefined. We now consider our embryo dataset (124 embryos), and consider the conditional probability of clinical pregnancy given an embryo is OB, for a range of $\Delta\mathcal{A}$. We plot the size of the sets IB and OB in Figure 7.9a, and the conditional probabilities of positive clinical pregnancy in Figure 7.9b, both against the choice of $\Delta\mathcal{A}$.



(a) The number of embryos classified as IB or OB for each $\Delta\mathcal{A}$.

(b) The percentage of embryos that generated clinical pregnancy, for each classification IB or OB.

Figure 7.9: (7.9a) The number of embryos labelled IB or OB as a function of $\Delta\mathcal{A}$. (7.9b) the percentage of embryos that led to clinical pregnancy for embryos classified as IB and OB, as a function of $\Delta\mathcal{A}$. When $\Delta\mathcal{A} > 0.42\%$, IB embryos are more likely to generate a clinical pregnancy than OB embryos, and beyond $\Delta\mathcal{A} = 2\%$, all OB embryos had a negative clinical pregnancy outcome.

From Figure 7.9b, we highlight two key values of $\Delta\mathcal{A}$. The first occurs at $\Delta\mathcal{A} > 0.42\%$, for which the likelihood of clinical pregnancy is higher for IB

embryos than OB. This represents the smallest value of $\Delta\mathcal{A}$ for which OB embryos have a lower probability of positive clinical pregnancy than IB embryos. We note that the point at which the dataset is split exactly in half, between IB to OB, is around $\Delta\mathcal{A} = 0.8\%$, where the clinical pregnancy rate is around 40% for OB embryos and 50% for IB. When $\Delta\mathcal{A} = 0.8\%$, we therefore have a new metric which partitions the dataset of 124 embryos approximately equally, in such a way that one partition is more likely to result in a clinical pregnancy than the other. Whilst OB embryos appear to have a higher clinical pregnancy rate than IB for $\Delta\mathcal{A} < 0.42$, the sample size of IB embryos is less than 15, and thus is too small to be certain of the validity of these results.

The second value of $\Delta\mathcal{A}$ to note is $\Delta\mathcal{A} > 2\%$. When $\Delta\mathcal{A} > 2\%$, all remaining OB embryos did not lead to clinical pregnancy. Whilst this extreme choice of $\Delta\mathcal{A}$ could be a way of identifying embryos which are guaranteed to not result in a clinical pregnancy, this considers a sample size of only 4 embryos classified as OB, which is not a large enough group to allow a confident conclusion. Therefore, the interval of interest is $0.42 < \Delta\mathcal{A} < 2\%$. When the largest change in area for an embryo in this dataset between two consecutive frames is greater than $\Delta\mathcal{A}$ in this interval, the embryo is less likely to result in a clinical pregnancy than an embryo which does not change by more than $\Delta\mathcal{A}$. We now explore the statistical certainty of this result.

One way in which we can quantify the statistical certainty that the observed results reflect a genuine pattern in the data, and are not a result of chance, is to apply a χ^2 test to determine whether IB embryos have a greater clinical pregnancy rate than OB embryos, for some $\Delta\mathcal{A}$. The null hypothesis of this test is that the probability of clinical pregnancy does not depend on the embryo classification, IB or OB. That is, under the null hypothesis, we assume that the clinical pregnancy rate for the IB group will be equal to the clinical pregnancy rate of the whole set of 124 embryos, found by the number of clinical pregnancies over the number of embryos i.e. $53/124 = 0.4274$ (see Figure 7.2). We, thus, define the null and alternate hypothesis as follows:

- H_0 : The probability of clinical pregnancy for the IB group is $p = 0.4274$.
- H_1 : The probability of clinical pregnancy for the IB group is $p \neq 0.4274$.

For each $\Delta\mathcal{A}$, we assign each embryo to one of four groups, defined by the pair of characteristics of being either positive or negative clinical pregnancy, and IB or OB. We can also predict the expected number of embryos in each of these groups, by multiplying the number of IB/OB embryos by the expected success

probability in H_0 . We may then apply the χ^2 test, using the test statistic

$$\chi^2 = \sum_{k=1}^n \frac{(O_k - E_k)^2}{E_k},$$

where O_k denotes the observed number of embryos in a group and E_k denotes the expected number of embryos in that group.

We have one degree of freedom in this system, since there is only one independent variable. We reject H_0 if the χ^2 statistic exceeds the critical value of 3.84, which corresponds to a p-value of 0.05. We plot the χ^2 test statistic in Figure 7.10, for each $\Delta\mathcal{A}$; this is greater than the critical value 3.84 when $0.86\% < \Delta\mathcal{A} < 1.44\%$. This means that the probability of clinical pregnancy given an embryo is IB is statistically significantly higher than that of the general cohort, for these $\Delta\mathcal{A}$. Within this region, selecting only IB embryos improves clinical pregnancy rates by between 0.6-3%.

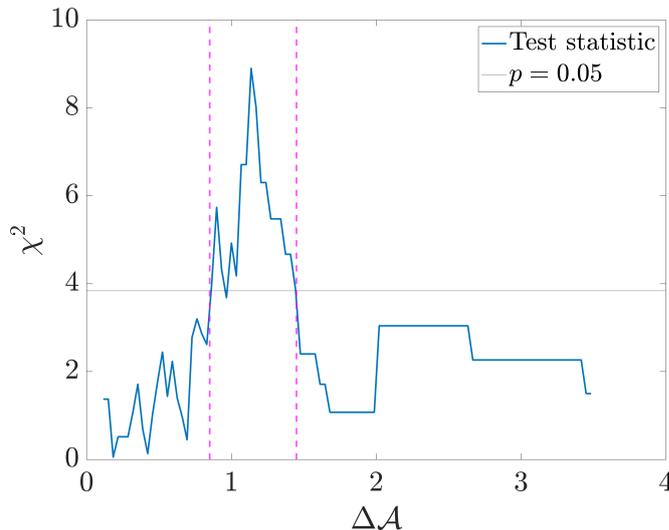


Figure 7.10: The χ^2 test statistic value as $\Delta\mathcal{A}$ varies from 0 to 3.5%, compared to the critical value 3.84 corresponding to the p-value $p = 0.05$. Any test statistic value over the critical value 3.84 is considered statistically significant.

Another measurement of the difference between groups is the odds ratio, defined as

$$\frac{\# \text{ Positive IB} \times \# \text{ Negative OB}}{\# \text{ Negative IB} \times \# \text{ Positive OB}}.$$

The odds ratio when $0.86\% < \Delta\mathcal{A} < 1.44\%$ region varies between 2.45-12.55, so picking IB embryos means you are 2-10 times more likely to get a clinical pregnancy compared to choosing OB embryos.

We have, therefore, found a new metric which, in this dataset of 124 embryos, appears to partition the cohort of embryos into two sets based on the size of the largest change in the blastocyst area, where the two sets have statistically

significant differences in clinical pregnancy probability. Our new metric considers large contractions and expansions to result in an embryo being assigned the OB label, but we wish to explore whether a combination of the two phenomena drives the partitioning of our data, or whether either contraction or expansion dominate the larger changes in the blastocyst area.

We, thus, plot histograms in Figure 7.11 that show the distribution of changes in blastocyst area in the 124 embryos, separately for embryos which resulted in clinical pregnancy and those which did not. Embryos which result in negative clinical pregnancy are skewed more towards negative changes in area than positive outcome embryos, with a heavy tail that suggests larger contraction predicts negative outcome. This suggests that it is primarily contraction which drives the classification of an embryo as OB. However, we also note that the range of blastocyst area change for positive outcome embryos is contained within the range observed for negative outcome embryos, suggesting that exceeding some upper threshold on expansion could still be indicative of negative outcome. In this case, it may be that some asymmetric threshold on the acceptable contraction and expansion rates would provide a better partition than our current symmetric approach.

7.4 Summary and discussion

In this chapter, we analysed time-lapse image data of post-thaw expanding embryos, and identified a possible metric which might indicate the viability of the embryo based on its re-expansion behaviour. The clinical embryo data are provided by the London Women’s Clinic, and are movies and additional data for human embryos thawed and used in Frozen Embryo Transfers (FET). This metric allows us to correlate morphokinetic parameters with clinical outcome, in a more complete manner than in previous studies.

First; we implemented a neural network to determine the blastocyst area in each time frame of the movies through semantic segmentation, which allowed us to quickly extract key morphokinetic parameters. By plotting the blastocyst area at each time point against the blastocyst area at the next time frame, we identified that large expansion or contraction correlates with negative clinical pregnancy outcomes. This led to a new metric which partitioned the dataset into In-Bounds (IB) or Out-of-Bounds (OB) embryos, depending on whether the embryo expansion or contraction was bounded above by some threshold, $\Delta\mathcal{A}$ (units % of the whole frame size), at each time step. When $0.86 < \Delta\mathcal{A} < 1.44\%$, IB embryos were statistically significantly more likely to result in a clinical pregnancy than OB embryos. We also showed in Figure 7.11 that classification as

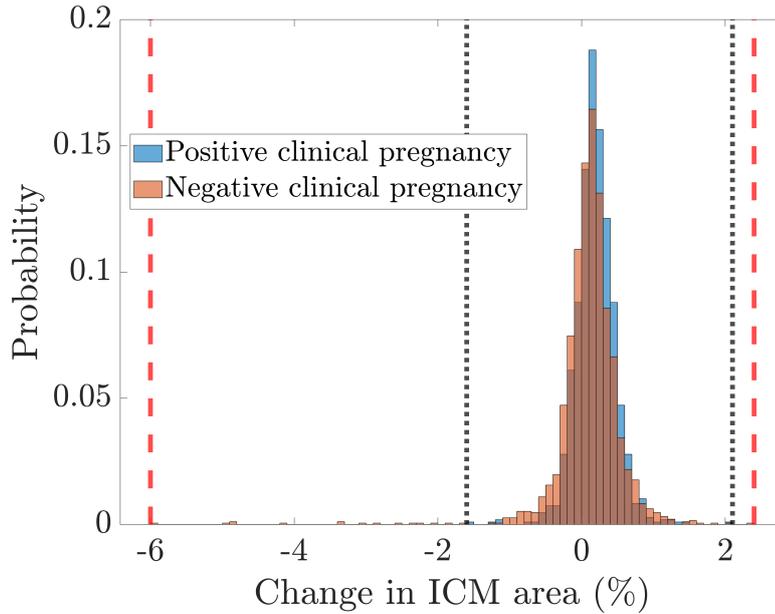


Figure 7.11: Histogram exhibiting the Probability Density Function (PDF) of the changes in the blastocyst area between each time frame in the time-lapse movies; plotted separately for embryos with a positive and negative clinical pregnancy outcome. Black dotted vertical lines denote the upper and lower bounds on the blastocyst area change for embryos with a positive clinical outcome, and red dashed lines show the upper and lower limits on the blastocyst area change for negative clinical outcome. Embryos resulting in negative clinical pregnancy have wider tails in PDF, and are skewed negatively compared to the embryos resulting in positive clinical pregnancy.

OB was predominantly driven by the observation of large contractions, although exceptionally large expansions also appeared to indicate negative outcome.

The metric identified here extends the existing approaches to using morphokinetic parameters in post-thaw embryos, by considering measurements of the embryo at each individual time points during expansion, rather than waiting for expansion to conclude [69, 249, 252, 253], or considering only the change between the beginning and end of expansion [78, 247, 250, 255]. Our approach complements existing work which studies the relationship between contraction during post-thaw expansion and clinical outcome [247, 250, 253, 254], since our metric is mostly based on identifying contraction, as proven above in Figure 7.11.

Whilst [253] explored the frequency of contraction, our approach is more similar to that of [254], which correlated the size of contraction with clinical outcome in samples of 30 mouse blastocysts. Although our study uses a much larger sample size than [254], our results agree with this work, by suggesting that there is some limit in the size of contraction, above which negative clinical outcome is more likely. Our results are more directly relevant to the clinic, since we use human embryo images instead of mouse images, and a larger sample size. Additionally,

we considered a threshold not only on contraction, but also expansion, such that our study captures more anomalous behaviour in the blastocyst re-expansion.

Although our work suggests that the proposed metric for classifying embryos could lead to increased clinical pregnancy rates, we have been careful to identify that without some biological explanation for how absolute blastocyst area change correlates with embryo health, this metric is not immediately clinically useable. For example, contraction frequency has been theorised to relate to the health of the embryo through the physiology of the trophoctoderm, with less dense cell arrangement resulting in fewer ion channels that drive osmotic expansion of the blastocoel [247]. The link between significant changes to the cell volume during re-expansion, and lower rates of clinical pregnancy, could be similarly related to the function of the ion channels. Re-expansion is the result of the movement of water into the cell, which can be accelerated through the intake of sodium ions. This is an active process, demanding the use of ATP, which could suggest that embryos with well-regulated re-expansion have larger ATP reserves than embryos which collapse, or re-expand too quickly. If true, this would imply that low variance in re-expansion is indicative of higher embryo viability. In order to explain and validate our proposed metric, future work is needed to design controlled experiments which identify biological health markers correlating with the rate of expansion and contraction.

Another direction for future work would be to consider additional features beyond the blastocyst area. In the movies in the considered dataset, different embryos exhibited different behaviours that could have the potential to discern between them, based on health. For example, some embryos demonstrated cell fragmentation, which has been identified as a possible sign of negative outcome [393]. Our approach focussed on only one measurement, the blastocyst area, but in reality, some combination of measurements could be better predictors of viability.

One limitation of our work is that the accuracy with which we measure the size of the blastocyst depends on the performance of the neural network used for semantic segmentation. Although we demonstrated that the network achieved over 98% accuracy during training, the labels themselves were provided by the author, and not a trained embryologist. Lack of expertise in label definition may limit the performance of the neural network compared to the ‘true value’, which would affect all subsequent analysis. As such, we recommend that this study is repeated with labelling by an embryologist, and with a larger annotated training set, in order to increase our confidence in the results.

Another limitation is that, as a retrospective study, our results suffer from survivorship bias, defined to be error produced by focussing on data from subjects

which have already satisfied some selection criteria, whilst ignoring those which did not [394]. We only have here the data from embryos which have been used in FET cycles; we do not know how well the alternative embryos might have performed. This means that even if our metric does have some biological motivation, we cannot predict how much benefit its implementation would have in the clinic. For example, suppose maternal genetics are the root cause of excessively large contraction or expansion, and therefore identifying the embryo as a poor candidate. Then it is possible that every embryo available in the cohort would have the same OB classification; in this scenario, classification has no benefit, because this new information cannot motivate improved embryo selection. Survivorship bias is a common problem in retrospective studies, and is mitigated only through control studies. The requirement that new morphokinetic parameters be subject to control studies is highlighted in [199], who observe that the many seemingly useful morphokinetic parameters which are identified in a retrospective study subsequently fail to correlate with clinical outcome in other datasets.

Even if we can find some biological explanation for the results observed in this work, and can prove through controlled studies that the selection criteria would have some benefit in clinical pregnancy rates, the observed benefit of using our metric in FET needs to be much higher than in fresh cycles to justify its use. In fresh cycles, we have some cohort of embryos or oocytes to choose from, and any unused embryos can be vitrified for use in subsequent cycles, if needed. Hence, any improvement in clinical pregnancy rate gained purely from optimised selection, regardless of magnitude, is worth implementing in the clinic, since selection has no disadvantages. In FETs, however, we only thaw one embryo at a time, and have to make the decision to either use the embryo, or thaw another and discard the current candidate. Making this decision is difficult for the embryologist; consider the scenario in which an embryo is thawed, and given a ‘bad’ classification that predicts a low probability of success compared to a ‘good’ embryo. The decision to choose another embryo is motivated by the probability that the next embryo might have a ‘good’ classification, the size of the potential increase in clinical pregnancy rate which a ‘good’ embryo could provide, and the number of embryos a patient has available.

Furthermore, the opinion of the patient is a vital consideration in this decision. Choosing to discard embryos is an emotionally challenging decision for patients [395, 396], so some patients may prefer to use an embryo with a slightly lower success rate, rather than destroy an embryo in the hope of improving their likelihood of getting pregnant. To provide the best outcome for patients in terms of maximising their chances of getting pregnant whilst also allowing them to retain autonomy in the decision making process, there must be some overwhelming

advantage to moving on to the next candidate in order to justify this decision. This typically means that there must be evidence that the current embryo is guaranteed to not result in pregnancy. As such, selection criteria that result in a raised probability of success may be less useful in FET than criteria which identify only a few embryos with exceptionally poor prospects.

In addition to identifying a range of blastocyst area changes for which the success rate for IB embryos was elevated over OB embryos, we also identified that for above an area change threshold of 2% of the frame size, all OB embryos resulted in a negative clinical pregnancy. We only had four embryos in this sample, which is insufficient for statistical certainty in our observation, but this type of exclusion criteria would have much more power than cohort-selection criteria in a FET setting. Identifying and discarding embryos which are almost certain to fail will improve clinical pregnancy rates, although by very small amounts. However, if multiple exclusion criteria can be identified, the sum of these marginal gains could be clinically significant. As such, we propose that future research into the morphokinetics of post-thaw embryos should focus on identifying these incredibly poor embryos with very low success rates.

Despite the outlined flaws and limitations of this work, the results presented here show that there is promise in morphokinetic analysis of time-lapse images of post-thaw embryo re-expansion. Time-lapse incubators have offered clinics a wealth of new imaging data, which may offer new insight into embryo health. Morphokinetics are a non-invasive tool which have the potential to improve selection processes and, hence, clinical pregnancy rates, at no additional cost to clinics already using time-lapse incubators. Furthermore, automated image analysis software, if designed in conjunction with embryologists, can allow advanced analysis techniques to be easily incorporated into clinical protocol, with immediate benefit to patients. Whilst we do not claim to have ‘solved’ post-thaw embryo health assessment, our results show promise in the 124 embryos studied, and we encourage further investigation into using measurements of re-expansion to assess the viability of embryos.

Chapter 8

Summary, discussion and future work

IVF, as a treatment for infertility, has revolutionised assisted reproductive technologies (ART) [5, 15, 16]. Development of cryopreservation technology has further advanced the use of IVF, offering more control over family planning [37, 38], protecting against fertility threats such as cancer treatment [23], and preserving leftover embryos from previous IVF cycles [32, 33]. The booming fertility industry has evolved to meet the increasing demand for IVF [29] and improve treatment standards, but there are still several challenges to overcome. The IVF success rate (# of live births per treatment cycle), summarised in Table 1.1, remains low, at only 19% world-wide [7, 83, 85–87]. Whilst success rates have improved since the 1980s, when IVF first became available, they have stagnated in recent years [79]. The high cost of IVF is also a barrier for many patients who have to seek treatment privately [89], when they are not eligible for state-funded treatment, or are faced with long waiting lists [92]. Improving IVF success rates would benefit patients, clinics and society at large. Mathematics can play a central role in tackling important IVF challenges [101], through mathematical modelling, data science and the use of machine learning. In this thesis, we have used a combination of mathematical methodologies to tackle some of these IVF challenges, in collaboration with the London Women’s Clinic (LWC) [10], as well as experimental biologists and imaging specialists, respectively, from the Schools of Biosciences and Physics and Astronomy, Cardiff University. In particular, we focus on challenges related to non-invasive oocyte assessment using time-lapse microscopy, and cryopreservation of oocytes and embryos. Below, we summarise the main results and conclusions obtained in this thesis and outline future work directions.

In the Introduction (Section 1.2), we described a number of important mathematical challenges faced in the IVF clinic. In particular, we highlighted that identifying which oocytes are most likely to develop into healthy embryos is a

key consideration for improving IVF success rates. Current oocyte assessment methods mostly focus on identifying abnormalities in the oocyte morphology which may be indicative of poor health [109, 119, 122], but these methods remain unreliable, as many seemingly normal oocytes do not result in a healthy pregnancy [109, 122]. In Section 1.2.2.1, we discussed how measurements of cytoplasmic movement, extracted from time-lapse imaging of oocytes, may predict oocyte health. Of the available image analysis techniques, summarised in Section 2.1.2, we chose to investigate Differential Dynamic Microscopy (DDM) as an oocyte health assessment technique. Specifically, we use DDM to describe the randomly distributed, spatially incoherent advection-diffusion behaviours which drive cytoplasmic movement (see Figure 1.6).

In Section 2.2, we first provided an in-depth review of the theory of DDM; this is, to our knowledge, the most comprehensive overview of DDM to date. We consider this a very useful contribution to the field given the often sparse details provided in previous literature [164, 266, 267]; we hope that this will make DDM more accessible to other researchers in the various fields in which it has been applied, including soft matter [397], chemical physics [311] and mathematical biology [142, 163, 270]. In this review, we examined the numerous assumptions necessary for DDM, and identified a set of practical considerations that can improve the algorithm’s performance with respect to accuracy and computational workload. These considerations can be applied to a large variety of datasets, both biological and inorganic.

In Section 2.3, we applied DDM analysis to a sample mouse oocyte dataset (provided by the Swann lab, Cardiff University), and drew several conclusions about the potential use of DDM in oocyte assessment in IVF. First, we demonstrated in Section 2.3.1 that cytoplasmic movements in FCCP-poisoned oocytes could be described using the intermediate scattering function (ISF, see Equation (2.2.15)) for Brownian motion, allowing us to extract a diffusion coefficient for each oocyte which is characteristic of the distribution of the speed of its cytoplasmic movement. This demonstrated that DDM can be used to analyse time-lapse microscopy of oocyte data. However, we also identified several challenges faced when applying DDM to oocyte data. Firstly, we need to derive the ISF describing cytoplasmic movement in healthy oocytes. In Sections 2.3.1-2.3.3, we show that the two ISF models available, namely Brownian motion (Equation (2.2.22)) and an undirected advection-diffusion motion (Equation (2.2.24)), fail to describe the cytoplasmic movement in healthy oocytes. Thus, even though the DDM matrix may contain information about the health of the oocytes, we currently lack a suitable model to extract this information. Despite this limitation, when applying the ‘incorrect’ Brownian model to healthy oocytes as well as to poisoned ones,

we were able to assign a diffusion coefficient D to each oocyte, and show that D decreases following poisoning, as expected (see Figure 2.16). This demonstrates that DDM has the potential to inform on underlying oocyte health, but we identified three limitations which prevent us from forming any statistically significant inference of oocyte health using DDM: (i) the lack of model for the ISF means we cannot guarantee our fitted parameters correspond to the observable cytoplasmic movements, (ii) with weak parameter fitting approaches, we cannot validate whether our fitted movement parameters match the ‘true’ value for each oocyte, and (iii) the DDM assumption that there are no boundary effects (Section 2.2.8.4) requires large amounts of image data to be discarded. These three challenges may be found in other image datasets, and are therefore are not limited to our application of DDM to oocyte images. There is, hence, wider value in resolving these challenges. Identifying methods to overcome these limitations forms the work in Chapters 3 and 4.

There is limited image data to analyse when tackling these limitations; therefore, in Chapter 3, we introduced synthetic data as an alternative. We, first, reviewed and improved current synthetic data practice in the context of DDM. Whilst Equation (3.1.2) provides the commonly accepted approach to generating such data, it introduces several parameters for which there is little guidance to govern their selection. We, hence, developed new guidelines which optimise the generation of synthetic images for use in DDM, through a mixture of heuristic arguments and simulated experiments. These simulations were constructed by creating particle trajectories, which are converted into images as per Equation (3.1.2). In particular, we demonstrated the existence of an upper and lower bound on the term controlling the width of each particle, $0.6 < \sigma < 5$ pixels (see Figures 3.2 and 3.4); this was numerically validated in Figures 3.3 and 3.5. Additionally, we studied the importance of the relationship between the number of particles, N_p , and the number of pairs of image differences, M . For a fixed value of $M = 500$, we identified a minimum required particle number, $N_t \approx 150$, that guarantees a sufficient statistical sample of the particle ensemble. This value of M was chosen to limit the computational cost of the generated simulations, but it is highly likely that such a choice is not optimal. For larger values of M , fewer particles will be required, and conversely, more particles may be needed for smaller values of M .

Together, the guidelines we proposed ensure the validity of our subsequent analysis of synthetic data. Further investigation of the relationship between imaging parameters, such as N_p and M , and DDM performance may be an important step towards defining conditions and tests to ensure the appropriate use of DDM on real image datasets. Given a set of biological time-lapse images, we can always

apply DDM to extract some set of parameters. When images are ‘poorly conditioned’, such as when the particle number is excessively high, or too few images are contained within the movie, these parameters will have no physical meaning. As such, it is critical in future work to develop tests which can be applied to new datasets, to indicate whether they are suitable for DDM.

In Chapter 4, using the simulation methodology and parameters derived in Chapter 3, we tackled limitations and challenges faced when applying DDM to oocyte image data, such as the bright-field oocyte data obtained in the Swann lab. First, in Section 4.2 we introduced a new methodology for modelling boundary effects, through numerically approximating the ISF, which only requires to describe the behaviour of a single particle in the system. This is, thus, much less restrictive than current DDM analysis, which requires us to also analytically express the average behaviour of the ensemble. Through the ‘toy’ problem of Brownian motion confined to a circular domain with reflective boundaries, we demonstrated in Section 4.2 that our new method yielded improved accuracy, compared to the existing alternative approach of cropping the image to subdomains away from the boundary, as seen in other works [142, 280, 302] (Figures 4.4 and 4.5). This work, aside to the benefit of enabling the inclusion of boundaries in images, offers a new mechanism for the generation of the ISF for datasets where an analytic expression is challenging or impossible, which is most often the case.

Despite the aforementioned new methodology we developed for numerically constructing the ISF, which has a clear advantage over current DDM approaches, we demonstrated in Section 4.3 that we are still limited by flawed parameter fitting approaches. We, thus, created a virtual experiment with a synthetic directed advection-diffusion problem (with velocity \mathbf{v} , assumed constant) and highlighted the inadequacy of the current fitting approach taken in DDM. We showed in Figure 4.9 that the accuracy with which we approximate the diffusion coefficient and velocity terms depends on the non-dimensional Péclet number. The Péclet number, which is the ratio of the velocity magnitude $|\mathbf{v}|$ to the diffusion coefficient D , is particularly important in the context of oocyte assessment, because it describes the relative magnitudes of the key driving forces in cytoplasmic movement, namely active diffusion and cytoplasmic streaming. Our hypothesis is that oocyte quality determines the rate of cytoplasmic movement, which may imply that oocytes of different quality have different Péclet numbers. If Péclet number correlates with both the accuracy of DDM analysis, and the quality of the oocyte, then comparison of DDM-derived parameters between oocytes becomes unreliable.

In Section 4.4, to overcome the poor performance of current fitting algorithms in DDM, we developed a new methodology for fitting parameters from image data, based on machine learning. We introduce a pipeline (Figure 4.13) which uses a neural network (NN) to fit parameters from DDM matrices. By careful preconditioning of the DDM matrix, we allow the NN to be trained on purely synthetic image data. The algorithm performs well on real image data, as evidenced by application to time-series images of colloidal dispersions supplied by the Swann lab. We have successfully approximated the diffusion coefficients characterising these colloidal dispersions; these agreed with the Stokes-Einstein relation [299], and achieved similar accuracy to existing fitting approaches. Although the accuracy of our proposed pipeline did not surpass current approaches, we propose that the benefit of our methodology is that it can be generalised to other forms of motion more easily than current DDM fitting. Given a set of rules which determine the trajectory of each particle, we can train the NN to fit key movement parameters without needing to determine the ISF.

The proposed machine learning pipeline is an exciting first step towards applying DDM analysis in non-invasive oocyte assessment, although there is still a great deal of work to be done. In order to accurately generate the synthetic data to train the NN, we require live imaging experiments to inform our understanding of oocyte cytoplasmic movement. The cost of generating such data, as well as the ethical approvals required, hinder its production. Additionally, the NN we developed is basic; whilst it has proved sufficient for approximating diffusion coefficients in colloidal dispersions, more complex forms of particle movement will likely require more advanced architecture. An avenue for future work is to explore different types of NN, as well as optimising NN meta-parameters such as the number of, and size of, layers, to increase fitting accuracy and ensure the proposed pipeline can be applied to different image datasets.

Our final contribution to DDM, in Chapter 5, involves a previously unexplored interaction with phase-contrast imaging. Despite previous implementation of DDM in differential interference contrast (DIC) images, we observe that image shadowing, a known feature of phase contrast images [268], induces an artificial anisotropic deformation of the DDM tensor, which corresponds to a previously unknown effect on DDM analysis. In the context of oocyte assessment, this anisotropy is important to consider since a popular bench-top incubator in the IVF clinic, the Embryoscope[®], uses phase contrast microscopy [333]. In any imaged dataset, assuming the observed motion is isotropic, we derive an expression in Equation (5.2.21) which interprets this artificial anisotropy as a frequency-dependent scaling, and conclude that parameter fitting is invariant under phase-contrast shadowing. Critically, this means that DDM can be applied

in phase-contrast images of isotropic movement, without any consideration for the shadowing effect. We extended this analysis to a directed advection-diffusion process, as a specific form of anisotropic motion. For this problem, we describe the shape of the DDM matrix in Equation (5.3.14), showing that the induced artificial anisotropy cannot be expressed by a scaling term as in the isotropic case. This means that in order to fit the velocity magnitude and diffusion coefficient from the DDM matrix, an adapted ISF must be determined, which is not an easy task.

Seeking an adapted ISF to tackle phase contrast images is not ideal, since it can be difficult to describe analytically. In the advection-diffusion problem, a ‘natural’ anisotropy problem which approximates combined active diffusion and cytoplasmic streaming in the oocyte [147, 154], the adapted ISF also contains a number of non-linear dependencies on velocity, the shear distance and the angle between the velocity field and shear. These non-linear dependencies result in less stable parameter fitting (see Figure 5.13). In Section 5.3.2, we explored whether accurate parameter fitting can be undertaken without adapting the ISF, as in the isotropic case. We identified a condition for this to be true in Equation (5.3.19), namely that the radial average of the product of the velocity and shear correlation functions is approximately equal to the product of the radial average of each. This assumption implies the shear distance in DIC is much smaller than the displacement observed in the images, which is realistic. This suggests that we may be able to ignore shadowing in DIC imaging.

However, we advise caution, since in Section 5.4, we show that failure to use the adapted ISF when the small-shear approximation does not hold causes parameter fitting algorithms to converge on incorrect estimates of the velocity (see Figure 5.18). Furthermore, we show that for some combinations of velocity and shear, we can predict how the magnitude of this error (see Figure 5.19). With more analysis, we might be able to propose post-processing steps which restore the correct estimate of the velocity, but for now, we advise that the adapted form of the ISF is used in anisotropic motion to prevent unexpected fitting error.

Our work is important in clinical applications, since we have shown, for anisotropic motions such as cytoplasmic streaming, how DDM analysis must be adapted to consider the shadowing effect and obtain more accurate estimates of the streaming velocity. Important future work must be conducted to further explore the interplay between anisotropic motion and phase-contrast shadowing in other behaviours than the advection-diffusion problem we consider.

The second part of this thesis, Chapters 6-7, focuses on challenges in cryo-preservation and thawing, directly addressing questions posed to us by LWC. In Chapter 6, we developed a new computational thermal model of the Open

Cryotop[®] device in COMSOL Multiphysics, and determined the temperatures and cooling rates during oocyte or embryo freezing. This allowed us to virtually experiment with different numbers and arrangements of embryos or oocytes on a single device. Comparing the derived cooling rates against a ‘Worst Case’ viable benchmark (Figure 6.13), we concluded that the cooling rates are approximately equal in all considered arrangements. Our model could be easily adapted to other cryopreservation devices, increasing the value of the proposed framework across the wider IVF community. This work has direct clinical application, as it provides confidence and time savings for embryologists, who can choose the number and any arrangement of oocytes or embryos within the limits of protocol, without risking cryogenic damage.

Our computational model adopts similar assumptions with other modelling of cryopreservation and thawing processes in IVF [240, 242, 244, 245]. However, the model can be improved by relaxing some of these assumptions. In particular, we considered liquid nitrogen to be a stationary, isothermal medium, as in other works [240, 242], whereas it is an actively boiling liquid [357]. Another assumption which could be relaxed is that there is no moving boundary which models phase change from the liquid to the solid state of the cryoprotectant solution. Whilst relaxing this assumption would make our model more realistic, we noted that previous models making the same assumption [240, 242, 244, 245] demonstrated cooling rates which agree with experimental results [229, 246]. This is likely a result of the very small volume device, oocyte/embryo and CPA droplet, which results in very fast cooling and, therefore, a very fast moving boundary. The cooling rates predicted by our model would therefore be likely to remain unchanged given the inclusion of a moving boundary problem.

In Chapter 7, we introduced a new metric for non-invasive health assessment of embryos, using time-lapse images of 124 cryopreserved embryos re-expanding after being thawed. The challenge posed by LWC was to determine whether the rate of re-expansion of thawing embryos can be used to predict clinical pregnancy. LWC provided us with a dataset describing the outcome of 147 Frozen Embryo Transfers (FET), in addition to time-lapse imaging of the re-expansion of each embryo following warming. Excluding 23 embryos from the study due to missing data, or lack of survival during cryopreservation, 67 (54.0%) registered a positive pregnancy test, 53 (42.7%) registered a positive clinical pregnancy, and 37 (29.8%) resulted in a live birth (see Figure 7.2). We applied a machine learning semantic segmentation approach to the time-lapse images provided by LWC, to measure the area of the expanding cell mass in the embryo at each time step. We identified that embryos with non-linear re-expansion patterns were less likely to result in clinical pregnancy than embryos with linear re-expansion (Figure 7.7). In Figure 7.9, we

determined a threshold on the absolute area change, which identifies a subset of 57 out of 124 embryos with lower clinical pregnancy rates, at a statistically significant level ($p < 0.05$). Based on LWC dataset, this threshold could increase clinical pregnancy rates by up to 3% by identifying poor quality embryos, which could be replaced by a second, more viable candidate from the pool of vitrified embryos.

Whilst the latter threshold appears promising, it has been previously reported in [199] that morphokinetic parameters identified in one dataset frequently fail to yield results when applied to another. As such, a repeat study is needed, with the sole intention of investigating and validating this hypothesis; any new morphokinetic parameter should be tested with randomised controlled trials before being used in the clinic. Additionally, it is important to investigate the cause of the relationship between expansion rate and the underlying embryo health. We identified (Figure 7.11) that the majority of large area changes observed in embryos which did not lead to clinical pregnancy were due to sudden contractions, which have already been considered as indicators of poor embryo quality [247, 250]. However, it is not known whether contraction is the cause, caused by, or even related to embryo quality [72, 253, 254].

The work in this thesis is an exciting step forward in using mathematics to tackle important challenges in IVF, in the non-invasive assessment of oocytes using advanced image processing algorithms, in optimisation of cryopreservation protocols, and in thawed embryo selection during Frozen Embryo Transfers. The implications of our results range in their applicability and scope; our thermal modelling work has direct clinical applications, allowing embryologists to save time and be more confident as cooling rates are approximately constant as the number and arrangement of oocytes or embryos vitrified on a single Open Cryoptop[®] device changes. For the wider biological imaging community, our work on explaining the effect of phase-contrast shadowing on DDM opens the door to applying the DDM technique in a range of new datasets. The rest of our work provides a foundation for future development. In particular, we consider the use of DDM as a non-invasive oocyte (and early 1-cell embryo) assessment technique to be a highly exciting prospect; however, there are steps to take before this work can reach the clinic. We, first, require improved models that describe the cytoplasmic movements in oocytes, which will likely only be available through experiments with tracer particles, similar to experiments in [275, 398–400]. These experiments may act as a ‘ground truth’ against which DDM can be compared.

Secondly, potentially invasive experiments which relate cytoplasmic movements to measurements of health in oocytes are required. For example, in [135],

the relationship between cytoplasmic contractions, and the health of the cytoskeleton, was established by injecting the oocytes with messenger RNAs which encode two proteins, each binding to either myosin or actin. These proteins were tagged with Green Fluorescent Protein (GFP) or Enhanced GFP, to highlight myosin and actin activity respectively. Spikes in the cytoplasm speed could therefore be correlated with spikes in fluorescence, which in turn indicate enhanced actin or myosin activity. Hence, invasive experiments provided data which could bridge the gap between a clinically safe measurement, and key biological measurements which are indicative of embryo viability, but not clinically accessible through direct measurement. Currently, similar data which may correlate output statistics from DDM to oocyte viability are lacking. Pairing the machine learning pipeline outlined in Figure 4.13 (Section 4.4) with such oocyte data is a vital step towards employing DDM as a predictive tool in the IVF clinic.

Whilst improving oocyte and embryo selection has significant potential to improve IVF success rates, there is almost certainly no ‘silver bullet’ that will guarantee the desired outcome of one live birth per treatment. IVF is a complex, multistage process [64], with clinical outcomes being the product of a plethora of variables including patient age [8], environmental influence [263], clinical protocol [401], oocyte quality [132] and more. The fertility community has made great strides to identify the areas of IVF which need the most improvement [50, 100, 101], but we echo ESHRE in that ‘interdisciplinary collaborations involving the use of mathematics are bringing innovation to reproductive medicine’.

A recent review [402] discussed fields where the greatest open questions in mathematical biology can be found, identifying parameter sensitivity, model selection and development, the consideration of multi-scale, hierarchical, and spatio-temporal models, stochastic non-linear dynamics and hybrid, data-driven modelling. All these challenges and methodologies have featured in this thesis, and various avenues for future work have been outlined.

A critical direction for future work is the development of machine learning and artificial intelligence (AI) approaches in the IVF clinic. IVF treatment is accompanied by a vast amount of data, including patient age and BMI [403], environmental factors such as patient occupation [261], morphokinetics of the embryo [57, 404] and more. Understanding how these many variables determine IVF outcomes is critical to improving success rates, and is a challenge well suited to AI [405]. Integrating AI into IVF is a recent, but highly active, research field [395, 405]. However, whilst AI is being investigated as a predictive tool in IVF treatments, and has begun to see some clinical applications [406, 407], AI integration into most clinics is still limited. This is partly attributed to poorly conditioned training datasets which contain bias, noise and imbalanced classes,

in addition to a lack of standardised approaches to choosing the optimal network architecture [405]. With the help of AI experts, these problems can be mitigated, and AI has the potential to become a part of regular day-to-day operation for IVF clinics.

However, before the mathematical challenges of AI can be tackled in the fertility industry, the foundations of good working principles of AI must first be established. Healthcare is highly regulated, which means any technology adopted for healthcare applications must also be subjected to stringent regulatory requirements. The importance of regulatory oversight in AI healthcare applications is driven by a number safety and ethical concerns, as well as privacy concerns surrounding training data, the possibility of bias, and a lack of transparency in its application [408]. The emergence of Large Language Models (LLMs) such as ChatGPT and Bard, which are vastly complex and have frequently been misused or misunderstood, have accelerated these concerns [409]. Regulation of AI is not a concern unique to the healthcare setting, so lessons can be learned from similarly heavily regulated industries, such as banking and the financial sector [410]. For the fertility industry in particular, however, the setting of policy and guidelines that ensure safe and effective application of AI in healthcare settings is a task demanding multidisciplinary contribution, from clinicians, government regulators, academics and AI industrial experts.

Our primary recommendation for enhancing the use of mathematics in fertility treatments is therefore to pursue interdisciplinary collaborations, like [411], and the project this thesis forms a part of. With the combined efforts of modellers, statisticians, operations researchers, imaging specialists, AI experts, experimental biologists and clinical embryologists, the many barriers to improving success rates in IVF may be overcome.

Table of Notation

Notation	Meaning
\mathbf{r}	Cartesian coordinates (x, y) , units μm .
t	Time, units s.
C	Local concentration of some chemical, units $\text{mol}/\mu\text{m}^3$.
$c(\mathbf{r}, z, t)$	Local particle density.
D	Diffusion coefficient, units $\mu\text{m}^2/\text{s}$.
\mathbf{v}	Velocity vector field, units $\mu\text{m s}^{-1}$.
η	Dynamic viscosity, units Pa s.
p	pressure, units Pa.
$I(\mathbf{r}, t)$	Image stack.
N_p	Number of particles within the imaged system.
N_t	Number of images in the image stack $I(\mathbf{r}, t)$.
L, W	Length and width of each image $I(\mathbf{r}, t)$, respectively.
ν	Physical length of a pixel, units μm .
T	The set of times corresponding to each frame in the image stack $I(\mathbf{r}, t)$, such that $t \in T$.
i_0	Background image brightness.
$K(\mathbf{r} - \mathbf{r}', z)$	Point Scattering Function (PSF).
\star	Convolution operator.
$G(\Delta\mathbf{r}, \Delta t)$	The van Hove function.
V	Volume, units μm^3 .
Δt	The time lag between two frames of the image stack, $I(\mathbf{r}, t)$.
Δt_{\min}	The time lag between two <i>consecutive</i> frames of the image stack, $I(\mathbf{r}, t)$.
\mathbf{q}	The spatial frequency vector in two dimensions, with length q μm^{-1} and angle θ .
$\mathcal{D}(\mathbf{q}, \Delta t)$	The Differential Dynamic Microscopy (DDM) tensor evaluated at the spatial frequency vector \mathbf{q} and time lag Δt .
Γ	The circular contour of length q with origin $(0, 0)$.
$\mathcal{D}_R(q, \Delta t)$	The DDM Matrix, equal to the radial average of $\mathcal{D}(\mathbf{q}, \Delta t)$ over the circular contour Γ .
$f(q, \Delta t)$	The Intermediate Scattering Function (ISF) at spatial frequency q and time lag Δt , equal to the real part of the Fourier transform of the van Hove function.
$W(x, y)$	The two-dimensional Blackman filter.
M	The number of image pairs sampled during DDM.
L_s	Length of the simulation window in pixels, where $L_s \leq L$.
α	Motility fraction.
a	Particle brightness.
σ	The standard deviation of the Gaussian beam formed by each simulated particle in the corresponding synthetic image, units pixels.

Notation	Meaning
$\mathcal{P}(D, \Delta t) _q$	Polynomial interpolation of the ISF, $f(q, \Delta t)$, for a fixed spatial frequency and variable diffusion coefficient.
K_B	The Boltzmann constant, units J/K.
$X(\mathbf{r}, t)$	The phase image formed by one of the pairs of beams from a DIC image stack.
$\mathcal{D}_I(\mathbf{q}, \Delta t)$	DDM tensor as defined above, where the subscript I specifies the image stack on which the DDM tensor is evaluated.
$\mathcal{D}_{I,R}(q, \Delta t)$	DDM matrix as defined above, where the subscript I specifies the image stack on which the DDM tensor is evaluated, and the R denotes the application of the radial average.
$J_n(x)$	The Bessel function of order n of the first kind.
$E(q)$	The residual fitting error between the DDM matrix and the model fitting equation.
γ	The ratio of the fitted velocity v_{out} to the input velocity v_{in} .
$j_{n,k}$	The k^{th} positive zero of the Bessel function of order n of the first kind, $J_n(x)$.
Ln	The domain occupied by liquid nitrogen.
F	The domain occupied by the cryoprotectant solution.
S	The domain occupied by the embryo or oocyte.
P	The domain occupied by the Cryotop [®] plastic plate.
$u(x, y, z, t)$	The temperature at Cartesian position (x, y, z) and time t , units K.
$c(u)$	Specific heat capacity at temperature u , units J/kg K.
$k(u)$	Thermal conductivity at temperature u , units W/mK.
$\rho(u)$	Density at temperature u , units kg/m ³ .
h	The heat transfer coefficient, units W/m ² K.
$\varkappa_{A,B}$	The tension force between the two domains A and B , units N.
ℓ	The latent heat of freezing, units J/kg.
t_D	The characteristic timescale of diffusion, units s.
$\Delta\mathcal{A}$	The change in area of the cell mass in an image of an embryo re-expanding after thawing from cryopreservation, units pixels ² .
χ^2	The χ^2 test statistic.
IB	In Bounds, a binary classifier embryo denoting that the frame-wise absolute change in cell mass area, $ \Delta\mathcal{A} $, is below a fixed tolerance.
IB	Out of Bounds, a binary classifier embryo denoting that the frame-wise absolute change in cell mass area, $ \Delta\mathcal{A} $, is above a fixed tolerance.

Appendix A

Experimental details regarding mouse oocyte collection

This appendix provides detail on the setup of the experiments used to gather oocyte data for this body of work. This appendix applies to the following datasets:

- Bright-field images of mouse oocytes introduced in Section 2.1.1
- DIC images of mouse oocytes introduced in Section 5.1.

Eight-week old CD1 mice were intraperitoneally injected with 5IU pregnant mare's serum gonadotrophin (PMSG) to induce ovarian follicle development. They were again injected with 10IU human chorionic gonadotrophin (hCG) approximately 48hrs later to induce ovulation. Ovulated mature (MII) eggs were collected from oviducts approximately 15hrs later. All animals were handled according to UK Home Office regulations, and procedures were carried out under a UK Home Office Project License, with the approval of Cardiff University Animal Ethics Committee. Live imaging experiments were carried out with mature eggs cultured in HKSOM, as described previously [412].

Appendix B

Single Particle Tracking & Particle Image Velocimetry

B.1 Single Particle Tracking

When dealing with a finite number of distinct particles, the exact trajectory of each particle over time is the most information we can possibly have. This is the idea behind the first method of image analysis, Single Particle Tracking (SPT), which is reviewed in [274].

Consider a finite collection of N particles, depicted in an image stack $I(\mathbf{r}, t)$ where \mathbf{r} denotes the Cartesian coordinates $\mathbf{r} = (x, y)$. We denote by $\mathbf{r}_n(t)$ the position of the centroid of particle $n \in \{1, 2, \dots, N\}$ at time t . The first step of SPT is *particle localisation*, the process of identifying the set of particle centroids for each frame. For optical microscopes, single point objects have the appearance of concentric rings with decaying intensity moving away from the center. The middle ring is known as an Airy disc, which for a particle located at (x_0, y_0) , can be approximated by [274]

$$I(\mathbf{r}, t) \approx i_0 e^{-\frac{(x-x_0)^2}{2\sigma^2}} e^{-\frac{(y-y_0)^2}{2\sigma^2}},$$

where i_0 is the central intensity and σ the standard deviation of the intensity profile. Fitting this profile to the image can determine the centroid of the particle below even the optical resolution (sub-pixel), or averaging methods can be used to locate the centroid [277]. The next step is *particle linking*, whereby trajectories are constructed by determining which centroids over time originate from the same particle.

Whilst knowing the trajectory of each particle allows for in depth statistical analysis, there are a number of physical and experimental constraints that often limit the use of SPT, at both the localisation and linking stages. Stochastic variance in the illumination and unavoidable detection and readout noise mean

that the location of a particle is not exact, but estimated, which may affect subsequent analysis [274]. More importantly, SPT encounters problems with high particle density. In ultra low densities, particle tracking is a simple ‘connect the dots’ exercise, but in higher densities or where each particle moves a significant distance, trajectories can overlap, precluding the possibility of incorrect or fragmented localisations, as in Figure B.1 [274].

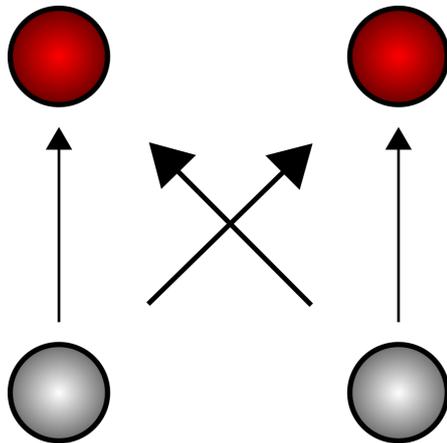


Figure B.1: An example of linking error in SPT, where two particles have moved from the grey to the red location over a fixed time step. Based on only observation of the initial and final state of the system, the trajectory of each particle is uncertain.

Trajectory mapping in dense particle populations is essentially a combinatorial optimisation problem [274]. A huge variety of SPT algorithms have arisen to tackle the computational complexity of tracking and improve accuracy [413], but in general, lack of sparsity will inhibit the ability of SPT. Additionally, SPT requires distinct, visible particles to function. When classifying a fluid flow in which no such particles exist, a common solution is to introduce tracer particles such as fluorescent beads or gold nanoparticles, or through tracking genetically modified GFP-expressing proteins [275, 398–400]. Clearly, these are invasive interventions and are unsuitable in the clinic.

SPT has been previously successfully applied to track the movement of endogenous particles in *Xenopus laevis* oocytes [160], since the large size of these particles means that they are much larger than the cytoskeleton mesh, and thus move as tracers in a fluid, as well as being visually well contrasted from the background. Similar particles exist in mouse oocytes, which could be tracked [161], but looking at the bright-field image in Figures 2.1, and the DIC image in Figure 5.2, we can see that although visible particles exist, the surrounding medium is visually messy and the particles themselves are sparsely integrated into the cytoplasm. This means only a few tracers are considered, and much of the cytoplasmic data, making SPT less favourable.

B.2 Particle Image Velocimetry

The second technique we discuss is Particle Image Velocimetry (PIV), which refers to a class of techniques for measuring local instantaneous vector velocity fields by tracking the motion of numerous fine tracer particles [277]. First, we partition the image domain into a grid of subdomains, \mathcal{W}_i , known as ‘interrogation windows’. On each window, for each pair of consecutive images I_1 and I_2 , we estimate the cross correlation $R(\mathbf{s})$, defined as [277]

$$R(\mathbf{s}) = \int_{\mathcal{W}_1} I_1(\mathbf{r})I_2(\mathbf{r} + \mathbf{s})d\mathbf{r}.$$

The cross correlation can be considered equivalent to sliding I_2 across I_1 by a vector \mathbf{s} , and measuring how much overlap exists at each shift. When some local linear flow exists with velocity \mathbf{v} , the majority of particles will move by the same vector $\Delta\mathbf{r} = \mathbf{v}\Delta t$ during a time lag of Δt . Hence, a peak in the cross correlation occurs at $\mathbf{s} = \Delta\mathbf{r}$, where the most overlap occurs, allowing us to infer the local velocity. In practice, cross correlation is calculated using Fourier transforms to improve the speed of the algorithm.

An important subsequent step is to validate the results of interrogation, because anomalous detection can be caused either by the presence of boundaries in the domain, or groups of particles creating a peak in the noise stronger than that of the bulk flow [277]. Compared to SPT, PIV is more statistically averaged, working under the assumption that particles follow local flow. This means when particles move independently, and not under some bulk flow, the cross correlation may not exhibit a dominant peak, resulting in anomalous local velocity classification. This means classifying diffusion using PIV can be challenging in optically dense samples [279].

There is a wealth of literature reporting the use of PIV in biological application. PIV was used to classify spasms in mouse oocytes in relation to Ca^{2+} oscillations [152], which can act as a predictor of viability with respect to potential to fertilise [135]. PIV-detected flows in the cytoplasm have also been used to reflect the architecture of the cytoskeleton, as well as changes in Kinesin activity [150]. PIV has also served to verify numerical models of cytoplasmic streaming in starfish oocytes [162].

Appendix C

DDM applied to bright-field oocyte images

In this Appendix, we display the fitting results from applying DDM ‘out of the box’ to bright-field images of mouse oocytes, as discussed in Section 2.3. Each oocyte is imaged before and after poisoning with FCCP, a metabolic inhibitor. To the DDM matrices generated for each oocyte in each experimental condition, we fit movement parameters using two different fitting functions; one assumes a diffusive-type movement in the cytoplasm, referred to as the Brownian model here, and the second assumes that some ballistic velocity is superimposed on the diffusive movement of each particle in the cytoplasm, referred to as the bacterial model here.

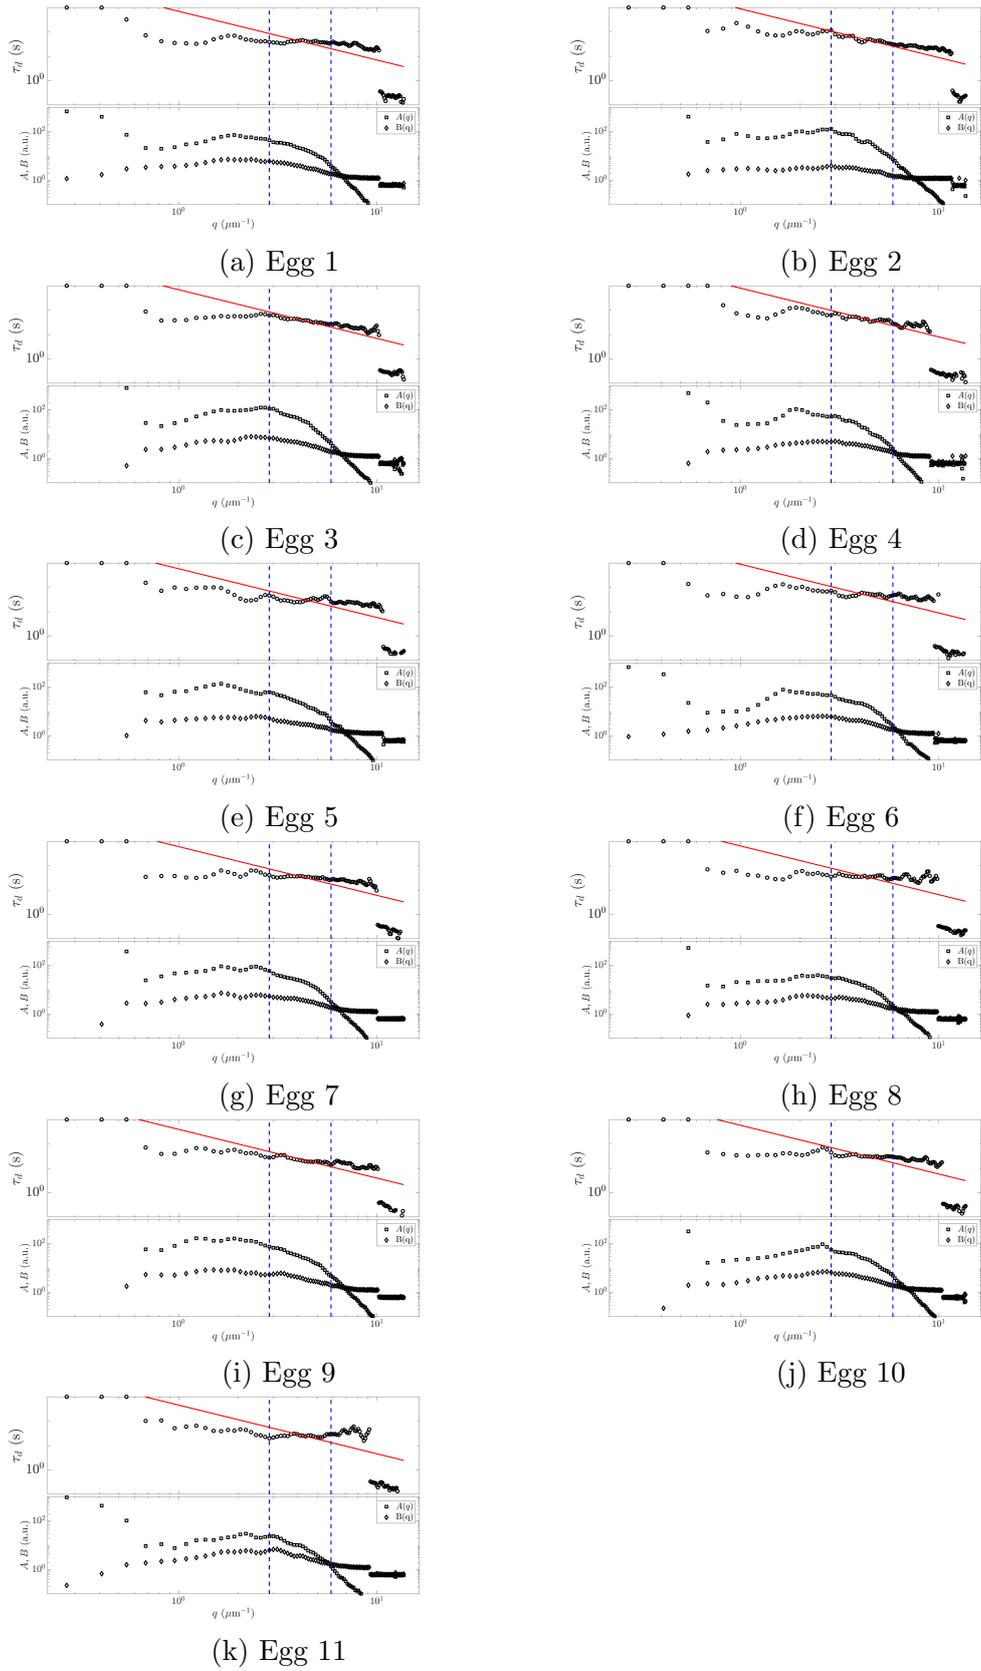


Figure C.1: Fitting results performing DDM analysis for brightfield imaged mouse oocytes, under control conditions, where the Brownian fitting model is used.

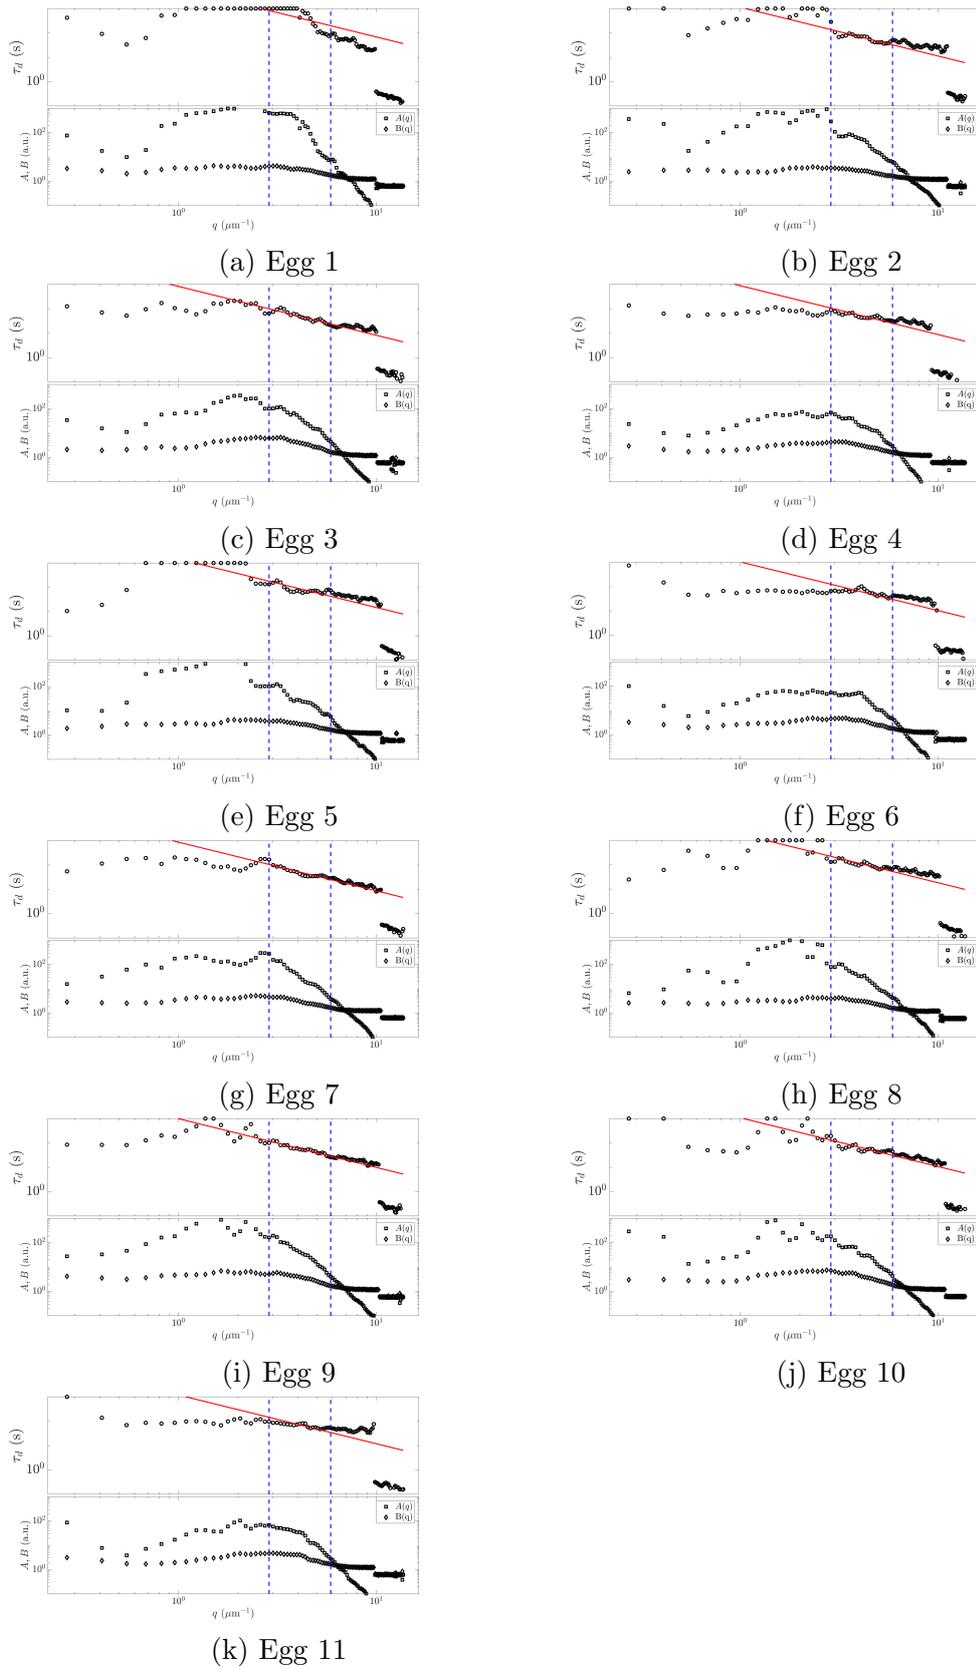


Figure C.2: Fitting results performing DDM analysis for brightfield imaged mouse oocytes, in the FCCP condition, where the Brownian fitting model is used.

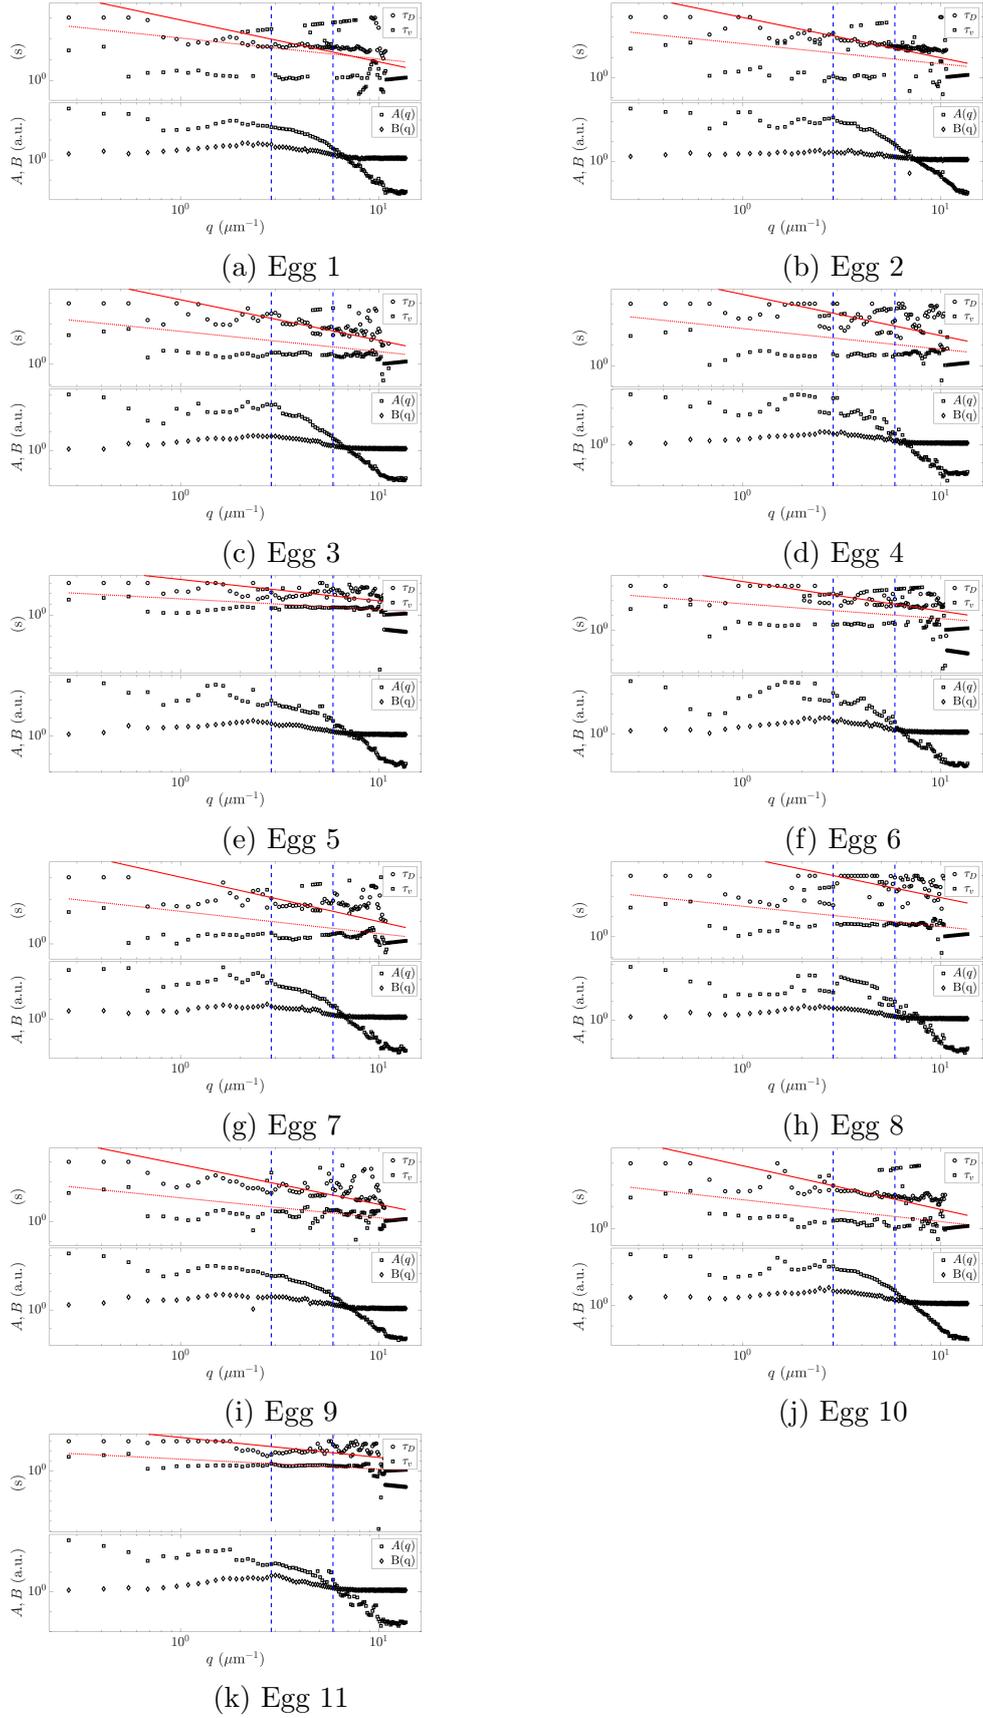


Figure C.3: Fitting results performing DDM analysis for brightfield imaged mouse oocytes, under control conditions, where the bacteria fitting model is used.

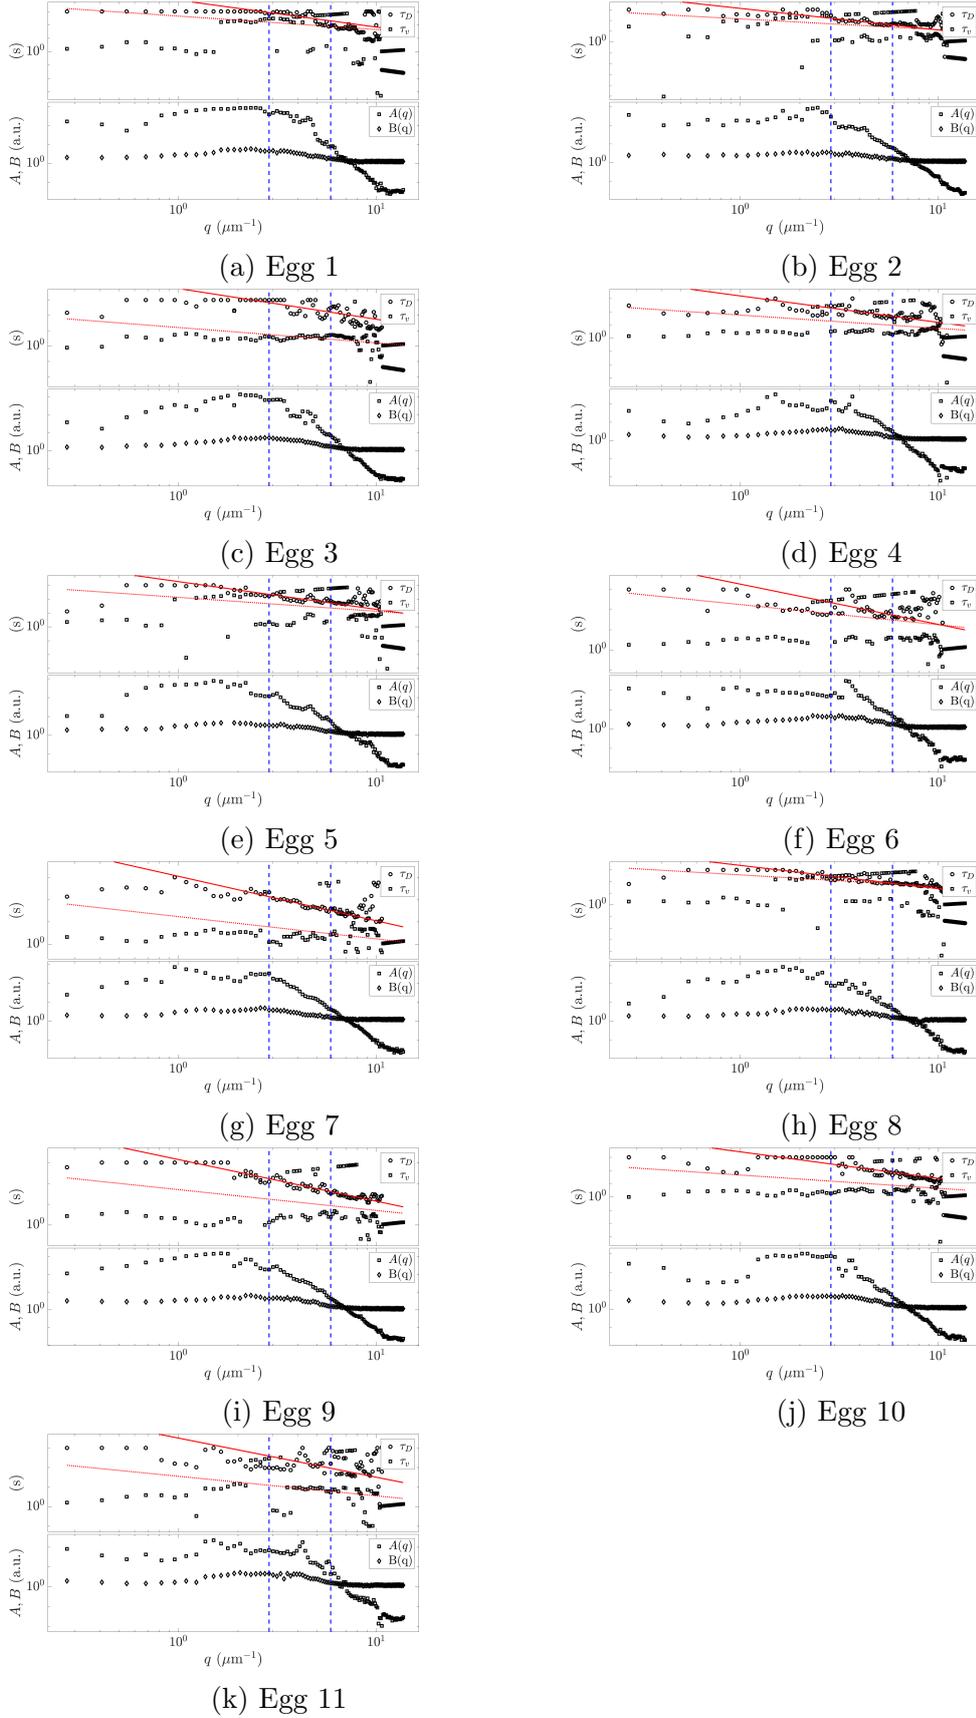


Figure C.4: Fitting results performing DDM analysis for bright-field imaged mouse oocytes, under FCCP conditions, where the bacteria fitting model is used.

C.1 The theory of the formation of a Differential Interference Contrast (DIC) image

This appendix gives a more thorough explanation for the underlying physics describing the formation of an image taken using DIC. The following explanation was provided by Wolfgang Langbein, a collaborator in the School of Physics and Astronomy, Cardiff University, as part of a submitted paper [335].

A sketch of a DIC microscope with de-Sénarmont compensation and its workings is given in Figure C.5. Starting from a bright-field transmission microscope in Köhler illumination, the illumination is linearly polarised at $45^\circ + \theta$, and then transmitted through a quarter-wave plate with the fast axis at 45° , resulting in electric field components E_x and E_y in the x and y directions, respectively, of equal amplitude and a relative phase shift of $\psi = 2\theta$. A Wollaston prism splits the propagation direction of E_x and E_y in the condenser back focal plane, resulting in a relative displacement of the corresponding focal points at the sample by the shear \mathbf{s} along the x direction. The fields probe the sample properties accordingly at two positions displaced by \mathbf{s} , and acquire a spatially dependent phase difference $\delta(\mathbf{r})$, with the in-plane sample position $\mathbf{r} = (x, y)$. After transmission through the sample, the two fields are recombined in direction at the back focal plane of the objective by a Wollaston prism matched to the first one (a Nomarski prism can be used instead, which allows for the effective recombination position to be displaced from the prism and thus to enabling the positioning of the prism at a accessible place in the beam path after the objective). A second polariser, called analyser, orientated at -45° , orthogonal to the first one, projects the recombined fields along its axis to provide an intensity interference. The resulting transmitted intensity I_t has a dependence on the phase shift $\delta(\mathbf{r})$ given by

$$I_t(\mathbf{r}, \psi) = \frac{I_e}{2} [1 - \cos(\psi - \delta(\mathbf{r}))], \quad (\text{C.1.1})$$

with the excitation intensity I_e , the position in the sample plane \mathbf{r} , the phase offset ψ , and the difference $\delta(\mathbf{r})$ of the optical phase shift ϕ for the two beams that pass through the sample in two adjacent points separated by the shear vector \mathbf{s} .

Adjusting the phase offset ψ , the contrast can be changed from a dark-field type at $\psi = 0$, where the intensity is proportional to $\delta^2(\mathbf{r})$, to a bright field image modulated by a term linear with the phase shift, with the largest linear range for $\psi = 90^\circ$, for a polariser angle of $\theta = 45^\circ$. More details on the theory of DIC can be found in Ref. [414] for two-dimensional samples, in Ref. [415] for three-dimensional samples, and a fully vectorial treatment of the fields is given in Ref. [416].

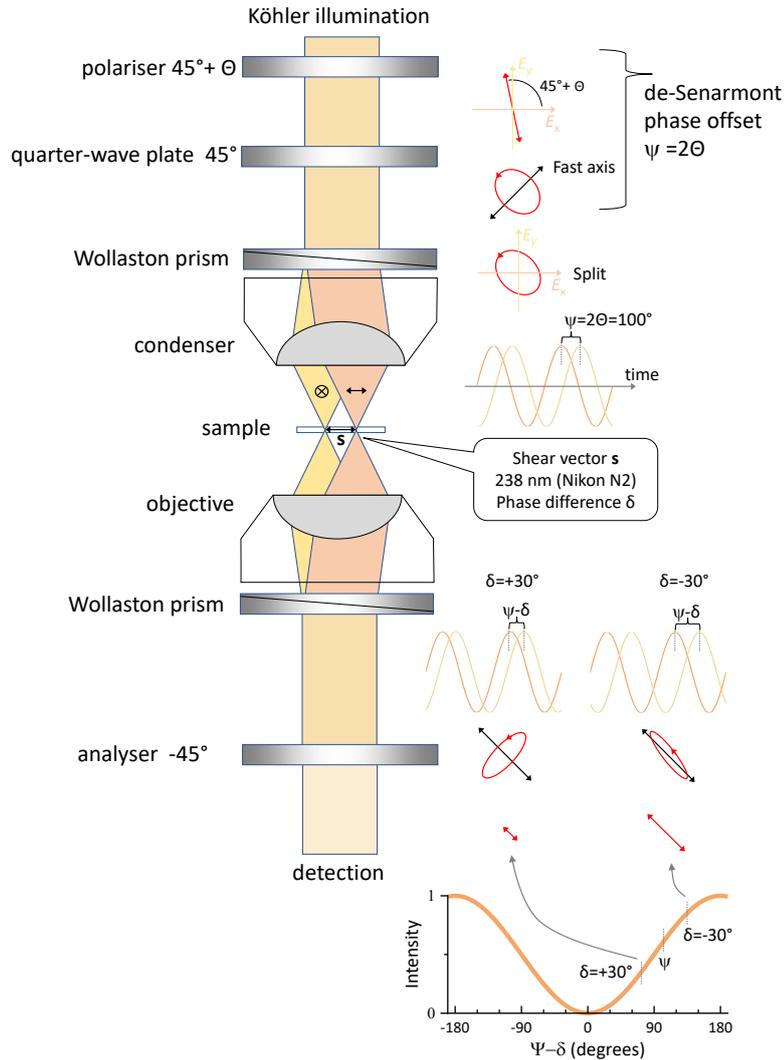


Figure C.5: A sketch of the principle of operation of a DIC microscope with de-Sénarmont compensation. After adjusting a transmission microscope for Köhler illumination, a linear polariser at $45^\circ + \theta$ to the sample x (horizontal) and y (out of plane) directions creates a linear polarisation (see electric field sketches on the right), which is converted by a quarter wave plate to a relative phase shift of 2θ of the x and y polarised field components E_x and E_y of equal amplitude (see sketch of temporal field oscillations on the right). A Wollaston prism splits the two fields, creating two focal points separated by the shear \mathbf{s} . The transmitted fields, which experienced a relative phase shift $\delta(\mathbf{r})$ as function of sample position $\mathbf{r} = (x, y)$, are collimated by the objective, recombined in direction by a second Wollaston prism, and projected by the analyser of -45° orientation, providing an intensity proportional to $1 - \cos(\psi - \delta(\mathbf{r}))$. Using $\psi = 90^\circ$, a nearly linear dependence of the intensity on the phase shift $\delta(\mathbf{r})$ is obtained, leading to the relief-type contrast.

C.2 Approximation of the ISF for DIC images as a weighted average of parallel and orthogonal components

This appendix demonstrates that in DIC images of advection-diffusion processes, where the displacement from the velocity is much greater than the size of the DIC shear, the ISF outlined in Equation (5.3.15) can be bounded above and below by the ‘extreme’ cases where $\Phi = 0$ or $\Phi = \pi/2$. Although this observation does not feature in our analysis in this work, it is nonetheless interesting to present, as it enhances our understanding of the behaviour of the ISF. In particular, it allows us to demonstrate that in DIC images, the dependence of the ISF on Φ is weak, since $f(q, \Delta t)$ is bounded between two functions which converge towards the same limit as DIC shear distance decreases.

We begin by recalling the final term in Equation (5.3.15), which was derived from the term T_4 in Equation (5.3.6). Using the non-dimensional parameter groups

$$\lambda = qv\Delta t, \quad \xi = qs, \quad \kappa = q^2D\Delta t,$$

we can redefine T_4 by

$$T_4(\lambda, \xi, \Phi) = \frac{e^{-\kappa}}{2} \left(J_0 \left(\sqrt{\lambda^2 + \xi^2 + 2\lambda\xi \cos \Phi} \right) + J_0 \left(\sqrt{\lambda^2 + \xi^2 - 2\lambda\xi \cos \Phi} \right) \right). \quad (\text{C.2.1})$$

We are interested in the behaviour of T_4 when $\Phi = 0$ and $\Phi = \pi/2$, which form the most ‘extreme’ conditions for when the DIC shear and advection velocity are parallel and perpendicular, respectively. When $\Phi = 0$, Equation (C.2.1) simplifies to

$$T_4(\lambda, \xi, 0) = \frac{e^{-\kappa}}{2} (J_0(\lambda + \xi) + J_0(\lambda - \xi)). \quad (\text{C.2.2})$$

Similarly, when $\Phi = \pi/2$

$$T_4 \left(\lambda, \xi, \frac{\pi}{2} \right) = e^{-\kappa} J_0 \left(\sqrt{\lambda^2 + \xi^2} \right). \quad (\text{C.2.3})$$

Although Equations (C.2.2) and (C.2.3) describe the function T_4 at the most extreme angles, they do not give upper or lower bounds for the T_4 (see Figure 5.14 when $s = 5$). However, we make the assumption that for DIC images, $v\Delta t \gg s$. Equivalently, this also means that $\lambda \gg \xi$. We may therefore take the Taylor expansion of $T_4(\lambda, \xi, \Phi)$ around $\xi = 0$,

$$T_4(\lambda, \xi, \Phi) = e^{-\kappa} \left(J_0(\lambda) - \frac{2 \sin^2(\Phi) J_1(\lambda) + \lambda \cos^2(\Phi) (J_0(\lambda) - J_2(\lambda))}{4\lambda} \xi^2 + \mathcal{O}(\xi^4) \right),$$

which by the recurrence relation on the Bessel Functions [324],

$$\frac{2a}{x}J_a(x) = J_{a-1}(x) + J_{a+1}(x), \quad (\text{C.2.4})$$

can be expressed as

$$T_4(\lambda, \xi, \Phi) = e^{-\kappa} \left(J_0(\lambda) + \frac{J_2(\lambda) (\cos^2(\Phi) - \sin^2(\Phi)) - J_0(\lambda)}{4} \xi^2 + \mathcal{O}(\xi^4) \right). \quad (\text{C.2.5})$$

We also take the Taylor expansion of $T_4(\lambda, \xi, 0)$ and $T_4(\lambda, \xi, \pi/2)$ around $\xi = 0$,

$$T_4(\lambda, \xi, 0) = e^{-\kappa} \left(J_0(\lambda) + \frac{1}{4} (J_2(\lambda) - J_0(\lambda)) \xi^2 + \mathcal{O}(\xi^4) \right), \quad (\text{C.2.6})$$

$$T_4\left(\lambda, \xi, \frac{\pi}{2}\right) = e^{-\kappa} \left(J_0(\lambda) - \frac{1}{4} (J_2(\lambda) + J_0(\lambda)) \xi^2 + \mathcal{O}(\xi^4) \right). \quad (\text{C.2.7})$$

Let us suppose that $T_4(\lambda, \xi, 0)$ and $T_4(\lambda, \xi, \pi/2)$ form a linearly independent basis, from which we can construct $f(q, \Delta t)$. Under this assumption, there exists some scalar values $\alpha(\Phi)$ and $\beta(\Phi)$ such that

$$T_4(\lambda, \xi, \Phi) \approx \alpha(\Phi)T_4(\lambda, \xi, 0) + \beta(\Phi)T_4(\lambda, \xi, \pi/2). \quad (\text{C.2.8})$$

Substituting (C.2.5)-(C.2.7) into (C.2.8), we aim to solve for α and β by equating terms of matching order. The $\mathcal{O}(1)$ terms give us

$$\alpha + \beta = 1,$$

and the $\mathcal{O}(\xi^2)$ terms give

$$\frac{1}{4}J_0(\lambda) (\alpha + \beta - 2 \cos^2(\Phi) + 2 \cos^2(\Phi) - 1) + \frac{1}{4}J_2(\lambda) (\alpha - \beta - \sin^2(\Phi) + \cos^2(\Phi)) = 0,$$

implying

$$\begin{aligned} \alpha + \beta &= 1, \\ \alpha - \beta &= \cos^2(\Phi) - \sin^2(\Phi). \end{aligned}$$

The conditions on both the $\mathcal{O}(1)$ and $\mathcal{O}(\xi^2)$ terms are satisfied by

$$\alpha = \cos^2(\Phi), \quad \beta = \sin^2(\Phi), \quad (\text{C.2.9})$$

which informs us that the linearly independent basis in Expression (C.2.8) is in fact a weighted average of the parallel and perpendicular behaviours of T_4 .

The error from the approximation in Equation (C.2.8) is determined by subtracting the right hand side from the left. The leading order term, which is order $\mathcal{O}(\xi^4)$, is equal to

$$\frac{1}{24}e^{-\kappa}J_4(\lambda) \cos^2(\Phi)(\cos^2(\Phi) - 1).$$

which is equal to 0 when $\Phi = 0$ or $\Phi = \pi/2$, and has an absolute upper bound of $0.0032e^{-\kappa}$ when $\Phi = \pi/4$ and the Bessel function of fourth order attains its maximum of 0.3102 at $\lambda = 4.2012$ [346]. The maximum value of ξ is attained at the maximum spatial frequency, which from Equation (2.2.17) is given by

$$\max q = \frac{\pi}{\nu},$$

where ν is the length of a pixel. Under our assumption that the DIC shear is approximately the same size as ν , the maximal value of ξ is therefore

$$\max \xi = \frac{\pi s}{\nu} \approx \pi,$$

which would mean the $\mathcal{O}(\xi^4)$ term is much smaller than 1. Hence, the Taylor series converges. This means that, in DIC images, the ISF is sandwiched between the extreme behaviours for when $\Phi = 0$ and $\Phi = \pi/2$. Furthermore,

$$\lim_{\xi \rightarrow 0} T_4(\lambda, \xi, 0) = \lim_{\xi \rightarrow 0} T_4\left(\lambda, \xi, \frac{\pi}{2}\right) = J_0\lambda.$$

If $T_4(\lambda, \xi, \Phi)$ is the weighted average of both $T_4(\lambda, \xi, 0)$ and $T_4(\lambda, \xi, \pi/2)$, which are pointwise convergent to $J_0(\lambda)$, then it can be said that

$$\lim_{\xi \rightarrow 0} \frac{\partial T_4(\lambda, \xi, \Phi)}{\partial \Phi} = 0,$$

or in other words, varying Φ has less effect on the value of the ISF as ξ approaches zero. This means, in DIC, that fitting is insensitive to Φ .

References

1. Zegers-Hochschild, F., Adamson, G. D., de Mouzon, J., Ishihara, O. *et al.* The International Committee for Monitoring Assisted Reproductive Technology (ICMART) and the World Health Organization (WHO) Revised Glossary on ART Terminology, 2009. *Hum. Reprod.* **24**, 2683–2687 (2009).
2. Jain, M. & Singh, M. Assisted Reproductive Technology (ART) Techniques. *StatPearls* (2023). ppublish.
3. Geyter, C. D. Assisted reproductive technology: Impact on society and need for surveillance. *Best Pract. Res. Clin. Endocrinol.* **33**, 3–8 (2019).
4. London Women’s Clinic. *In Vitro Fertilisation (IVF)* London Women’s Clinic. <https://www.londonwomensclinic.com/fertility-treatments/ivf/> (2023).
5. Steptoe, P. C. & Edwards, R. G. Birth after the reimplantation of a human embryo. *Lancet* **312**, 366 (1978).
6. European Society of Human Reproduction and Embryology (ESHRE). *More than 8 million babies born from IVF since the world’s first in 1978* <https://www.eshre.eu/Annual-Meeting/Barcelona-2018/ESHRE-2018-Press-releases/De-Geyter>.
7. Human Fertilisation & Embryology Authority (HFEA). *Fertility treatment 2019: trends and figures* <https://www.hfea.gov.uk/about-us/publications/research-and-data/fertility-treatment-2019-trends-and-figures/>. 2022. <https://www.hfea.gov.uk/about-us/publications/research-and-data/fertility-treatment-2019-trends-and-figures/>.
8. Aflatoonian, A., Eftekhari, M., Mohammadian, F. & Yousefnejad, F. Outcome of assisted reproductive technology in women aged 40 years and older. *Iran. J. Reprod. Med.* **9**, 281–284. ISSN: 1680-6433 (4 2011). ppublish.
9. Wyns, C., Geyter, C. D., Calhaz-Jorge, C., Kupka, M. S. *et al.* ART in Europe, 2018: results generated from European registries by ESHRE. *Hum. Reprod. Open* **2022** (2022).
10. London Women’s Clinic. *Welcome to London Women’s Clinic* <https://www.londonwomensclinic.com> (2023).
11. Bosch, E. & Ezcurra, D. Individualised controlled ovarian stimulation (iCOS): maximising success rates for assisted reproductive technology patients. *Reprod. Biol. Endocrinol.* **9**, 82 (2011).
12. Human Fertilisation & Embryology Authority (HFEA). *Intracytoplasmic sperm injection (ICSI)* 2023. <https://www.hfea.gov.uk/treatments/explore-all-treatments/intracytoplasmic-sperm-injection-icsi/>.

13. Garbhini, P. G., Suardika, A., Anantasika, A. A. N., Adnyana, I. B. P. *et al.* Day-3 vs. Day-5 fresh embryo transfer. *JBRA Assist. Reprod.* (2022).
14. Borghot, M. V. & Wyns, C. Fertility and infertility: Definition and epidemiology. *Clin. Biochem.* **62**, 2–10 (2018).
15. McLaren, J. F. Infertility Evaluation. *Obstet. Gynecol. Clin. North Am.* **39**, 453–463 (2012).
16. Organisation, W. H. *Infertility* <https://www.who.int/news-room/fact-sheets/detail/infertility> (2023).
17. Cousineau, T. M. & Domar, A. D. Psychological impact of infertility. *Best Pract. Res. Clin. Obstet. Gynaecol.* **21**, 293–308 (2007).
18. Petraglia, F., Serour, G. I. & Chapron, C. The changing prevalence of infertility. *International Journal of Gynecology Obstetrics* **123**, S4–S8 (2013).
19. Anwar, S. & Anwar, A. Infertility: A Review on Causes, Treatment and Management. *J. Womens Health Gynecol.* ISSN: 2369-307X (2016).
20. Wasser, S. K., Sewall, G. & Soules, M. R. Psychosocial stress as a cause of infertility. *Fertil. Steril.* **59**, 685–689 (1993).
21. Tsevat, D. G., Wiesenfeld, H. C., Parks, C. & Peipert, J. F. Sexually transmitted diseases and infertility. *Am. J. Obstet. Gynecol.* **216**, 1–9 (2017).
22. Hart, R. PCOS and infertility. *Panminerva Med.* **50**, 305–314. ISSN: 0031-0808 (4 2008). ppublish.
23. Bahadur, G. Fertility issues for cancer patients. *Mol. Cell. Endocrinol.* **169**, 117–122 (2000).
24. Henriksson, H. W. & Bergnehr, D. in *Single Parents: Representations and Resistance in an International Context* (ed Åström, B.) 215–234 (Springer International Publishing, 2021).
25. Pettersson, M. L., Sydsjö, G., Lampic, C., Svanberg, A. S. *et al.* Perceived social support in solo women seeking treatment with donor gametes and in women in heterosexual couples seeking IVF-treatment with own gametes. *Sci. Rep.* **13** (2023).
26. Marina, S., Marina, D., Marina, F., Fosas, N. *et al.* Sharing motherhood: biological lesbian co-mothers, a new IVF indication. *Hum. Reprod.* **25**, 938–941 (2010).
27. Yeshua, A., Lee, J. A., Witkin, G. & Copperman, A. B. Female Couples Undergoing IVF with Partner Eggs (Co-IVF): Pathways to Parenthood. *LGBT Health* **2**, 135–139 (2015).
28. Lindheim, S. R., Madeira, J. L., Ludwin, A., Kemner, E. *et al.* Societal pressures and procreative preferences for gay fathers successfully pursuing parenthood through IVF and gestational carriers. *Reprod. Biomed. Soc. Online* **9**, 1–10 (2019).
29. Patrizio, P., Albertini, D. F., Gleicher, N. & Caplan, A. The changing world of IVF: the pros and cons of new business models offering assisted reproductive technologies. *J. Assist. Reprod. Genet.* **39**, 305–313 (2022).

30. Wyns, C., Geyter, C. D., Calhaz-Jorge, C., Kupka, M. S. *et al.* ART in Europe, 2017: results generated from European registries by ESHRE. *Hum. Reprod. Open* **2021** (2021).
31. Ishizaki, T., Takeuchi, Y., Ishibashi, K., Gotoh, N. *et al.* Cryopreservation of tissues by slow-freezing using an emerging zwitterionic cryoprotectant. *Sci. Rep.* **13** (2023).
32. Abreu, C. W. D. P. A., Abreu, M. L. A., Abreu, M. M. A., Abreu, J. P. A. *et al.* Final destination of surplus cryopreserved embryos. What decision should be made? *JBRA Assist. Reprod.* **25**, 276–281 (2021).
33. Christianson, M. S., Stern, J. E., Sun, F., Zhang, H. *et al.* Embryo cryopreservation and utilization in the United States from 2004–2013. *F&S Reports* **1**, 71–77 (2020).
34. Moomjy, M., Cholst, I., Davis, O., Applegarth, L. *et al.* Donor Oocytes in Assisted Reproduction - An Overview. *Semin. Reprod. Med.* **13**, 173–186 (1995).
35. Kawwass, J. F. Trends and Outcomes for Donor Oocyte Cycles in the United States, 2000-2010. *JAMA* (2013).
36. Priddle, H. *Personal communications* 2023.
37. Goold, I. & Savulescu, J. In favour of freezing eggs for non-medical reasons. *Bioethics* **23**, 47–58 (2009).
38. European Society of Human Reproduction and Embryology Capri Workshop Group. Fertility and ageing. *Hum. Reprod. Update* **11**, 261–276 (2005).
39. Schmidt, L., Sobotka, T., Bentzen, J. G. & Andersen, A. N. Demographic and medical consequences of the postponement of parenthood. *Hum. Reprod. Update* **18**, 29–43 (2011).
40. For National Statistics, O. *Birth Characteristics in England and Wales: 2021* research rep. (2023). <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/livebirths/bulletins/birthcharacteristicsinenglandandwales/2021>.
41. Deatsman, S., Vasilopoulos, T. & Rhoton-Vlasak, A. Age and Fertility: A Study on Patient Awareness. *JBRA Assist. Reprod.* **20**, 99–106. ISSN: 1518-0557 (3 2016). epublish.
42. Eisenberg, M. L. & Meldrum, D. Effects of age on fertility and sexual function. *Fertil. Steril.* **107**, 301–304 (2017).
43. Human Fertilisation & Embryology Authority (HFEA). *Egg freezing infertility treatment* 2016. <https://www.hfea.gov.uk/media/2656/egg-freezing-in-fertility-treatment-trends-and-figures-2010-2016-final.pdf> (2023).
44. Human Fertilisation & Embryology Authority (HFEA). *Impact of COVID-19 on fertility treatment 2020* 2022. <https://www.hfea.gov.uk/about-us/publications/research-and-data/impact-of-covid-19-on-fertility-treatment-2020/>.

45. Human Fertilisation & Embryology Authority (HFEA). *Embryo Freezing* <https://www.hfea.gov.uk/treatments/fertility-preservation/embryo-freezing/> (2021).
46. United Kingdom Government. *Storage limit for frozen eggs, sperm and embryos to be reconsidered* 2021. <https://www.gov.uk/government/news/storage-limit-for-frozen-eggs-sperm-and-embryos-to-be-reconsidered> (2023).
47. Bishop, K. *How fertility became a workplace perk* (ed British Broadcasting Corporation (BBC)) 2021. <https://www.bbc.com/worklife/article/20211008-how-fertility-became-a-workplace-perk> (2023).
48. Lensen, S., Chen, S., Goodman, L., Rombauts, L. *et al.* IVF add-ons in Australia and New Zealand: A systematic assessment of IVF clinic websites. *Aust. N. Z. J. Obstet. Gynaecol.* **61**, 430–438 (2021).
49. Lensen, S., Hammarberg, K., Polyakov, A., Wilkinson, J. *et al.* How common is add-on use and how do patients decide whether to use them? A national survey of IVF patients. *Hum. Reprod.* **36**, 1854–1861 (2021).
50. Macklon, N. S., Ahuja, K. K. & Fauser, B. C. J. M. Building an evidence base for IVF ‘add-ons’. *Reprod. BioMed. Online* **38**, 853–856 (2019).
51. Perrotta, M. & Hamper, J. The crafting of hope: Contextualising add-ons in the treatment trajectories of IVF patients. *Social Science & Medicine* **287**, 114317 (2021).
52. Glatthorn, H. N. & Decherney, A. The efficacy of add-ons: selected IVF “add-on” procedures and future directions. *J. Assist. Reprod. Genet.* **39**, 581–589 (2022).
53. Massip, A. & Mulnard, J. Time-lapse cinematographic analysis of hatching of normal and frozen-thawed cow blastocysts. *Reprod.* **58**, 475–478. ISSN: 1470-1626 (1980).
54. Harrity, C., Arce, H., Ghaei, R. E., Keane, D. *et al.* Effect of the Geri incubator on clinical outcomes in fresh transfer ART cycles. *Fertil. Steril.* **114**, e144–e145 (2020).
55. Armstrong, S., Vail, A., Mastenbroek, S., Jordan, V. *et al.* Time-lapse in the IVF-lab: how should we assess potential benefit? *Human Reprod.* **30**, 3–8 (2014).
56. Racowsky, C., Kovacs, P. & Martins, W. P. A critical appraisal of time-lapse imaging for embryo selection: where are we and where do we need to go? *J. Assist. Reprod. Genet.* **32**, 1025–1030 (2015).
57. Meseguer, M., Herrero, J., Tejera, A., Hilligsoe, K. M. *et al.* The use of morphokinetics as a predictor of embryo implantation. *Hum. Reprod.* **26**, 2658–2671 (2011).
58. Basile, N., Carbajosa, A. R. & Meseguer, M. in Gardner, D. K., Weissman, A., Howles, C. M. & Shoham, Z. *Textbook of Assisted Reproductive Techniques, Fifth Edition : Volume 1 Laboratory Perspectives* 168–176 (Taylor & Francis Group, 2018). ISBN: 9781498740142.

59. Kirkegaard, K., Agerholm, I. E. & Ingerslev, H. J. Time-lapse monitoring as a tool for clinical embryo assessment. *Hum. Reprod.* **27**, 1277–1285 (2012).
60. Kovacs, P. Embryo selection: the role of time-lapse monitoring. *Reprod. Biol. Endocrinol.* **12**, 124 (2014).
61. Wong, C. C., Loewke, K. E., Bossert, N. L., Behr, B. *et al.* Non-invasive imaging of human embryos before embryonic genome activation predicts development to the blastocyst stage. *Nat. Biotechnol.* **28**, 1115–1121 (2010).
62. Kirkegaard, K., Ahlström, A., Ingerslev, H. J. & Hardarson, T. Choosing the best embryo by time lapse versus standard morphology. *Fertil. Steril.* **103**, 323–332 (2015).
63. Rienzi, L., Cimadomo, D., Vaiarelli, A., Gennarelli, G. *et al.* Measuring success in IVF is a complex multidisciplinary task: time for a consensus? *Reprod. BioMed. Online* **43**, 775–778 (2021).
64. Wilkinson, J., Roberts, S. A., Showell, M., Brison, D. R. *et al.* No common denominator: a review of outcome measures in IVF RCTs. *Hum. Reprod.* **31**, 2714–2722 (2016).
65. Ledger, W. L., Anumba, D., Marlow, N., Thomas, C. M. *et al.* Fertility and assisted reproduction: The costs to the NHS of multiple births after IVF treatment in the UK. *BJOG: Int. J. Obstet. Gynaecol.* **113**, 21–25 (2005).
66. Giunco, H., Connerney, M., Boylan, C., Koelper, N. *et al.* Embryo re-expansion does not affect clinical pregnancy rates in frozen embryo transfer cycles: a retrospective study. *J. Assist. Reprod. Genet.* **38**, 2933–2939 (2021).
67. Zhao, J., Yan, Y., Huang, X., Sun, L. *et al.* Blastocoele expansion: an important parameter for predicting clinical success pregnancy after frozen-warmed blastocysts transfer. *Reprod. Biol. Endocrinol.* **17** (2019).
68. Li, L., Zhou, L., Feng, T., Hao, G. *et al.* Ambient air pollution exposed during preantral-antral follicle transition stage was sensitive to associate with clinical pregnancy for women receiving IVF. *Environ. Pollut.* **265**, 114973 (2020).
69. Ebner, T., Oppelt, P., Radler, E., Allerstorfer, C. *et al.* Morphokinetics of vitrified and warmed blastocysts predicts implantation potential. *J. Assist. Reprod. Genet.* **34**, 239–244 (2016).
70. Tran, D., Cooke, S., Illingworth, P. J. & Gardner, D. K. Deep learning as a predictive tool for fetal heart pregnancy following time-lapse incubation and blastocyst transfer. *Hum. Reprod.* **34**, 1011–1018 (2019).
71. Du, Q., Wang, E., Huang, Y., Guo X. and Xiong, Y. *et al.* Blastocoele expansion degree predicts live birth after single blastocyst transfer for fresh and vitrified/warmed single blastocyst transfer cycles. *Fertil. Steril.* **105**, 910–919 (2016).
72. Bodri, D., Sugimoto, T., Serna, J. Y., Kawachiya, S. *et al.* Blastocyst collapse is not an independent predictor of reduced live birth: a time-lapse study. *Fertil. Steril.* **105**, 1476–1483 (2016).

73. Ahlstrom, A., Westin, C., Wikland, M. & Hardarson, T. Prediction of live birth in frozen-thawed single blastocyst transfer cycles by pre-freeze and post-thaw morphology. *Hum. Reprod.* **28**, 1199–1209 (2013).
74. Maheshwari, A., McLernon, D. & Bhattacharya, S. Cumulative live birth rate: time for a consensus? *Hum. Reprod.*, 2703–2707 (2015).
75. Griesinger, G. Beware of the ‘implantation rate’! Why the outcome parameter ‘implantation rate’ should be abandoned from infertility research. *Hum. Reprod.*, dev322 (2016).
76. National Health Service. *Ultrasound scans in pregnancy* <https://www.nhs.uk/pregnancy/your-pregnancy-care/ultrasound-scans/> (2023).
77. Balaban, B., Yakin, K. & Urman, B. Randomized comparison of two different blastocyst grading systems. *Fertil. Steril.* **85**, 559–563 (2006).
78. Hershko-Klement, A., Raviv, S., Nemerovsky, L., Rom, T. *et al.* Standardization of Post-Vitrification Human Blastocyst Expansion as a Tool for Implantation Prediction. *J. Clin. Med.* **11**, 2673 (2022).
79. Geyter, C. D., Wyns, C., Calhaz-Jorge, C., de Mouzon, J. *et al.* 20 years of the European IVF-monitoring Consortium registry: what have we learned? A comparison with registries from two other regions. *Hum. Reprod.* **35**, 2832–2849 (2020).
80. Busnelli, A., Reschini, M., Cardellicchio, L., Vegetti, W. *et al.* How common is real repeated implantation failure? An indirect estimate of the prevalence. *Reprod. Biomed. Online* **40**, 91–97 (2020).
81. Malina, A. & Pooley, J. Psychological consequences of IVF fertilization – Review of research. *Ann. Agric. Environ. Med.* **24**, 554–558 (2017).
82. Hynes, G. J., Callan, V. J., Terry, D. J. & Gallois, C. The psychological well-being of infertile women after a failed IVF attempt: The effects of coping. *Br. J. Clin. Psychol.* **65**, 269–278 (1992).
83. Archary, P., Potgieter, L., Elgindy, E., Adageba, R. K. *et al.* Assisted reproductive technologies in Africa: The African Network and Registry for ART, 2018 and 2019. *Reprod. Biomed. Online* **46**, 835–845 (2023).
84. Fauser, B. C. J. M. Towards the global coverage of a unified registry of IVF outcomes. *Reprod. Biomed. Online* **38**, 133–137 (2019).
85. European Society of Human Reproduction and Embryology (ESHRE). *ART fact sheet Report* (2023). <https://www.eshre.eu/Europe/Factsheets-and-infographics>.
86. Zegers-Hochschild, F., Crosby, J. A., Musri, C., Petermann-Rocha, F. *et al.* ART in Latin America: the Latin American Registry, 2020. *Reprod. BioMed. Online*, 103195 (2023).
87. Centers for Disease Control (CDC). *ART Success Rates* <https://www.cdc.gov/art/artdata/index.html> (2023).
88. Kushnir, V. A., Barad, D. H., Albertini, D. F., Darmon, S. K. *et al.* Systematic review of worldwide trends in assisted reproductive technology 2004–2013. *Reprod. Biol. Endocrinol.* **15** (2017).

89. Smith, J. F., Eisenberg, M. L., Glidden, D., Millstein, S. G. *et al.* Socioeconomic disparities in the use and success of fertility treatments: analysis of data from a prospective cohort in the United States. *Fertil. Steril.* **96**, 95–101 (2011).
90. Mladovsky, P. & Sorenson, C. Public Financing of IVF: A Review of Policy Rationales. *Health Care Anal.* **18**, 113–128 (2009).
91. Service, N. H. *IVF Availability* <https://www.nhs.uk/conditions/ivf/availability/> (2023).
92. Taylor, R. *Long NHS waiting lists cutting chances of pregnancy, fertility chief warns* (ed The Guardian) <https://www.theguardian.com/society/2023/jan/22/long-nhs-waiting-lists-cutting-chances-of-pregnancy-fertility-chief-warns> (2023).
93. Haimes, E. ‘Time waits for no one’: IVF patients’ evaluations of their experiences of the National Health Service waiting list. *Hum. Fert.* **2**, 122–126 (1999).
94. Human Fertilisation & Embryology Authority (HFEA). *How do clinics decide how many embryos to put back in the womb?* 2022. <https://www.hfea.gov.uk/treatments/explore-all-treatments/decisions-to-make-about-your-embryos/> (2022).
95. Messinis, I. E., Messini, C. I., Daponte, A., Garas, A. *et al.* The current situation of infertility services provision in Europe. *Eur. J. Obstet. Gynecol. Reprod. Biol.* **207**, 200–204 (2016).
96. Connolly, M. P., Hoorens, S. & Chambers, G. M. The costs and consequences of assisted reproductive technology: an economic perspective. *Hum. Reprod. Update* **16**, 603–613 (2010).
97. European Society of Human Reproduction and Embryology (ESHRE). *The funding of IVF treatment* 2017. <https://www.eshre.eu/-/media/sitecore-files/Press-room/Resources/4-Reimbursement.pdf?la=en&hash=C714A51012C21CCD5C475982A8BC18DDD0586806> (2022).
98. Of the American Society for Reproductive Medicine, T. E. C. Disparities in access to effective treatment for infertility in the United States: an Ethics Committee opinion. *Fertil. Steril.* **116**, 54–63 (2021).
99. Human Fertilisation & Embryology Authority (HFEA). *In Vitro Fertilisation (IVF)* <https://www.hfea.gov.uk/treatments/explore-all-treatments/in-vitro-fertilisation-ivf/> (2023).
100. Duffy, J. M. N., Adamson, G. D., Benson, E., Bhattacharya, S. *et al.* Top 10 priorities for future infertility research: an international consensus development study. *Hum. Reprod.* **35**, 2715–2724 (2020).
101. Ammar, O. F., Massarotti, C., Sharma, K., Fraire-Zamora, J. J. *et al.* ESHREjc report: non-invasive aneuploidy testing as an example of interdisciplinary innovation in ART. *Hum. Reprod.* **37**, 1694–1696 (2022).
102. Williams, C. J. & Erikson, G. F. Morphology and Physiology of the Ovary. *Endotext* (eds Feingold, K. R., Anawalt, B. & Blackman, M. R.) (2000). ppublish.

103. Hardy, K., Wright, C. S., Franks, S. & Winston, R. M. L. In Vitro maturation of oocytes. *Br. Med. Bull.* **56**, 588–602 (2000).
104. Thiyagarajan, D. K., Basit, H. & Jeanmonod., R. Physiology, Menstrual Cycle. *StatPearls* (2023). ppublish.
105. Santos, M. A., Kuijk, E. W. & Macklon, N. S. The impact of ovarian stimulation for IVF on the developing embryo. *Reproduction* **139**, 23–34. ISSN: 1741-7899 (2010).
106. Verberg, M. F. G., Eijkemans, M. J. C., Macklon, N. S., Heijnen, E. M. E. W. *et al.* The clinical significance of the retrieval of a low number of oocytes following mild ovarian stimulation for IVF: a meta-analysis. *Hum. Reprod. Update* **15**, 5–12 (2008).
107. Meyer, M., Broughton, D. & Jungheim, E. Best Practices for Controlled Ovarian Stimulation in In Vitro Fertilization. *Semin. Reprod. Med.* **33**, 077–082 (2015).
108. Hillier, S. G. The Parkes lecture: controlled ovarian stimulation in women. *Reproduction*, 201–210 (2000).
109. Rienzi, L. F. & Ubaldi, F. M. in Gardner, D. K., Weissman, A., Howles, C. M. & Shoham, Z. *Textbook of Assisted Reproductive Techniques, Fifth Edition : Volume 1 Laboratory Perspectives* 88–107 (Taylor & Francis Group, 2018). ISBN: 9781498740142.
110. Haahr, T., Esteves, S. C. & Humaidan, P. Individualized controlled ovarian stimulation in expected poor-responders: an update. *Reprod. Biol. Endocrinol.* **16** (2018).
111. Yovich, J., Alsbjerg, B., Conceicao, J., Hinchliffe, P. *et al.* PIVET rFSH dosing algorithms for individualized controlled ovarian stimulation enables optimized pregnancy productivity rates and avoidance of ovarian hyperstimulation syndrome. *Drug Des. Devel. Ther.* **Volume 10**, 2561–2573 (2016).
112. Palomba, S., Falbo, A., Orio, F., Manguso, F. *et al.* A randomized controlled trial evaluating metformin pre-treatment and co-administration in non-obese insulin-resistant women with polycystic ovary syndrome treated with controlled ovarian stimulation plus timed intercourse or intrauterine insemination. *Hum. Reprod.* **20**, 2879–2886 (2005).
113. De Castro, F., Morón, F. J., Montoro, L., Galán, J. J. *et al.* Human controlled ovarian hyperstimulation outcome is a polygenic trait. *Pharmacogenetics* **14**, 285–293 (2004).
114. Clark, L. H. Multiple Stable Periodic Solutions in a Model for Hormonal Control of the Menstrual Cycle. *Bull. Math. Biol.* **65**, 157–173 (2003).
115. Panza, N. M., Wright, A. A. & Selgrade, J. F. A delay differential equation model of follicle waves in women. *J. Biol. Dyn.* **10**, 200–221 (2015).
116. Reinecke, I. & Deuffhard, P. A complex mathematical model of the human menstrual cycle. *J. Theor. Biol.* **247**, 303–330 (2007).

117. Röblitz, S., Stötzel, C., Deuffhard, P., Jones, H. M. *et al.* A mathematical model of the human menstrual cycle for the administration of GnRH analogues. *J. Theor. Biol.* **321**, 8–27 (2013).
118. Fischer, S., Ehrig, R., Schäfer, S., Tronci, E. *et al.* Mathematical Modeling and Simulation Provides Evidence for New Strategies of Ovarian Stimulation. *Front. Endocrinol.* **12** (2021).
119. Sciorio, R., Miranian, D. & Smith, G. D. Non-invasive oocyte quality assessment. *Biol. Reprod.* **106**, 274–290 (2022).
120. Wang, Q. & Sun, Q.-Y. Evaluation of oocyte quality: morphological, cellular and molecular predictors. *Reprod. Fertil. Dev.* **19**, 1 (2007).
121. Keefe, D., Kumar, M. & Kalmbach, K. Oocyte competency is the key to embryo potential. *Fertil. Steril.* **103**, 317–322 (2015).
122. Lemseffer, Y., Terret, M.-E., Campillo, C. & Labrune, E. Methods for Assessing Oocyte Quality: A Review of Literature. *Biomedicines* **10**, 2184 (2022).
123. Kumar, S., Padubidri, V. G. & Daftary, S. N. *Howkins and Bourne. Shaw's Textbook of Gynaecology, 17 Edition* ISBN: 9788131254110 (Elsevier India, 2018).
124. Targosz, A., Przystałka, P., Wiaderkiewicz, R. & Mrugacz, G. Semantic segmentation of human oocyte images using deep neural networks. *Biomed. Eng. Online* **20** (2021).
125. Patrizio, P., Fragouli, E., Bianchi, V., Borini, A. *et al.* Molecular methods for selection of the ideal oocyte. *Reprod. BioMed. Online* **15**, 346–353 (2007).
126. Ebner, T., Moser, M. & Tews, G. Is oocyte morphology prognostic of embryo developmental potential after ICSI? *Reprod. BioMed. Online* **12**, 507–512 (2006).
127. Balaban, B. & Urman, B. Effect of oocyte morphology on embryo development and implantation. *Reprod. BioMed. Online* **12**, 608–615 (2006).
128. Sutter, P. D., Dozortsev, D., Qian, C. & Dhont, M. Oocyte morphology does not correlate with fertilization rate and embryo quality after intracytoplasmic sperm injection. *Hum. Reprod.* **11**, 595–597 (1996).
129. Ten, J., Mendiola, J., Vioque, J., de Juan, J. *et al.* Donor oocyte dysmorphisms and their influence on fertilization and embryo quality. *Reprod. BioMed. Online* **14**, 40–48 (2007).
130. Esfandiari, N., Burjaq, H., Gotlieb, L. & Casper, R. F. Brown oocytes: Implications for assisted reproductive technology. *Fertil. Steril.* **86**, 1522–1525 (2006).
131. Balaban, B., Ata, B., Isiklar, A., Yakin, K. *et al.* Severe cytoplasmic abnormalities of the oocyte decrease cryosurvival and subsequent embryonic development of cryopreserved embryos. *Hum. Reprod.* **23**, 1778–1785 (2008).
132. Loutradis, D., Drakakis, P., Kallianidis, K., Milingos, S. *et al.* Oocyte morphology correlates with embryo quality and pregnancy rate after intracytoplasmic sperm injection. *Fertil. Steril.* **72**, 240–244 (1999).

133. Sanchez, T., Venturas, M., Aghvami, S. A., Yang, X. *et al.* Combined non-invasive metabolic and spindle imaging as potential tools for embryo and oocyte assessment. *Hum. Reprod.* **34**, 2349–2361 (2019).
134. Bogliolo, L., Ledda, S., Innocenzi, P., Ariu, F. *et al.* Raman microspectroscopy as a non-invasive tool to assess the vitrification-induced changes of ovine oocyte zona pellucida. *Cryobiology* **64**, 267–272 (2012).
135. Ajduk, A., Ilozue, T., Windsor, S., Yu, Y. *et al.* Rhythmic actomyosin-driven contractions induced by sperm entry predict mammalian embryo viability. *Nat. Commun.* **2** (2011).
136. Alberts, B., Bray, D., Hopkin, K., Johnson, A. D. *et al.* *Essential Cell Biology 4th Edition* ISBN: 978-0-8153-4454-4 (Garland Publishing, 2013).
137. Bui, T. T. H., Belli, M., Fassina, L., Vigone, G. *et al.* Cytoplasmic movement profiles of mouse surrounding nucleolus and not-surrounding nucleolus antral oocytes during meiotic resumption. *Mol. Reprod. Dev.* **84**, 356–362 (2017).
138. Brangwynne, C. P., Koenderink, G. H., MacKintosh, F. C. & Weitz, D. A. Cytoplasmic diffusion: molecular motors mix it up. *J. Cell Biol.* **183**, 583–587 (2008).
139. Brown, R. XXVII. A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies. *Philos. Mag.* **4**, 161–173 (1828).
140. Brangwynne, C. P., Koenderink, G. H., MacKintosh, F. C. & Weitz, D. A. Intracellular transport by active diffusion. *Trends Cell Biol.* **19**, 423–427 (2009).
141. Guo, M., Ehrlicher, A. J., Jensen, M. H., Renz, M. *et al.* Probing the Stochastic, Motor-Driven Properties of the Cytoplasm Using Force Spectrum Microscopy. *Cell* **158**, 822–832 (2014).
142. Drechsler, M., Giavazzi, F., Cerbino, R. & Palacios, I. M. Active diffusion and advection in *Drosophila* oocytes result from the interplay of actin and microtubules. *Nat. Commun.* **8**, 1–11 (2017).
143. Almonacid, M., Ahmed, W. W., Bussonnier, M., Maily, P. *et al.* Active diffusion positions the nucleus in mouse oocytes. *Nat. Cell Biol.* **17**, 470–479 (2015).
144. Quinlan, M. E. Cytoplasmic Streaming in the *Drosophila* Oocyte. *Annu. Rev. Cell Dev.* **32**, 173–195 (2016).
145. Huang, W. Y. C., Cheng, X. & Ferrell, J. E. Cytoplasmic organization promotes protein diffusion in *Xenopus* extracts. *Nat. Commun.* **13** (2022).
146. Corti, B. *Osservazioni Microscopiche Sulla Tremella E Sulla Circolazione Del Fluido In Una Pianta Acquajoula, Dell' Abate Bonaventura Corti* (1774).
147. Goldstein, R. E., Tuval, I. & van de Meent, J.-W. Microfluidics of cytoplasmic streaming and its implications for intracellular transport. *Proc. Natl. Acad. Sci.* **105**, 3663–3667 (2008).

148. Pickard, W. F. The role of cytoplasmic streaming in symplastic transport. *Plant Cell Environ.* **26**, 1–15 (2003).
149. Allen, N. S. & Allen, R. D. Cytoplasmic streaming in green plants. *Annu. Rev. Biophys. Eng.* **7**, 497–526 (1978).
150. Ganguly, S., Williams, L. S., Palacios, I. M. & Goldstein, R. E. Cytoplasmic streaming in *Drosophila* oocytes varies with kinesin activity and correlates with the microtubule cytoskeleton architecture. *Proc. Natl. Acad. Sci.* **109**, 15109–15114 (2012).
151. Niwayama, R., Nagao, H., Kitajima, T. S., Hufnagel, L. *et al.* Bayesian Inference of Forces Causing Cytoplasmic Streaming in *Caenorhabditis elegans* Embryos and Mouse Oocytes. *PLOS ONE* **11** (ed Chiam, K.-H.) e0159917 (2016).
152. Swann, K., Windsor, S., Campbell, K., Elgmati, K. *et al.* Phospholipase C- ζ -induced Ca^{2+} oscillations cause coincident cytoplasmic movements in human oocytes that failed to fertilize after intracytoplasmic sperm injection. *Fertil. Steril.* **97**, 742–747 (2012).
153. Wang, H. & Dey, S. K. Roadmap to embryo implantation: clues from mouse models. *Nat. Rev. Genet.* **7**, 185–199 (2006).
154. Goldstein, R. E. & van de Meent, J.-W. A physical perspective on cytoplasmic streaming. *Interface Focus* **5**, 20150030 (2015).
155. Yi, K., Unruh, J. R., Deng, M., Slaughter, B. D. *et al.* Dynamic maintenance of asymmetric meiotic spindle position through Arp2/3-complex-driven cytoplasmic streaming in mouse oocytes. *Nat. Cell Biol.* **13**, 1252–1258 (2011).
156. Van de Meent, J.-W., Tuval, I. & Goldstein, R. E. Nature’s Microfluidic Transporter: Rotational Cytoplasmic Streaming at High Péclet Numbers. *Phys. Rev. Lett.* **101**, 178102 (2008).
157. Woodhouse, F. G. & Goldstein, R. E. Cytoplasmic streaming in plant cells emerges naturally by microfilament self-organization. *Proc. Natl. Acad. Sci.* **110**, 14132–14137 (2013).
158. Hu, J., Jafari, S., Han, Y., Grodzinsky, A. J. *et al.* Size- and speed-dependent mechanical behavior in living mammalian cytoplasm. *Proc. Natl. Acad. Sci.* **114**, 9529–9534. ISSN: 1091-6490 (2017).
159. Zabielski, L. & Mestel, A. J. Steady flow in a helically symmetric pipe. *J. Fluid Mech.* **370**, 297–320 (1998).
160. Selvaggi, L., Salemme, M., Vaccaro, C., Pesce, G. *et al.* Multiple-Particle-Tracking to investigate viscoelastic properties in living cells. *Methods* **51**, 20–26 (2010).
161. Ahmed, W. W., Fodor, É., Almonacid, M., Bussonnier, M. *et al.* Active Mechanics Reveal Molecular-Scale Force Kinetics in Living Oocytes. *Biophys. J.* **114**, 1667–1679 (2018).
162. Klughammer, N., Bischof, J., Schnellbacher, N. D., Callegari, A. *et al.* Cytoplasmic flows in starfish oocytes are fully determined by cortical contractions. *PLOS Comput. Biol.* **14** (ed Asthagiri, A. R.) e1006588 (2018).

163. Wilson, L. G., Martinez, V. A., Schwarz-Linek, J., Tailleur, J. *et al.* Differential Dynamic Microscopy of Bacterial Motility. *Physical Review Letters* **106**, 018101 (2011).
164. Germain, D., Leocmach, M. & Gibaud, T. Differential dynamic microscopy to characterize Brownian motion and bacteria motility. *Am. J. Phys.* **84**, 202–210. ISSN: 1943-2909. <http://dx.doi.org/10.1119/1.4939516> (2016).
165. Silverberg, K. M. & Turner, T. in Gardner, D. K., Weissman, A., Howles, C. M. & Shoham, Z. *Textbook of Assisted Reproductive Techniques, Fifth Edition : Volume 1 Laboratory Perspectives* 36–49 (Taylor & Francis Group, 2018). ISBN: 9781498740142.
166. Alberts, B., Johnson, A., Lewis, J., Raff, M. *et al.* *Molecular Biology of the Cell, Fourth Edition* 1616. ISBN: 0-8153-3218-1 (Garland, 2002).
167. Millan, N. M., Lau, P., Hann, M., Ioannou, D. *et al.* Hierarchical radial and polar organisation of chromosomes in human sperm. *Chromosome Res.* **20**, 875–887 (2012).
168. Liu, D. Y. & Baker, H. W. G. Evaluation and assessment of semen for IVF/ICSI. *Asian J. Androl.* **4**, 281–285. ISSN: 1008-682X (4 2002). ppublish.
169. Gaffney, E. A., Ishimoto, K. & Walker, B. J. Modelling Motility: The Mathematics of Spermatozoa. *Front. Cell Dev. Bio.* **9** (2021).
170. Fauci, L. J. & McDonald, A. Sperm motility in the presence of boundaries. *Bull. Math. Biol.* **57**, 679–699 (1995).
171. Dillon, R. H., Fauci, L. J. & Yang, X. Sperm Motility and Multiciliary Beating: An Integrative Mechanical Model. *Comput. Math. Appl.* **52**, 749–758 (2006).
172. Gaffney, E. A., Gadêlha, H., Smith, D. J., Blake, J. R. *et al.* Mammalian Sperm Motility: Observation and Theory. *Annu. Rev. Fluid Mech.* **43**, 501–528 (2011).
173. Finelli, R., Leisegang, K., Tumallapalli, S., Henkel, R. *et al.* The validity and reliability of computer-aided semen analyzers in performing semen analysis: a systematic review. *Transl. Androl. Urol.* **10**, 3069–3079. ISSN: 2223-4691 (7 2021). ppublish.
174. Agarwal, A., Henkel, R., Huang, C. & Lee, M. Automation of human semen analysis using a novel artificial intelligence optical microscopic technology. *Andrologia* **51** (2019).
175. Agarwal, A., Selvam, M. K. P. & Ambar, R. F. Validation of LensHooke® X1 PRO and Computer-Assisted Semen Analyzer Compared with Laboratory-Based Manual Semen Analysis. *World J. Men's Health* **39**, 496 (2021).
176. Inge, G. B., Brinsden, P. R. & Elder, K. T. Oocyte number per live birth in IVF: were Steptoe and Edwards less wasteful? *Hum. Reprod.* **20**, 588–592 (2005).

177. Magnusson, Å., Källén, K., Thurin-Kjellberg, A. & Bergh, C. The number of oocytes retrieved during IVF: a balance between efficacy and safety. *Hum. Reprod.* **33**, 58–64 (2017).
178. Walker, Z., Lanes, A. & Ginsburg, E. Oocyte cryopreservation review: outcomes of medical oocyte cryopreservation and planned oocyte cryopreservation. *Reprod. Biol. Endocrinol.* **20** (2022).
179. Thurin, A., Hausken, J., Hillensjö, T., Jablonowska, B. *et al.* Elective Single-Embryo Transfer versus Double-Embryo Transfer in in Vitro Fertilization. *New Engl. J. Med.* **351**, 2392–2402 (2004).
180. Sazonova, A., Källén, K., Thurin-Kjellberg, A., Wennerholm, U.-B. *et al.* Neonatal and maternal outcomes comparing women undergoing two in vitro fertilization (IVF) singleton pregnancies and women undergoing one IVF twin pregnancy. *Fertil. Steril.* **99**, 731–737 (2013).
181. Gerris, J. M. R. Single embryo transfer and IVF/ICSI outcome: a balanced appraisal. *Hum. Reprod. Update* **11**, 105–121 (2005).
182. Human Fertilisation & Embryology Authority (HFEA). *Our campaign to reduce multiple births* <https://www.hfea.gov.uk/about-us/our-campaign-to-reduce-multiple-births/>.
183. Montag, M., Toth, B. & Strowitzki, T. New approaches to embryo selection. *Reprod. BioMed. Online* **27**, 539–546 (2013).
184. Mastenbroek, S., van der Veen, F., Aflatoonian, A., Shapiro, B. *et al.* Embryo selection in IVF. *Hum. Reprod.* **26**, 964–966 (2011).
185. Montag, M., Liebenthron, J. & Köster, M. Which morphological scoring system is relevant in human embryo development? *Placenta* **32**, S252–S256 (2011).
186. Wong, K., Repping, S. & Mastenbroek, S. Limitations of Embryo Selection Methods. *Semin. Reprod. Med.* **32**, 127–133 (2014).
187. Sadler, T. W. *Langman's medical embryology* ISBN: 9780781790697 (Lippincott William & Wilkins, 2009).
188. Wassarman, P. M. & Litscher, E. S. Female fertility and the zona pellucida. *eLife* **11** (2022).
189. Cecchele, A., Cermisoni, G. C., Giacomini, E., Pinna, M. *et al.* Cellular and Molecular Nature of Fragmentation of Human Embryos. *Int. J. Mol. Sci.* **23**, 1349 (2022).
190. Khosravi, P., Kazemi, E., Zhan, Q., Malmsten, J. E. *et al.* Deep learning enables robust assessment and selection of human blastocysts after in vitro fertilization. *NPJ Digit. Med.* **2** (2019).
191. Gardner, D. K., Surrey, E., Minjarez, D., Leitz, A. *et al.* Single blastocyst transfer: a prospective randomized trial. *Fertil. Steril.* **81**, 551–555 (2004).
192. Boiso, I., Veiga, A. & Edwards, R. G. Fundamentals of human embryonic growth in vitro and the selection of high-quality embryos for transfer. *Reprod. Biomed. Online* **5**, 328–350 (2002).

193. Weitzman, V. N., Schnee-Riesz, J., Benadiva, C., Nulsen, J. *et al.* Predictive value of embryo grading for embryos with known outcomes. *Fertil. Steril.* **93**, 658–662 (2010).
194. Cimadomo, D., Fernandez, L. S., Soscia, D., Fabozzi, G. *et al.* Inter-centre reliability in embryo grading across several IVF clinics is limited: implications for embryo selection. *Reprod. Biomed. Online* **44**, 39–48 (2022).
195. Martínez-Granados, L., Serrano, M., González-Utor, A., Ortiz, N. *et al.* Reliability and agreement on embryo assessment: 5 years of an external quality control programme. *Reprod. Biomed. Online* **36**, 259–268 (2018).
196. Payne, D., Flaherty, S. P., Barry, M. F. & Matthews, C. D. Preliminary observations on polar body extrusion and pronuclear formation in human oocytes using time-lapse video cinematography. *Hum. Reprod.* **12**, 532–541 (1997).
197. Niakan, K. K., Han, J., Pedersen, R. A., Simon, C. *et al.* Human pre-implantation embryo development. *Development* **139**, 829–841 (2012).
198. Mandawala, A. A., Harvey, S. C., Roy, T. K. & Fowler, K. E. Time-lapse embryo imaging and morphokinetic profiling: Towards a general characterisation of embryogenesis. *Anim. Reprod. Sci.* **174**, 2–10 (2016).
199. Kramer, Y. G., Kofinas, J. D., Melzer, K., Noyes, N. *et al.* Assessing morphokinetic parameters via time lapse microscopy (TLM) to predict euploidy: are aneuploidy risk classification models universal? *J. Assist. Reprod. Genet.* **31**, 1231–1242 (2014).
200. Chavez, S. L., Loewke, K. E., Han, J., Moussavi, F. *et al.* Dynamic blastomere behaviour reflects human embryo ploidy by the four-cell stage. *Nat. Commun.* **3** (2012).
201. Fragouli, E., Alfarawati, S., Spath, K., Jaroudi, S. *et al.* The origin and impact of embryonic aneuploidy. *Hum. Genet.* **132**, 1001–1013 (2013).
202. Hassold, T. & Hunt, P. To err (meiotically) is human: the genesis of human aneuploidy. *Nat. Rev. Genet.* **2**, 280–291 (2001).
203. Castelló, D., Motato, Y., Basile, N., Remohí, J. *et al.* How much have we learned from time-lapse in clinical IVF? *Mol. Hum. Reprod.* **22**, 719–727 (2016).
204. Chen, A. A., Tan, L., Suraj, V., Pera, R. R. *et al.* Biomarkers identified with time-lapse imaging: discovery, validation, and practical application. *Fertil. Steril.* **99**, 1035–1043 (2013).
205. Lundin, K. & Park, H. Time-lapse technology for embryo culture and selection. *Ups. J. Med. Sci.* **125**, 77–84 (2020).
206. Conaghan, J., Chen, A. A., Willman, S. P., Ivani, K. *et al.* Improving embryo selection using a computer-automated time-lapse image analysis test plus day 3 morphology: results from a prospective multicenter trial. *Fertil. Steril.* **100**, 412–419 (2013).

207. Diamond, M. P., Suraj, V., Behnke, E. J., Yang X. and Angle, M. J. *et al.* Using the Eeva Test™ adjunctively to traditional day 3 morphology is informative for consistent embryo assessment within a panel of embryologists with diverse experience. *J. Assist. Reprod. Genet.* **32**, 61–68 (2014).
208. Glatstein, I., Chavez-Badiola, A. & Curchoe, C. L. New frontiers in embryo selection. *J. Assist. Reprod. Genet.* **40**, 223–234 (2023).
209. Raudonis, V., Paulauskaite-Taraseviciene, A., Sutiene, K. & Jonaitis, D. Towards the automation of early-stage human embryo development detection. *Biomed. Eng. Online* **18** (2019).
210. Chavez-Badiola, A., Saiffe-Farías, A. F., Mendizabal-Ruiz, G., Drakeley, A. J. *et al.* Embryo Ranking Intelligent Classification Algorithm (ERICA): artificial intelligence clinical assistant predicting embryo ploidy and implantation. *Reprod. Biomed. Online* **41**, 585–593 (2020).
211. Bori, L., Paya, E., Alegre, L., Vilorio, T. A. *et al.* Novel and conventional embryo parameters as input data for artificial neural networks: an artificial intelligence model applied for prediction of the implantation potential. *Fertil. Steril.* **114**, 1232–1241 (2020).
212. Kragh, M. F. & Karstoft, H. Embryo selection with artificial intelligence: how to evaluate and compare methods? *J. Assist. Reprod. Genet.* **38**, 1675–1689 (2021).
213. Bormann, C. L., Thirumalaraju, P., Kanakasabapathy, M. K., Kandula, H. *et al.* Consistency and objectivity of automated embryo assessments using deep neural networks. *Fertil. Steril.* **113**, 781–787 (2020).
214. Munck, N. D., Verheyen, G., Landuyt, L. V., Stoop, D. *et al.* Survival and post-warming in vitro competence of human oocytes after high security closed system vitrification. *J. Assist. Reprod. Genet.* **30**, 361–369 (2013).
215. Valojerdi, M. R., Eftekhari-Yazdi, P., Karimian, L., Hassani, F. *et al.* Vitrification versus slow freezing gives excellent survival, post warming embryo morphology and pregnancy outcomes for human cleaved embryos. *J. Assist. Reprod. Genet.* **26**, 347–354 (2009).
216. Evans, J., Hannan, N. J., Edgell, T. A., Vollenhoven, B. J. *et al.* Fresh versus frozen embryo transfer: backing clinical decisions with scientific and clinical evidence. *Hum. Reprod. Update* **20**, 808–821 (2014).
217. Mazur, P. Cryobiology: The Freezing of Biological Systems. *Science* **168**, 939–949 (1970).
218. Gao, D. & Critser, J. K. Mechanisms of Cryoinjury in Living Cells. *ILAR J.* **41**, 187–196 (2000).
219. Akhoondi, M., Oldenhof, H., Sieme, H. & Wolkers, W. F. Freezing-induced cellular and membrane dehydration in the presence of cryoprotective agents. *Mol. Membr. Biol.* **29**, 197–206 (2012).
220. Mazur, P. Kinetics of Water Loss from Cells at Subzero Temperatures and the Likelihood of Intracellular Freezing. *J. Gen. Physiol.* **47**, 347–369 (1963).

221. Fang, L., Jin, L., Li, E., Cui, L. *et al.* Clinical evaluation of two formulations of slow-freezing solutions for cleavage stage embryos. *J. Assist. Reprod. Genet.* **33**, 1389–1393 (2016).
222. Son, W.-Y. & Tan, S. L. Comparison between slow freezing and vitrification for human embryos. *Expert Rev. Med. Devices* **6**, 1–7 (2009).
223. Arav, A. & Natan, Y. in *Oocyte Biology in Fertility Preservation* (ed Kim, S. S.) 69–84 (Springer, 2015).
224. Kuwayama, M., Vajta, G., Kato, O. & Leibo, S. P. Highly efficient vitrification method for cryopreservation of human oocytes. *Reprod. Biomed. Online* **11**, 300–308 (2005).
225. Rall, W. F. & Fahy, G. M. Ice-free cryopreservation of mouse embryos at -196°C by vitrification. *Nature* **313**, 573–575 (1985).
226. Vajta, G., Rienzi, L. & Ubaldi, F. M. Open versus closed systems for vitrification of human oocytes and embryos. *Reprod. Biomed. Online* **30**, 325–333 (2015).
227. Li, W. J., Zhou, X. L., Wang, H. S., Liu, B. L. *et al.* Heat transfer coefficient of cryotop during freezing. *Cryo-Lett.* **34**, 255–260. ISSN: 0143-2044. <http://europepmc.org/abstract/MED/23812315> (3 2013). ppublish.
228. Mazur, P. & Seki, S. 92. The dominance of warming rate over cooling rate in the survival of mouse oocytes subjected to a vitrification procedure. *Cryobiology* **59**, 396 (2009).
229. Seki, S. & Mazur, P. Ultra-Rapid Warming Yields High Survival of Mouse Oocytes Cooled to -196°C in Dilutions of a Standard Vitrification Solution. *PLOS ONE* **7** (ed Baltz, J. M.) e36058 (2012).
230. Zhang, X., Catalano, P. N., Gurkan, U. A., Khimji, I. *et al.* Emerging technologies in medical applications of minimum volume vitrification. *Nanomed.* **6**, 1115–1129 (2011).
231. Hochi, S. Cryodevices developed for minimum volume cooling vitrification of bovine oocytes. *Anim. Sci. J.* **93** (2022).
232. Jain, J. K. & J.Paulson, R. Oocyte cryopreservation. *Fertil. Steril.* **86**, 1037–1046 (2006).
233. Vajta, G., Holm, P., Kuwayama, M., Booth, P. J. *et al.* Open pulled straw (OPS) vitrification: A new way to reduce cryoinjuries of bovine ova and embryos. *Mol. Reprod. Dev.* **51**, 53–58 (1998).
234. Mukaida, T. Vitrification of human blastocysts using cryoloops: clinical outcome of 223 cycles. *Hum. Reprod.* **18**, 384–391 (2003).
235. Park, S.-P. Ultra-rapid freezing of human multipronuclear zygotes using electron microscope grids. *Hum. Reprod.* **15**, 1787–1790 (2000).
236. Kitazato Corporation. *Cryotop vitrification brochure* https://www.kitazato-ivf.com/wp-content/uploads/2021/06/Brochure_Vitrification_v2021.pdf (2023).
237. Kitazato Corporation. *A device for manipulating biological materials in a process of cryopreservation and a use of such a device* 2021. <https://patents.google.com/patent/AU2011254894A1/en>.

238. Aizer, A., Noach-Hirsh, M., Dratviman-Storobinsky, O., Haas, J. *et al.* Does the number of embryos loaded on a single cryo-carrier affect post-vitrification survival rate? *Zygote* **29**, 87–91 (2020).
239. Tain, D. *Oocyte reading* (ed LinkedIn) https://www.linkedin.com/posts/diana-tain-ab7155159_embryology-embryologist-embryologistsarefunpeopletofollow-activity-7042157097966338048-K17W?utm_source=share&utm_medium=member_desktop.
240. Sansinena, M., Santos, M. V., Zaritzky, N. & Chirife, J. Numerical simulation of cooling rates in vitrification systems used for oocyte cryopreservation. *Cryobiology* **63**, 32–37. ISSN: 0011-2240. <https://doi.org/10.1016/j.cryobiol.2011.04.006> (2011).
241. Sansinena, M., Santos, M. V., Taminelli, G. & Zaritzky, N. Implications of storage and handling conditions on glass transition and potential devitrification of oocytes and embryos. *Theriogenology* **82**, 373–378 (2014).
242. Santos, M. V., Sansinena, M., Zaritzky, N. & Chirife, J. Assessment of external heat transfer coefficient during oocyte vitrification in liquid and slush nitrogen using numerical simulations to determine cooling rates. *Cryo letters* **33**, 31–40. ISSN: 0143-2044 (1 2012). ppublish.
243. Santos, M. V., Sansinena, M., Chirife, J. & Zaritzky, N. Mathematical modeling of the heat transfer process in vitrification devices used for oocyte cryopreservation. *Cryobiology* **65**, 353. ISSN: 2591-3522 (2012).
244. Santos, M. V., Sansinena, M., Chirife, J. & Zaritzky, N. Determination of heat transfer coefficients in plastic French straws plunged in liquid nitrogen. *Cryobiology* **69**, 488–495. ISSN: 0011-2240. <https://www.sciencedirect.com/science/article/pii/S0011224014006348> (2014).
245. Santos, M. V., Sansinena, M., Chirife, J. & Zaritzky, N. Convective heat transfer coefficients of open and closed Cryotop® systems under different warming conditions. *Cryobiology* **84**, 20–26. ISSN: 0011-2240. <http://www.sciencedirect.com/science/article/pii/S001122401830172X> (2018).
246. Mazur, P. & Seki, S. Survival of mouse oocytes after being cooled in a vitrification solution to -196°C at 95° to 70,000°C/min and warmed at 610° to 118,000°C/min: A new paradigm for cryopreservation by vitrification. *Cryobiology* **62**, 1–7 (2011).
247. Iwasawa, T., Takahashi, K., Goto, M., Anzai, M. *et al.* Human frozen-thawed blastocyst morphokinetics observed using time-lapse cinematography reflects the number of trophectoderm cells. *PLOS ONE* **14** (ed Yu, Y.) e0210992 (2019).
248. Kitazato Corporation. *Vitrification Protocol for Cryotop method* https://www.kitazato-dibimed.com/wp-content/uploads/CRYOTOP_SAFETYKIT_A5_2019_LOW.pdf (2021).
249. Shu, Y., Watt, J., Gebhardt, J., Dasig, J. *et al.* The value of fast blastocoele re-expansion in the selection of a viable thawed blastocyst for transfer. *Fertil. Steril.* **91**, 401–406 (2009).

250. Coello, A., Meseguer, M., Galán, A., Alegre, L. *et al.* Analysis of the morphological dynamics of blastocysts after vitrification/warming: defining new predictive variables of implantation. *Fertil. Steril.* **108**, 659–666 (2017).
251. Coello, A., Meseguer, M., Galán, A., Alegre, L. *et al.* Dynamics of blastocyst re-expansion after vitrification: defining new morphokinetic parameters and their relationship with implantation. *Fertil. Steril.* **104**, e193–e194 (2015).
252. Desai, N. & Goldfarb, J. Examination of frozen cycles with replacement of a single thawed blastocyst. *Reprod. Biomed. Online* **11**, 349–354 (2005).
253. Maezawa, T., Yamanaka, M., Hashimoto, S., Amo, A. *et al.* Possible selection of viable human blastocysts after vitrification by monitoring morphological changes. *J. Assist. Reprod. Genet.* **31**, 1099–1104 (2014).
254. Niimura, S. Time-Lapse Videomicrographic Analyses of Contractions in Mouse Blastocysts. *J. Reprod. Dev.* **49**, 413–423 (2003).
255. *Time-lapse analysis of re-expansion of warmed blastocysts: is re-expansion related to pregnancy outcome?* 145 (Informa UK Limited, 2022).
256. Muharam, R. & Firman, F. Lean Management Improves the Process Efficiency of Controlled Ovarian Stimulation Monitoring in IVF Treatment. *J. Healthc. Eng.* **2022** (ed Improta, G.) 1–10 (2022).
257. Parmegiani, L., Maccarini, A. M., Rastellini, A., Bernadi, S. *et al.* Oocyte vitrification/storage/handling/transportation/warming, effect on survival and clinical results in donation programmes. *Curr. Trends Clin. Embryol.* **4**, 34 (2017).
258. Ageron, B., Benzidia, S. & Boulakis, M. Healthcare logistics and supply chain – issues and future challenges. *Supply Chain Forum* **19**, 1–3 (2018).
259. Brandeau, M. L., Sainfort, F. & Pierskalla, W. P. *Operations Research and Health Care* (eds Brandeau, M. L., Sainfort, F. & Pierskalla, W. P.) (Springer US, 2004).
260. Denton, B. T. *Handbook of Healthcare Operations Management. Methods and Applications* (ed Denton, B. T.) 536. ISBN: 978-1-4614-5884-5 (Springer New York, 2013).
261. Younglai, E. V., Holloway, A. C. & Foster, W. G. Environmental and occupational factors affecting fertility and IVF success. *Hum. Reprod. Update* **11**, 43–57 (2005).
262. Lefebvre, T., Fréour, T., Ploteau, S., Bizec, B. L. *et al.* Associations between human internal chemical exposure to Persistent Organic Pollutants (POPs) and In Vitro Fertilization (IVF) outcomes: Systematic review and evidence map of human epidemiological evidence. *Reprod. Toxicol.* **105**, 184–197 (2021).
263. Perin, P. M., Maluf, M., Czeresnia, C. E., Januário, D. A. N. F. *et al.* Impact of short-term preconceptional exposure to particulate air pollution on treatment outcome in couples undergoing in vitro fertilization and embryo transfer (IVF/ET). *J. Assist. Reprod. Genet.* **27**, 371–382 (2010).

264. Hanna, C. W., Bloom, M. S., Robinson, W. P., Kim, D. *et al.* DNA methylation changes in whole blood is associated with exposure to the environmental contaminants, mercury, lead, cadmium and bisphenol A, in women undergoing ovarian stimulation for IVF. *Hum. Reprod.* **27**, 1401–1410 (2012).
265. O’Raghallaigh, R. *How green is your IVF?* (International IVF Initiative (i3), 2022). <https://ivfmeeting.com/products/how-green-is-your-ivf> (2023).
266. Cerbino, R. & Trappe, V. Differential dynamic microscopy: probing wave vector dependent dynamics with a microscope. *Phys. Rev. Lett.* **100**, 188102 (2008).
267. Giavazzi, F. & Cerbino, R. Digital Fourier microscopy for soft matter dynamics. *J. Opt.* **16**, 083001 (2014).
268. Murphy, D. B. & Davidson, M. W. *Fundamentals of Light Microscopy and Electronic Imaging* (John Wiley & Sons, Inc., 2012).
269. Colin, R., Zhang, R. & Wilson, L. G. Fast, high-throughput measurement of collective behaviour in a bacterial population. *J. R. Soc. Interface* **11**, 20140486 (2014).
270. Jepson, A., Arlt, J., Statham, J., Spilman, M. *et al.* High-throughput characterisation of bull semen motility using differential dynamic microscopy. *PloS one* **14** (ed Gaffney, E. A.) e0202720 (2019).
271. He, S., Pascucci, D. R., Caggioni, M., Lindberg, S. *et al.* Rheological properties of phase transitions in polydisperse and monodisperse colloidal rod systems. *AIChE J.* **67** (2021).
272. Laurati, M., Sentjabrskaja, T., Ruiz-Franco, J., Egelhaaf, S. U. *et al.* Different scenarios of dynamic coupling in glassy colloidal mixtures. *Phys. Chem. Chem. Phys.* **20**, 18630–18638 (2018).
273. Vasudevan, S. A., Rauh, A., Kröger, M., Karg, M. *et al.* Dynamics and Wetting Behavior of Core-Shell Soft Particles at a Fluid-Fluid Interface. *Langmuir* **34**, 15370–15382 (2018).
274. Manzo, C. & Garcia-Parajo, M. F. A review of progress in single particle tracking: from methods to biophysical insights. *Rep. Prog. Phys.* **78**, 124601 (2015).
275. Cognet, L., Leduc, C. & Lounis, B. Advances in live-cell single-particle tracking and dynamic super-resolution imaging. *Curr. Opin. Chem. Biol.* **20**, 78–85 (2014).
276. Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J. *et al.* *Particle image velocimetry: a practical guide* (Springer, 2018).
277. Adrian, R. J. & Westerwell, J. *Particle Image Velocimetry* 586 pp. ISBN: 0521440084. https://www.ebook.de/de/product/12359208/ronald_j_adrian_adrian_ronald_j_particle_image_velocimetry.html (Cambridge University Press, 2010).

278. Westerweel, J., Elsinga, G. E. & Adrian, R. J. Particle Image Velocimetry for Complex and Turbulent Flows. *Annu. Rev. Fluid Mech.* **45**, 409–436 (2013).
279. Cerbino, R. & Cicuta, P. Perspective: Differential dynamic microscopy extracts multi-scale activity in complex fluids and biological systems. *J. Chem. Phys.* **147**, 110901 (2017).
280. Reufer, M., Martinez, V. A., Schurtenberger, P. & Poon, W. C. K. Differential Dynamic Microscopy for Anisotropic Colloidal Dynamics. *Langmuir* **28**, 4618–4624 (2012).
281. Martinez, V. A., Besseling, R., Croze, O. A., Tailleur, J. *et al.* Differential Dynamic Microscopy: A High-Throughput Method for Characterizing the Motility of Microorganisms. *Biophys. J.* **103**, 1637–1647 (2012).
282. Brizioli, M., Sentjabrskaja, T., Egelhaaf, S. U., Laurati, M. *et al.* Reciprocal Space Study of Brownian Yet Non-Gaussian Diffusion of Small Tracers in a Hard-Sphere Glass. *Front. Phys.* **10** (2022).
283. Karlin, S. & Taylor, H. E. *First Course in Stochastic Processes* 557. ISBN: 9780080570419 (Elsevier Science & Technology Books, 2012).
284. Cramer, H. On the Theory of Stationary Random Processes. *Ann. Math.* **41**, 215 (1940).
285. Hansen, J.-P. & McDonald, I. R. *Theory of Simple Liquids With Applications to Soft Matter. With Applications to Soft Matter* 636. ISBN: 9780123870322 (Elsevier Science & Technology Books, 2013).
286. Giavazzi, F., Brogioli, D., Trappe, V., Bellini, T. *et al.* Scattering information obtained by optical microscopy: Differential dynamic microscopy and beyond. *Phys. Rev. E* **80**, 031403 (2009).
287. Van Dijk, G. *Distribution theory convolution, fourier transform, and laplace transform. convolution, fourier transform, and laplace transform* 105. ISBN: 9783110295917 (De Gruyter, 2013).
288. Srivastava, H. M. & Buschman, R. G. *Theory and Applications of Convolution Integral Equations* (Springer Netherlands, 1992).
289. McLean, I. S. *Electronic Imaging in Astronomy* 2nd, 241–275. ISBN: 9783540765837 (Springer Berlin Heidelberg, 2008).
290. Van Hove, L. Correlations in space and time and Born approximation scattering in systems of interacting particles. *Phys. Rev.* **95**, 249 (1954).
291. Vineyard, G. H. Scattering of slow neutrons by a liquid. *Phys. Rev.* **110**, 999 (1958).
292. Bayles, A. V., Squires, T. M. & Helgeson, M. E. Dark-field differential dynamic microscopy. *Soft Matter* **12**, 2440–2452 (2016).
293. Rahman, M. *Applications of Fourier transforms to generalized functions* 170. ISBN: 1845645642 (WIT Press, 2011).
294. McGillem, C. D. *Continuous and discrete signal and system analysis* 418. ISBN: 0030617030 (Holt, Rinehart, and Winston, 1984).

295. Giavazzi, F., Edera, P., Lu, P. J. & Cerbino, R. Image windowing mitigates edge effects in Differential Dynamic Microscopy. *Eur. Phys. J. E* **40**, 1–9 (2017).
296. Goldberg, W. I. Dynamic light scattering. *Am. J. Phys.* **67**, 1152–1160 (1999).
297. Struntz, P. & Weiss, M. The hitchhiker’s guide to quantitative diffusion measurements. *Phys. Chem. Chem. Phys.* **20**, 28910–28919 (2018).
298. Perrin, J. B. *Nobel Lecture, Discontinuous Structure of Matter* 1926. <https://www.nobelprize.org/prizes/physics/1926/perrin/lecture/> (2022).
299. Einstein, A. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Ann. Phys.* **4** (1905).
300. Pecora, R. *Dynamic Light Scattering. Applications of Photon Correlation Spectroscopy* (ed Pecora, R.) (Springer US, 1985).
301. Hanasaki, I. & Ooi, Y. Particle image diffusometry: Resolving diffusion coefficient field from microscopy movie data without particle tracking. *AIP Adv.* **8**, 065014 (2018).
302. Nixon-Luke, R., Arlt, J., Poon, W. C. K., Bryant, G. *et al.* Probing the dynamics of turbid colloidal suspensions using differential dynamic microscopy. *Soft Matter* **18**, 1858–1867 (2022).
303. Chang, C.-C., Zhang, B., Li, C.-Y., Hsieh, C.-C. *et al.* Exploring cytoplasmic dynamics in zebrafish yolk cells by single particle tracking of fluorescent nanodiamonds in *SPIE Proceedings. Advances in Photonics of Quantum Computing, Memory, and Communication* (eds Hasan, Z. U., Hemmer, P. R., Lee, H. & Santori, C. M.) (SPIE, 2012).
304. Liu, M., Li, Q., Liang, L., Li, J. *et al.* Real-time visualization of clustering and intracellular transport of gold nanoparticles by correlative imaging. *Nat. Commun.* **8** (2017).
305. Gu, M., Luo, Y., He, Y., Helgeson, M. E. *et al.* Uncertainty quantification and estimation in differential dynamic microscopy. *Phys. Rev. E* **104**, 034610 (2021).
306. Croze, O. A., Martinez, V. A., Jakuszeit, T., Dell’Arciprete, D. *et al.* Helical and oscillatory microswimmer motility statistics from differential dynamic microscopy. *New J. Phys.* **21**, 063012 (2019).
307. Larobina, M. & Murino, L. Medical Image File Formats. *J. Digit. Imaging* **27**, 200–206 (2013).
308. Moler, C. & Eddins, S. *Faster Finite Fourier Transforms MATLAB* (ed Mathworks) <https://uk.mathworks.com/company/newsletters/articles/faster-finite-fourier-transforms-matlab.html> (2023).
309. Lintusaari, J., Gutmann, M. U., Dutta, R., Kaski, S. *et al.* Fundamentals and Recent Developments in Approximate Bayesian Computation. *Syst. Biol.*, syw077 (2016).

310. Martineau, R. L., Bayles, A. V., Hung, C.-S., Reyes, K. G. *et al.* Engineering Gelation Kinetics in Living Silk Hydrogels by Differential Dynamic Microscopy Microrheology and Machine Learning. *Adv. Biol.* **6**, 2101070 (2021).
311. Cerbino, R., Giavazzi, F. & Helgeson, M. E. Differential dynamic microscopy for the characterization of polymer systems. *J. Polym. Sci.* **60**, 1079–1089 (2021).
312. Jacob, J. D. C., He, K., Retterer, S. T., Krishnamoorti, R. *et al.* Diffusive dynamics of nanoparticles in ultra-confined media. *Soft Matter* **11**, 7515–7524 (2015).
313. Jepson, A., Martinez, V. A., Schwarz-Linek, J., Morozov, A. *et al.* Enhanced diffusion of nonswimmers in a three-dimensional bath of motile bacteria. *Phys. Rev. E.* **88**, 041002 (2013).
314. Chen, X., Liu, D., Cai, D., Qiu, J. *et al.* Coaxial differential dynamic microscopy for measurement of Brownian motion in weak optical field. *Opt. Express* **26**, 32083 (2018).
315. Bayles, A. V., Squires, T. M. & Helgeson, M. E. Probe microrheology without particle tracking by differential dynamic microscopy. *Rheol. Acta* **56**, 863–869 (2017).
316. Giavazzi, F., Haro-Pérez, C. & Cerbino, R. Simultaneous characterization of rotational and translational diffusion of optically anisotropic particles by optical microscopy. *J. Phys. Condens. Matter* **28**, 195201 (2016).
317. Dieball, C., Krapf, D., Weiss, M. & Godec, A. Scattering fingerprints of two-state dynamics. *New J. Phys.* **24**, 023004 (2022).
318. He, K., Khorasani, F. B., Retterer, S. T., Thomas, D. K. *et al.* Diffusive dynamics of nanoparticles in arrays of nanoposts. *ACS Nano* **7**, 5122–5130 (2013).
319. Sentjabrskaja, T., Zaccarelli, E., Michele, C. D., Sciortino, F. *et al.* Anomalous dynamics of intruders in a crowded environment of mobile obstacles. *Nat. Commun.* **7** (2016).
320. Gregorin, Ž., Sebastián, N., Osterman, N., Boštjančič, P. H. *et al.* Dynamics of domain formation in a ferromagnetic fluid. *J. Mol. Liq.* **366**, 120308 (2022).
321. Pantoja G. C. F. and Elias, W. S. Green’s functions and method of images: an interdisciplinary topic usually cast aside in physics textbooks. *arXiv* (2020).
322. Sadegh, S. *2D random walk confined in a circular domain with reflective boundaries* 2016. <https://uk.mathworks.com/matlabcentral/fileexchange/60673-2d-random-walk-confined-in-a-circular-domain-with-reflective-boundaries> (2023).
323. Arko, M. & Petelin, A. Cross-differential dynamic microscopy. *Soft Matter* **15**, 2791–2797 (2019).

324. Abramowitz, M. & Stegun, I. A. *Handbook of mathematical functions with formulas, graphs, and mathematical tables* 5th ed., 1046. ISBN: 0-486-61272-4 (US Government printing office, 1964).
325. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* <http://www.deeplearningbook.org> (MIT Press, 2016).
326. Higham, C. F. & Higham, D. J. Deep Learning: An Introduction for Applied Mathematicians. *arXiv* (2018).
327. Aggarwal, C. C. *Neural Networks and Deep Learning : A Textbook. A Textbook* 497. ISBN: 3319944622 (Springer, 2018).
328. Uhlmann, V., Donati, L. & Sage, D. A Practical Guide to Supervised Deep Learning for Bioimage Analysis: Challenges and good practices. *IEEE Signal Process. Mag.* **39**, 73–86 (2022).
329. Brownlee, J. *A Gentle Introduction to the Rectified Linear Unit (ReLU)* (ed Mastery, M. L.) 2019. <https://machinelearningmastery.com/rectified-linear-activation-function-for-deep-learning-neural-networks/>.
330. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* (2014).
331. He, K., Spannuth, M., Conrad, J. C. & Krishnamoorti, R. Diffusive dynamics of nanoparticles in aqueous dispersions. *Soft Matter* **8**, 11933–11938 (2012).
332. Lu, P. J., Giavazzi, F., Angelini, T. E., Zaccarelli, E. *et al.* Characterizing Concentrated, Multiply Scattering, and Actively Driven Fluorescent Systems with Confocal Differential Dynamic Microscopy. *Phys. Rev. Lett.* **108**, 218103 (2012).
333. Leahy, B. D., Jang, W.-D., Yang, H. Y., Struyven, R. *et al.* *Automated Measurements of Key Morphological Features of Human Embryos for IVF in Medical Image Computing and Computer Assisted Intervention – MICCAI 2020* (Springer International Publishing, 2020), 25–35.
334. Hoffman, R. & Gross, L. Modulation Contrast Microscope. *Appl. Opt.* **14**, 1169 (1975).
335. Ostler, T., Langbein, W., Borri, P., Lewis, E. *et al.* Differential Dynamic Microscopy can be applied to Differential Interference Contrast images despite shadowing effects. *arXiv* (2022).
336. Hamilton, S., Regan, D., Payne, L., Langbein, W. *et al.* Sizing individual dielectric nanoparticles with quantitative differential interference contrast microscopy. *Analyst* **147**, 1567–1580 (2022).
337. Jenkins, F. A. & White, H. E. Fundamentals of optics. *Indian J. Phys.* **25**, 265–266 (1957).
338. Allen, R. D. & David, G. B. The Zeiss-Nomarski differential interference equipment for transmitted-light microscopy. *Z. Wiss. Mikrosk.* **69**, 193–221 (1969).

339. Poon, W. C. K. in Terentjev, E. M. & Weitz, D. A. *The Oxford Handbook of Soft Condensed Matter* (eds Terentjev, E. M. & Weitz, D. A.) (Oxford University Press, 2015). ISBN: 9780191044069.
340. Wittmeier, A., Leeth Holterhoff, A., Johnson, J. & Gibbs, J. G. Rotational analysis of spherical, optically anisotropic janus particles by dynamic microscopy. *Langmuir* **31**, 10402–10410 (2015).
341. Shribak, M. & Inoué, S. Orientation-independent differential interference contrast microscopy. *Appl. Opt.* **45**, 460 (2006).
342. Shribak, M. Quantitative orientation-independent differential interference contrast microscope with fast switching shear direction and bias modulation. *J. Opt. Soc. Am. A* **30**, 769 (2013).
343. Shribak, M., Larkin, K. G. & Biggs, D. Mapping optical path length and image enhancement using quantitative orientation-independent differential interference contrast microscopy. *J. Biomed. Opt.* **22**, 016006 (2017).
344. Regan, D., Williams, J., Borri, P. & Langbein, W. Lipid bilayer thickness measured by quantitative DIC reveals phase transitions and effects of substrate hydrophilicity. *Langmuir* **35**, 13805–13814 (2019).
345. De Castro, E. & Morandi, C. J. I. T. Registration of translated and rotated images using finite Fourier transforms. *IEEE Trans. Pattern Anal. Machine Intell.*, 700–703 (1987).
346. Mecholsky, N. A., Akhbarifar, S., Lutze, W., Brandys, M. *et al.* Dataset of Bessel function J_n maxima and minima to 600 orders and 10000 extrema. *Data in Brief* **39**, 107508 (2021).
347. Watson, G. N. *A treatise on the theory of Bessel functions* 804. ISBN: 9780521483919 (Cambridge University Press, 1995).
348. Ostler, T., Woolley, T. E., Swann, K., Thomson, A. *et al.* Vitrifying multiple embryos in different arrangements does not alter the cooling rate. *Cryobiology* **103**, 22–31 (2021).
349. Kitazato Corporation. *Best Practice in Vitrification* <https://www.kitazato-ivf.com/best-practice-in-vitrification/> (2023).
350. Kitazato Corporation. *Cryotop vitrification webpage* 2023. <https://www.kitazato-ivf.com/vitrification/cryotop/> (2023).
351. Davidson, A. F., Benson, J. D. & Higgins, A. Z. Mathematically optimized cryoprotectant equilibration procedures for cryopreservation of human oocytes. *Theor. Biol. Medical Model.* **11** (2014).
352. Cairo Consensus Group. ‘There is only one thing that is truly important in an IVF laboratory: everything’ Cairo Consensus Guidelines on IVF Culture Conditions. *Reprod. Biomed. Online* **40**, 33–60 (2020).
353. Vajta, G. & Kuwayama, M. Improving cryopreservation systems. *Theriogenology* **65**. IETS 2006 Pre- and Post-Conference Symposia, 236–244. ISSN: 0093-691X. <https://www.sciencedirect.com/science/article/pii/S0093691X05004164> (2006).

354. Carslaw, H. S. & Jaeger, J. C. *Conduction of heat in solids* 2nd ed. ISBN: 9780198533689. <https://books.google.co.uk/books?id=y20sAAAAYAAJ> (Clarendon Press, Oxford, 1959).
355. COMSOL Inc., Stockholm, Sweden. *COMSOL Multiphysics 5.5 Reference Manual* https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/COMSOL_ReferenceManual.pdf, last accessed on 14/10/2021. COMSOL Inc. https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/COMSOL_ReferenceManual.pdf (2021).
356. Kleinhans, F., Seki, S. & Mazur, P. Simple, inexpensive attainment and measurement of very high cooling and warming rates. *Cryobiology* **61**, 231–233. ISSN: 0011-2240. <https://www.sciencedirect.com/science/article/pii/S0011224010001094> (2020).
357. Roth, E. *Nucleation and Heat Transfer in Liquid Nitrogen* PhD thesis (Portland State University, 2000).
358. Griffin, J., Emery, B. R., Huang, I., Peterson, C. M. *et al.* Comparative analysis of follicle morphology and oocyte diameter in four mammalian species (mouse, hamster, pig, and human). *Journal of Experimental & Clinical Assisted Reproduction* **3**, 1–9 (2006).
359. Veeck, L. *An atlas of human gametes and conceptuses: an illustrated reference for assisted reproductive technology* (Taylor & Francis, 1999).
360. Halim, B., Lubis, H. P., Novia, D. & Thaharuddin, M. Does oval oocyte have an impact on embryo development in in vitro fertilization? *JBRA Assisted Reproduction* **21** (2017).
361. Mullen, S. F., Li, M., Li, Y., Chen, Z.-J. *et al.* Human oocyte vitrification: the permeability of metaphase II oocytes to water and ethylene glycol and the appliance toward vitrification. *Fertil. Steril.* **89**, 1812–1825 (2008).
362. Boyard, J., Reignier, A., Chtourou, S., Lefebvre, T. *et al.* Should artificial shrinkage be performed prior to blastocyst vitrification? A systematic review of the literature and meta-analysis. *Hum. Fert.* **25**, 1–9 (2020).
363. Lagalla, C., Barberi, M., Orlando, G., Sciajno, R. *et al.* A quantitative approach to blastocyst quality evaluation: morphometric analysis and related IVF outcomes. *J. Assist. Reprod. Genet.* **32**, 705–712 (2015).
364. Baudot, A. & Odagescu, V. Thermal properties of ethylene glycol aqueous solutions. *Cryobiology* **48**, 283–294 (2004).
365. Choi, J. & Bischof, J. C. Review of biomaterial thermal property measurements in the cryogenic regime and their use for prediction of equilibrium and non-equilibrium freezing applications in cryobiology. *Cryobiology* **60**, 52–70 (2010).
366. Amini, M. & Benson, J. D. Technologies for Vitrification Based Cryopreservation. *Bioengineering* **10**, 508. ISSN: 2306-5354 (2023).
367. Bald, W. B. & Fraser, J. Cryogenic surgery. *Rep. Prog. Phys.* **45**, 1381–1434 (1982).
368. Lide, D. R. *CRC Handbook of Chemistry and Physics, 86th Edition* (ed Lide, D. R.) 2544. ISBN: 0849304865 (CRC, 2005).

369. Wang, T., Zhao, G., Tang, H. & Jiang, Z. Determination of convective heat transfer coefficient at the outer surface of a cryovial being plunged into liquid nitrogen. *CryoLetters* **36**, 285–288 (2015).
370. The Engineering Toolbox. *Specific heat of common substances* https://www.engineeringtoolbox.com/specific-heat-capacity-d_391.html (2021).
371. Priddle, H. *Personal communications* Personal Communication. 2021.
372. Kleiner, B. An isoperimetric comparison theorem. *Invent Math* **108**, 37–47. <http://eudml.org/doc/143987> (1992).
373. Adam, N. K. Use of the Term ‘Young’s Equation’ for Contact Angles. *Nature* **180**, 809–810 (1957).
374. Fuller, B. & Paynter, S. Fundamentals of cryobiology in reproductive medicine. *Reprod. Biomed. Online* **9**, 680–691 (2004).
375. Alexiades, V. & Solomon, A. D. *Mathematical modeling of melting and freezing processes* 323. ISBN: 1560321253 (Hemisphere Pub. Corp., 1993).
376. Crank, J. *Free and moving boundary problems* 425. ISBN: 0198533578 (Clarendon Press, 1987).
377. Astarita, G. & Kenny, J. M. The Stefan and Deborah numbers in polymer crystallization. *Chem. Eng. Commun.* **53**, 69–84 (1987).
378. Cash, J. R. in *Encyclopedia of Applied and Computational Mathematics* (ed Engquist, B.) (Springer Berlin Heidelberg, 2015).
379. Ascher, U. M. & Petzold, L. R. *Computer methods for ordinary differential equations and differential-algebraic equations* 314. ISBN: 0898714125 (Society for Industrial and Applied Mathematics, 1998).
380. Süli, E. & Mayers, D. *An introduction to numerical analysis* 433. ISBN: 0-521-00794-1 (Cambridge University Press, 2006).
381. Söderlind, G. *Numerical Algorithms* **1/4**, 281–310 (Springer Science and Business Media LLC, 2002).
382. Wilkinson, J. H. *Rounding errors in algebraic processes* 161. ISBN: 0486679993 (Dover, 1994).
383. Kharangate, C. R. & Mudawar, I. Review of computational studies on boiling and condensation. *International Journal of Heat and Mass Transfer* **108**, 1164–1196 (2017).
384. Benson, J. D. in *Cryopreservation and Freeze-Drying Protocols* (eds Walkers, W. & Oldenhof, H.) 129–172 (Springer US, Totowa, N.J, 1995). ISBN: 978-1-0716-0783-1. https://doi.org/10.1007/978-1-0716-0783-1_4.
385. Roque, M., Lattes, K., Serra, S., Solà, I. *et al.* Fresh embryo transfer versus frozen embryo transfer in in vitro fertilization cycles: a systematic review and meta-analysis. *Fertil. Steril.* **99**, 156–162 (2013).
386. Regan, L. & Rai, R. Epidemiology and the medical causes of miscarriage. *Best Pract. Res. Clin. Obstet. Gynaecol.* **14**, 839–854 (2000).
387. Filho, E. S., Noble, J. A. & Wells, D. A Review on Automatic Analysis of Human Embryo Microscope Images. *Open Biomed. Eng.* **4**, 170–177 (2010).

388. He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. *arXiv* (2015).
389. Yu, X. & Wang, S.-H. Abnormality Diagnosis in Mammograms by Transfer Learning Based on ResNet18. *Fundam. Inform.* **168** (eds Skarbek, W. & Zhang, Y.-D.) 219–230 (2019).
390. Odusami, M., Maskeliūnas, R., Damaševičius, R. & Krilavičius, T. Analysis of Features of Alzheimer’s Disease: Detection of Early Stage from Functional Brain Changes in Magnetic Resonance Images Using a Fine-tuned ResNet18 Network. *Diagnostics* **11**, 1071 (2021).
391. Al-Falluji, R. A., Katheeth, Z. D. & Alathari, B. Automatic Detection of COVID-19 Using Chest X-Ray Images and Modified ResNet18-Based Convolution Neural Networks. *Comput. Mater. Contin.* **66**, 1301–1313 (2021).
392. Megahed, F. M., Chen, Y.-J., Megahed, A., Ong, Y. *et al.* The class imbalance problem. *Nat. Methods* **18**, 1270–1272 (2021).
393. Abeyta, M. & Behr, B. Morphological Assessment of Embryo Viability. *Semin. Reprod. Med.* **32**, 114–126 (2014).
394. Elston, D. M. Survivorship bias. *J. Am. Acad. Dermatol.* (2021).
395. Simopoulou, M., Sfakianoudis, K., Giannelou, P., Rapani, A. *et al.* Discarding IVF embryos: reporting on global practices. *J. Assist. Reprod. Genet.* **36**, 2447–2457 (2019).
396. De Lacey, S. Decisions for the fate of frozen embryos: Fresh insights into patients’ thinking and their rationales for donating or discarding embryos. *Hum. Reprod.* **22**, 1751–1758 (2007).
397. Cerbino, R., Villa, S., Palamidessi, A., Frittoli, E. *et al.* Disentangling collective motion and local rearrangements in 2D and 3D cell assemblies. *Soft Matter* **17**, 3550–3559 (2021).
398. Dupont, A. & Lamb, D. C. Nanoscale three-dimensional single particle tracking. *Nanoscale* **3**, 4532 (2011).
399. Soni, G. V., Ali, B. M. J., Hatwalne, Y. & Shivashankar, G. V. Single Particle Tracking of Correlated Bacterial Dynamics. *Biophys. J.* **84**, 2634–2637 (2003).
400. Elmore, S., Müller, M., Vischer, N., Odijk, T. *et al.* Single-particle tracking of oriC-GFP fluorescent spots during chromosome segregation in *Escherichia coli*. *J. Struct. Biol.* **151**, 275–287 (2005).
401. Elizur, S. E., Lerner-Geva, L., Levron, J., Shulman, A. *et al.* Factors predicting IVF treatment outcome: a multivariate analysis of 5310 cycles. *Reprod. Biomed. Online* **10**, 645–649 (2005).
402. Vittadello, S. T. & Stumpf, M. P. H. Open problems in mathematical biology. *Math. Biosci.* **354**, 108926 (2022).
403. Rittenberg, V., Seshadri, S., Sunkara, S. K., Sobaleva, S. *et al.* Effect of body mass index on IVF treatment outcome: an updated systematic review and meta-analysis. *Reprod. BioMed. Online* **23**, 421–439 (2011).

404. Campbell, A., Fishel, S., Bowman, N., Duffy, S. *et al.* Modelling a risk classification of aneuploidy in human embryos using non-invasive morphokinetics. *Reprod. Biomed. Online* **26**, 477–485 (2013).
405. Fernandez, E. I., Ferreira, A. S., Cecílio, M. H. M., Chéles, D. S. *et al.* Artificial intelligence in the IVF laboratory: overview through the application of different types of algorithms for the classification of reproductive data. *J. Assist. Reprod. Genet.* **37**, 2359–2376 (2020).
406. Apricity. *AI to aid fertility – Apricity Fertility Predictor® launches* <https://www.apricity.life/fertility-hub/ai-aids-fertility-apricity-fertility-predictor> (2023).
407. Concievable Life Sciences. *Concievable Life Sciences* <https://www.conceivable.life/> (2023).
408. Reddy, S., Allan, S., Coghlan, S. & Cooper, P. A governance model for the application of AI in health care. *J. Am. Med. Inform. Assoc.* **27**, 491–497 (2019).
409. Meskó, B. & Topol, E. J. The imperative for regulatory oversight of large language models (or generative AI) in healthcare. *NPJ Digit. Med.* **6** (2023).
410. Rathi, N. *Our emerging regulatory approach to Big Tech and Artificial Intelligence* (ed Financial Conduct Authority) <https://www.fca.org.uk/news/speeches/our-emerging-regulatory-approach-big-tech-and-artificial-intelligence> (2023).
411. GW4. *Fertility: In Vitro, in Silico, In Clinico* <https://gw4.ac.uk/community/fertility-in-vitro-in-silico-in-clinico/> (2023).
412. Wang, Y., Pope, I., Brennan-Craddock, H., Poole, E. *et al.* A primary effect of palmitic acid on mouse oocytes is the disruption of the structure of the endoplasmic reticulum. *Reproduction* **163**, 45–56 (2022).
413. Chenouard, N., Smal, I., de Chaumont, F., Maška, M. *et al.* Objective comparison of particle tracking methods. *Nat. Methods* **11**, 281–289 (2014).
414. Holmes, T. J. & Levy, W. J. Signal-processing characteristics of differential-interference-contrast microscopy. *Appl. Opt.* **26**, 3929 (1987).
415. Preza, C., Snyder, D. L. & Conchello, J.-A. Theoretical development and experimental evaluation of imaging models for differential-interference-contrast microscopy. *J. Opt. Soc. Am. A* **16**, 2185–2199 (1999).
416. Munro, P. R. T. & Török, P. Vectorial, high numerical aperture study of Nomarski’s differential interference contrast microscope. *Opt. Express* **13**, 6833 (2005).