



Data-Driven multi-Contrast spectral microstructure imaging with InSpect: INtegrated SPECTral component estimation and mapping

Paddy J. Slator^{a,*}, Jana Hutter^{b,c}, Razvan V. Marinescu^a, Marco Palombo^a,
Laurence H. Jackson^{b,c}, Alison Ho^d, Lucy C. Chappell^d, Mary Rutherford^b,
Joseph V. Hajnal^{b,c}, Daniel C. Alexander^a

^a Centre for Medical Image Computing, Department of Computer Science, University College London, UK

^b Centre for the Developing Brain, Kings College London, London, UK

^c Biomedical Engineering Department, Kings College London, London, UK

^d Women's Health Department, King's College London, London, UK

ARTICLE INFO

Article history:

Received 31 July 2020

Revised 8 February 2021

Accepted 16 March 2021

Available online 20 April 2021

Keywords:

MRI

Microstructure imaging

Diffusion-relaxation MRI

Inverse Laplace transform

Unsupervised learning

Quantitative MRI

Placenta MRI

ABSTRACT

We introduce and demonstrate an unsupervised machine learning technique for spectroscopic analysis of quantitative MRI experiments. Our algorithm supports estimation of one-dimensional spectra from single-contrast data, and multidimensional correlation spectra from simultaneous multi-contrast data. These spectrum-based approaches allow model-free investigation of tissue properties, but require regularised inversion of a Laplace transform or Fredholm integral, which is an ill-posed calculation. Here we present a method that addresses this limitation in a data-driven way. The algorithm simultaneously estimates a canonical basis of spectral components and voxelwise maps of their weightings, thereby pooling information across whole images to regularise the ill-posed problem. We show in simulations that our algorithm substantially outperforms current voxelwise spectral approaches. We demonstrate the method on multi-contrast diffusion-relaxometry placental MRI scans, revealing anatomically-relevant sub-structures, and identifying dysfunctional placentas. Our algorithm vastly reduces the data required to reliably estimate spectra, opening up the possibility of quantitative MRI spectroscopy in a wide range of new applications. Our InSpect code is available at github.com/paddyslator/inspect.

© 2021 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Quantitative MRI can measure and map physical and chemical quantities that strongly relate to underlying tissue structure and function. A convenient, and data-driven, way to analyse quantitative MRI data is to assume a population of spins with a distribution of quantities (e.g. relaxivity, diffusivity) that are encoded in a one-dimensional, or multidimensional correlation, spectrum. By estimating such distributions, multiple microstructural components can be distinguished without making a-priori modelling assumptions, such as fixing the number of tissue compartments. This approach, which we refer to as quantitative MRI spectroscopy in this paper, has the potential to provide novel biomarkers (Bai et al., 2014). Quantitative MRI spectroscopy has been demonstrated for single-contrast approaches, including T2 component analysis of multi-echo relaxometry data, as has been used to image myelin

(Alonso-Ortiz et al., 2015; Nagtegaal et al., 2020) and luminal water (Sabouri et al., 2017; Devine et al., 2019).

Quantitative MRI spectroscopy is also an attractive technique for analysing MRI experiments that concurrently measure multiple MR properties such as T1, T2, and diffusivity. By providing information on correlations and couplings between complementary MR properties, this approach can resolve distinct microstructural compartments that are indistinguishable with a single contrast. Various types of multi-contrast correlation spectroscopy have been demonstrated, such as relaxometry-relaxometry (English et al., 1991) and diffusion-relaxometry (Van Dusschoten et al., 1996). Several papers have leveraged recent advances in scanner hardware to extend these ideas into imaging, in the T1-diffusion (De Santis et al., 2016b; 2016a), T2-diffusion (Veraart et al., 2018; Kim et al., 2017a; Melbourne et al., 2018; De Almeida Martins and Topgaard, 2018; Lampinen et al., 2020; De Almeida Martins et al., 2020; Reymbaut et al., 2020; Gong et al., 2020), T1-T2-diffusion (Benjamini and Basser, 2017), T2*-diffusion (Slator et al., 2019b), and T1-T2*-diffusion (Hutter et al., 2018) domains.

* Corresponding author.

E-mail address: p.slator@ucl.ac.uk (P.J. Slator).

A general continuum model for quantitative MRI spectroscopy gives a Fredholm integral equation on the MR signal (Benjamini and Basser, 2020), or a Laplace transform in the specific exponential decay case, e.g. for diffusivity or relaxivity. The spectrum can be estimated using regularised inversion of this integral (English et al., 1991), although this is highly ill-posed. Estimating spectra in each image voxel independently therefore requires unrealistically high signal-to-noise ratio (SNR). Moreover, to derive meaningful quantitative maps from voxelwise spectra typically involves a procedure known as *spectral integration*; see top panel Fig. 1 for a two-dimensional example. In spectral integration, the user first manually identifies regions of the spectrum (termed spectral regions of interest, sROI). These sROIs are typically found by identifying prominent features in spectra estimated from the signal averaged over image regions of interest (ROIs). Scalar indices are then calculated by numerical integration of each voxel's reconstructed spectrum over these sROIs (Mackay et al., 1994; Kim et al., 2017a; Benjamini and Basser, 2017).

Recently, methods have been proposed for increasing the robustness of voxelwise spectral fits, utilising marginal distributions (Benjamini and Basser, 2016) or spatial regularisation (Kim et al., 2017a). These methods can improve inversion stability and thus give more meaningful quantitative maps. However, inherent limitations remain. In particular, they rely on ad-hoc choices of regularisation terms and manually defined sROIs. A recent technique automatically identifies these spectral integration regions (Pas et al., 2020), but restricts to rectangular and non-overlapping sROIs, depends on a user-defined threshold value, and still requires voxelwise spectral estimation.

Here we present a method which addresses these limitations in a data-driven way. It simultaneously estimates a canonical basis of spectral components for a whole image (or a data set comprising multiple images), and the voxelwise weighting factors of each component. We build on previous work that exploits data redundancies in nuclear magnetic resonance (NMR) datasets, including (a) COmponent REsolved (CORE) NMR (Stilbs et al., 1996) - a technique for analysing Fourier Transform Pulsed-Gradient Spin-Echo (FT-PGSE) data sets that has also been applied to NMR images (Stilbs, 1998) - and (b) our previously published discrete InSpect algorithm (Slator et al., 2019a). We hence derive a technique that estimates a set of underlying spectra and corresponding maps of their relative contributions in each voxel. This allows us to capture smooth changes in assignment to spectral components across the image, rather than forcing hard categorisation of pixels into a small set of components, which discards subtle variation. Our approach is inspired by the fact that the MR signal comes from a multitude of microstructural environments, motivating the estimation of a set of canonical bases disentangling these environments. Moreover, the spectra and maps outputted by our algorithm are inherently interpretable, unlike many AI - particularly deep learning - techniques (Castelvecchi, 2016; Lundervold and Lundervold, 2019).

Unlike standard inversion approaches, the InSpect algorithm introduced here exploits the large amount of information shared among voxels, dramatically reducing the SNR required for stable inversion, hence enabling quantitative MRI spectroscopy in a wide variety of new situations. The method provides a natural lower-dimensional representation enabling standard downstream analysis of images without manual division of the spectral domain (into sROIs) or image segmentation (into image ROIs). Fig. 1 visualises the existing voxelwise method (top panel) and our proposed InSpect algorithm (bottom panel) side-by-side. In this paper we demonstrate InSpect by estimating multidimensional correlation spectra from combined T2*-diffusion data, but we emphasise that our algorithm is applicable to data of any dimension.

The paper proceeds as follows: we first describe the assumptions underlying the continuous InSpect algorithm, then derive the

iterative optimisation algorithm. We demonstrate InSpect first in simulations, then on in-vivo diffusion-relaxometry placental MRI data. Our Matlab (The MathWorks, Natick, MA) implementation of InSpect is available at github.com/paddyslator/inspect.

2. Methods

2.1. Related methods

InSpect is based on a continuum model, which assumes that single voxels contain spins with a spectrum of MR properties. For a general n -dimensional multi-contrast MRI experiment the voxel signal is

$$S(\mathbf{t}) = \int \dots \int F(\boldsymbol{\omega})K(\mathbf{t}, \boldsymbol{\omega}) d\omega_1 \dots d\omega_n \quad (1)$$

where \mathbf{t} is a vector of experimental parameters which are varied to yield contrast in intrinsic MR properties $\boldsymbol{\omega}$, via the specific form of the kernel $K(\mathbf{t}, \boldsymbol{\omega})$. F is the n -dimensional spectrum over $\boldsymbol{\omega}$, i.e. the distribution of these values across all spins. For example, in two-dimensional T2*-diffusivity (or T2-diffusivity) imaging, $\mathbf{t} = (b, T_E)$, the b-value and echo time (TE); $\boldsymbol{\omega} = (ADC, T_2^*)$, the apparent diffusion coefficient and T2*; $K(\mathbf{t}, \boldsymbol{\omega}) = \exp(-T_E/T_2^*) \exp(-bADC)$ is the kernel; and F is the T2*-ADC spectrum.

The standard approach for estimating the spectrum from a general n -dimensional quantitative MRI experiment, following (Menon and Allen, 1991; English et al., 1991; Ronen et al., 2006), proceeds as follows. Equation (1) is first discretised onto an n -dimensional grid, with lengths defined by the user-defined vector $N_{\boldsymbol{\omega}} = (N_{\omega_1}, \dots, N_{\omega_n})$. This gives the following signal expression

$$S(\mathbf{t}) = \sum_{l_1=1}^{N_{\omega_1}} \dots \sum_{l_n=1}^{N_{\omega_n}} F(\boldsymbol{\omega}(l_1, \dots, l_n))K(\mathbf{t}, \boldsymbol{\omega}(l_1, \dots, l_n)). \quad (2)$$

By choosing an ordering of the elements of the n -dimensional grid coordinates, the signal for all MR encodings in the experiment can thus be written in matrix form

$$\mathbf{S} = \mathbf{K}\mathbf{F} \quad (3)$$

where \mathbf{S} is a column vector, length N_s , of the signals at each MR encoding, \mathbf{K} is an N_s by $\prod_{l=1}^n N_{\omega_l}$ matrix of discretised kernel values, and \mathbf{F} is an $\prod_{l=1}^n N_{\omega_l}$ length column vector of spectrum values. The spectrum F is then calculated with regularised non-negative least squares

$$\mathbf{F} = \underset{\mathbf{F} \geq 0}{\arg \min} \|\mathbf{K}\mathbf{F} - \mathbf{S}\|_2^2 + \alpha \|\mathbf{F}\|_2^2. \quad (4)$$

By solving Equation (4) in each voxel, the spectrum can be estimated across a whole image. Volume fraction maps are then produced by numerically integrating voxelwise spectra over user-defined sROIs, (Kim et al., 2017b; Benjamini and Basser, 2017; Mackay et al., 1994), see top panel Fig. 1. However low SNR can lead to noisy spectrum estimates and hence poor spatial maps.

CORE NMR is an existing method that produces spatial maps by sharing information across voxels and hence increases the effective SNR. CORE assumes a fixed number of components across an image (Stilbs, 1998) (or equivalently an NMR spectra (Stilbs et al., 1996)). In particular, for the two-dimensional combined T2-D imaging application (Stilbs, 1998), assuming M components, the signal for a single MR encoding in voxel n is

$$S_n(b, T_E) = \sum_{m=1}^M z_{nm} \exp(-T_E/T_{2m}) \exp(-bADC_m) \quad (5)$$

where z_{nm} is the weighting of component m in voxel n , and T_{2m} , ADC_m are the single-valued MR properties (i.e. not spectra) for each

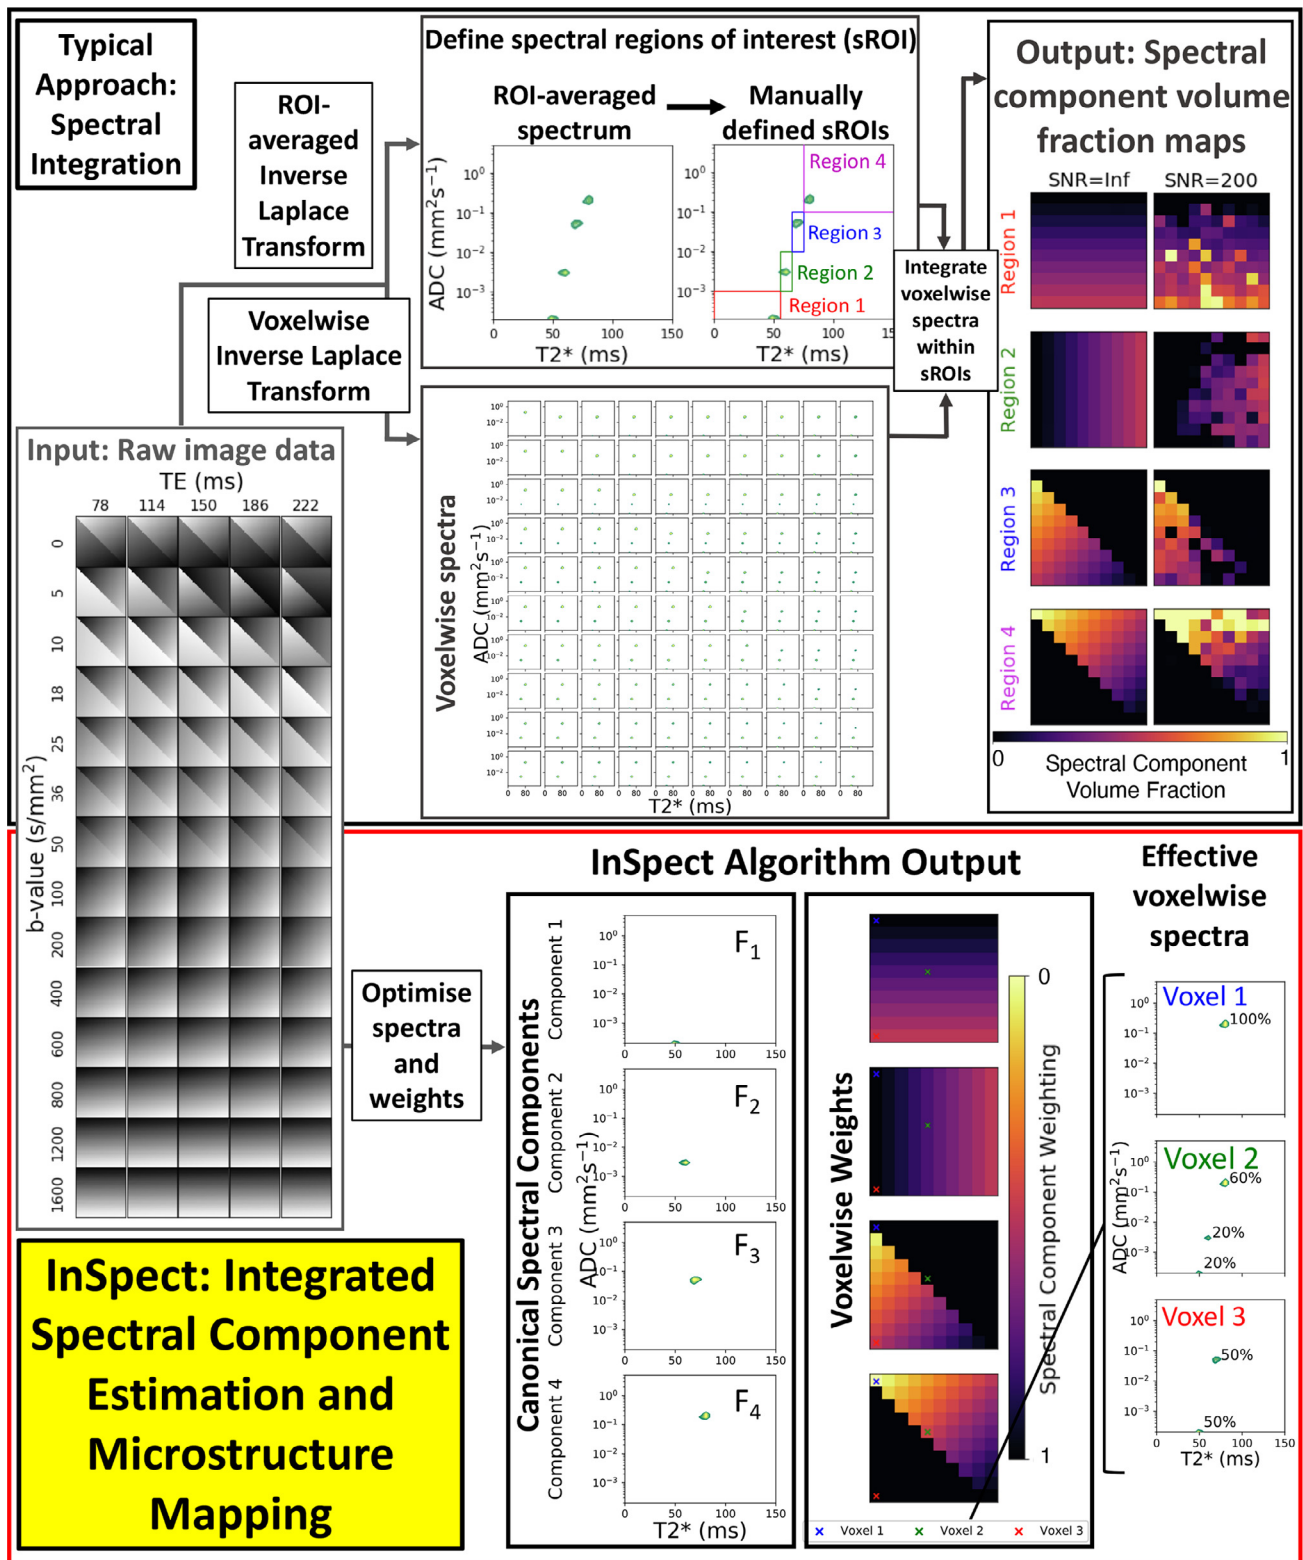


Fig. 1. Schematic comparison of voxelwise integration and InSpect for spectral estimation from combined T2*-diffusion data. Top panel: derivation of volume fraction maps from multi-contrast MRI using spectral integration. In spectral integration, spectral regions of interest (sROI) are manually identified, typically by examining spectra calculated from the signal averaged across image ROIs. Spectra are also calculated for every image voxel. Volume fraction maps are thus calculated by integrating the voxelwise spectra in the sROIs. Bottom panel: Illustration of InSpect applied to multi-contrast MRI. The algorithm simultaneously infers a set of M canonical spectral components (e.g. F₁ to F₄, these examples here have single modes but they can equally be multimodal) and their corresponding voxelwise weightings.

component. Similarly to Equation (3) the signal for all MR encodings in voxel n can be written in vector form as

$$\mathbf{S}_n = \sum_{m=1}^M z_{nm} \exp(-\mathbf{T}_E/T_{2m}) \exp(-\mathbf{b}ADC_m) \quad (6)$$

where \mathbf{T}_E and \mathbf{b} are column vectors, length N_s , of the MR encoding parameters. The CORE method optimises the component weightings for all voxels, $\{z_{n1}, z_{n2}, \dots, z_{nM}\}_{n=1}^N$, and global MR properties, T_{2m} and ADC_m in this case, using least squares fitting.

2.2. InSpect model

Our InSpect algorithm automates spectral mapping, and undertakes data-driven inversion of the Fredholm integral (or Laplace transform). Rather than naively fitting spectra to each voxel independently, the algorithm learns a data-driven low-dimensional representation consistent with the whole image, see bottom panel of Fig. 1 for a visualisation.

The first element of the representation is a pre-specified number, M , of canonical spectral components, $\{F_1, F_2, \dots, F_M\}$, i.e. InSpect infers *spectra* of MR properties, rather than the single values of CORE. These spectral components can consist of a single mode (or *peak*) as in the examples in Fig. 1, or can have multiple modes. Each spectral component has a corresponding voxelwise weight in each N image voxels. Again the weighting of component m in voxel n is denoted z_{nm} , so that the full set of voxelwise weights is

$$\mathbf{z}_n = \{z_{n1}, z_{n2}, \dots, z_{nM}\}_{n=1}^N, \quad \text{where } 0 < z_{nm} < 1. \quad (7)$$

Note that we relax the condition that these spectral component weights must sum to 1 (as used in (Slator et al., 2020)) to allow for overall changes in scale across voxels.

The signal from each voxel, \mathbf{S}_n , is described by the continuum model of Equation (1), with the effective spectrum in each voxel n given by the weighted sum of the canonical spectral components

$$F(\mathbf{z}_n) = \sum_{m=1}^M z_{nm} \mathbf{F}_m \quad (8)$$

where $\mathbf{z}_n = \{z_{nm}\}_{m=1}^M$ are the component weights for voxel n . The discrete model for a single voxel is therefore given by rewriting Equation (3)

$$\mathbf{S}_n = \mathbf{K}F(\mathbf{z}_n) \quad (9)$$

and we consider Gaussian noise on the observed signal

$$\mathbf{S}_n \sim N(\mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I). \quad (10)$$

The log-likelihood across the whole image is therefore

$$\log \pi(\{\mathbf{S}_n\}_{n=1}^N | \mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_M, \mathbf{z}_n, \{\sigma_n^2\}_{n=1}^N) = \sum_{n=1}^N \log \mathcal{N}(\mathbf{S}_n; \mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I) \quad (11)$$

where $\mathcal{N}(x; \mu, \Sigma)$ is the multivariate normal PDF with mean μ and covariance matrix Σ . Note that we assume all observations in a voxel have the same variance, i.e. the covariance matrix is $\sigma_n^2 I$. For simplicity we denote the log-likelihood $\log \pi(\mathbf{D}|\theta)$, where $\mathbf{D} = \{\mathbf{S}_n\}_{n=1}^N$ is the full single- or multi-contrast quantitative MRI dataset (either a single scan or set of scans), and $\theta = \{\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_M, \{z_{n1}, z_{n2}, \dots, z_{nM}\}_{n=1}^N, \{\sigma_n^2\}_{n=1}^N\}$ are the model parameters to be inferred by InSpect.

2.3. InSpect inference algorithm

Here we present the InSpect inference algorithm and provide some intuition behind it. The full derivation is given in Supplementary Material. The inference algorithm seeks the parameters

θ that maximise $\log \pi(\mathbf{D}|\theta)$. In practice, we estimate σ_n^2 directly from the data, e.g. for a T2*-diffusion experiment we estimate by calculating the empirical variance of the volumes with $b=0$ and the lowest TE, and we use the same approach for lower or higher dimensional datasets. The canonical spectral components, \mathbf{F}_m , and voxelwise maps, $\{z_{n1}, z_{n2}, \dots, z_{nM}\}_{n=1}^N$ are iteratively optimised by sequentially maximising as follows

1. Initialise the canonical spectral components $\{F_1, F_2, \dots, F_M\}$, by calculating the mean spectrum across the whole image, then assigning the M biggest separate peaks as the initial spectral components. (Naturally, other initialisation strategies are possible.)
2. Initialise the spectral weights for all voxels $\{z_{nm}\}_{n=1}^N$, given $\{F_1, F_2, \dots, F_M\}$
3. Update \mathbf{F}_m for some m by solving Equation (12) below
4. Update $\{z_{nm}\}_{n=1}^N$, subject to $\sum_{m=1}^M z_{nm} = 1$ by solving Equation (13) for all voxels
5. Repeat steps 3 and 4 until all voxelwise weights have converged according to some tolerance value, or the maximum iteration number is reached.

The update equations are

$$\mathbf{F}_m = \arg \min_{\mathbf{F}_m \geq 0} \left\| \sum_{n=1}^N (\mathbf{K}z_{nm}) \mathbf{F}_m - \sum_{n=1}^N \left(\mathbf{S}_n - \mathbf{K} \left(\sum_{m \neq n} z_{nm} \mathbf{F}_m \right) \right) \right\|_2^2 \quad (12)$$

for the spectral components $m = 1, \dots, M$ and

$$\{z_{nm}\}_{m=1}^M = \arg \max_{0 < z_{nm} < 1} \log \mathcal{N}(\mathbf{S}_n; \mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I) \quad (13)$$

for the weights in voxels $n = 1, \dots, N$. Equations (12) and (13) are solved with non-negative least squares and interior-point algorithms respectively.

2.3.1. Model selection to determine the number of canonical spectral components

Finally, we determine the number of canonical spectral components, M . In this paper, we use the Bayesian information criterion (BIC) to guide the choice of M , with a view to automating this process in future iterations of InSpect. For an InSpect run with fixed M , the BIC is

$$\text{BIC} = k \log n - 2 \log \pi(\mathbf{D}|\hat{\theta}) \quad (14)$$

where $k = (M-1)N + M \prod_{m=1}^M N_{\omega_m}$ is the number of parameters, $n = NN_s$ is the number of observations, and $\log \pi(\mathbf{D}|\hat{\theta})$ is the maximised value of the log likelihood function, i.e. at the last step of the inference algorithm. The value of M which produces the lowest BIC best explains the data, although we emphasise that other considerations, such as interpretability and prior microstructural knowledge, are also important when choosing the number of canonical spectral components.

3. Experiments

3.1. Placenta diffusion-relaxometry data

We demonstrate InSpect using previously published placental T2*-diffusion data (Slator et al., 2019b). Whilst this is two-dimensional multi-contrast data, we again emphasise that InSpect is a general method that is also applicable to single- and higher-dimensional datasets. The protocol has 66 diffusion-weightings (ranging from $b = 5$ to 1600 s mm^{-2} , including six $b = 0$ volumes) and 5 TEs (78, 114, 150, 186, 222 ms) for a total of 330 contrast-encodings. Other acquisition parameters were FOV = $300 \times 320 \times 84 \text{ mm}$, TR = 7 s, SENSE = 2.5, halfscan = 0.6, resolution = 3 mm^3 . We considered 13 scans from 12 women, of whom

Table 1
Participant details.

Participant ID	GA at scan (weeks)	Cohort
1	23.72	Control
2	23.86	Control
3	25.43	Control
4	25.72	Control
5	26.14	Control
6	26.72	Control
7	27.14	Control
8	28.29	Control
9	28.86	Control
10	34.43	CH
11	27.7	PE+FGR
12 (scan 1)	30.71	CH
12 (scan 2)	34.14	CH+PE

GA - gestational age, PE - pre-eclampsia,
CH - chronic hypertensive, FGR - fetal growth restriction

9 were categorised as healthy controls, two had chronic hypertension in pregnancy, and one had pre-eclampsia (PE) with additional fetal growth restriction (FGR). One participant with chronic hypertension was scanned two times, four weeks apart, and developed superimposed pre-eclampsia by the second scan (see Table 1 for full participant details). An ROI comprising the whole placenta and the adjacent section of uterine wall was manually segmented on all images.

There are two main approaches possible when applying InSpect to a dataset with multiple scans: run on each MRI scan independently and find separate canonical spectra and mappings for each; or run on all scans simultaneously and estimate a common set of spectra. The best approach will depend on the specific application. One important consideration is the extent to which one wants to probe within-image heterogeneity, as opposed to across-image differences. In this paper our aim is to gain an initial idea of typical placental T2*-ADC spectra, and their spatial distributions, across healthy and unhealthy tissue. We therefore run InSpect on the data for all patients simultaneously - i.e. fitting to all images jointly, estimating a common set of spectral components across all participants. We ran InSpect, including calculating the BIC, for all choices of the number of components from $M = 2$ to $M = 10$. We terminated the algorithm when voxelwise spectral weight maps had converged to within a tolerance of 10^{-3} , or after five iterations were completed.

3.2. Simulations

We also tested InSpect on simulated diffusion-relaxometry data. Four synthetic canonical spectral components - informed by observed placental spectra (Slator et al., 2019b) - were first defined (first columns, Fig. 2). We next defined ground truth voxelwise weights on a 50-by-50 image (i.e. 2500 voxels total, see third columns in Fig. 2). Given these we simulated diffusion-relaxometry scans as

$$S_n(b, T_E) = \mathbf{K}F(\mathbf{z}_n) = \mathbf{K} \left(\sum_{m=1}^M z_{nm} F_m \right) \quad (15)$$

using the same b-values and TEs as the placental data. We normalised each voxel's data so that the mean of the volumes with $b=0$ and the lowest TE equalled 1. We added Rician noise with SNR from 50 to 400 - this is comparable to the placenta data where we calculated SNRs ranging from 50 to 200. We ran InSpect for all choices of the number of components from $M = 2$ to $M = 8$ on these scans. Equivalently to the placenta case, we terminated the algorithm when the weights had converged to within a tolerance

of 10^{-3} , which typically took around three iterations for the simulated data. We also fit voxelwise T2*-ADC spectra to all scans by solving Equation (4), with α set to 0.01 using the L-curve method (Hansen, 1992). We hence derived spectral volume fraction maps by integrating voxelwise spectra in four regions of T2*-ADC space (i.e. sROIs) as in (Mackay et al., 1994; Kim et al., 2017a; Benjamini and Basser, 2017). We chose these sROIs based on the ground truth components, which we note would not be possible on real data.

4. Results

Figs. 2 and 3 demonstrate InSpect applied to simulated diffusion-relaxometry images, and Figs. 4, 5 and 6 present the simultaneous InSpect run on all participants' placental MR images. Fig. 7 displays how InSpect parameters in the placenta change over gestation.

Figs. 2 and S1 demonstrate that InSpect dramatically outperforms the voxelwise approach on simulated data. At all noise levels InSpect captures the core features of the maps, and more accurately recovers the ground truth than voxelwise maps. There is a modest diagonal artifact in component 1 and 2 InSpect maps, most noticeable at lower SNRs, which arises because of the large difference in component 3 and 4 weightings on either side of this line slightly biasing component 1 and 2 weight estimation. At SNR=400 (Fig. S1) InSpect maps are almost identical to the ground truth, whereas the voxelwise spectral integration maps remain somewhat noisy. InSpect also accurately recovers ground truth spectral components (first columns, Fig. 2), we emphasise that it is not possible to automatically recover these components using the standard voxelwise approach - rather corresponding regions of the spectral domain are manually defined.

Fig. 3 shows that the BIC correctly suggests that choosing $M = 4$ components best explains the simulated data at all noise levels, except for one anomaly at SNR=50 where $M = 6$ has a very slightly lower BIC value.

Fig. 4 shows that InSpect with $M = 7$ components had the lowest BIC value for the placenta scans, demonstrating that this choice is optimal in an information theoretic sense. We emphasise that this is not necessarily the "best" choice - a smaller choice of M may make the InSpect algorithm output more readily interpretable in practice.

Figs. 5 and 6 display the seven-component InSpect run on all participants' placental MR images simultaneously. The first column in Fig. 5 shows the canonical spectral components (which are shared across all participants), and the remaining columns show the corresponding weighting maps. The seven canonical spectral components have distinct characteristics and - even though the algorithm imposes no direct anatomical analogue for any of the components - the corresponding maps identify clear anatomical structures which are consistent across control placentas, and show clear differences in dysfunctional placentas. This suggests that the canonical spectral components are potentially identifying distinct tissue environments, and that those tissue environments could be salient to placental dysfunction.

Given the observed spatial patterns and corresponding canonical spectral component characteristics in placental data, we can make initial speculations about the microstructural and microcirculatory environments associated with each component, although we emphasise that whilst some signal components appear to latch on to specific tissue structures, they cannot be fully interpreted in this way as nothing in the InSpect algorithm guarantees this specific mapping:

- Components one, two and three all consist of a single spectral peak with ADC close to free water. In weighting maps, component one appears more prominent in the uterine wall, com-

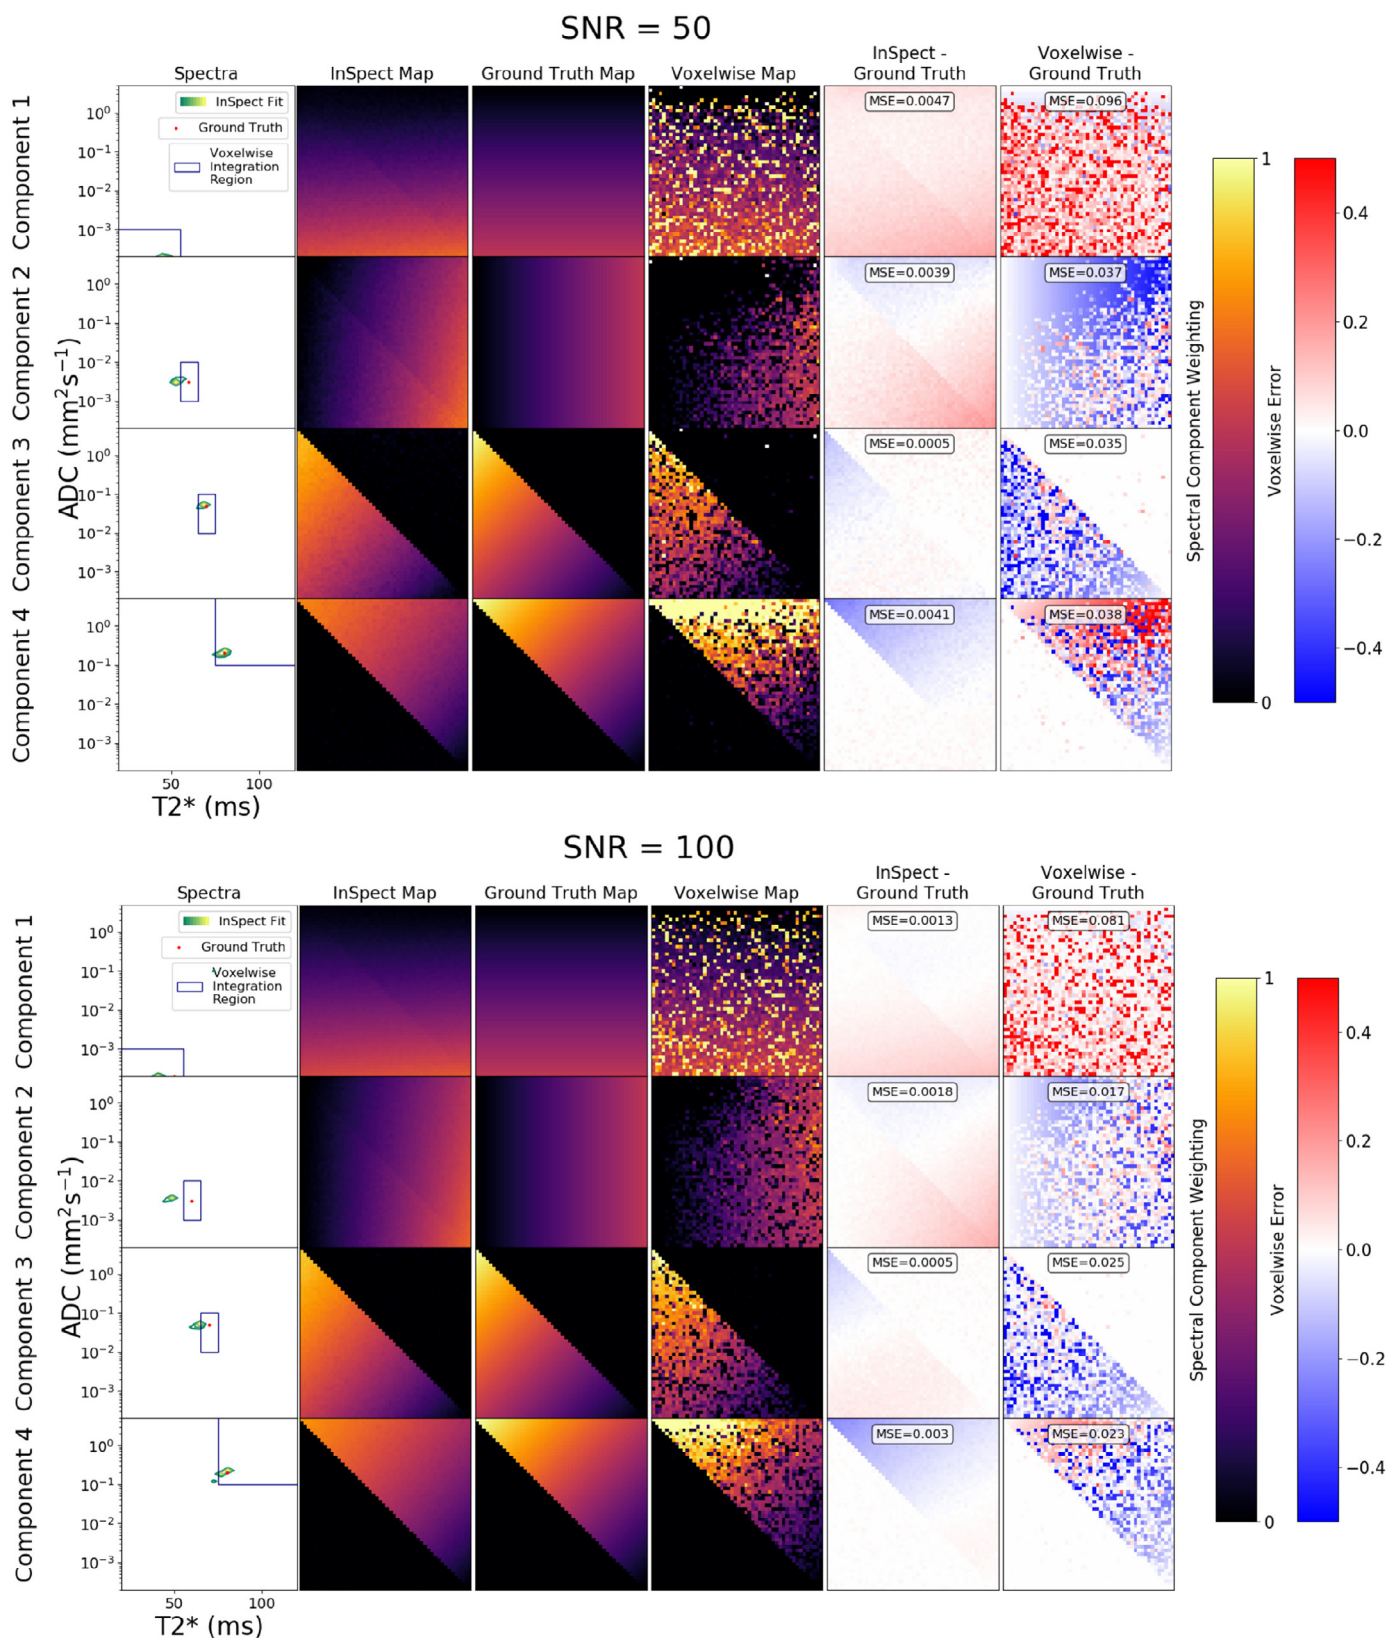


Fig. 2. InSpect applied to simulated images for two noise levels. For each SNR, the InSpect output is presented in columns 1–2: inferred spectral components (green-yellow peaks, column 1), and corresponding voxelwise weights (column 2). Column 3: ground truth voxelwise weights; ground truth spectral components (red dots in column 1) have fixed $T2^*$ -ADC values of (0.05, 0.0002), (0.06, 0.003), (0.07, 0.05) and (0.08 ms, 0.2 mm²/s). Column 4: maps obtained by numerical integration (within blue regions of column 1) of voxelwise spectral fits (we set $\alpha = 0.01$ for these fits, following (Hansen, 1992)). Columns 5 and 6: difference and mean square errors between ground truth and inferred maps. See Fig. S1 for corresponding maps for simulations with higher SNR. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

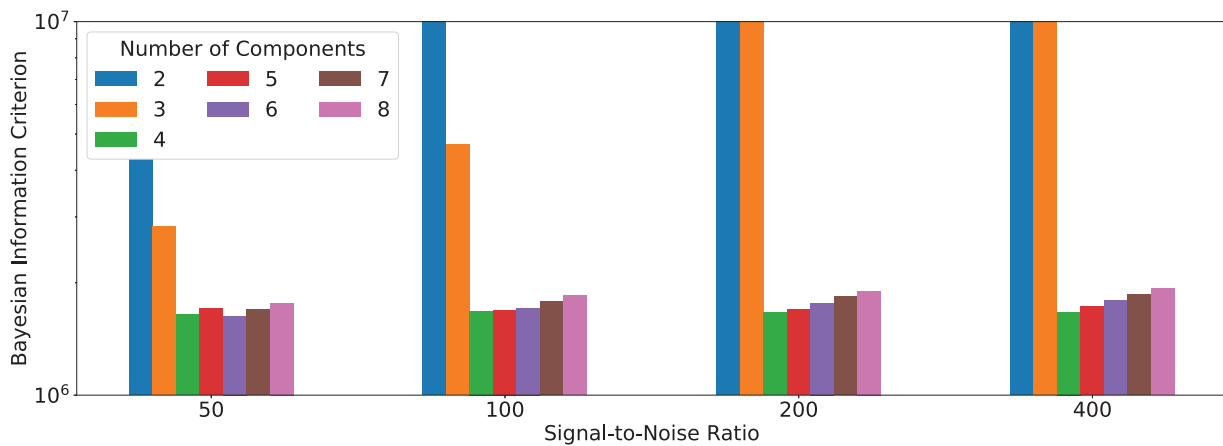


Fig. 3. Bayesian information criterion for InSpect run on simulations for various choices of M , the number of canonical spectral components. For clarity the y-axis is limited to between 10^6 and 10^7 . Simulation parameters as Fig. 2.

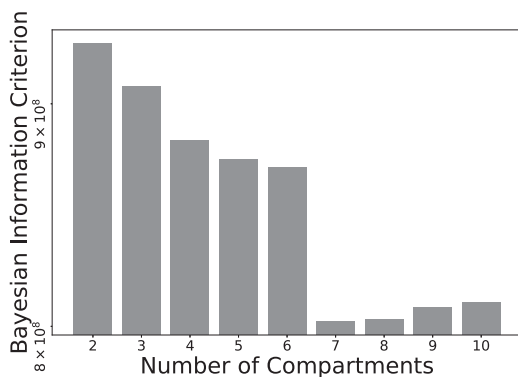


Fig. 4. Bayesian Information Criterion for InSpect run on all placenta scans for various choices of M , the number of canonical spectral components. Fits with lower BIC values explain the data better.

ponent two appears in the areas surrounding placental lobules, and component three appears in the centre of placental lobules. These observations are consistent with these components representing freely diffusing water in the uterine wall, septa, and lobules respectively.

- Component four has a single peak with low ADC and low T2*, and its maps show high weighting in areas of maternal tissue. This is consistent with this component representing water with lower oxygen saturation and diffusion hindered by tissue structures, potentially the maternal venous system.
- Component five has a single peak with low ADC and high T2*, and is spatially prominent in lobular patterns within the placenta. These observations are consistent with this component representing water with high oxygen saturation whose diffusion is hindered, such as oxygenated maternal blood percolating through convoluted intervillous space (e.g. (Dellschaft et al., 2020)).
- Component six has two peaks with ADC much higher than free water, and reasonably high T2*. It is spatially prominent in the uterine wall. This suggests that this component reflects maternal blood perfusing in maternal vasculature.
- Component seven has two peaks with ADC higher than free water, and is spatially prominent in lobular patterns in the placenta. This suggests that this component reflects fetal blood perfusing in fetal vasculature.

To make an initial assessment of InSpect’s ability to quantify and predict changes during development and disease, we plotted

the mean canonical spectral component weightings in the placenta and uterine wall ROI for all participants across gestation (Fig. 7).

Given our speculative assignments of components to tissue environments, the trends we observe are consistent with known placental structure changes across gestation and in disease:

- Component three (“free water in placental lobules”) and component five (“maternal blood in intervillous space”) decrease over gestational age; this is consistent with terminal villi growth shrinking the size of maternal blood pools (Jackson et al., 1992).
- The perfusion-associated components six and seven both decrease over gestational age. This is consistent with the observed decrease in diffusion MRI-derived perfusion fraction over gestational age (Slator et al., 2018b; Anderson et al., 2020).
- Component four (“maternal venous system”) is significantly higher in dysfunctional placentas compared to controls, potentially reflecting lower blood oxygenation (Derwig et al., 2013; Sinding et al., 2016; 2017; 2018; Ingram et al., 2017; Hutter et al., 2019; Ho et al., 2020).
- Component six (“perfusing maternal blood”) is significantly lower in dysfunctional placentas compared to controls, potentially reflecting reduced perfusion (Derwig et al., 2013; Sohlberg et al., 2014; 2015; Anderson et al., 2020).

5. Discussion

We introduce and demonstrate InSpect, a data-driven algorithm for continuous mapping of spectral components in single- or multi-contrast quantitative MRI experiments. The approach addresses key limitations of traditional approaches, which are unstable with standard MRI noise levels and require manual spectral labelling to obtain parametric maps. Specifically, InSpect exploits the redundancy in spectral variation across individual (or groups of) images, and automates mapping of spectral components.

On simulated data we show that InSpect dramatically outperforms the standard voxelwise approach, even when the total number of voxels is relatively small compared to a typical clinical scan (Fig. 2). On placental diffusion-relaxometry MRI data InSpect maps clearly show anatomical structures (Figs. 5 and 6). These maps and the properties of their corresponding spectral components reveal some association of InSpect components with known tissue structures. These associations are consistent with the observed patterns across gestation (Fig. 7), and reveal insights that are consistent with known placenta microstructure and microcirculation properties, in healthy controls and participants with pregnancy complications.

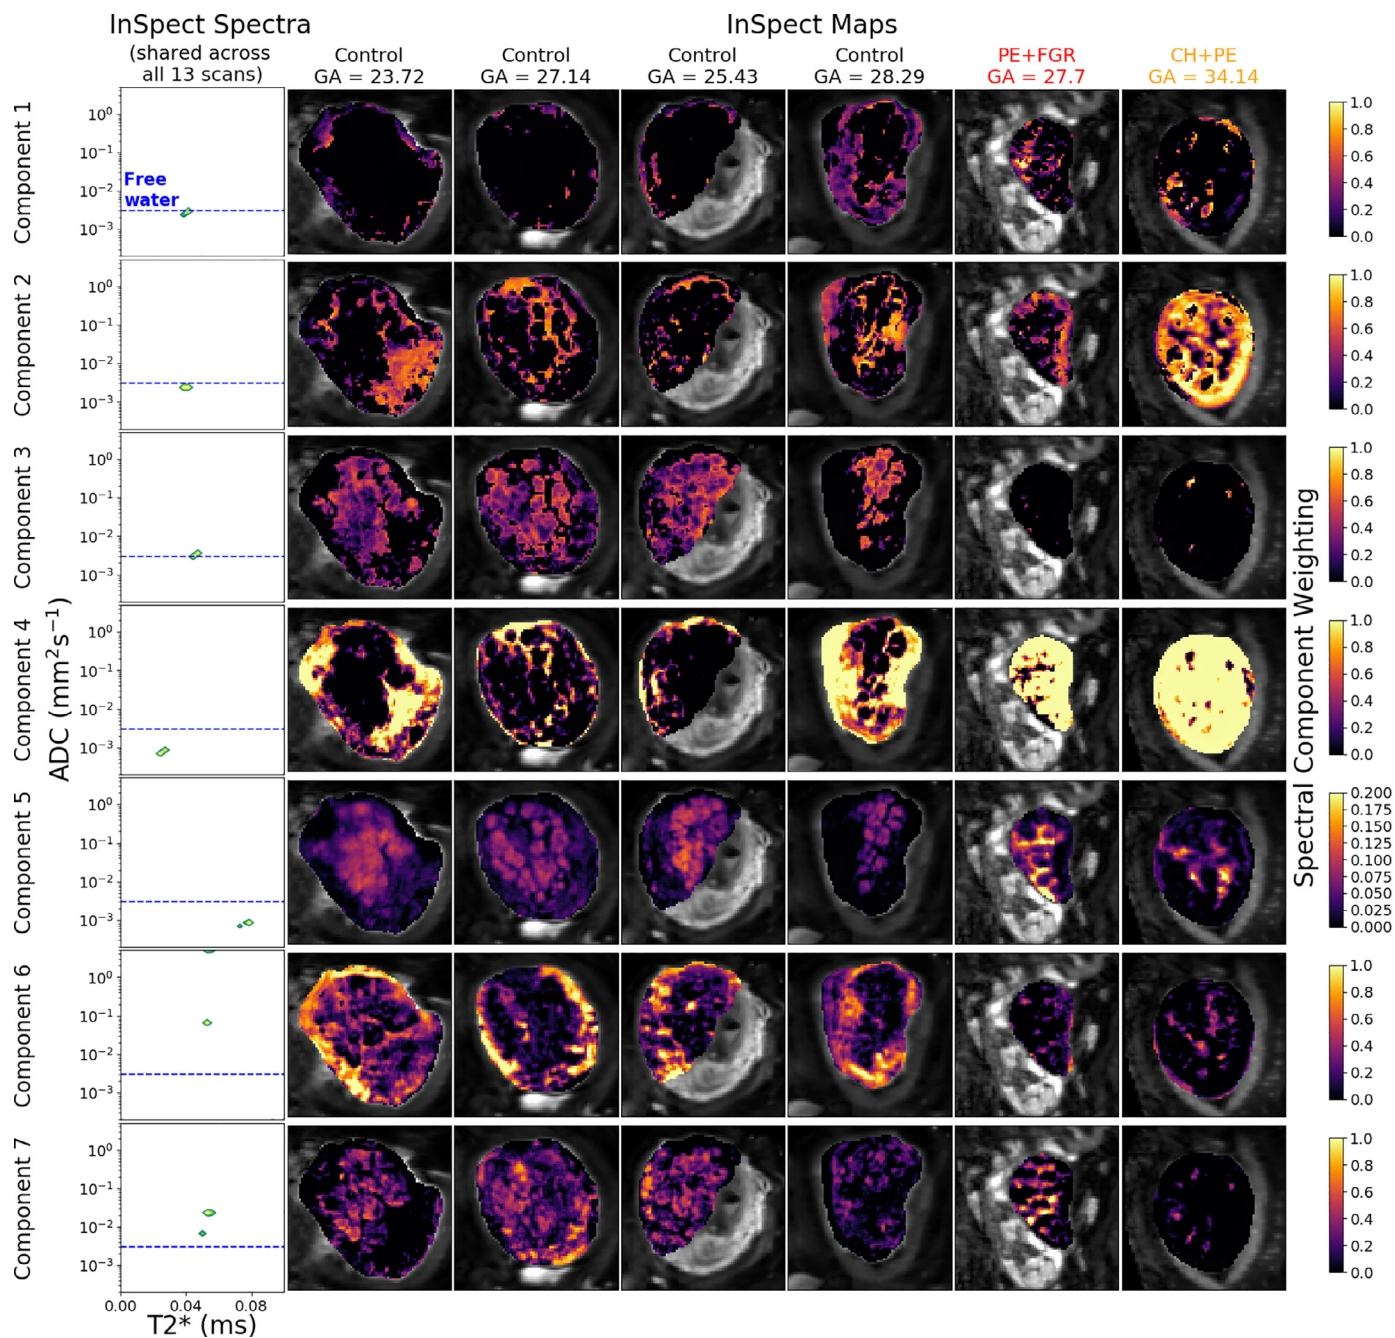


Fig. 5. Seven-component InSpect run on 13 placenta diffusion-relaxometry scans. The leftmost boxes display the canonical spectral components, shared across all 13 participants. The remaining boxes show the corresponding component weighting maps for 6 of the 13 scans, with each column displaying a single participant's maps (see Fig. 6 for the remaining 7 maps). Two participants had been diagnosed with a pregnancy complication at scan time (red outline). Note that the final columns of this Figure and Fig. 6 display maps for the same participant, scanned twice, four weeks apart. See (Slator et al., 2019b) for corresponding maps calculated with a voxelwise spectral approach. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Components one, two, and three have very similar spectral properties - a single peak with ADC close to free water - but maps that are prominent in the uterine wall, septa, and lobules respectively. Compared to component one, component two has a slightly lower ADC, this is consistent with the higher ADC values typically seen in the uterine wall (e.g. (Slator et al., 2019b)). Component three has a slightly higher $T2^*$ than component one, this is consistent with the higher $T2^*$ values seen in the centre of the lobule-like structures (e.g. (Hutter et al., 2019)). These observations suggest that although their spectra appear very similar, these three components could encode subtle microstructural changes affecting the diffusion of free water across the uterine wall, septa, and lobules.

Components six (“perfusing maternal blood”) and seven (“perfusing fetal blood”) both have two peaks with ADC above free water, although component six peaks have markedly higher ADCs than seven. Component six is prominent in the uterine wall, whereas component seven is prominent in the placenta. These observations are consistent with the higher values seen in the uterine wall in ADC maps (Slator et al., 2018a; 2019b; Hutter et al., 2019). The possibility that these maps could separately quantify maternal and fetal circulations, similarly to (Melbourne et al., 2018; Aughwane et al., 2020), has potential applications in identifying placental insufficiency.

Components three (“lobule free water”), four (“maternal venous system”), and six (“perfusing maternal blood”) show the clearest




Fig. 6. InSpect maps for 7 participants from the seven-component InSpect run on 13 placenta diffusion-relaxometry scans. Each row displays maps for a single canonical spectral component - see first column of Fig. 5 for the corresponding spectra. Columns display the maps for a single scan. Note that the final column in Figs. 5 and 6 display maps for the same participant, scanned twice.

differences in dysfunctional placentas, likely indicative of pathology, and are hence promising for predicting and diagnosing pregnancy complications. However, we need to increase participant numbers, particularly by scanning healthy controls across a wider distribution of gestational ages, in order to fully evaluate the predictive and diagnostic power of these metrics.

Whilst it is tempting, as we have speculated above for the placenta, to think of InSpect components as representing distinct “tissue types”, there are many potential inaccuracies and uncertainties. The relationship between the two will likely be disrupted by various effects, and so we emphasise again the need for caution in lieu of any direct validation of these links. We also emphasise that we cannot definitively know the spatial scales of these structures. They could be bulk “tissue types” that contain multiple microstructural compartments, e.g. maternal placenta, fetal placenta; or more

specific “tissue compartments”, e.g. intracellular, extracellular, fetal blood, maternal blood. The scale of the mapped structures will depend on the number of InSpect components - less components will tend to output maps of tissue types consisting of multiple tissue compartments, whereas more components will disentangle these tissue types into single tissue compartments. The long-term hope is that InSpect components are close enough to specific tissue structures that they can reveal pathology with greater specificity than indices that do not attempt to disentangle tissue types.

5.1. Relationship to other methods

InSpect offers a new way of mapping single- or multi-dimensional quantitative MRI data that is complementary to existing methods based on spatial regularisation (Kim et al., 2017a) or

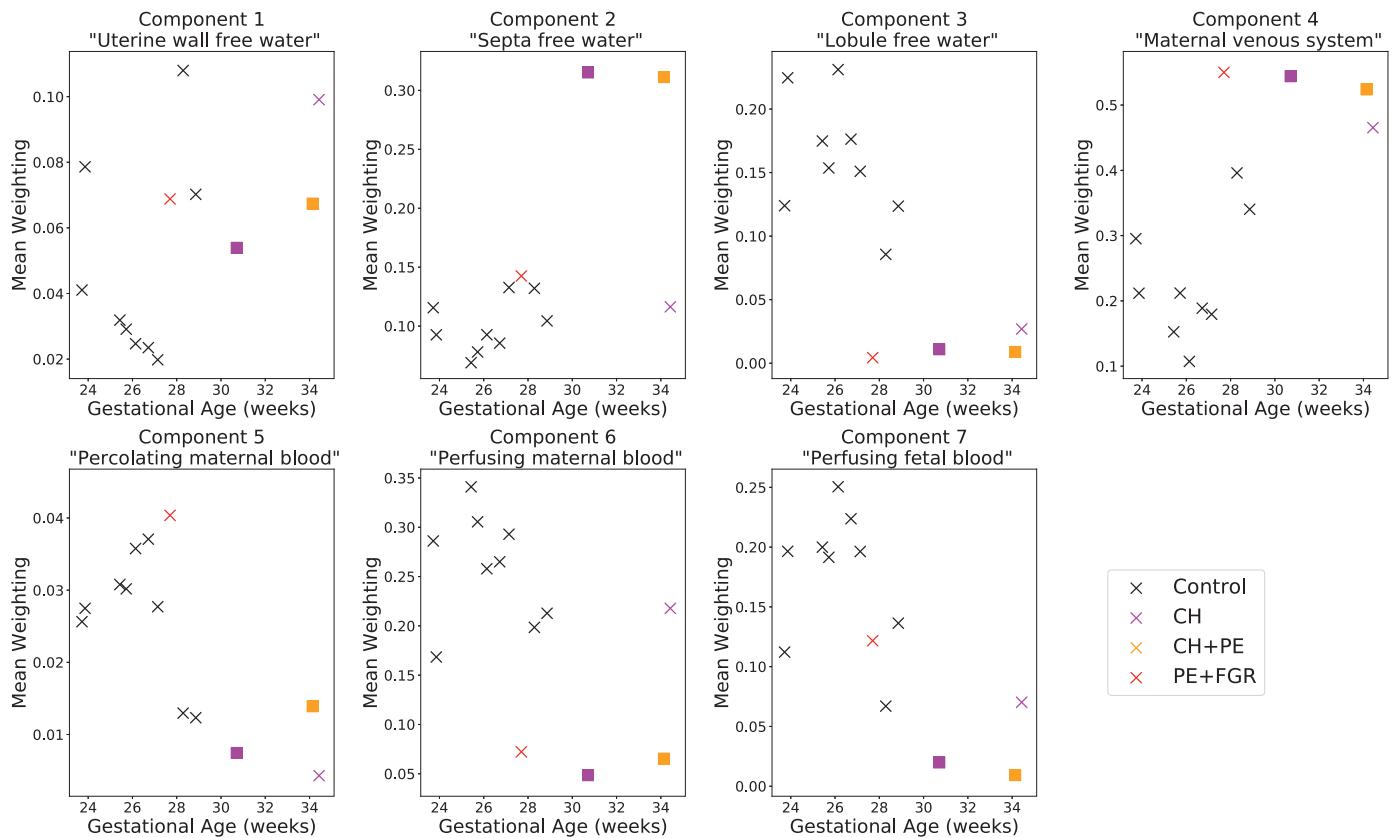


Fig. 7. Scatterplot of mean values of spectral component weightings in placenta and uterine wall ROI against gestational age. Square markers indicate the single participant who was scanned twice.

marginal distributions (Benjamini and Bassler, 2016). Unlike these methods, InSpect does not require manual identification of sROIs to obtain parametric maps, avoiding this problem through data-driven regularisation. CORE NMR (Stilbs et al., 1996; Stilbs, 1998) was a successful first attempt at regularisation of similar inverse problems in a data-driven way. Whilst InSpect and CORE share conceptual similarities, we note the following advances and novelties:

- CORE infers single fixed values for the underlying MR properties of each component, InSpect infers spectra across these properties. This allows InSpect to capture components with more complex diffusion, relaxation, or diffusion-relaxation profiles than a single fixed value, which can reflect more intricate tissue environments.
- InSpect can automatically select the number of spectral components, whereas CORE uses a fixed number of components.
- As opposed to CORE, InSpect uses a general log-likelihood formulation. This framework accommodates more accurate MRI noise models, e.g. Rician, in future.
- The differences in model formulation lead to different optimisation procedures. InSpect uses a modified EM algorithm to infer component spectra (M-step) and their associated weightings (E-step). On the other hand, CORE uses a multi-level least squares optimisation to infer the global MR properties (top level) and component weightings (lower level).

Our method is also comparable to blind source separation (BSS) techniques (e.g. Kim et al. (2015) and Molina-Romero et al. (2018)). The main advantage of InSpect over BSS is that we incorporate a well-defined basic MRI model, allowing us to explicitly reconstruct signal components that we can associate with distinct tissue compartments. However, BSS would be more appropriate when explicit signal models are unknown or inaccurate.

Moreover, InSpect is complementary to image denoising techniques, such as Marchenko-Pastur distribution-based principal component analysis (MP-PCA) (Veraart et al., 2016), which has been shown to improve voxelwise mapping of multi-exponential relaxometry data (Does et al., 2019). This PCA-based denoising technique could be adapted in future to spectral mapping of quantitative MRI data by using an InSpect-like kernel instead of MP.

In future, we will undertake an extensive comparison of all these techniques, whilst carefully considering how to compare methods fairly as this is not a clearly defined standalone task - so the optimal technique will depend on the specific application.

5.2. Limitations and future work

There are limitations to our algorithm in its current form that motivate potential improvements. Firstly, although we demonstrated choosing the number of canonical spectral components by running InSpect with several values of M , then determining the final M with the BIC, this approach may not be optimal in all cases. Running InSpect with fewer components is less computationally expensive and post-processing will be simpler, although it is not immediately clear which approach is better for interpretation. To explore this, we plotted output from all InSpect runs on the placental data with $M = 2$ to $M = 6$ components (Figs. S2 to S6 in Supplementary Material). We note that for InSpect runs with less than six components (i.e. Figs. S2 to S5, and also (Slator et al., 2020) where we used four components) there are not two clearly separated perfusion components. Therefore, for these fits, it is not possible to putatively assign two components to maternal and fetal perfusion, as for components six and seven in the seven-component fit. Moreover, in runs with lower M we typically recover spectral components which have multiple peaks, so likely

consist of multiple tissue microenvironments, whereas for $M = 7$ most components have a single peak, so are more likely to represent a single microstructural environment. In this sense the seven-component run is easier to interpret. However, it may be the case that less than seven components are sufficient to make the necessary diagnostic or prognostic observations from the data - we will test this hypothesis extensively in future. We will also explore other ways of selecting the number of components, such as quantifying the distance between estimated maps (or spectra) with an appropriate image (or distribution) similarity measure then thresholding in order to set an upper limit on M , and model selection via cross validation. Users may also have a fixed number of target microstructural compartments in mind, so it may be beneficial to leave the possibility of manually defining M .

In this study, we show that InSpect spectral components appear to be associated with distinct anatomical features in the placenta. In future, depending on the specific application, we could constrain the spectra to ensure sensitivity and specificity to anatomical features of interest. Another option is to incorporate a Markov random field (similar to (Marinescu et al., 2019)) into the weighting maps to make it more likely that anatomical features are picked up.

Here we ran InSpect on groups of placental MRI scans simultaneously. This is the obvious approach for group studies since the components need to be consistent across images/individuals, but may average over some important within-individual features. In future we will compare to individual scan fits, with focus on which approach best differentiates controls from disease. This choice also impacts on downstream analysis, for example if a dataset is extended, it may not be straightforward to compare InSpect results for the extended dataset with the original dataset (in the group fitting case). There is hence scope for substantial further work to assess the consistency of components and model selections across different instances of datasets - a bootstrap or cross-validation analysis would likely be very informative.

Since it assumes a high level of redundancy across images, InSpect will likely latch onto the strongest spectral components and may miss important components that appear in only a small number of voxels. This could potentially be addressed by running InSpect with a varying number of spectral components - inferring aggregate information from the fits with few components and finer details from the many component fits, or alternatively by optimising a different objective function that strongly emphasises regions of the image that are poorly reconstructed.

Our algorithm assumes that underlying spectral components are fixed. We tested the robustness of InSpect to voxelwise variations in ground truth spectra by running simulations with additional noise. Specifically, we introduced voxelwise normally-distributed perturbations to the ADC and $T2^*$ values of the four synthetic canonical spectral components seen in Figs. 2 and S1. Figs. S7 and S8 display the resulting maps, which are noisier but still recover the salient structures, and show that InSpect still dramatically outperforms the voxelwise approach. In future we will explore more general approaches for increasing InSpect's robustness to voxel-by-voxel noise, such as hierarchical modelling that explicitly accounts for variation in the MR properties of the underlying spectra.

Although we applied to two-dimensional $T2^*$ -diffusion data in this paper, our InSpect formulation is agnostic to the dimensionality. Applying the framework to single-contrast quantitative MRI is particularly attractive, since there are applications in many tissue types and imaging modalities. Our approach could enable significantly improved mappings compared to voxelwise model fitting methods. For example, the framework is immediately applicable to spectral analysis of multi-echo $T2$ relaxometry for myelin water imaging in the brain (Alonso-Ortiz et al., 2015; Nagtegaal et al., 2020), and luminal water imaging in the prostate

(Sabouri et al., 2017; Devine et al., 2019). Whilst we focused on diffusion-relaxometry experiments, InSpect also generalises to other multidimensional spectral techniques, such as the range of "diffusion-X" spectral techniques (Callaghan, 1991), relaxometry-relaxometry (English et al., 1991), and diffusion exchange spectroscopy (Benjamini et al., 2017; Breen-Norris et al., 2020).

InSpect reduces the SNR necessary to estimate and map spectral features using continuum models and quantitative MRI data. The minimum SNR limit for data reconstruction could be further lowered by accounting for Rician noise. The method can hence underpin novel scanning methods for obtaining new data, by leveraging the SNR gain to reduce scanning times and/or improve spatial resolution, ultimately increasing the likelihood of translating quantitative MRI spectroscopy to clinical applications.

6. Conclusion

We present a data-driven approach for quantitative MRI data analysis, and demonstrate its ability to quantify placental dysfunction. InSpect exploits within-image redundancies to simultaneously estimate a set of canonical spectral components and their mapping across images, offering significant advantages over typical single- and multidimensional spectrum estimation methods. The method makes quantitative MRI spectroscopy possible in a wide range of new application areas across tissue types and imaging modalities.

Funding

This work was supported by the NIH Human Placenta Project grant 1U01HD087202-01 (Placenta Imaging Project [PiP]); Wellcome Trust (201374/Z/16/Z); EPSRC (N018702, M020533, EP/N018702/1); NIHR (RP-2014-05-019); the National Institute for Health Research (NIHR) Biomedical Research Centre at University College London Hospitals NHS Foundation Trust and University College London; the Wellcome EPSRC Centre for Medical Engineering at Kings College London (WT 203148/Z/16/Z) and by the NIHR Biomedical Research Centre based at Guys and St Thomas NHS Foundation Trust and Kings College London. PJS and DCA were funded by the European Unions Horizon 2020 research and innovation programme under grant agreement No. 666992. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Paddy J. Slator: Writing - original draft, Writing - review & editing, Conceptualization, Methodology, Software, Formal analysis. **Jana Hutter:** Methodology, Conceptualization, Data curation, Writing - review & editing. **Razvan V. Marinescu:** Conceptualization, Writing - review & editing. **Marco Palombo:** Writing - review & editing, Conceptualization. **Laurence H. Jackson:** Methodology, Data curation. **Alison Ho:** Data curation, Resources. **Lucy C. Chappell:** Data curation, Resources, Funding acquisition. **Mary Rutherford:** Conceptualization, Funding acquisition. **Joseph V. Hajnal:** Conceptualization, Funding acquisition. **Daniel C. Alexander:** Conceptualization, Writing - review & editing, Funding acquisition.

Acknowledgments

We thank all mothers, midwives, obstetricians, and radiographers who played a key role in obtaining the datasets.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.media.2021.102045](https://doi.org/10.1016/j.media.2021.102045).

References

- Alonso-Ortiz, E., Levesque, I.R., Pike, G.B., 2015. MRI-Based myelin water imaging: a technical review. *Magn. Reson. Med.* 73, 70–81. doi:[10.1002/mrm.25198](https://doi.org/10.1002/mrm.25198).
- Andersen, K.B., Hansen, D.N., Haals, C., Sinding, M., Petersen, A., Frøkjær, J.B., Peters, D.A., Sørensen, A., 2020. Placental diffusion-weighted MRI in normal pregnancies and those complicated by placental dysfunction due to vascular malperfusion. *Placenta* 91, 52–58. doi:[10.1016/j.placenta.2020.01.009](https://doi.org/10.1016/j.placenta.2020.01.009).
- Aughwane, R., Mufti, N., Flouri, D., Maksym, K., Spencer, R., Sokolska, M., Kendall, G., Atkinson, D., Bainbridge, A., Deprest, J., Vercauteren, T., Ourselin, S., David, A.L., Melbourne, A., 2020. MRI Measurement of placental perfusion and oxygen saturation in early onset fetal growth restriction. *BJOG: An International Journal of Obstetrics & Gynaecology* 1471–0528.16387. doi:[10.1111/1471-0528.16387](https://doi.org/10.1111/1471-0528.16387).
- Bai, R., Koay, C.G., Hutchinson, E., Basser, P.J., 2014. A framework for accurate determination of the t2distribution from multiple echo magnitude MRI images. *J. Magn. Reson.* 244, 53–63. doi:[10.1016/j.jmr.2014.04.016](https://doi.org/10.1016/j.jmr.2014.04.016).
- Benjamini, D., Basser, P.J., 2016. Use of marginal distributions constrained optimization (MADCO) for accelerated 2d MRI relaxometry and diffusometry. *J. Magn. Reson.* 271, 40–45. doi:[10.1016/j.jmr.2016.08.004](https://doi.org/10.1016/j.jmr.2016.08.004).
- Benjamini, D., Basser, P.J., 2017. Magnetic resonance microdynamic imaging reveals distinct tissue microenvironments. *Neuroimage* 163, 183–196. doi:[10.1016/j.neuroimage.2017.09.033](https://doi.org/10.1016/j.neuroimage.2017.09.033). <http://www.sciencedirect.com/science/article/pii/S1053811917307814>.
- Benjamini, D., Basser, P.J., 2020. Multidimensional correlation MRI. *NMR Biomed* doi:[10.1002/nbm.4226](https://doi.org/10.1002/nbm.4226).
- Benjamini, D., Komlos, M.E., Basser, P.J., 2017. Imaging local diffusive dynamics using diffusion exchange spectroscopy MRI. *Phys. Rev. Lett.* 118, 158003. doi:[10.1103/PhysRevLett.118.158003](https://doi.org/10.1103/PhysRevLett.118.158003).
- Breen-Norris, J.O., Siow, B., Walsh, C., Hipwell, B., Hill, I., Roberts, T., Hall, M.G., Lythgoe, M.F., Ianus, A., Alexander, D.C., Walker-Samuel, S., 2020. Measuring diffusion exchange across the cell membrane with DEXSY (diffusion exchange spectroscopy). *Magn. Reson. Med.* 1–9. doi:[10.1002/mrm.28207](https://doi.org/10.1002/mrm.28207).
- Callaghan, P.T., 1991. Principles of nuclear magnetic resonance microscopy. doi:[10.1016/0730-725x\(82\)90227-2](https://doi.org/10.1016/0730-725x(82)90227-2).
- Castelvecchi, D., 2016. Can we open the black box of AI? *Nature* 538, 20–23. doi:[10.1038/538020a](https://doi.org/10.1038/538020a). <http://www.nature.com/news/can-we-open-the-black-box-of-ai-1.20731>.
- De Almeida Martins, J.P., Tax, C.M., Reymbaut, A., Szczepankiewicz, F., Jones, D.K., Topgaard, D., 2020. Visualizing orientation-specific relaxation-diffusion features mapped onto orientation distribution functions estimated via nonparametric monte carlo MRI signal inversion. *bioRxiv* 1–27.
- De Almeida Martins, J.P., Topgaard, D., 2018. Multidimensional correlation of nuclear relaxation rates and diffusion tensors for model-free investigations of heterogeneous anisotropic porous materials. *Sci. Rep.* 8, 2488. doi:[10.1038/s41598-018-19826-9](https://doi.org/10.1038/s41598-018-19826-9).
- De Santis, S., Assaf, Y., Jeurissen, B., Jones, D.K., Roebroeck, A., 2016. T1 Relaxometry of crossing fibres in the human brain. *Neuroimage* 141, 133–142. doi:[10.1016/j.neuroimage.2016.07.037](https://doi.org/10.1016/j.neuroimage.2016.07.037). <https://www.sciencedirect.com/science/article/pii/S1053811916303445>.
- De Santis, S., Barazany, D., Jones, D.K., Assaf, Y., 2016. Resolving relaxometry and diffusion properties within the same voxel in the presence of crossing fibres by combining inversion recovery and diffusion-weighted acquisitions. *Magn. Reson. Med.* 75, 372–380. doi:[10.1002/mrm.25644](https://doi.org/10.1002/mrm.25644).
- Dellschaft, N.S., Hutchinson, G., Shah, S., Jones, N.W., Bradley, C., Leach, L., Platt, C., Bowtell, R., Gowland, P.A., 2020. The haemodynamics of the human placenta in utero. *PLoS Biol.* 18, e3000676. doi:[10.1371/journal.pbio.3000676](https://doi.org/10.1371/journal.pbio.3000676).
- Derwig, I., Lythgoe, D., Barker, G., Poon, L., Gowland, P., Yeung, R., Zelaya, F., Nicolaidis, K., 2013. Association of placental perfusion, as assessed by magnetic resonance imaging and uterine artery doppler ultrasound, and its relationship to pregnancy outcome. *Placenta* 34, 885–891. doi:[10.1016/j.placenta.2013.07.006](https://doi.org/10.1016/j.placenta.2013.07.006). <http://linkinghub.elsevier.com/retrieve/pii/S0143400413005936>.
- Devine, W., Giganti, F., Johnston, E.W., Sidhu, H.S., Panagiotaki, E., Punwani, S., Alexander, D.C., Atkinson, D., 2019. Simplified luminal water imaging for the detection of prostate cancer from multiple t2 MR images. *J. Magn. Reson. Imaging* 50, 910–917. doi:[10.1002/jmri.26608](https://doi.org/10.1002/jmri.26608).
- Does, M.D., Olesen, J.L., Harkins, K.D., Serradas-Duarte, T., Gochberg, D.F., Jespersen, S.N., Shemesh, N., 2019. Evaluation of principal component analysis image denoising on multi-exponential MRI relaxometry. *Magn. Reson. Med.* 27658. doi:[10.1002/mrm.27658](https://doi.org/10.1002/mrm.27658).
- English, A.E., Whittall, K.P., Joy, M.L., Henkelman, R.M., 1991. Quantitative two-dimensional time correlation relaxometry. *Magn. Reson. Med.* 22, 425–434. doi:[10.1002/mrm.1910220250](https://doi.org/10.1002/mrm.1910220250).
- Gong, T., Tong, Q., He, H., Sun, Y., Zhong, J., Zhang, H., 2020. MTE-NODDI: Multi-TE NODDI for disentangling non-t2-weighted signal fractions from compartment-specific t2 relaxation times. *Neuroimage* 217, 116906. doi:[10.1016/j.neuroimage.2020.116906](https://doi.org/10.1016/j.neuroimage.2020.116906).
- Hansen, P.C., 1992. Analysis of discrete ill-posed problems by means of the l-curve. *SIAM Rev.* 34, 561–580. doi:[10.1137/1034115](https://doi.org/10.1137/1034115), arXiv preprint: [1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- Ho, A.E., Hutter, J., Jackson, L.H., Seed, P.T., McCabe, L., Al-Adnani, M., Marnierides, A., George, S., Story, L., Hajnal, J.V., Rutherford, M.A., Chappell, L.C., 2020. T2* Placental magnetic resonance imaging in preterm preeclampsia: an observational cohort study. *Hypertension* (Dallas, Tex. : 1979) 75, 1523–1531. doi:[10.1161/HYPERTENSIONAHA.120.14701](https://doi.org/10.1161/HYPERTENSIONAHA.120.14701).
- Hutter, J., Slator, P.J., Christiaens, D., Teixeira, R.P.A.G., Roberts, T., Jackson, L., Price, A.N., Malik, S., Hajnal, J.V., 2018. Integrated and efficient diffusion-relaxometry using ZEBRA. *Sci. Rep.* 8, 15138. doi:[10.1038/s41598-018-33463-2arXiv preprint: 1806.03874](https://doi.org/10.1038/s41598-018-33463-2arXiv%20preprint%201806.03874).
- Hutter, J., Slator, P.J., Jackson, L., Gomes, A.D.S., Ho, A., Story, L., O'Muirheartaigh, J., Teixeira, R.P.A.G., Chappell, L.C., Alexander, D.C., Rutherford, M.A., Hajnal, J.V., 2019. Multimodal functional MRI to explore placental function over gestation. *Magn. Reson. Med.* 81, 1191–1204. doi:[10.1002/mrm.27447](https://doi.org/10.1002/mrm.27447).
- Ingram, E., Morris, D., Naish, J., Myers, J., Johnstone, E., 2017. MR Imaging measurements of altered placental oxygenation in pregnancies complicated by fetal growth restriction. *Radiology* 285, 953–960. doi:[10.1148/radiol.2017162385](https://doi.org/10.1148/radiol.2017162385).
- Jackson, M.R., Mayhew, T.M., Boyd, P.A., 1992. Quantitative description of the elaboration and maturation of villi from 10 weeks of gestation to term. *Placenta* 13, 357–370. doi:[10.1016/0143-4004\(92\)90060-7](https://doi.org/10.1016/0143-4004(92)90060-7).
- Kim, D., Doyle, E.K., Wisnowski, J.L., Kim, J.H., Haldar, J.P., 2017. Diffusion-relaxation correlation spectroscopic imaging: a multidimensional approach for probing microstructure. *Magn. Reson. Med.* 78, 2236–2249. doi:[10.1002/mrm.26629](https://doi.org/10.1002/mrm.26629).
- Kim, D., Kim, J.H., Haldar, J.P., 2015. Automatic Tissue Decomposition Using Non-negative Matrix Factorization for Noisy MR Magnitude Images. In: *ISMRM 2015*, p. 3701.
- Kim, D., Wisnowski, J.L., Haldar, J.P., 2017. Improved Efficiency for Microstructure Imaging Using High-dimensional MR Correlation Spectroscopic Imaging. In: *2017 51st Asilomar Conference on Signals, Systems, and Computers*. IEEE, pp. 1264–1268. doi:[10.1109/ACSSC.2017.8335555](https://doi.org/10.1109/ACSSC.2017.8335555).
- Lampinen, B., Szczepankiewicz, F., Mårtensson, J., van Westen, D., Hansson, O., Westin, C.F., Nilsson, M., 2020. Towards unconstrained compartment modeling in white matter using diffusion-relaxation MRI with tensor-valued diffusion encoding. *Magn. Reson. Med.* 2020. doi:[10.1002/mrm.28216](https://doi.org/10.1002/mrm.28216).
- Lundervold, A.S., Lundervold, A., 2019. An overview of deep learning in medical imaging focusing on MRI. *Z Med Phys* 29, 102–127. doi:[10.1016/j.zemedi.2018.11.002](https://doi.org/10.1016/j.zemedi.2018.11.002), arXiv preprint: [1811.10052](https://arxiv.org/abs/1811.10052).
- Mackay, A., Whittall, K., Adler, J., Li, D., Paty, D., Graeb, D., 1994. In vivo visualization of myelin water in brain by magnetic resonance. *Magn. Reson. Med.* 31, 673–677. doi:[10.1002/mrm.1910310614](https://doi.org/10.1002/mrm.1910310614).
- Marinescu, R.V., Eshaghi, A., Lorenzi, M., Young, A.L., Oxtoby, N.P., Garbarino, S., Crutch, S.J., Alexander, D.C., 2019. DIVE: A spatiotemporal progression model of brain pathology in neurodegenerative disorders. *Neuroimage* 192, 166–177. doi:[10.1016/j.neuroimage.2019.02.053](https://doi.org/10.1016/j.neuroimage.2019.02.053), arXiv preprint: [1901.03553](https://arxiv.org/abs/1901.03553).
- Melbourne, A., Aughwane, R., Sokolska, M., Owen, D., Kendall, G., Flouri, D., Bainbridge, A., Atkinson, D., Deprest, J., Vercauteren, T., David, A., Ourselin, S., 2018. Separating fetal and maternal placenta circulations using multiparametric MRI. *10.1002/mrm.27406*
- Menon, R.S., Allen, P.S., 1991. Application of continuous relaxation time distributions to the fitting of data from model systems and excised tissue. *Magn. Reson. Med.* 20, 214–227. doi:[10.1002/mrm.1910200205](https://doi.org/10.1002/mrm.1910200205).
- Molina-Romero, M., Gómez, P.A., Sperl, J.I., Czisch, M., Sämann, P.G., Jones, D.K., Menzel, M.L., Menze, B.H., 2018. A diffusion model-free framework with echo time dependence for free-water elimination and brain tissue microstructure characterization. *Magn. Reson. Med.* 80, 2155–2172. doi:[10.1002/mrm.27181](https://doi.org/10.1002/mrm.27181).
- Nagtegaal, M., Koken, P., Amthor, T., Bresser, J.D., Vos, F., Doneva, M., 2020. Myelin water imaging from multi-echo t2 MR relaxometry data using a joint sparsity constraint. *Neuroimage* 117014. doi:[10.1016/j.neuroimage.2020.117014](https://doi.org/10.1016/j.neuroimage.2020.117014).
- Pas, K., Komlos, M.E., Perl, D.P., Basser, P.J., Benjamini, D., 2020. Retaining information from multidimensional correlation MRI using a spectral regions of interest generator. *Sci. Rep.* 1–10. doi:[10.1038/s41598-020-60092-5](https://doi.org/10.1038/s41598-020-60092-5).
- Reymbaut, A., Martins, J.P.d.A., Tax, C.M.W., Szczepankiewicz, F., Jones, D., Topgaard, D., 2020. Resolving orientation-specific diffusion-relaxation features via monte-carlo density-peak clustering in heterogeneous brain tissue. 45–48, arXiv preprint: [2004.08626](https://arxiv.org/abs/2004.08626).
- Ronen, I., Moeller, S., Ugurbil, K., Kim, D.S., 2006. Analysis of the distribution of diffusion coefficients in cat brain at 9.4 t using the inverse laplace transformation. *Magn. Reson. Imaging* 24, 61–68. doi:[10.1016/j.mri.2005.10.023](https://doi.org/10.1016/j.mri.2005.10.023). <https://www.sciencedirect.com/science/article/pii/S0730725X05003279?via%3Dihub> <http://linkinghub.elsevier.com/retrieve/pii/S0730725X05003279>.
- Sabouri, S., Chang, S.D., Savdie, R., Zhang, J., Jones, E.C., Goldenberg, S.L., Black, P.C., Kozlowski, P., 2017. Luminal water imaging: a new MR imaging t2 mapping technique for prostate cancer diagnosis. *Radiology* 284, 451–459. doi:[10.1148/radiol.2017161687](https://doi.org/10.1148/radiol.2017161687).
- Sinding, M., Peters, D.A., Frøkjær, J.B., Christiansen, O.B., Petersen, A., Uldbjerg, N., Sørensen, A., 2016. Placental magnetic resonance imaging t2* measurements in normal pregnancies and in those complicated by fetal growth restriction. *Ultrasound in Obstetrics & Gynecology* 47, 748–754. doi:[10.1002/uo.14917](https://doi.org/10.1002/uo.14917).
- Sinding, M., Peters, D.A., Frøkjær, J.B., Christiansen, O.B., Petersen, A., Uldbjerg, N., Sørensen, A., 2017. Prediction of low birth weight: comparison of placental t2* estimated by MRI and uterine artery pulsatility index. *Placenta* 49, 48–54. doi:[10.1016/j.placenta.2016.11.009](https://doi.org/10.1016/j.placenta.2016.11.009). <http://www.ncbi.nlm.nih.gov/pubmed/28012454>.
- Sinding, M., Peters, D.A., Poulsen, S.S., Frøkjær, J.B., Christiansen, O.B., Petersen, A., Uldbjerg, N., Sørensen, A., 2018. Placental baseline conditions modulate the hyperoxic BOLD-MRI response. *Placenta* 61, 17–23. doi:[10.1016/j.placenta.2017.11.002](https://doi.org/10.1016/j.placenta.2017.11.002). <https://www.sciencedirect.com/science/article/pii/S0143400417311724?via%3Dihub>.
- Slator, P.J., Hutter, J., Marinescu, R.V., Palombo, M., Jackson, L.H., Ho, A., Chappell, L.C.,

- Rutherford, M., Hajnal, J.V., Alexander, D.C., 2020. Data-driven multi-contrast spectral microstructure imaging with inspect. MICCAI 2020 12266 LNCS. doi:[10.1007/978-3-030-59725-2_36](https://doi.org/10.1007/978-3-030-59725-2_36).
- Slator, P.J., Hutter, J., Marinescu, R.V., Palombo, M., Young, A.L., Jackson, L.H., Ho, A., Chappell, L.C., Rutherford, M., Hajnal, J.V., Alexander, D.C., 2019. Inspect: Integrated SPECTral Component Estimation and Mapping for Multi-contrast Microstructural MRI. In: IPMI 2019. Springer, Cham, pp. 755–766. doi:[10.1007/978-3-030-20351-1_59](https://doi.org/10.1007/978-3-030-20351-1_59).
- Slator, P.J., Hutter, J., McCabe, L., Gomes, A.D.S., Price, A.N., Panagiotaki, E., Rutherford, M.A., Hajnal, J.V., Alexander, D.C., 2018. Placenta microstructure and microcirculation imaging with diffusion MRI. *Magn. Reson. Med.* 80, 756–766. doi:[10.1002/mrm.27036](https://doi.org/10.1002/mrm.27036).
- Slator, P.J., Hutter, J., Palombo, M., Jackson, L.H., Ho, A., Panagiotaki, E., Chappell, L.C., Rutherford, M.A., Hajnal, J.V., Alexander, D.C., 2019b. Combined diffusion-relaxometry MRI to identify dysfunction in the human placenta. *Magn. Reson. Med.*, 1–22. doi:[10.1002/mrm.27733](https://doi.org/10.1002/mrm.27733), arXiv preprint: [1810.04156](https://arxiv.org/abs/1810.04156).
- Slator, P.J., Hutter, J., Panagiotaki, E., Rutherford, M.A., Hajnal, J.V., Alexander, D.C., Hutter, J., Panagiotaki, E., Rutherford, M.A., Hajnal, J.V., Alexander, D.C., 2018. IVIM MRI of the Placenta. In: Le Bihan, D., Iima, M., Federau, C., Sigmund, E.E. (Eds.), *Intravoxel Incoherent Motion (IVIM) MRI*. Pan Stanford, pp. 345–366. doi:[10.1201/9780429427275-16](https://doi.org/10.1201/9780429427275-16).
- Sohlberg, S., Mulic-Lutvica, A., Lindgren, P., Ortiz-Nieto, F., Wikström, A.K., Wikström, J., 2014. Placental perfusion in normal pregnancy and early and late preeclampsia: a magnetic resonance imaging study. *Placenta* 35, 202–206. doi:[10.1016/j.placenta.2014.01.008](https://doi.org/10.1016/j.placenta.2014.01.008).
- Sohlberg, S., Mulic-Lutvica, A., Olovsson, M., Weis, J., Axelsson, O., Wikström, J., Wikström, A.K., 2015. Magnetic resonance imaging-estimated placental perfusion in fetal growth assessment. *Ultrasound in Obstetrics & Gynecology* 46, 700–705. doi:[10.1002/uog.14786](https://doi.org/10.1002/uog.14786).
- Stilbs, P., 1998. Component separation in NMR imaging and multidimensional spectroscopy through global least-squares analysis, based on prior knowledge. doi:[10.1006/jmre.1998.1566](https://doi.org/10.1006/jmre.1998.1566).
- Stilbs, P., Paulsen, K., Griffiths, P.C., 1996. Global least-squares analysis of large, correlated spectral data sets: application to component-resolved FT-PGSE NMR spectroscopy. *J. Phys. Chem.* 100, 8180–8189. doi:[10.1021/jp9535607](https://doi.org/10.1021/jp9535607). <https://pubs.acs.org/sharingguidelines>.
- Van Dusschoten, D., Moonen, C.T., Adrie De Jager, P., Van As, H., 1996. Unraveling diffusion constants in biological tissue by combining carr-purcell-meiboom-gill imaging and pulsed field gradient NMR. *Magn. Reson. Med.* 36, 907–913. doi:[10.1002/mrm.1910360613](https://doi.org/10.1002/mrm.1910360613).
- Veraart, J., Novikov, D.S., Christiaens, D., Ades-aron, B., Sijbers, J., Fieremans, E., 2016. Denoising of diffusion MRI using random matrix theory. *Neuroimage* 142, 394–406. doi:[10.1016/j.neuroimage.2016.08.016](https://doi.org/10.1016/j.neuroimage.2016.08.016), arXiv preprint: [15334406](https://arxiv.org/abs/15334406).
- Veraart, J., Novikov, D.S., Fieremans, E., 2018. TE Dependent diffusion imaging (TEdDI) distinguishes between compartmental t2 relaxation times. *Neuroimage* 182, 360–369. doi:[10.1016/j.neuroimage.2017.09.030](https://doi.org/10.1016/j.neuroimage.2017.09.030). <http://linkinghub.elsevier.com/retrieve/pii/S1053811917307784>.