Automated Extraction of Low Surface Brightness Sources

&

The Physical Nature of Ultra-Diffuse Galaxies

by

Daniel James Prole

A THESIS SUBMITTED TO CARDIFF UNIVERSITY FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

September, 2019

Acknowledgements

I am extraordinarily grateful to my primary supervisor, Jonathan Davies, who has supported me throughout both my PhD and undergraduate studies and has provided me with countless opportunities that have greatly enriched my time as a student.

I would also like to sincerely thank Michael Hilker, who supported my application to the European Southern Observatory (ESO) Studentship Programme and was a fantastic and supremely knowledgable supervisor throughout my short time in Garching.

In addition, I am exceptionally grateful to my co-supervisor and fellow mentor at ESO, Remco van der Burg, who I have learned a great deal from in a very short space of time and whose grasp of the field has inspired a considerable chunk of my thesis.

I am immensely grateful for the funding I have received from Cardiff University that has enabled this work; in particular to Haley Gomez and Jonathan Davies. I additionally acknowledge that this work has been partially funded by ESO as part of the ESO Studentship Programme.

Chapter 2 is based upon work published in Prole et al. (2018). I would like to thank the co-authors, Jonathan Davies, Olivia Keenan and Luke Davies for their valuable advise and suggestions that contributed to this chapter. I would like to sincerely thank the anonymous reviewer of the paper for the significant contribution they made to the quality of the work. I am also grateful for additional comments provided by Michael Hilker.

Chapter 3 is based upon work published in Prole et al. (2019a). I would like to thank the co-authors, Remco van der Burg, Michael Hilker and Jonathan Davies for their valuable advise and suggestions that contributed to this chapter. I am additionally thankful to Javier Román, Alister Graham, Ivan Baldry and Benne Holwerda for providing insightful comments and suggestions regarding this chapter, and to Javier Román, Yu Rong and Michael G. Jones (and collaborators) for kindly sharing their data. I am also thankful to Caroline Haigh and the MTObjects team for providing their software.

Chapter 4 is based upon work published in Prole et al. (2019b). I would like to thank the co-authors, Michael Hilker, Remco van der Burg, Michele Cantiello, Aku Venhola, Enrica Iodice, Glenn van de Ven, Carolin Wittman, Reynier Peletier, Steffen Mieske, Massimo Capaccioli, Nicola Napolitano, Maurizio Paolillo, Marilena Spavone and Edwin Valentijn for their valuable advise and suggestions that contributed to this chapter. I am also very grateful for the comments provided by Adriano Agnello, Arianna Di Cintio and Michael Beasley regarding this chapter.

Parts of this thesis are based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA/DAPNIA, at the Canada-France-Hawaii Telescope (CFHT) which is operated by the National Research Council (NRC) of Canada, the Institut National des Science de l'Univers of the Centre National de la Recherche Scientifique (CNRS) of France, and the University of Hawaii.

Based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs 177.A-3016, 177.A-3017 and 177.A-3018, and on data products produced by Target/OmegaCEN, INAF-OACN, INAF-OAPD and the KiDS production team, on behalf of the KiDS consortium. OmegaCEN and the KiDS production team acknowledge support by NOVA and NWO-M grants. Members of INAF-OAPD and INAF-OACN also acknowledge the support from the Department of Physics & Astronomy of the University of Padova, and of the Department of Physics of Univ. Federico II (Naples).

This work is based on visitor mode observations collected at the European Organisation for Astronomical Research in the Southern Hemisphere under the following VST GTO programs: 094.B-0512(B), 094.B- 0496(A), 096.B- 0501(B), 096.B-0582(A).

GAMA is a joint European-Australasian project based around a spectroscopic campaign using the Anglo-Australian Telescope. The GAMA input catalogue is based on data taken from the Sloan Digital Sky Survey and the UKIRT Infrared Deep Sky Survey. Complementary imaging of the GAMA regions is being obtained by a number of independent survey programmes including GALEX MIS, VST KiDS, VISTA VIKING, WISE, Herschel-ATLAS, GMRT and ASKAP providing UV to radio coverage. GAMA is funded by the STFC (UK), the ARC (Australia), the AAO, and the participating institutions. The GAMA website is http://www.gama-survey.org/.

This research made use of Astropy,[1] a community-developed core Python package for Astronomy (Astropy Collaboration et al., 2013, 2018).

This work was supported by the Advanced Research Computing@Cardiff (ARCCA) Division, Cardiff University.

This research has made use of the VizieR catalogue access tool, CDS, Strasbourg, France (DOI : 10.26093/cds/vizier). The original description of the VizieR service was published in A&AS 143, 23.

This research has made use of NASA's Astrophysics Data System.

I acknowledge support from the European Research Council (ERC) in the form of Consolidator Grant CosmicDust (PI H L Gomez).

[1] http://www.astropy.org

Abstract

As the size and depth of astronomical data improves, so must the software with which we use to extract sources from it. In this thesis, I describe a new software package, DeepScan, which I have written & designed to offer improvements over contemporary software packages such as SExtractor for detecting extended low surface brightness (LSB) features, particularly to avoid issues of excessive segment fragmentation. I demonstrate the technique by applying it over a 5 degree^2 portion of the Next-Generation Virgo Survey data to reveal 53 LSB galaxies that are candidate cluster members based on their sizes and colours; 30 of which are new detections. The final sample have extremely low surface brightness ($26.0 \leq \bar{\mu}_e \leq 28.5$) and low stellar masses ($10^{6.3\pm0.5} M_\odot$), making them some of the faintest known dwarf galaxies in the Virgo Cluster.

Improved detection and measurement techniques are crucial for identifying large samples of LSB galaxies, which are important for characterising the formation mechanisms of ultra-diffuse galaxies (UDGs): large, LSB galaxies with mysterious origins. In this thesis I have used the Kilo-Degree Survey together with the Hyper Suprime-Cam Subaru Strategic Program to measure the colours and abundances of UDGs in the field, where they are expected to form preferentially via secular mechanisms compared with those in denser environments. I show that a scenario in which cluster-like red sequence UDGs occupy a large fraction of field galaxies is unlikely, with most field UDGs being significantly bluer. I also estimate an upper-limit on the total field abundance of UDGs of $8\pm3\times10^{-3}$cMpc^{-3} within my selection range; this is the first observational constraint on the total field abundance of UDGs. The implied mass formation efficiency of UDGs is high enough that I cannot rule out the possibility that secular formation mechanisms produce a significant proportion of UDGs.

The halo masses M_{halo} of LSB galaxies are also critical measurements for understanding their formation processes. In this thesis, I have estimated M_{halo} for 175 LSB ($23 \leq \bar{\mu}_e[\text{mag arcsec}^{-2}] \leq 28$) galaxies in the core of the Fornax cluster by measuring the number of globular clusters associated with each galaxy. This is the largest sample of low mass galaxies so-far analysed in this kind of study. The sample includes 12 UDGs, with projected r-band half-light radii greater than 1.5 kpc. The UDGs are consistent with having dwarf sized halos, with typical masses between 10^{10} and 10^{11}M$_\odot$. The UDG sample is indistinguishable from smaller LSB galaxies in the same magnitude range. I do not find any candidates likely to be as rich in globular clusters as some of those found in the Coma cluster.

Contents

Chapter 1

Introduction

1.1 A Brief History of Extra-Galactic Astronomy

Looking into the night sky, our eyes are sufficient only to observe a miniscule proportion of the closest objects in our galaxy, the Milky Way. Our ability to observe objects in the cosmos is not only limited by their distance, but also by their intrinsic brightness (luminosity). Thus, what we actually see is a mixture of nearby and bright objects, the most common of which being nearby stars.

It was not until the early seventeenth century that Galileo Galilei's telescope first allowed humans to improve upon the natural limitations of their eyes and peer much deeper into the cosmos, thus allowing astronomers to probe the nature of the Milky Way. Advances in the design of the telescope later enabled Sir William Herschel to estimate the shape of our galaxy. He achieved this by counting the number of stars at different apparent (i.e. observed) brightnesses; under the assumption that all stars have the same luminosity, he estimated the relative distances to each star he observed and concluded that the galaxy was roughly disk-like in shape (Herschel, 1785). While today we know that stars have a great range of intrinsic brightnesses, the disk analogy is still consistent with today's understanding of the Milky Way.

Yet there were certain "nebulous" objects that Herschel could not easily explain; faint patches of light that could not be resolved into individual stars by the instruments of the day. Such objects had two possible explanations: Either they were not made from stars and perhaps relatively nearby, or they were made from stars but so far away that telescopes could not resolve them.

Much of astronomy in the early twentieth century was devoted to measuring the shape of our galaxy, its size, and the distance to the mysterious nebulous objects. Distances to far-away stars could now be accurately measured thanks to the observed

relationship between period and luminosity of Cepheid variable stars that was discovered by Leavitt & Pickering (1912). This was a crucial discovery that enabled accurate measurements of the the size of the Milky Way and to the revelation that the Sun is not situated at the centre of the galaxy, but is instead located in the disk outskirts (Shapley, 1919).

Using the Cepheid variable distance estimation technique in combination with the then-largest telescope in the world, Hubble (1929b) was able to show that a relatively bright nebulous object in the vicinity of the Andromeda constellation was much further away than any of the stars in the galaxy; roughly a hundred times the distance from Earth to the centre of the Milky Way! This was a profound discovery as it confirmed the existence of galaxies other than our own and gave birth to extragalactic astronomy.

Hubble (1929a) was not only able to measure distances to other galaxies, but also their line-of-sight velocities relative to the Earth, made possible by observations of Doppler-shifts in the galaxy spectra. This led to a startling observation: more distant galaxies appeared to be receding at higher velocities. A profound implication was that the Universe is expanding, something that is still believed and observationally supported today. Moreover, this idea was compatible with Einstein's theory of general relativity (Friedmann, 1922). The isotropic expansion of the Universe suggests it had a beginning, where all matter and energy is condensed into a single point in space. This is known as the big bang theory, first proposed by Lemaître (1927).

It was not long after this discovery that Zwicky (1933) discovered something odd about the velocity dispersion of galaxies in the Coma galaxy cluster. In particular, his measurements indicated that galaxies in the cluster were moving much too quickly to be easily explained by the luminous matter that he could observe. He therefore argued that "dark matter", i.e. non-luminous matter, had to constitute the majority of the mass in the cluster.

Several decades later, Rubin & Ford (1970) showed that the stars in the Andromeda Galaxy were orbiting too quickly as to be explained by the mass inferred from its normal baryonic matter. Later, several more galaxies were shown to have the same property (Rubin et al., 1980), providing substantial evidence for the existence of dark matter.

However, the second half of the twentieth century had another major surprise for astronomers. Observations of distant supernovae, exploding stars of predictable luminosity, revealed that the expansion rate of the Universe as measured by Hubble (1929a) was actually increasing over time (Riess et al., 1998; Perlmutter et al., 1999). Cosmologists invoked a constant energy density term (known as the cosmological constant

Λ) in the Einstein field equations to account for this effect, which became known as dark energy, now thought to comprise around 70% of the mass-energy of the Universe (Planck Collaboration et al., 2018).

1.2 The Early Universe

The current understanding of galaxy formation is remarkably simple: large galaxies form from smaller ones that merge together under mutual gravity. This is known as the hierarchical model of galaxy formation (White & Rees, 1978; Kauffmann et al., 1993; Cole et al., 2000), widely accepted by science. Yet this is not the full picture; galaxies show a myriad of variation in their observable properties that are not simply direct consequences of hierarchical merging.

Galaxy formation is the natural consequence of the initial distribution of mass-energy in the Universe developing over time according to physical law[1]. In order to fully understand galaxy formation, it is therefore necessary to begin with an account of where the matter that constitutes galaxies came from. I note that the primary sources of reference for the following section are Coles & Lucchin (1995) & Weinberg (2008).

It is generally believed that we live in a spatially flat Universe (at least as far out to what we can measure) dominated by dark energy, that has expanded outwards from a state of infinite density known as the big bang. The relative proportions of each form of energy-density are typically represented by Ω_m for matter, Ω_r for radiation, Ω_Λ for dark energy and Ω_κ for the spatial curvature density. These quantities change over time. Of these, $\Omega_{0,\kappa}$ and $\Omega_{0,r}$ are thought to be very small, with $\Omega_{0,m} \sim 0.3$ and $\Omega_{0,\Lambda} \sim 0.7$. The "0" in the subscript indicates that the Ω parameters are quoted with their current values (Planck Collaboration et al., 2018).

It is thought that very quickly after the big bang the Universe went through a period of exponential expansion known as inflation (Guth, 1981). During this time, quantum fluctuations in energy-density that existed prior to inflation could have been enlarged to cosmic scales and become "frozen-in" because they could no longer vary after they grew larger than quantum scales.

Despite being insignificant now, it is thought that radiation was the dominant component in the energy density of the Universe early after inflation. During this period, the density of matter in the Universe fluctuated periodically - the effect of gravity balancing the pressure of radiation that was coupled to ionised matter in thermodynamic equilibrium. The origin of these inhomogeneities are theorised to

[1] In fact, this is true of almost everything!

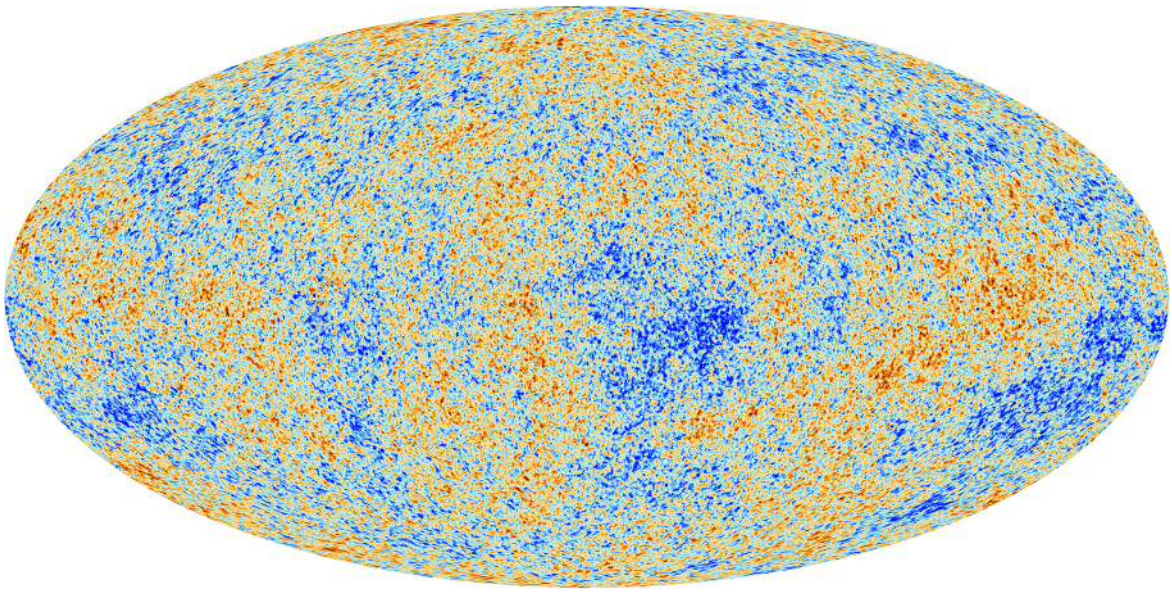

Figure 1.1: Temperature fluctuations (at the ∼3mK level) in the cosmic microwave radiation (CMB), over the whole sky. The temperature fluctuations correspond to density perturbations that are the seeds structure in the Universe. Credit: ESA and the Planck Collaboration.

have been the "frozen-in" quantum mechanical fluctuations from before inflation.

At some point in cosmic history, the density of radiation became low enough that it could no-longer keep the charged matter from recombining into neutral atoms. Following this, thermal equilibrium between the matter radiation was lost and the individual components were separated in an event referred to as "decoupling". After decoupling, the radiation pressure was no longer sufficient to resist the gravitational collapse of the matter over-densities; this was the basis for the first structure in the Universe. The radiation that decoupled from matter was then able to "free stream" through the Universe and the remnant of this light, stretched and dimmed by cosmic expansion, is still visible today as the cosmic-microwave background (CMB). The angular scale over-which the density perturbations exist are directly measurable from observations of the CMB. The CMB was theoretically predicted by Alpher & Herman (1949) and subsequently detected by Penzias & Wilson (1965). In figure 1.1, it is possible to see the density fluctuations of mass in the early Universe that seeded the structures in our Universe today.

1.3 Dark Matter & Structure Formation

Dark matter is thought to be the dominant matter component driving structure formation. Its basic property in many models is that it interacts gravitationally with "normal" baryonic matter, but only exceptionally weakly, if at all, through other mechanisms such as the electromagnetic interaction[2]. Dark matter is therefore able to collapse into large structures much more quickly than ordinary baryonic matter because it does not experience resistive radiation pressure. Individual over-densities (also known as "haloes") of dark matter can accrete baryonic matter (i.e. that which can form stars) because of their large gravitational potentials; this is the basis of galaxy formation.

Since the discovery of dark matter in the twentieth century, numerous theoretical models have been proposed to explain it. Dark matter theories can be classified as being "cold", "warm" or "hot", referring to the thermal velocity of the dark matter particles from low to high, respectively. Correspondingly, dark matter particle candidates range in mass, with colder models having higher particle masses. In hot dark matter models, the large free streaming length of the dark matter particle causes the suppression of energy-density perturbations on small spatial scales in the early Universe, but does not affect fluctuations on larger scales. This means that large structures are able to form before smaller ones in a "top-down" manner. In cold dark matter models, this suppressive effect is much less prolific, allowing small structures to form before larger ones. This is known as "bottom-up" or hierarchical growth of structure, as smaller structures combine over time to form larger ones.

Astronomical observations provide some of the strongest constraints on the nature of dark matter because it is possible to observe structure over a very wide range of spatial scales. This range spans Milky Way globular clusters (a few parsecs in size), to ultra-faint dwarf galaxies in the Local Group (a few tens of parsecs), to clusters of galaxies (megaparsecs), to the whole observable Universe (gigaparsecs), and lots more in between. Modern astronomers use computer simulations in order to make observational predictions from the multitude of dark matter theories. The first simulations only included dark matter (e.g. Frenk et al., 1983; Davis et al., 1985; Springel et al., 2005) because baryonic physics is far more complex than that of simple dark matter models, which only evolve under gravity.

Since dark matter dominates galaxy formation, the expected spatial distribution of

[2]Perhaps because of this weak interaction with baryonic matter, the proposed "dark matter particle" has yet to be directly detected (e.g. Bertone et al., 2005).

Figure 1.2: Predicted distribution of CDM (left) vs. WDM (right) for haloes at redshift z=0 in the Aquarius simulations (Springel et al., 2008). Brightness indicates density, whereas colour indicates velocity dispersion, yellow being high dispersion. Each box is 1.5×1.5 Mpc. Credit: Lovell et al. (2012).

galaxies from such simulations can be compared against redshift surveys (e.g. Geller & Huchra, 1989; Springel et al., 2006) to probe the nature of the dark matter particle. Such comparisons immediately revealed that hot dark matter was strongly disfavoured, with models preferring either cold or warm dark matter. However, distinguishing between these two models is difficult because both provide similar observational predictions on large scales.

However, there initially appeared to be several problems associated with cold dark matter on small scales. One example is the "missing satellites" problem (e.g. Kauffmann et al., 1993; Moore et al., 1999), whereby the simulations seemed to overproduce low mass dark matter satellite haloes compared to observations of the local group. Since warm dark matter suppresses the formation of smaller haloes (see figure 1.2), it was argued that this was evidence for the model. A second problem with CDM that appeared to be resolved by WDM is known as the "too big to fail" problem (Klypin et al., 1999) whereby there are not enough observed dark matter satellites with high enough mass in comparison to theoretical predictions.

One way to probe the nature of dark matter is to observe the column density of redshifted Lyman-α absorption along the line of sight to a distant, bright source such as a quasar (this is known as the Lyman-α forest). Warmer dark matter models predict lower densities of absorbers. Despite their potential to fix the problems with

the simulations, constraints from this method are sufficient to disfavour warm dark matter models (e.g. Viel et al., 2013; Kennedy et al., 2014).

It soon became apparent that some of the issues stemmed from problems in the simulations themselves. Of particular importance was the inclusion of baryon physics in the simulations (e.g. Sawala et al., 2016). This makes it difficult for small dark matter haloes to produce enough stars to form a galaxy because forces other than gravity can play a role and tend to resist star formation. In particular, stellar feedback mechanisms like supernovae are sufficient to remove gas from small haloes, effectively preventing them from forming more stars. It is also likely that the ionizing background radiation in the early Universe was able to delay gas collapse (and therefore star formation) in low mass haloes.

1.4 Star Formation & Evolution

Stars are complex objects but can be broadly classified into spectral types that correspond to surface temperatures and luminosities. Stars can be modelled as black-body radiators, meaning that hotter stars are bluer and cooler stars are redder. Typically, the coolest stars are known as M-type, while the hottest are labelled O-type. It is common for astronomers to visualise the spectral type of stars stars on the Hertzsprung-Russell (HR) diagram[3] (figure 1.4).

In the study of distant galaxies, it is difficult to study individual stars because of limitations on the angular resolution of our telescopes. As such, astronomers are forced to model populations of stars, consisting of a range of masses and possibly other properties like age and metallicity (i.e. the fraction of elements heavier than Helium). The most massive stars are typically the brightest, and therefore dominate the total stellar luminosity of a galaxy.

Massive stars evolve much faster than smaller ones due to their increased rate of fuel consumption. This means that as the population ages, the bright, blue stars disappear much more quickly than the less massive, cooler ones. As stellar populations age (i.e. as the fraction of older stars increases), their integrated colours therefore become redder. The evolution of stars from the main-sequence into the red giant branch also drives this effect. It is therefore possible to estimate the age of a stellar population based on its colour. However, age is not the only factor that can affect a star's colour. The presence of metals in the stellar envelope can absorb energy, reducing the effective surface temperature and therefore causing stars to appear more red.

[3]similar to an observational colour-magnitude diagram.

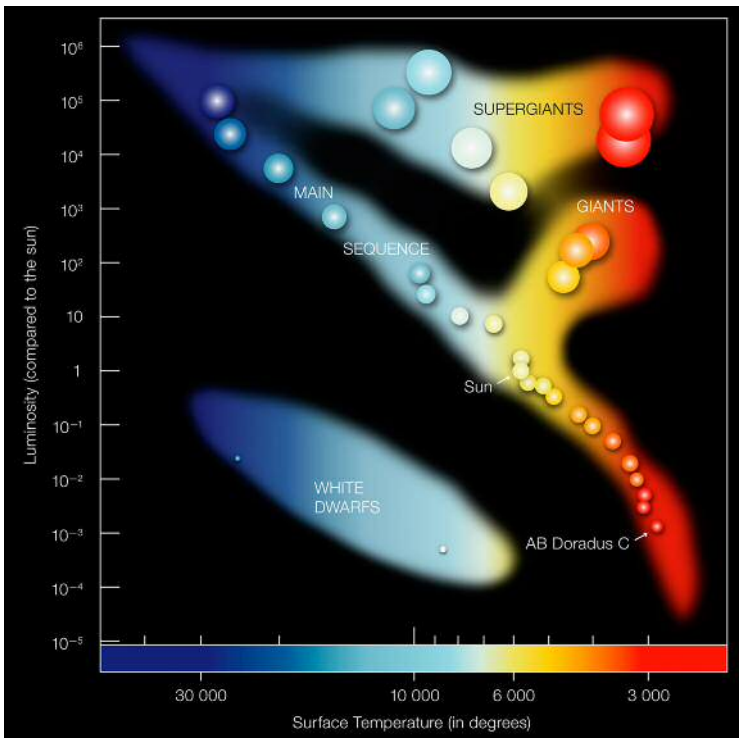

Figure 1.3: Schematic Hertzsprung-Russell (HR) diagram. The surface temperature of a star determines its colour. Hotter stars towards the left of the figure tend to use up their fuel quicker than cooler ones, meaning that older stellar populations appear redder. Credit: European Southern Observatory.

I note that the main source of reference for this section has been Prialnik (2000).

1.5 Galaxy Evolution

Edwin Hubble classified galaxies based mainly on their morphology (i.e. shape). The scheme is commonly represented on the "tuning fork diagram" (figure 1.4), progressing from elliptical galaxies (labelled E) towards either spirals (S) or barred spirals (SB). There are irregular types that are not easily classified. It was suggested that this progression from E types to S or Sb types was representative of galaxy evolution, where elliptical galaxies turn in to spiral galaxies over time. Thus, elliptical galaxies are known as "early type" and spiral galaxies are known as "late type". We know now that this theory is likely incorrect, but the nomenclature remains.

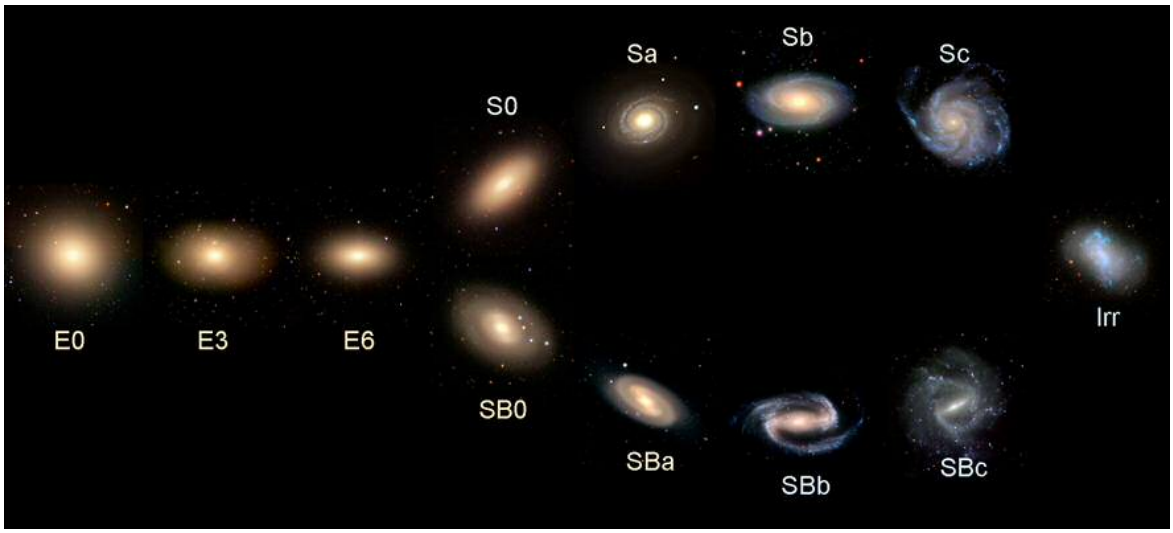

Figure 1.4: The Hubble galaxy classification scheme, or "tuning fork diagram". Credit: Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey (CANDLES).

Galaxy evolution refers to the physical changes that occur to a galaxy over time. The generally accepted modern theory of galaxy formation is hierarchical growth, a consequence of the ΛCDM paradigm[4], wherein smaller galaxies coalesce into successively larger systems through merger events. The mechanisms that cause physical changes in a galaxy can either be secular in nature, whereby events within the galaxy itself are responsible, or environmentally driven, in which an external body or bodies (e.g. another galaxy) affect it.

A fundamental property of galaxies is their ability to form stars. In turn, this corresponds to an ability to turn available gas into dense molecular clouds, the birthplaces of stars. Of course, this is, heavily dependent on the physical characteristics of the galaxy in question.

It is believed that majority of massive galaxies formed the bulk of their stars relatively early in cosmic time. This is the idea of "cosmic downsizing" (Cowie et al., 1996); the observation that smaller galaxies are typically more star-forming compared to massive ones at the present time. This can be explained using the hierarchical model of galaxy formation if the net effect of galaxy mergers is to shut down star formation (e.g. by quickly using up the available gas to form stars) more quickly than for isolated galaxies. Another reason behind this is that low-mass systems have smaller gravitational potential with which to accrete gas, so star formation naturally occurs later in smaller systems. Observationally, this is supported by the fact that some dwarf galaxies harbour higher than average amounts of unused gas by total mass fraction

[4]i.e. the cosmological model combining cold dark matter with the cosmological constant

(Schombert et al., 2001). Additionally, they tend to have relatively low metallicities in comparison to more massive galaxies. This implies that such galaxies have produced relatively fewer supernovae relative to their total baryonic mass and therefore that star formation has been relatively less efficient over time.

A quiescent galaxy is one that is not forming significant numbers of stars. The process of ending star formation in a galaxy is known as "quenching". The origin of quenching in galaxies is a major concern of astronomy because it directly affects the results of cosmological simulations. Once a galaxy is quenched, the luminous, blue stars (e.g. OB stars) that significantly contribute to a galaxy's spectral energy distribution (SED) die-off first owing to the fast energy production rate. Thus, quiescent galaxies appear much redder than star forming ones; this is an observable property and can be directly compared with the results of cosmological simulations. An effect of star formation is chemical enrichment of the ISM from supernovae and stellar winds, also causing a gradual reddening of a galaxy over time. Thus there are two competing mechanisms to redden a galaxy, resulting in the well-known age-metallicity degeneracy.

Because quiescent galaxies are red in colour they belong to the so-called "red sequence". Low mass quiescent galaxies (i.e. quiescent dwarf galaxies) are typically bluer than those of higher mass. Conversely, star-forming galaxies occupy the "blue cloud". Typically for massive galaxies (but not exclusively), red sequence galaxies tend to have elliptical morphologies, while blue cloud galaxies are generally of late-type. (e.g. Taylor et al., 2015) Unless they can find an alternative supply of gas, or become able to form stars out of gas that they could not previously, galaxies typically evolve from the blue cloud to the red sequence.

The mechanisms that cause quenching vary as a function of galaxy mass and environmental density. For example: AGN in high mass galaxies can stop gas cooling and forming stars, sometimes permanently ejecting it from the host galaxy (e.g. Teyssier et al., 2011). In low mass systems, supernovae winds can be sufficient to permanently expel gas. Sometimes, this can also be enough to change the morphological appearance of the host galaxy (e.g. Di Cintio et al., 2017). Since these events are coming from within the host galaxies, they are categorised as secular evolution mechanisms. In low mass systems, small amounts of stellar feedback are more significant owing to the small gravitational potentials that must be overcome to prevent star formation. Because of their relatively small sizes, such mechanisms can be sufficient to temporarily prevent any star formation in the system. Star formation in dwarf galaxies can therefore often be "bursty", with short periods of star formation followed by longer quiescent periods (e.g. Gerola et al., 1980).

There are also environmentally-driven mechanisms that can cause quenching and/or morphological change in galaxies. One particularly dramatic example is major mergers, i.e. the merger of two galaxies of similar masses. It has been suggested that such mergers could be capable of quenching the resulting larger galaxy, but there is some evidence to suggest that this is not true (Weigel et al., 2017). However, it is certainly true that mergers can modify a galaxy's morphology, generally enlarging the stellar halo and making it appear more spheroidal or elliptical.

Galaxies in groups or clusters with others are subject to quenching mechanisms that galaxies in less "dense" environments are not; as such galaxies in these environments are typically quiescent. The space in between galaxies in such places is occupied mainly by hot gas known as the intra-cluster medium (ICM). As galaxies move through this medium, they experience a drag force that is proportional to the ICM density and relative velocity. "Ram pressure stripping" occurs if the resulting pressure is sufficient to overcome the gravitational potential of a galaxy and remove its cold gas. The removal of cold gas restricts star formation and therefore quenches a galaxy over time.

Galaxies in groups or clusters can have high velocity dispersions and can therefore pass each other at very high velocities. The interaction between such galaxies is known as galaxy "harassment" and can cause significant morphological change. The gravitational potential of the galaxy cluster and its DM halo itself can apply tidal forces to nearby galaxies that can cause them to lose their gas to the ICM, perhaps in combination with ram-pressure stripping. This process is referred to as galaxy "strangulation" and can result in quenching. Gravitational interactions between galaxies and/or the cluster potential can cause the production of tidal tails; extended and diffuse features that extend away from their parent galaxy. In addition, gas that is ejected from galaxies in tidal interactions can coalesce to form galaxies of their own, known as tidal dwarf galaxies, which are typically thought to be devoid of dark matter.

1.6 The Low Surface Brightness Universe

In astronomy, surface brightness is a measure of brightness per unit area projected on the sky, typically measured in magnitudes per square arc-second. The optical surface brightness of a galaxy therefore scales with the column density of its stars, as viewed by the observer[5]. Low surface brightness (LSB) refers to surface brightness levels that are fainter than that of the night sky. While this is a rather loose definition, traditionally the LSB regime is thought of as $\mu_{0,B} \geq 23$, where μ denotes a quantity measured in

[5]provided that any obscuring ISM (dust) is optically thin.

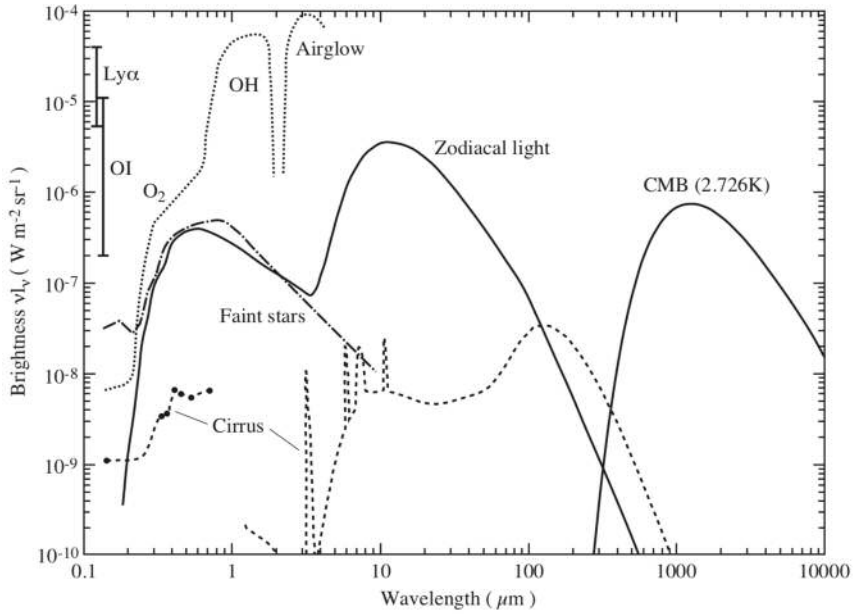

Figure 1.5: Contributions to the brightness of the sky from various sources. The optical regime extends from 0.4-0.8μm. In this region, airglow, integrated light faint stars and zodiacal light play a significant role. Credit: Leinert et al. (1998).

magnitudes per square arc-second and $\mu_{0,B}$ is the central surface brightness of a galaxy as observed in the B-band.

The surface brightness level of the sky has several components which vary in intensity as a function of wavelength (Leinert et al., 1998), as shown in figure 1.6. This includes: *Zodiacal light*: Light originating in the solar system that scatters off dust particles towards the observer, most prominent along the plane of the ecliptic; *Airglow*: light produced by chemical reactions in the upper atmosphere such as the recombination of ions with electrons that were ionised by the Sun during the day; *Integrated starlight*: The combined luminescence of stars in the Milky Way that are too faint to be individually detected, most prominent at low galactic latitudes; *Diffuse galactic light*: Starlight that is scattered off particles in interstellar space towards the observer; *Extragalactic background light*: undetected light originating from outside of the Milky Way; and also artificial light pollution caused by humans.

As the brightness of the sky increases, so does the amount of random fluctuation in its brightness. This is an unavoidable consequence of photon shot noise, which has natural fluctuations in intensity proportional to the square-root of the mean intensity (see chapter 2). As such, components of galaxies with surface brightnesses comparable

to or lower than that of the sky become shrouded in noise and are consequently very difficult or impossible to detect. This systematic bias was first made explicit by Disney (1976), who likened the visibility of galaxies to icebergs; often it is only the "tip" (i.e. bright regions) of a galaxy that can be observed, while the remainder lies undetected below the detection limit of the data.

The work of Disney (1976) was later expanded upon by Disney & Phillipps (1983). At the time, it was observed that the central surface brightness of disc galaxies had values of μ_0=21.65±0.35 through the B filter; this was known as the "Freeman law" (Freeman, 1970). This is close to the surface brightness of the sky. Disney (1976) suggested that this was no coincidence, but in fact was a bias in the way that galaxies were being selected. From Disney & Phillipps (1983): "*The likelihood of a galaxy with given intrinsic profile appearing in a photographic catalogue with limiting criteria on apparent magnitude and apparent size will depend on the maximum distance at which such a galaxy can lie and still obey both criteria.*" They then demonstrated that this maximum distance at which a galaxy of given surface brightness profile is detectable is highly dependant on its central surface brightness.

After accounting for the selection effects and using improved techniques to detect LSB galaxies such as automated surface photometry (Davies et al., 1988), photographic amplification (e.g. Impey et al., 1988; Irwin et al., 1990) and other advances in digital photometry (such as the advent of the charge-coupled device or CCD), it was eventually observed that the Freeman law was indeed a consequence of selection bias (see figure 1.6). In particular, it was found that the number density of low surface brightness galaxies was much higher than expected from the Freeman law (e.g. Davies et al., 1989; McGaugh et al., 1995; Dalcanton et al., 1997), vindicating the work of Disney (1976).

So-called "LSB" galaxies span a wide variety of Hubble types. In groups and clusters, it is common to find LSB dwarf elliptical (dE) and dwarf spheroidal (dSph) galaxies. In less dense environments, late-type LSB irregulars (Irr) and loose spiral galaxies (e.g. SBc) are prevalent. While LSB galaxies tend to have relatively low stellar masses akin to dwarf galaxies, $M_* \leq \sim 10^9 M_\odot$, some are much more massive than dwarfs (e.g. McGaugh et al., 1995). The classic example of such a galaxy is Malin I, a giant LSB spiral galaxy approximately three times the size of the Milky Way (Bothun et al., 1987). Such extreme environments are important but challenging for any astrophysical model to explain (e.g. Lelli et al., 2010); it is now thought that such galaxies could perhaps form in co-planar, co-rotating mergers (Di Cintio et al., 2019) of spiral galaxies.

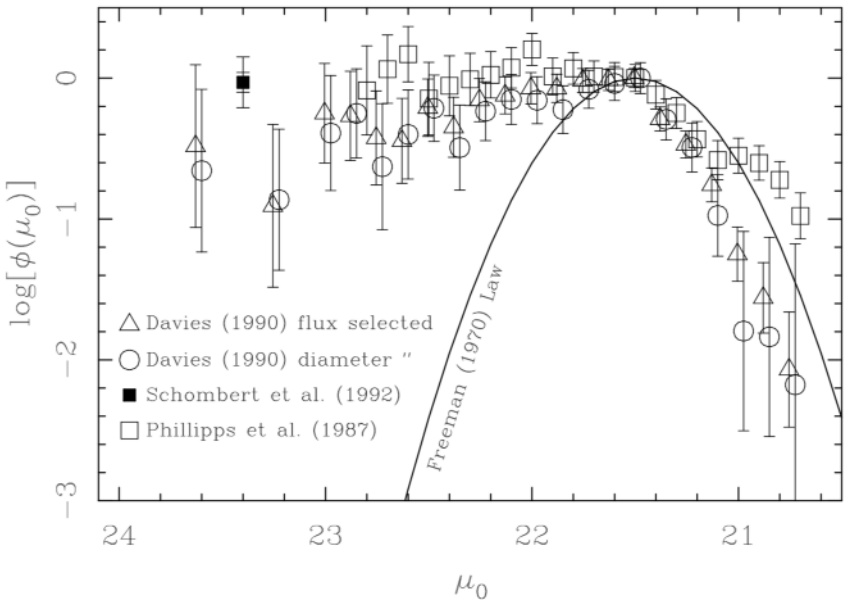

Figure 1.6: The number density of galaxies as a function of their central surface brightness. The Freeman law is obviously a result of selection bias after properly accounting for LSB galaxies. Credit: McGaugh et al. (1995).

Many LSB galaxies are metal poor, having systematically bluer colours than brighter ones. This means that they have not undergone much chemical enrichment from supernovae events. This implies that LSB galaxies have been relatively slow in forming stars. It is thought that most LSB galaxies follow similar evolutionary paths to brighter ones, but slower. This is because they cannot form stars efficiently, despite having very high fractions of their baryonic mass occupied by gas (Schombert et al., 2001). An obvious reason for this is that their gravitational potentials are much weaker than other galaxies, so baryonic feedback mechanisms can more easily prevent their gas from cooling into star-forming regions. A possible implication is therefore that LSB galaxies are in fact dwarf galaxies in terms of their halo mass, as otherwise they might be expected to form stars more efficiently.

Some of the faintest and lowest stellar mass LSB galaxies we know of are in and around the Local Group, many of which being ultra-faint dSph galaxies. This is an observational bias, but the observations still provide important test beds for many physical models. The reason that we are able to detect such galaxies is that we are able to resolve many of their individual stars, over-densities of which stand out above foreground stars in the Milky Way. Limits on angular resolution and depth prevent

the detection of more distant galaxies by this technique. Currently there are around 100 known Local Group satellites (McConnachie, 2012), and advances in observational capabilities e.g. from the Gaia survey (Torrealba et al., 2019) and Hyper-Suprime Cam Strategic Survey Program (Homma et al., 2019) continue to add to this list.

The abundance of LSB galaxies has implications for the galaxy luminosity (or almost equivalently, stellar mass) function, recording the number-density of galaxies per luminosity bin, an important measurement for constraining galaxy formation and dark matter models. As automated detection procedures have improved, more and more LSB galaxies have been discovered, steepening the faint-end slope of the luminosity function (e.g. Sandage & Binggeli, 1984; Bothun et al., 1991; Sabatini et al., 2003). This was particularly relevant because of the "missing satellites problem" (e.g. Klypin et al., 1999; Moore et al., 1999), where there were not enough satellite galaxies observed in the Local Group compared to theoretical predictions from dark matter simulations. Nowadays, this problem is thought to be largely alleviated by the inclusion of baryon physics in computer simulations as this limits star formation in smaller haloes (e.g. Wetzel et al., 2016) and can prevent them from being easily observed. This leads to the idea of "dark galaxies" - dark matter haloes devoid of any detectable starlight. It is possible that such objects can be detected by gravitational interactions with other bodies, or by associated gas clouds that have yet to condense into stars (e.g. Taylor et al., 2017). However, there is currently only limited observational evidence to support the existence of dark galaxies.

LSB galaxies have low densities of stars and other baryonic material compared to other galaxies. This means that they can be used to probe the inner-structures of dark matter haloes without excessive disturbance from baryon physics in an effort to address the "core-cusp" problem (de Blok, 2010). This refers to the observation that the density profiles of dark matter in the central regions of low mass galaxies are not observed to be as steep, or cuspy, as might be expected from dark matter simulations; either there is a problem with the simulations or our understanding of dark matter. However, even in the faintest LSB galaxies, baryon physics may play a non-negligible role in shaping the inner DM profile (Contenta et al., 2018; Genina et al., 2018), gravitationally heating the dark matter and causing it to form shallower, cored central density profiles.

Low surface brightness galaxies are not the only scientifically interesting LSB sources in the Universe. Unbound stars in galaxy clusters can produce a faint but detectable signal known as the intra-cluster light (ICL). It is thought that these stars are removed from their host galaxy in interactions with other galaxies and the cluster

potential. Thus, older and more massive galaxy clusters might be expected to harbour increased ICL. This test is one example of a prediction that can be tested by observing ICL across different galaxy clusters in order to validate galaxy formation models. Another test that can be made is to compare the colours of the stars in the ICL. It has been found that the colour of the ICL in the Virgo Cluster is consistent with the outer regions of massive galaxies, but likely has a variety of progenitors including dwarf galaxies (Mihos et al., 2015).

It has been shown that observations of very low surface brightness tidal features, caused by interactions between passing or merging galaxies, may provide constraints on dark matter models (e.g. Cooper et al., 2010). While it is currently very challenging to reach the necessary depths to test predictions (i.e. $>32\mu_V$), some authors have used a combination of state of the art telescopes, observing strategies and data reduction techniques to reach surface brightness levels approaching this level (e.g. Mihos et al., 2015; Trujillo & Fliri, 2016; Mihos et al., 2017; Román & Trujillo, 2018).

1.6.1 Observational Techniques

There are several important properties of an optical system that can limit its performance in the LSB regime. Of particular importance is the point spread function (PSF), which characterises the appearance of a point source in the focal plane. Many telescopes are optimised to have a narrow PSF in which high fractions of incident light from a point source are tightly focussed, but this often comes at the expense of extended "wings" of the PSF that cause the remainder of the light to be spread widely over the focal plane. This can lead to artefacts in the data that are hard to account for and can appear as genuine astrophysical sources (e.g. Sandin, 2014). There are now efforts to subtract the extended PSF wings from stars in order to improve the accuracy of sky estimates in their vicinities (e.g. Slater et al., 2009).

Also important is the ability to minimise scattered light reaching the focal plane, for example that which reflects off the internal optics of the telescope such as spiders that obscure the field of view (FOV). A significant contribution to scattered light reaching the detector are internal reflections of bright stars which cause bright "ghost rings" in the focal plane, not necessarily centred on the stars themselves. Such artefacts may at best degrade the quality of the data by artificially increasing the sky level or at worst be misinterpreted as genuine LSB sources. The serious problems caused by internal reflections mean that recent advances in LSB imaging have used relatively simple optical designs that reduce the surface area that light can reflect off towards

the focal plane within the telescope. Prominently, the Dragonfly array (Abraham & van Dokkum, 2014) uses an array of telephoto lenses that have no support structures obscuring the FOV. Together with the anti-reflection coatings on the lenses, the resulting telescope is well-optimized for the detection of extended LSB structure. The success of the Burrel-Schmidt telescope for wide field LSB imaging (e.g. Mihos et al., 2017) can be similarly explained.

Also importantly to LSB imaging is the ability to measure and control the sky. Ground-based telescopes suffer from atmospheric effects, that can cause variations in the sky level as a function of time and direction. Additionally, scattering and diffraction of light in the atmosphere can widen the PSF. This leads to the idea that space-based telescopes are the future for LSB science (e.g. Valls-Gabaud, 2016).

The observing strategy can also be optimised for LSB imaging. It is common for "dithering" to be used when observing; this refers to slightly changing the pointing of the telescope so that the source(s) appear at different points in the field of view. This enables systematic effects such as hot pixels or internal reflections to be more easily accounted for. Trujillo & Fliri (2016) detailed an approach whereby a rotation of the telescope is used in addition to normal dithering. The point of this is that internally scattered light always falls in a different part of the focal plane, so the maximum brightness caused by internal reflections is reduced.

Aside from the challenges involved in designing appropriate instrumentation and observing strategies, there are also issues surrounding the software with which we use to detect and measure LSB sources. At present, by far the most popular source extraction package is SExtractor (Bertin & Arnouts, 1996). However, the software is not optimal for the detection of LSB structures. In particular, it is well known than SExtractor can "shred" (i.e. detect as multiple separate sources) LSB sources, leading to incompleteness in observational samples. In order to fully utilise the depth of the data, there are now several efforts to improve upon SExtractor in the LSB regime; this is the subject of chapter 2.

1.6.2 Ultra Diffuse Galaxies

Recently, there has been a resurgence of interest in LSB galaxies thanks to the deep imaging of van Dokkum et al. (2015) with the Dragonfly array, who found a surprisingly high abundance of large LSB galaxies in the Coma Cluster (figure 1.7). The authors coined the term "Ultra-Diffuse Galaxy" or UDG for such objects, a name which has been widely adopted throughout the literature. Specifically, UDGs are de-

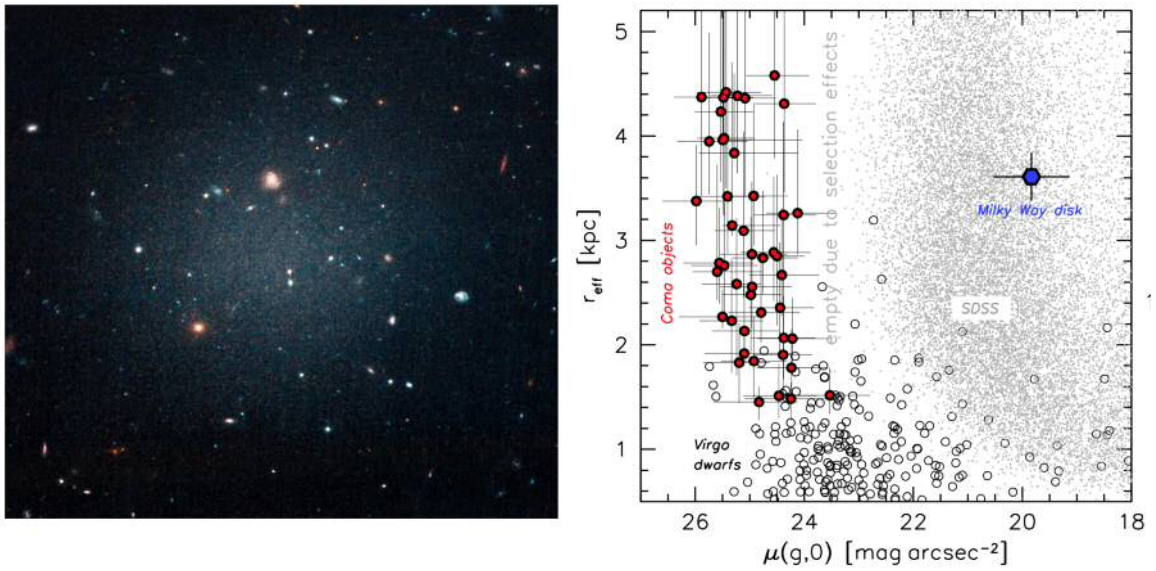

Figure 1.7: *Left*: The ultra diffuse galaxy NGC 1052-DF2, potentially lacking dark matter, taken with the Hubble Space Telescope. Credit: Pieter van Dokkum. *Right* Ultra diffuse galaxies in the Coma galaxy cluster (red) detected using the Dragonfly array, in comparison to more typical galaxies. Credit: van Dokkum et al. (2015).

fined as objects comparable in effective (half-light) radii to the Milky-Way ($r_{e,r}$>1.5 kpc) but are of much lower surface brightness ($\bar{\mu}_{e,r}$>24, where $\bar{\mu}_e$ is the average surface brightness within the effective radius), although the precise definition varies amongst the literature.

UDGs are not new to science. In fact, large LSB galaxies have been known to exist for several decades (e.g Bothun et al., 1987; Impey et al., 1988)[6]. The recent interest has been sparked by the discovery of their high abundance and ubiquity across a variety of environments. There has been much debate over the significance of UDGs and as to whether they make up a different population (in terms of their formation mechanism and thus intrinsic properties) to other, smaller low surface brightness galaxies. There is a growing consensus that UDGs share a continuum of properties with less extreme galaxies, in terms of star formation rate (Leisman et al., 2017), size and luminosity (Conselice, 2018; Danieli & van Dokkum, 2018), metallicity (Fensch et al., 2018) and perhaps also mass to light ratios (chapter 4), however some UDGs may be genuine outliers in the stellar mass - halo mass plane and are devoid of dark matter (e.g. van

[6]Although the giant LSB disk galaxy Malin 1 is not typically considered a UDG because of its central high surface brightness component and its spiral structure, but is nonetheless much larger than most LSB galaxies and indeed UDGs.

Dokkum et al., 2017, 2018, 2019, but see also Trujillo et al. (2018)), perhaps suggesting separate formation mechanisms. Whatever the case, the low mass, diffuse nature of UDGs makes them interesting and challenging objects to explain theoretically.

van Dokkum et al. (2015) initally suggested that UDGs could be "failed L⋆" galaxies, having Milky Way-sized halo masses but with a truncated star formation history. There have been several attempts to constrain the halo masses of UDGs with a variety of measurement techniques used, mainly focussing on UDGs in groups and clusters. Metrics include weak lensing (Sifón et al., 2018), prevalence of tidal features as a function of cluster radius (Mowla et al., 2017), comparisons of their spatial distribution with that of dwarf and massive galaxies (van der Burg et al., 2016; Román & Trujillo, 2017a), richness of their globular cluster systems (Beasley & Trujillo, 2016; Amorisco et al., 2018; van Dokkum et al., 2017; Lim et al., 2018, chapter 4) as well as direct measurements of the velocity dispersions of stellar populations (van Dokkum et al., 2016) and globular cluster systems (Beasley et al., 2016; Toloba et al., 2018). These studies generally conclude that UDGs are dwarf galaxies in terms of their stellar and halo masses, but are much larger than typical dwarfs.

It is likely that the popularity of UDGs among the literature is thanks in-part to their large sizes; this makes them easier to identify against background objects in groups and clusters. Indeed, this property has been exploited by several authors in their studies of UDGs in such environments (e.g. Koda et al., 2015; Muñoz et al., 2015; Yagi et al., 2016; van der Burg et al., 2016, 2017; Janssens et al., 2017; Venhola et al., 2017; Zaritsky et al., 2019; Mancera Piña et al., 2019). This bias towards studies in dense environments is exacerbated by the difficulty of obtaining spectroscopic redshifts (and therefore distances) for large samples of LSB galaxies. However, one recent development suggests that it may be possible to estimate distances to such galaxies by exploiting the ubiquity of the globular cluster luminosity function (Román et al., 2019a), although it is not currently clear whether the globular cluster luminosity function is ubiquitous for UDGs (e.g. van Dokkum et al., 2018, 2019).

UDGs in clusters are typically quiescent and on the red sequence (Koda et al., 2015; van der Burg et al., 2016), showing very few signs of star formation, and show little evidence for tidal interaction even close to the cluster centres (Mowla et al., 2017). This observation suggests relatively high mass-to-light (ML) ratios in order to survive unperturbed from tidal interactions with the cluster potential and other galaxies. Of course, this may not apply to the global population of UDGs; if UDGs exist in the field with lower ML ratios then it is possible that these can be destroyed quickly in a group or cluster environment, meaning that we would be more likely to observe the

ones that survive with higher ML ratios in such places.

There is tentative evidence that UDGs tend to be bluer towards the outskirts of galaxy groups (Román & Trujillo, 2017b), suggesting that interactions with the environment during in-fall can diminish star formation in the UDGs. While this is not surprising if UDGs are quenched during the in-fall, a separate analysis by Román & Trujillo (2017a) did not show a significant trend between environmental density and colour. While UDGs do tend to be somewhat bluer in the field (Leisman et al., 2017; Zaritsky et al., 2018), specific star formation rates seem to be quite low compared to other galaxies (e.g. Román et al., 2019a; RS et al., 2019).

One outstanding question regarding UDGs is whether they are able to form more efficiently in dense environments like groups and clusters, or whether density plays a detrimental role in UDG formation/survival efficiency. van der Burg et al. (2017) find that UDGs are relatively more common in higher-mass environments, but this is in tension with similar studies (Román & Trujillo, 2017b; Mancera Piña et al., 2018) that came to the opposite conclusion.

Whatever the role of environment in UDG production, several authors (e.g. van der Burg et al., 2016; Wittmann et al., 2017; Mancera Piña et al., 2018) have observed a relative dearth in their number density towards the centres of massive clusters. This suggests that in very high density regions, either UDGs are destroyed quickly or do not form as efficiently. One scenario suggested by Janssens et al. (2017) is that UDGs dissolve in cluster cores, possibly depositing ultra-compact dwarf galaxies in the process.

UDGs can also be understood from a theoretical point of view. UDG formation scenarios can be broadly classified as "in-situ" (i.e. secular formation in the absence of interactions with an exterior body or bodies), or environmentally driven. Di Cintio et al. (2017) showed through zoom-in cosmological simulations that gas outflows caused by internal feedback processes can produce UDGs within dwarf-sized halos. Further, Amorisco & Loeb (2016) argue that UDGs can form in-situ both in the field and in cluster environments, and should be expected to do so as the high angular momentum tail of the dwarf galaxy distribution. The importance of high-spin halos for UDG production was also noted by Rong et al. (2017) in their simulations.

However, there are several other feasible formation mechanisms that involve the transformation of normal dwarf galaxies to UDGs though environmental effects. One example is tidal heating, whereby galaxy-galaxy interactions cause an expansion of the dwarfs (Collins et al., 2013; Carleton et al., 2018). There are several pieces of observational evidence showing that some UDGs are associated with tidal interactions

(Wittmann et al., 2017; Bennet et al., 2018). Continually, ram-pressure stripping from the dense intra-cluster medium in groups & clusters is perhaps able to produce UDGs by quenching early in-fall galaxies (Yozin & Bekki, 2015). Jiang et al. (2018) argue that ram-pressure stripping is the primary effect that causes UDGs to lose gas (and therefore shutdown star formation) in dense environments. This can account for the red colours observed for UDGs in clusters.

1.7 Summary

The are two major topics addressed within this thesis. The first is the technical challenge of designing new automated procedures for detecting and accurately measuring LSB galaxies efficiently in wide-field optical imaging surveys. This is especially prudent given the limitations of contemporary software such as SExtractor in the LSB regime and the recent availability of deep survey data such as that from the Kilo-degree survey (KiDS, de Jong et al., 2013) and Hyper-Suprime Cam Strategic Survey Program (Aihara et al., 2018). Methods tested on these surveys can be modified and improved for the next generation of instrumentation such as LSST and the Euclid mission, which will provide all-sky surveys at unprecedented depths. This is the subject of chapter 2. Much of the work presented in this chapter has already been published in Prole et al. (2018).

The second topic is to address the nature of UDGs. In particular, to explain whether they represent a distinct population compared to other dwarf or LSB galaxies. In chapter 4, I address this question by using the correlation between the number of globular clusters associated with galaxies and their total (halo) masses in order to compare UDGs with other LSB galaxies in the Fornax cluster in the stellar mass - halo mass plane. Using these measurements, it is possible to distinguish between formation mechanisms and test whether the Fornax UDGs are really "failed L* galaxies" or have halo masses in the dwarf galaxy regime. Much of the work presented in this chapter has already been published in Prole et al. (2019b).

While much is known about UDGs in dense environments, relatively little is known about the population in less dense environments like the field, where secular formation mechanisms may be more prolific. In chapter 3, I demonstrate that UDGs are systematically bluer in the field than in clusters, implying that they are more star forming. In addition, I demonstrate that the mass formation efficiency of UDGs in the field is similar to that in clusters, but perhaps slightly lower. Such work is a first step in disentangling the relative importances of secular vs. environmental formation

mechanisms for UDGs. Much of the work presented in this chapter has already been published in Prole et al. (2019a).

Chapter 2

Source Extraction

2.1 Instrumentation for Astronomy

Modern optical / near infrared astronomical images are often captured using charge-coupled devices (CCDs). Incident photons interact with a semiconducting material, de-localising one or more electrons through the photoelectric effect. The free electrons are confined in a potential well until the CCD is read and the number electrons are counted. A CCD array combines many individual CCDs, typically in a square or rectangular structure. Following an observation, the electron counts for each CCD are counted in order and recorded digitally; they can then be reconstructed to form an image where each pixel has a value determined by the number of electrons, multiplied by a factor called the *gain*. Such an image is know as "greyscale" because each pixel element has a single, scalar value. This differs from colour images where each pixel is represented by a vector of values, typically three.

CCDs generally have very high quantum efficiency, i.e. the fraction of incident light that is recorded. However, they are not perfect. There are several mechanisms by which noise is introduced into the data. One example is noise produced by electron excitation due to the temperature of the CCD, known as dark current. In addition, noise can be introduced while the electrons are bring counted. This is known as "readout noise", and results from the noise of the amplifier that converts a raw electron count into a voltage measurement that is eventually digitalised.

Before reaching the CCD, incident light first passes through a filter. The filter is designed to control the wavelength range, or "band" of light that reaches the detector. Filters can be broad or narrow-band, depending on the size of their wavelength window. However, filters are not perfect and have have their own quantum efficiencies that vary as a function of wavelength.

The main contribution to the noise in modern wide field optical surveys is typically the natural fluctuations in the number of photons received from a source at fixed intensity, commonly referred to as "photon shot noise". This uncertainty is not an instrumental effect; it has its roots in quantum mechanics and is naturally unavoidable. The probability $P(X = i)$ of observing X photons per unit time given a constant mean photon arrival rate α can be modelled by the Poisson distribution as:

$$P(X = i) = \frac{\alpha^i}{i!} e^{-\alpha}, \text{for } i = 0, 1, ... \tag{2.1}$$

The Poisson distribution has the property that $E[X] = Var(X) = \alpha$, where E is the expectation value operator and Var is the variance operator. The signal to noise ratio (SNR) for sources governed by Poisson statistics scales with $\sqrt{\alpha}$ and α increases linearly with t_{obs}, the integration time of the observation. Therefore longer observations yield a higher SNR in proportion to the square-root of the integration time:

$$\text{SNR} = \sqrt{\alpha} \propto \sqrt{t_{obs}} \tag{2.2}$$

When the total observed number of photons is high, as can be the case for long exposure times, large telescope apertures and bright sources, the Poisson distribution in equation 2.1 can be well-approximated by a Gaussian distribution. This is a consequence of the central limit theorem, which states that for any sequence of random variables $X_1, X_2, ...$ drawn independently from identical distributions (i.i.d), with $\mu = E[X_1]$ and finite variance $\sigma^2 = E[(X_1 - \mu)^2] > 0$ the resulting combined distribution will approach a Gaussian distribution for large sample sizes. In other words:

$$P(X) \to g(X|\mu, \sigma^2) \text{ for } t_{obs} \to \infty \tag{2.3}$$

where $g(x|\mu, \sigma^2)$ represents the Gaussian distribution:

$$g(x|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \tag{2.4}$$

The Gaussian distribution may be preferred over the Poisson distribution for large numbers as it is easier to handle mathematically and computationally.

Measuring the sky distribution is a non-trivial task and can often be error-prone. The sky background distribution is often estimated over a fixed spatial scale l that is larger than the typical source size. However, this technique is susceptible to inaccuracy

if there are sharp gradients in background brightness over l. This can lead to overestimates in the background variance. Another consideration is undetected sources in the region that the sky is being estimated; these sources can cause an overestimate in the background level and its variance if they are not accounted for. This is discussed further in §2.3.4.

2.2 Source Extraction

Prior to the invention of digital processing of CCD data, individual stars and galaxies were mainly identified by eye. Even now, it is not uncommon for astronomers to visually identify galaxies. This is because the human optical system does a remarkably good job at pattern recognition compared to computers. However, there are two major flaws with this methodology: 1) It does not scale well to larger datasets and is not feasible for current or future wide area surveys; and 2) It introduces human bias - we are prone to accidentally missing or miss-classifying sources in ways that computers are not. The present chapter is concerned with efforts to create automated, computerised methods to detect sources in digital datasets.

Source extraction refers to the process of extracting information about individual sources from data. In this case, I will mainly discuss source extraction from two-dimensional image data as is typical of modern optical astronomy. Source extraction is a macro process that contains may subprocesses. Typically, this consists of a detection stage, whereby pixels that belong to sources are identified, a segmentation stage, where detected pixels are assigned unique labels corresponding to individual sources; and a measurement stage, where a set of parameters are estimated for the sources. Each of these subprocesses are comprised of many different subroutines, and can vary in complexity. I will now describe these stages in detail.

2.2.1 Detection

Typically, detection refers to the identification of significant pixels above the background. As such, essentially all contemporary detection algorithms rely on some form of background estimation. The background is often measured from pixels that are not part of a source[1]. However, this definition poses a serious problem, because the defini-

[1]There is a conceptual distinction between an astrophysical source (for example a star or galaxy) and a source in an image processing sense. The definition of a "source" from this point of view is entirely dependent on the extraction pipeline; different techniques will always give different results.

tion of a source relies on the background estimate. Various methods of circumventing this issue will be discussed later in the chapter.

The "sky" or "background" distribution refers to the distribution of pixel brightness values that are labelled as background pixels. In the first pass of background estimation, the whole image may be treated as sky. As such, bright sources will cause the distribution will be biased to higher values. Pixels can be treated as not-sky if they can be shown to belong to a source; this is the process of detection. A common approach is to detect pixels if they are significantly bright (in a statistical sense) from the background distribution. The background distribution can then be re-estimated ignoring these values.

The limits of detection are, obviously, dependent on the depth of the data. If a source is too faint, it will not be detectable at a high significance level. Lowering the detection threshold allows for fainter sources to be detected, but also results in a greater chance of detecting noise fluctuations in the sky. An astronomical survey must balance the merits of a high detection threshold (low numbers of noise detections but also less detections of astrophysical sources) with a low one (high numbers of noise detections with more detections of faint astrophysical sources).

The background measurement is coupled to the depth of the data. To see this, consider a large low surface brightness galaxy that is too faint to be detected in the data. Such a source will contribute to the background estimate, making it systematically higher than in other regions of sky. Continually, the low surface brightness outskirts of brighter galaxies may not be detected, having similar effects.

2.2.2 Segmentation

Segmentation is the process of assigning unique labels to pixels belonging to individual sources. A segmentation image can be made from such pixel classifications, having the same dimensions as the original image, where each detected pixel in the original image has its source label in the segmentation image. Correspondingly, each source has its own segment defined by the set of pixels in the segmentation image with a specific label. There are several different segmentation techniques, each with their own advantages and disadvantages. Some contemporary approaches are discussed later in the chapter.

There is no such thing as an ideal segmentation because different sources can "blend" together. This means that individual pixels generally have contributions from multiple sources. One might naively suggest that the pixel labels should represent the source that contributes the most flux to that pixel. However, consider a LSB source

embedded in the halo of a bright galaxy; since the bright galaxy dominates over all pixels, the LSB source has no pixels in the segmentation image. Indeed, one major challenge of segmentation algorithms is dealing with "nested" (embedded) sources.

2.2.3 Measurement

The most efficient and crude source measurements can be made directly from the source pixels (i.e. the pixels) themselves. For example, one can approximate the total luminosity of a source by the sum of its constituent pixels, or its radius from the maximum distance from its luminosity-weighted centre.

However, measurements derived from segments are always biased, more so for low surface brightness galaxies. The reason for this is simple; the size of segments themselves are dependent on the depth of the data. This follows from the discussion in §2.2.1. For example, consider an extended LSB source with a central brightness that is just enough to be detected. Its corresponding segment would be much smaller than the true extent of the source, encompassing only the central region. A resulting estimate of the size based on the segment would be a significant underestimate, the same applies for the total luminosity. However, the average surface brightness for that segment would be significantly brighter than its true value. Catalogues produced by measurements from segmentation images should therefore be used as first order estimates only, especially in the search for low surface brightness galaxies.

In practice, it is common for an additional layer of measurement that can use first-order estimates as inputs. Typically, this includes fitting of sources using parametric models. It is possible to fit several nearby sources concurrently, helping to alleviate the issue of nested sources. Parametric fitting can also increase the measurement accuracy for sources with undetected LSB components. Details of such procedures will be given in subsequent chapters.

There is a multitude of source extraction software available freely to the scientific community, SExtractor (Bertin & Arnouts, 1996) being by far the most widely used for extragalactic Astronomy. In what follows, I will briefly describe some of these packages and give an indication of their comparative strengths and shortcomings. Each piece of software has, sometimes many, parameters that the user can modify. I will introduce relevant parameters in the description of each package, but for exhaustive lists the reader is referred to their respective documentations.

2.2.4 DAOPHOT

DAOPHOT (Stetson, 1987) is a source extraction package suited to the detection of point sources such as stars. It was created primarily to detect and measure stars in crowded fields (i.e. images with a high density of point sources). It uses convolution with fixed Gaussian filters to estimate the peak flux associated with a hypothetical point source at every pixel location. The Gaussian filters are modified such that convolution over an area of fixed brightness (or a linear gradient) yield no signal. Since the surface brightness profiles of point sources (PSFs) are typically well fit by Gaussian profiles, the convolution produces a strong signal in their presence. Additionally, this signal is generally sharper than in the raw image, making it easier to distinguish between nearby sources. Individual sources can the acquired by applying a significance threshold to the convolved image, requiring that individual sources are made from contiguous (i.e. adjacent) pixels.

2.2.5 Clumpfind

Clumpfind (Williams et al., 1994) uses a wholly different algorithm, designed to detect structures embedded in three-dimensional space (two spatial dimensions and one for velocity). The algorithm operates by identifying sources as isolated regions of contiguous pixels above a discrete set of detection thresholds, starting at the most significant and proceeding downwards. Naturally, all regions detected at higher significance levels are also detected at lower ones, forming a hierarchical structure. If a detection at a particular significance level encompasses two or more detections at higher levels, then that layer is said to be blended. The approach Clumpfind uses to segment these layers, i.e. to assign existing labels in the top layer to each blended pixel in the blended layer, is to simply *dilate* each top layer segment until each pixel in the blended segment has a label. Here, dilation refers to the process of expanding a labelled region iteratively by a kernel. In this case, the kernel is defined by the (three-dimensional) pixels that are adjacent to the one being dilated. Each unlabelled pixel connected by the kernel to the dilated pixel are assigned its label; this process can happen iteratively so that existing segments grow until all unlabelled, detected pixels are labelled.

2.2.6 Source Extractor

SExtractor (Bertin & Arnouts, 1996) works in a similar way as Clumpfind but

with a few key differences that make it well suited for blind galaxy surveys (among other things). SExtractor first measures an estimate of the background, before labelling segments using a multi-threshold technique that is similar to Clumpfind. The software also produces a detailed catalogue of structural and photometric parameters, measured using the segments. Below I will describe the individual processes.

Background Estimation: The background distribution is estimated in a mesh grid, i.e. a grid structure where each grid element (mesh) is comprised of the individual pixels it contains. In order to estimate an unbiased sky level, a sigma-clipping algorithm (at 3σ) is applied to the brightness distribution in each mesh. If the resulting mean estimate doesn't change by more than 20% from the non-clipped mean, then the mesh is considered not-crowded and the clipped mean is used as the background level estimate. If it does change by more than this amount, then the mesh is considered crowded and the background level is estimated as 2.5×Median-1.5×Mean, which reduces the bias.

The resulting mesh grid can be median filtered over in order to reduce potential sky over-estimates in individual meshes. A bi-cubic spline is then fit to the mesh grid in order to produce a smooth background map of the same dimension as the original image.

Detection & Segmentation: Before detection, SExtractor typically applies a convolution with a Gaussian filter to the data in order to suppress noise fluctuations on a pixel-to-pixel level, therefore increasing the SNR. A lower-threshold is applied that defines which pixels are detected. Individual segments are initially defined as contiguous detected regions, but are then passed through a de-blending algorithm to separate merged objects. This is done by performing multi-thresholding: A set of discrete brightness thresholds (typically 32) are applied to the segment, beginning with the brightest threshold level. This creates a hierarchical tree structure for each segment, similar to that in Clumpfind. For each segment, every node (contiguous set of pixels above the the threshold associated with its level in the tree) in the tree is tested to see whether it should be considered as a separate source, based on the following criteria (directly from Bertin & Arnouts, 1996): a) the integrated pixel intensity (above the threshold) of the branch is greater than a certain fraction of the total intensity of the composite object; b) the condition a) is verified for at least one more branch at the same level.

One way that SExtractor prevents excessive de-blending due to noise peaks is to limit the lowest threshold used for the de-blend (this is a user-specified parameter). However, this presents a problem because it is not clear which de-blended component to which pixels lower than this threshold should be assigned. SExtractor deals

with this by fitting a simple bivariate Gaussian model to each de-blended source in the original segment. Unassigned faint pixels are assigned to whichever component has the most significant contribution at that pixel based on the model.

SExtractor can attempt to remove spurious detections caused by noise contributions to the (otherwise undetected) LSB wings of detected sources in its "cleaning" algorithm. This process involves fitting a Moffat profile to each source neighbouring the one in question, and assessing whether the object would have still been detected in their absence. If not, then the detected pixels are assigned to the neighbour based on the fitted profile. In this way, the pixels corresponding to a finalised SExtractor segment need not be contiguous.

2.2.7 MTObjects

MTObjects (Teeninga et al., 2016) is a source extraction package designed to improve on the ability of SExtractor to identify extended LSB structure. Unlike other detection algorithms, it does not use a fixed threshold for detection. Instead, it uses statistical significance tests to identify significant nodes in a *max-tree* structure, containing all the pixels in the image. The overall source extraction process is described below.

Background Estimation: MTObjects begins with a single valued background estimation: The image is split-up into a mesh grid. In each mesh, a statistical test is performed to test how likely the data are to be drawn from a Gaussian distribution. The closer to a Gaussian distribution, the more an individual mesh is considered as "flat" (i.e. devoid of sources). The mean value of all meshes that are sufficiently flat is used as the initial background estimate, representative of the whole image. Any pixel above the background value is eligible to become part of a source, pixels below are ignored.

Avoiding the creation of a spatially varying background map as done by SExtractor means that MTObjects can partially circumvent issues like over-estimating the background due to a mesh size that is too small, for example. However this is not an ideal solution; consider a large (in relation to the mesh size) image that contains an extended, slowly spatially varying LSB halo. The pixels in a mesh centred on this halo would be approximately normally distributed. Now consider a mesh that is placed on a different part of the image, in which there are no sources. In this case, the pixels are also approximately normally distributed, but the background level is much lower than before. Thus, the mean of these two "flat" tiles is positively biased. It is therefore

important to define the mesh spacing length so that it is larger than any other object in the image. At the time of writing however, this mesh size is fixed at 64×64 pixels.

Identifying sources: Following the background estimate, the next step is to create a max-tree from the whole image. The root (bottom-level) node is made from the background pixels. The tree structure is very similar to that used in SExtractor for de-blending. A key difference is that MTObjects uses a "continuous" threshold, i.e. each unique pixel value represents a new threshold level. Each individual node is statistically tested to see whether it is significant compared to its closest significant ancestor node, by default the root node.

Multiple nested significant nodes may correspond to a single source. Therefore, MTObjects applies logical operations in order to merge significant nodes into individual sources. Significant nodes with no significant ancestor nodes are considered as individual sources. Each node has a corresponding "main branch", defined as the descendent node with the largest area. If the closest significant ancestor of a significant node has a main branch which is different from the significant node in question, then it is labelled as a new source.

2.2.8 Fellwalker

Fellwalker (Berry, 2015) uses a gradient-tracing approach to assign individual pixels to local maxima. The basic idea behind the algorithm is that each detected pixel is the start of a walk; each walk traverses through neighbouring pixels, following the path of steepest ascent. A peak is defined by any pixel that has no higher-valued neighbour pixel. Before the peak is confirmed, the algorithm checks if there is a higher-valued pixel in a wider neighbouring region. In the case that there is, the peak is rejected and the algorithm continues the walk from the new pixel. Once a peak is found, it is assigned a unique label. This label is also assigned to each pixel visited on the walk that led to the peak. Any walk that intersects another adopts its label and terminates. In this way, individual segments area created for each peak.

The outcome of such an approach is clearly dependent on the size of the pixel neighbourhood used to define peaks. Too small and noise peaks will dominate. Too high and multiple sources will be blended together in a single segment. Fellwalker accommodates merging of segments to address the first issue; neighbouring segments that have peak values below a certain tolerance, relative to the highest value on the common segment boundary, are merged. This process is applied to every segment iteratively, until no more segments can merge.

Another limitation is that isolated noise peaks may never merge. FELLWALKER can remove segments that do not meet its clean criteria, e.g. that are too faint or contain too few pixels.

2.2.9 NoiseChisel

NOISECHISEL (Akhlaghi & Ichikawa, 2015) is source extraction software that has a unique pixel detection algorithm designed to detect very faint sources. The algorithm is fairly complex and as such I refer the interested reader to the original paper. Briefly, NOISECHISEL uses a combination of dilations and erosions (morphological operations describing the expansion / retraction of a binary segment by a kernel) in combination with statistical testing to reliably detect faint features while reducing the number of false detections due to noise compared to SEXTRACTOR.

2.2.10 ProFound

PROFOUND (Robotham et al., 2018) is a source extraction package that uses a watershed segmentation technique followed by an iterative segment dilation process in order to measure sources. PROFOUND addresses the issue of confusion between background and faint source pixels by iteratively remeasuring the background, masking out successively larger regions of the data corresponding to sources. PROFOUND is designed to be modular in nature such that it is flexible for the user. However, PROFOUND has a default source extraction procedure that I will now describe in detail:
Background Estimation: PROFOUND treats the background estimation as an iterative process, making several passes over the data. It uses a mesh grid, similarly to SEXTRACTOR and other software. However, PROFOUND offers improvements by reducing the bias caused by non-sky pixels in each individual mesh element. In each tile, a sigma-clipping algorithm is applied to all non-detected pixels (more on this later) and unbiased estimators for the sky level and its RMS uncertainty are applied. A bi-linear spline is fit to the resulting sky and RMS maps. The overall process of sky estimation is designed to be run several times over the source of a PROFOUND source extraction run, each time increasing the number of detected pixels and reducing the bias.
Source Extraction: Following an initial background estimate, PROFOUND applies a single detection threshold (in SNR units) to initially identify pixels. Following this, a watershed segmentation algorithm is applied. The following description of the algorithm is taken directly from Robotham et al. (2018):

(i) Identify the brightest pixel in the image above the specified surface brightness level which is not already assigned to a segment.
(ii) Progressively search unassigned image pixels surrounding the current segment, for each pixel searched:

a) If a searched pixel has less flux than any neighbouring pixels already assigned to the segment, then assign to the current segment if no unassigned pixels neighbouring the pixel under consideration have more flux.
b) If a searched pixel has more flux than its neighbours already assigned to the segment above some tolerance level then do not assign it to the current segment.
c) If no more pixels can be assigned to the current segment then terminate the growing process.

(iii) Select the next brightest unassigned pixel remaining in the image and assign it to a new segment, then repeat the above segment growing process.
(iv) Once all pixels above the specified surface brightness level have been assigned to a segment terminate the watershed process.

The result of this algorithm is a segmentation map with unique segments assigned to unique sources. One limitation associated with the watershed segmentation is that it does not allow for nested sources. Following this, the sky is re-estimated and basic statistics are calculated for each segment, most importantly for the present discussion, the flux. ProFound then dilates each segment by a kernel (recall that dilation of a segmentation image means that every unlabelled pixel that can be connected by the kernel to any labelled pixel is assigned its label), before once again remeasuring the sky using the newly dilated segmentation image to mask out detected pixels from the estimate.

The unique feature of ProFound is that it applies the sky estimate, segment measurement and segment dilation iteratively, each time checking if some statistic (typically the total flux) of a segment has converged to within a user-specified tolerance. If so, the segment stops dilating and the total flux of the segment approximates that of the underlying source.

One of the major strengths of ProFound is also one of its greatest drawbacks. Its iterative nature means that it is comparatively much slower that other packages overall. It is very computationally intensive, both in terms of CPU time and memory requirements. However, due to its modular design one can circumvent excessive run times by creating their own scripts, albeit sacrificing its full potential.

2.3 DeepScan

DeepScan is a source extraction package designed to overcome several of the limitations of contemporary software in detecting extended LSB objects. I wrote the DeepScan code from scratch (excluding its external dependencies) as part of the work for this thesis. The code is freely available online[2], and is in a state of continual improvement. It was designed to overcome the issue of fragmentation, or "shredding" of large LSB galaxies encountered by several authors (e.g. Davies et al., 2016) using SExtractor. In the following section, I give an overview of the algorithms used in DeepScan to measure the sky, detect pixels and perform segmentation and deblending. At the end of the section, I show an example of DeepScan being applied to the Next Generation Virgo Survey (NGVS; Ferrarese et al., 2012) data.

2.3.1 The `DBSCAN` algorithm

`DBSCAN` (Density-based spatial clustering of applications with noise, Ester et al., 1996), is a two-parameter algorithm that is designed to identify regions of high density within an n-dimensional data set. The algorithm has found recent use in astronomy through the classification of eclipsing binary stars (Kochoska et al., 2017) and the morphological analysis of open clusters (Bhattacharya et al., 2017). Broadly speaking, my application of `DBSCAN` operates in a similar way to other detection algorithms in that it builds detections by clustering together nearby pixels above a brightness threshold.

The fundamental difference between my use of `DBSCAN` and contiguous pixel clustering algorithms like that used by SExtractor is that a `DBSCAN` detection is based on the density of pixels above a SNR threshold within its ϵ radius, not necessarily being contiguous. (We note that SExtractor can amalgamate non-contiguous sources in its "cleaning" stage, which attempts to remove noise peaks that have been detected in the halos of brighter objects).

The first parameter of `DBSCAN` is a clustering scale length (ϵ) and the second (η) is the minimum number of data points required within an ϵ radius for a cluster to form. The algorithm iterates over every input data point. If the number of points φ within a circular radius of length ϵ meets the condition $\varphi \geq \eta$, the point is marked as a *core* point; this is the basic building block of a cluster. Then, each point enclosed within ϵ (known as *secondary* points) is checked to see if they also meet the condition to be core points, and if so then they are added to the same cluster. Thus, a cluster can contain more than one core point. This process repeats until there are no more core

[2] https://github.com/danjampro/DeepScan

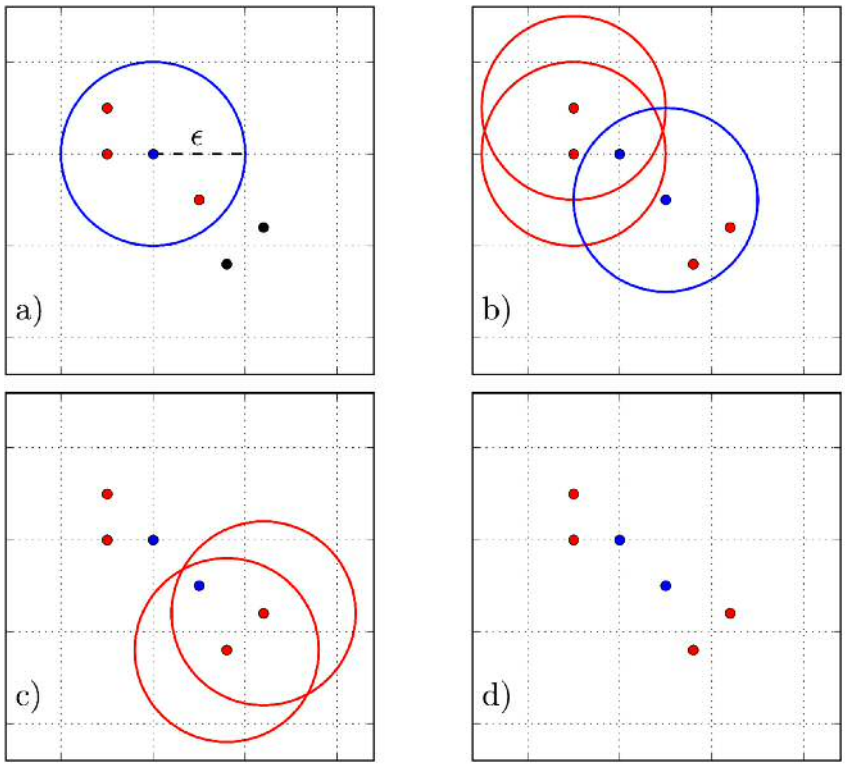

Figure 2.1: A visualisation of the DBSCAN algorithm with η=3. In a), the algorithm proceeds to identify a core point (blue). Then in b) it checks each of the secondary points (red) to see if they meet the criteria to be a core point - this is true of one of the secondary points, which is also shown in blue. The secondary points of the new core point are checked in c), and the final cluster, consisting of two core points, is shown in d).

points to add to the cluster and it is complete. The algorithm then repeats the process to identify separate clusters within the dataset, if they exist. The clustering process is illustrated in figure 2.1.

My approach is to use the spatial coordinates of pixels above a brightness threshold as inputs to DBSCAN, essentially identifying sources as over-densities of these pixels. This is analogous to how ultra faint local group dwarf galaxies are detected through over-densities of resolved stars against the foreground stars in the Milky Way. While I have not implemented usage of the upper detection threshold $I_{\max}$ in DeepScan I include it in the modelling for completeness.

Because DBSCAN uses a circular clustering region the algorithm has a "resolution"

determined by ϵ. This limitation does not rule out the detection of elongated structures if they are significant over scales similar to or larger than ϵ. The algorithm therefore performs poorly in identifying objects significantly smaller than the detection circle and in separating sources closer together than 2ϵ. While the former point is addressed by allowing the sensitivity of the algorithm to be set by the user, the latter can be remedied using an additional de-blending or segmentation algorithm.

The sensitivity of `DBSCAN` is set by η. To derive a value for η, a value of ϵ is assumed that remains a hyper-parameter of the algorithm (i.e. a parameter that is set by the user). η is estimated with the assumption that the noise brightness distribution is a zero-mean Gaussian of standard deviation σ_n, i.e.:

$$P(I_n) = g(I_n|0, \sigma_n^2) \tag{2.5}$$

for noise intensity I_n, mean μ and standard deviation σ. If a brightness threshold is applied with lower and upper boundaries $I_{\min}$ and $I_{\max}$ ($\mu_{\min}$ and $\mu_{\max}$ in magnitudes per square arc-second) respectively, the pixel-to-pixel noise distribution can be used to predict the probability P_{thresh} of a background pixel with a true brightness of I_b lying within the threshold:

$$P_{\text{thresh}} = P(I_{\min} \leq I \leq I_{\max}) = \int_{I_{\min}-I_b}^{I_{\max}-I_b} P(I'_n) dI'_n \tag{2.6}$$

Thus, an accurate model of the background is also assumed. As the amount of noise per pixel is modelled as an independent random variable, the binomial distribution can be used to calculate the number of pixels expected to lie within the brightness threshold within a circular region of radius ϵ. In the hunt for extended LSB objects, ϵ should be large so that it encapsulates a high number of pixels. Therefore the binomial distribution can be approximately represented by another Gaussian,

$$P(\varphi) = g(\varphi|\mu_\varphi, \sigma_\varphi^2) \tag{2.7}$$

with

$$\mu_\varphi = P_{\text{thresh}} N_{\text{pix}} \tag{2.8}$$

$$\sigma_\varphi = (P_{\text{thresh}}(1 - P_{\text{thresh}}) N_{\text{pix}})^{\frac{1}{2}} \tag{2.9}$$

where N_{pix} is the number of pixels enclosed by a circle of radius ϵ, and φ is the number of those pixels within the threshold. Equation 2.7 can be integrated between φ' and

$N_{\rm pix}$ to find the probability P_0 of φ' or more pixels in the circle lying within the threshold:

$$P_0 = \frac{1}{2}\left[\mathrm{erf}\left(\frac{N_{\rm pix}-\mu_\varphi}{\sigma_\varphi\sqrt{2}}\right) - \mathrm{erf}\left(\frac{\varphi'-\mu_\varphi}{\sigma_\varphi\sqrt{2}}\right)\right] \tag{2.10}$$

Setting φ'=η and rearranging for η, I obtain

$$\eta = \sigma_\varphi\sqrt{2}\,\mathrm{erf}^{-1}\left[\mathrm{erf}\left(\frac{N_{\rm pix}-\mu_\varphi}{\sigma_\varphi\sqrt{2}}\right) - 2P_0\right] + \mu_\varphi \tag{2.11}$$

I can replace the hyper-parameter η with a new parameter[3] κ , defined as the number of standard deviations (equation 2.9) above the expected number of points enclosed within ϵ. I can therefore write the somewhat simpler expression,

$$\eta = \mu_\varphi + \kappa\sigma_\varphi \tag{2.12}$$

where κ is in one-to-one correspondence with the probability P_0.

Equation 2.11 describes the number of pixels lying within the brightness threshold, within a circle of radius ϵ embedded within pure Gaussian noise. It is useful as it expresses η as a function of the probability of that many pixels occurring due to noise, a probability that can be set arbitrarily low by increasing κ. It also allows the prediction of what should be detected by the algorithm. For example, in the context of galaxy detection the detectable region on the central surface brightness (CSB), magnitude plane can be calculated. The brightness profiles of galaxies are often described by the Sérsic profile (Graham & Driver, 2005), which can be expressed as:

$$I(r) = I_0 \exp\left(\frac{-r}{h}\right)^{\frac{1}{n}} + I_b \tag{2.13}$$

for radius r, CSB I_0 (μ_0 in magnitudes per square arcsecond), scale size h and Sérsic index n. I_b is the background brightness. In analogy to equation 2.6, the probability of finding a pixel within the brightness threshold is:

$$P_{\rm thresh}(I(r)) = \int_{\Delta I_{\rm min}}^{\Delta I_{\rm max}} P(I'_n)dI' \tag{2.14}$$

where

$$\Delta I_{\substack{\rm min\\ \rm max}} = I_{\substack{\rm min\\ \rm max}} - I(r) \tag{2.15}$$

Equation 2.14 can be integrated over a circular region to obtain:

[3]Users of my software can still opt to specify η manually.

$$\varphi_s = \pi \int_0^{\epsilon} r \left[\mathrm{erf}\left(\frac{\Delta I_{\mathrm{max}}}{\sigma_n \sqrt{2}} \right) - \mathrm{erf}\left(\frac{\Delta I_{\mathrm{min}}}{\sigma_n \sqrt{2}} \right) \right] dr \tag{2.16}$$

The condition for the galaxy to be detected is then simply:

$$\varphi_s \geq \eta \tag{2.17}$$

Numerical approximations to this condition are shown for a variety of values of ϵ, κ and Sérsic index n in figure 2.2.

The effects of the PSF have not been modelled here. While the PSF is in general non-analytical and varies between datasets (and even in the same dataset), I probe the effects of a typical seeing PSF for ground based wide-area surveys in §2.3.2 and have found the effect to be negligible for my target sources.

2.3.2 Testing DBSCAN

In order to demonstrate the validity of the statistical modelling presented in section §2.3.1 I have performed artificial galaxy experiments, wherein sets of randomly generated circular Sérsic profiles (n=1) were generated using `ProFit`[4] (Robotham et al., 2017) and hidden in random noise of RMS=σ_n. n=1 was used because it is a fiducial value for dwarf galaxies (e.g. Koda et al., 2015) (but also see figure 2.2). I also include the effects of photon noise in the experiments and assume a gain of 1. A large grid of profiles was produced and embedded into random noise, using central surface brightnesses defined by their signal-to-noise (SNR) ratio (SNR logarithmically drawn between 0.1 and 100). The profiles have effective (half-light) radii between 1 and 15 pixels, where I have converted between Sérsic quantities using the prescriptions of Graham & Driver (2005).

Individual profiles were spaced by eight times the maximum effective radius of the sample, and were each truncated at 4 times this radius. This was to ensure that extended profiles could not contribute to their neighbour's detections.

My new `DBSCAN` implementation (see §2.3.1) was then applied to the synthetic image, and any synthetic source that had an object located (by the mean coordinate of the core points) within twice its effective radius was regarded as detected. No two `DBSCAN` detections could be assigned to the same source and it was asserted that there were no `DBSCAN` detections that did not have matches. This check was to ensure that large portions of the image had not been detected as one. Results from matching the detections with their profiles are shown in figure 2.3 for κ=10, ϵ=5 pixels and a lower

[4]https://github.com/ICRAR/ProFit

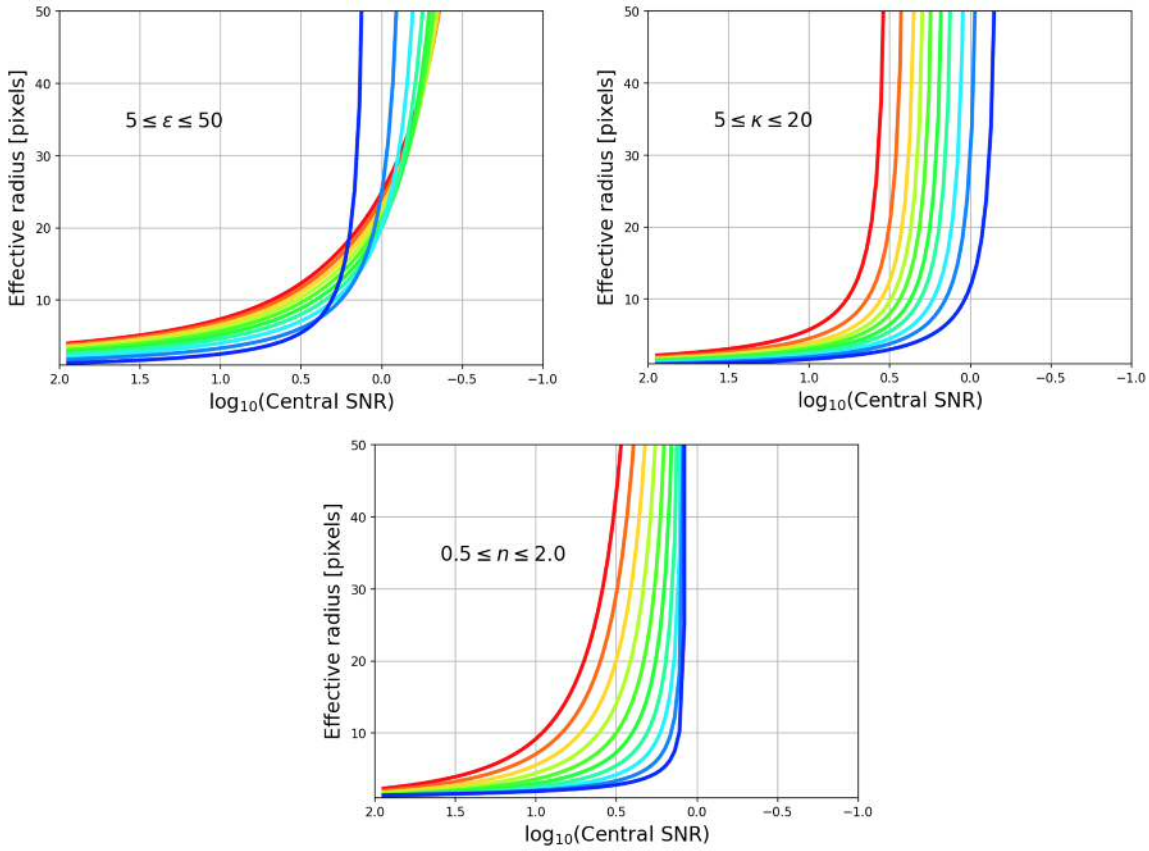

Figure 2.2: Variation of the limits of detection with ϵ, κ and Sérsic index n, ranging linearly between the limits shown from blue to red. I have assumed pure Gaussian noise. For the variation with ϵ, the results obtained with κ=10, $I_{\rm min}$=σ_n. For κ, the results were obtained with ϵ=5 pixels, $I_{\rm min}$=σ_n. For the variation with n, the results were obtained with κ=10, $I_{\rm min}$=σ_n, ϵ=5 pixels. Effects of the PSF are not considered here.

detection threshold of $1\sigma_n$. Also on the plot is shown a numerical approximation (using the Nelder-Mead minimisation algorithm to the condition in equation 2.17). Importantly, the detection boundary predicted through the modelling is in good agreement with the observation of the synthetic data.

I have performed the same experiment after convolving the Sérsic profiles with a mock Gaussian PSF with a full-width at half-maximum (FWHM) of 5 pixels (typical for 1″ seeing with a 0.2″ pixel size). The results were practically identical, even for the lower values of effective radii that I probed. This is easily explained from the fact that ϵ was also 5 pixels. However, DeepScan is intended for use on wide-area survey data (which typically has seeing and pixel scales of the order of what is probed here)

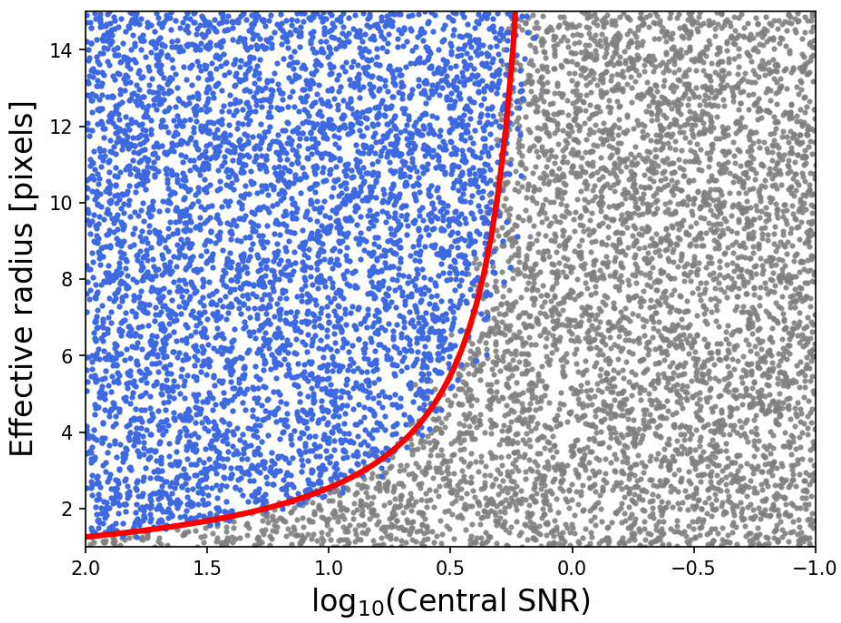

Figure 2.3: Results of the synthetic profile experiment for n=1 circular Sérsic profiles. Each point represents a synthetic source; blue ones have been detected and grey ones missed according to the criteria discussed in the text. The red line shows the boundary of detection predicted in equation 2.17. The sources presented here were not convolved with a PSF (see text).

and relatively large values of ϵ, so this effect is not considered important.

The preceding derivation of η also does not consider the effects of correlated noise. This is noise that is "clumpy", produced during image stacking (interpolation, drizzling etc.) and from sources such as faint background galaxies that have not been accounted for in the sky modelling. In particular, this correlation will tend to make the uncertainty in the number of points contained within ϵ due to the sky distribution (equation 2.9) an underestimate. The degree of noise correlation varies between datasets, making its effects difficult to quantify in general. For my current purpose of getting an estimate for η, the effects of underestimating σ_φ can be accounted for by using higher values of κ, as can be seen in equation 2.12. To estimate the degree to which κ should change to accommodate the correlation, I generated independent random noise and applied a Gaussian filter with σ=1 pixel to create noise correlated over scales of two pixels. In such a set-up, Monte-Carlo trials suggest the standard deviation in equation 2.9 is underestimated by a factor of $\sim$2.5 and thus κ would have

to be multiplied by this factor to obtain equivalent behaviour to the uncorrelated case in terms of robustness against the detection of noise peaks (this comes at the cost of sensitivity).

Despite this, the assignment of κ is likely to be done empirically rather than derived statistically because even with a source mask in place there will likely be unmasked sources contributing to a non-Gaussian background (see §2.5.3).

There are similarities between `DBSCAN` and conventional data smoothing techniques because of the size of the search radius ϵ. One major reason why very large smoothing kernels are not commonly used for detection is because nearby objects become confused with each other. Further, smoothing over bright, concentrated sources may produce detections that appear similar to LSB galaxies. By applying the source mask *before* the detection algorithm, this problem is alleviated and I can make use of the SNR obtained over larger areas without significant source confusion. `DBSCAN` is also more robust to the detection of small unmasked background objects because the input pixels are not flux-weighted; sources are forced to be significant over areas similar to the search area in order to be detected.

2.3.3 The DeepScan software

DeepScan is a Python package intended to identify regions of significant LSB light. The software uses a novel implementation of the `DBSCAN` algorithm that was created in order to operate much more efficiently than the standard. This efficiency is in-part due to many calls to integrated C code within `numpy`[5] and `scipy`[6]. One of the goals of DeepScan has been to be compatible with other pieces of software, and as such there is a lot of flexibility as to what can be input to the software in terms of, for example, user-generated background maps or object masks. Equally, the outputs of DeepScan such as segmentation maps and initial guesses on Sérsic parameters can be easily transferred and used by different tools. If however the user does not have ready-made background maps etc., the basic usage of DeepScan is as follows:

(i) Measurement of the sky distribution to produce sky and sky RMS maps (implicit source masking).

(ii) (Optional) Generation of a bright source mask (currently SExtractor is used to create masks).

(iii) Source detection on masked frames using `DBSCAN`.

(iv) (Optional) De-blend the segmentation image(s).

[5]http://www.numpy.org

[6]https://www.scipy.org

(v) Automatic measurement of detections based on segmentation image(s). Users are likely to want higher quality source measurements than the approximate ones provided with DeepScan. I envisage these measurements to serve as inputs into robust profile fitting algorithms like ProFit.

Novel `DBSCAN` implementation

An important feature of any detection algorithm is the runtime. Lower runtime helps users to fine-tune their parameters as well as enabling large data sets to be analysed over accessible (CPU) time-frames. Code optimisation usually proceeds by reducing the most significant time-consuming operation in the program; in the case of `DBSCAN` this is the region query, whereby the number of points within ϵ are counted for every data point. A simple implementation of `DBSCAN` may perform the region query by directly measuring the Euclidean distances from every input point to every other point and storing these distances in a symmetric distance matrix. This is inefficient in terms of memory as well as CPU time because every unique element of the matrix requires checking for every query. Indexing structures such as the R-tree (Guttman, 1984) are optimised for spatial queries and allow for a significant speed-up. Many `DBSCAN` implementations use this method, including those in `scikit-learn`[7] and R[8].

However, my implementation is done in a much different manner that obtains equivalent results in notably shorter time-frames through a convolution approach. The basic procedure is the following:

(i) Create a binary image where pixels above the detection threshold are assigned the value of 1 and all others 0. I term the detection threshold as `thresh`, which is quoted in units of the background RMS unless otherwise stated.

(ii) Convolve the result with a top-hat filter of unit height and radius equal to ϵ. This step essentially counts the number of thresholded pixels within an ϵ radius of every pixel.

(iii) Threshold the resulting image at η (derived from κ and `thresh`), creating a binary image with non-zero pixels being `DBSCAN` core points.

(iv) Convolve the result with the same top-hat filter as in step 2. This connects regions corresponding to unique `DBSCAN` clusters. Set all non-zero pixels to 1.

(v) Run a contiguous pixel clustering algorithm over the result. This assigns unique integer labels >0 to each `DBSCAN` cluster. The result of this is known as the `segmap_dilated`, and bounds the regions contained by all the `DBSCAN` core and

[7] http://scikit-learn.org/stable/

[8] https://cran.r-project.org/web/packages/dbscan/dbscan.pdf

secondary points.

(vi) Perform a binary erosion on each object in `segmap_dilated` to obtain the contiguous areas bounded by the core points. The result is simply called the `segmap`. A segmentation map of only the core points (`corepoints`) can also be retrieved.

We note that an "erosion" refers to contracting a source's segment with a kernel (in my case the top-hat filter of radius ϵ); For each pixel making up a source's segment, all pixels that are contained within the kernel's footprint centred on that pixel are removed from the segment. A dilation is the opposite transformation. Hence the name "`segmap_dilated`" is appropriate because it can be obtained by performing a dilation on the `segmap`. The only time an erosion is performed by DeepScan is to create the `segmap` from the `segmap_dilated`.

The speed-up from the above approach compared to standard implementations stems from the fact that a) fast-Fourier transform (FFT) techniques can be used for the convolution steps and b) The need for the `DBSCAN` region query is removed and is replaced by a much more efficient contiguous-pixel clustering algorithm.

I have tested DeepScan against the `scikit-learn` and `R` `DBSCAN` implementations (versions 0.18.2 and 1.1-1 respectively) on one processor, and find that this implementation is faster than both. The tests were performed on a mid-2013 MacBook Pro (2.5 GHz Intel Core i5) with 8GB of RAM, running OSX 10.12.6. I also note that I find the R implementation to be significantly faster than that in `scikit-learn`, but `scikit-learn` gives the option to run `DBSCAN` in parallel whereas R does not. For example, averaging over five runs for a 500×500 NGVS g-band cut-out, the times are DeepScan: 0.3s, `R`: 0.3s and `scikit-learn`: 0.8s. Enlarging the image to 1000×1000 pixels gives DeepScan: 3.5s, `R`: 6.0s and `scikit-learn`: 17.6s. Scaling up once again to 4000×4000 pixels, this time letting DeepScan and `scikit-learn` use four processors, the results are DeepScan: 12.8s `R`: 27.6s and `scikit-learn`: 89.5s. I tested whether there was a difference between the output of DeepScan compared to the other implementations and found that there was an exact match between the results.

2.3.4 Sky measurement

DeepScan can produce sky and sky RMS maps if required. To obtain an estimate of the sky I iteratively make measurements of the sky and the sky RMS, each iteration using `DBSCAN` with a low detection threshold (default `thresh`=0.5) to identify sources (including LSB components) which are masked from the sky calculation in the following iteration. Using suitable values for ϵ (default of 5 pixels - similar to a typical PSF

FWHM for wide field optical surveys) and κ (default a value of 5 - low enough to encapsulate LSB components), this iterative masking reduces the bias incurred from unmasked LSB components each time. The iterations terminate when the sky level has converged to a user-specified tolerance. I have provided figure 2.4 as an example of the sky-measurement algorithm, which was generated with default parameters.

The sky and sky RMS levels are estimated in meshes. The mesh size is a user-defined parameter. I maintain flexibility by allowing custom estimators for the measurements, but by default use the median for the sky and a lower-quantile estimate of the RMS (i.e. the level for which 15.9% of the data is enclosed below the median). These are computed for each iteration, ignoring any masked pixels. The meshes are then median filtered over a customisable scale, before being interpolated over using a bi-cubic spline to the full image resolution. Meshes with too-few pixels are ignored and interpolated over; by default at least 30% of the mesh must be unmasked to count. Following this, `DBSCAN` is run, and any pixel identified within the `segmap_dilated` is masked. The algorithm then checks for convergence on a mesh-by-mesh basis; individual meshes that have converged are ignored for further iterations and their converged values are used in the interpolation. This process repeats with the updated source mask, either until all the meshes have converged or a maximum number iterations has been reached (the default is 6).

I again emphasise that it is trivial to use sky and RMS maps generated externally from DeepScan. I also note that custom masks can be used as an input to the source masking routine, which can be combined with the mask generated with `DBSCAN` or even treated as the final mask, in which the iterative mask generation is not applied.

2.3.5 Masking bright sources

A crucial requirement of my detection method is a source mask. The mask must be created with the aim of eliminating all sources one does not wish to detect right out to their LSB halos. My approach here has been to use SExtractor to create the source mask. Measurements from the output catalogue such as the `FLUX_RADIUS` were found to do a poor job, underestimating the source sizes even with high values of `PHOT_FLUXFRAC` (the fraction of light contained within the flux radius). This prompted me to model each source with a Sérsic profile and to size the ellipse according to some isophotal radius below the DeepScan detection threshold.

I estimated the Sérsic index without performing any additional fitting from the ratio between the effective and Kron radii (Graham & Driver, 2005). This is use-

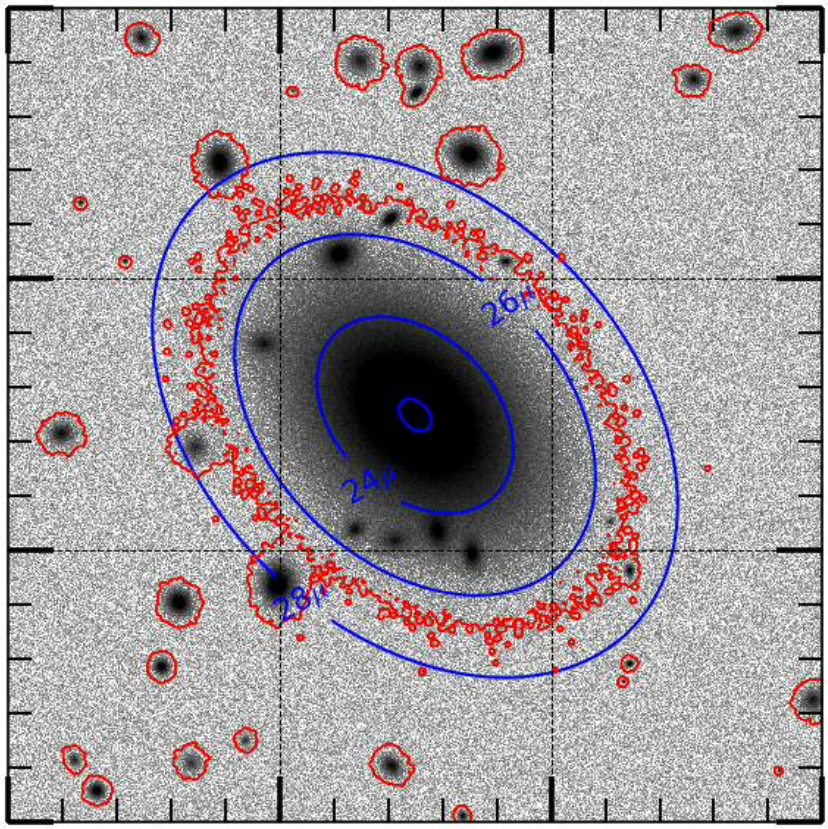

Figure 2.4: This shows a synthetic 1500×1500 pixel image, with an extended central source and some smaller "background" objects. Contours of the surface brightness profile are shown (blue) for the central profile, as well as a contour showing the masked region (red). Clearly the mask extends to very low surface brightness (the RMS of the image is 26.2μ). The 500×500 pixel meshes (black dashed lines) were median filtered over 3×3 regions. The actual sky level of the image was set to zero ADU. Despite the presence of the dominant LSB object in the frame, the maximum value of the sky was measured to be 28.4μ - well below the RMS level. The RMS was recovered with an error of less than 1%.

ful as both of these measurements can be efficiently retrieved by SExtractor with the `KRON_RADIUS` and the `FLUX_RADIUS` keywords and `PHOT_FLUXFRAC=0.5`. Combining these measurements with the total magnitude (also measured using the SExtractor`MAG_AUTO` parameter, by default with `PHOT_AUTOPARAMS` set to `2.5,3.5`), the profile is fully characterised. The source is then masked in an elliptical aperture

(based on SExtractorś elliptical parameters) to the derived isophotal radius. I note that SExtractor allows the possibility to perform Sérsic fitting by requiring the `SPHEROID_SERSICN` or `SPHEROID_REFF_IMAGE` columns in the output file. While I have not explored this in my current work, it is possible that this would improve the fits at the expense of some computation time.

A caveat of the above approach is that galaxies typically have non-elliptical LSB components and therefore may not be adequately masked. ProFound could potentially improve over SExtractor for the source masking because it offers non-parametric object masks and more reliable estimates of parameters such as the half-light radius (Robotham et al., 2018).

2.3.6 Source measurement

The goal of DeepScan is not to provide accurate profile fitting, but is rather to identify regions with significant LSB light. That said, I do provide a basic function for 1-dimensional Sérsic profile fitting in order to get initial estimates of parameters to input into robust 2D fitting programs such as `ProFit`[9]. The basic requirements for the fit are the data and a segment corresponding to the source.

The initial task that is performed is the estimation of the centroid position and elliptical parameters (axis ratio and position angle). This is done using the same method of flux-weighted moments as in SExtractor (see Bertin & Arnouts (1996) for detail), where the user has the choice to calculate the moments on either a masked or unmasked segmentation map. The user can choose which of the three segmentation maps provided (`segmap`, `segmap_dilated` or `corepoints`) to calculate these parameters.

The next step is to measure the average (default median) flux within concentric elliptical annuli of fixed width centred on the source, based on the measurements from the previous step. The annuli iteratively increase their radius until a user-defined isophotal surface brightness is reached (often a relatively robust method of measuring large LSB galaxies) or a maximum radius has been reached. If the isophotal level is reached before sufficient steps have been performed then more steps will be taken in order to achieve a minimum number of data points (default is 5).

I then proceed to fit the profile using `Scipy`'s `curve_fit` routine, which by default uses the Trust Region Reflective algorithm for parameter optimisation in the case of constrained problems (each Sérsic parameter is constrained by default to have positive

[9] `https://github.com/ICRAR/ProFit`

values). An initial parameter guess can be provided, but in its absence the parameters are estimated as following: The index n is assumed as 1; The effective radius is assumed as the semi-major axis of the ellipse that bounds the segmentation map; the surface brightness at the effective radius is estimated by measuring the average surface brightness within the segmentation map, rescaling to effective surface brightness using the default value of n.

I have not yet implemented methods to measure non-galaxy like objects. However, I suggest the segmentation images outputted by DeepScan can be used as inputs to non-parametric measurement tools such as that offered in `ProFound` to provide estimates on parameters like total magnitude etc.

2.4 DeepScan vs Source Extractor

During my testing of SExtractor I have found that it can perform fairly well in detecting LSB features provided specific input parameters are used. In this section I describe some observations about its usage and explain why DeepScan may be preferred to detect specifically highly extended LSB objects. My tests have consisted of using a cut-out from the publicly available NGVS g-band data, which has a pixel size of 0.186″ and typical RMS of $\sim$26.9$g\mu$ with various combinations of SExtractor settings. For the experiment, both SExtractor and DeepScan used the same background mesh of 50″ ($\sim$270 pixels) that was median filtered in 3×3 meshes (Using the `BACK_SIZE` and `BACK_FILTERSIZE` SExtractor keywords). The cut-out used had a size of 810×810 pixels so that the background estimation was realistic. For this experiment, I have used SExtractor to convolve the image with a Gaussian kernel of 5-pixels RMS, as in Greco et al. (2018b).

It is thought that the de-blending can routinely fragment large LSB features (Davies et al., 2016; Greco et al., 2018b) (aka "shredding"). Indeed I have observed the de-blending fragmentation (figure 2.5a), which was observed with SExtractor's de-blending contrast parameter `DEBLEND_MINCONT=0.005` (the default). A way around the problem is to deactivate the de-blending by setting `DEBLEND_MINCONT` to `1` - this value is used for the remainder of my tests. I have not experimented with different values of `DEBLEND_MINCONT` in this work. I also note that it is important for the `CLEAN` parameter to be switched on (default) in order to reduce LSB source fragmentation, so it is activated in all my tests here.

Figure 2.5b) shows the result of increasing the value of `DETECT_MINAREA` (the minimum number of pixels required for a detection to count) compared to the result shown

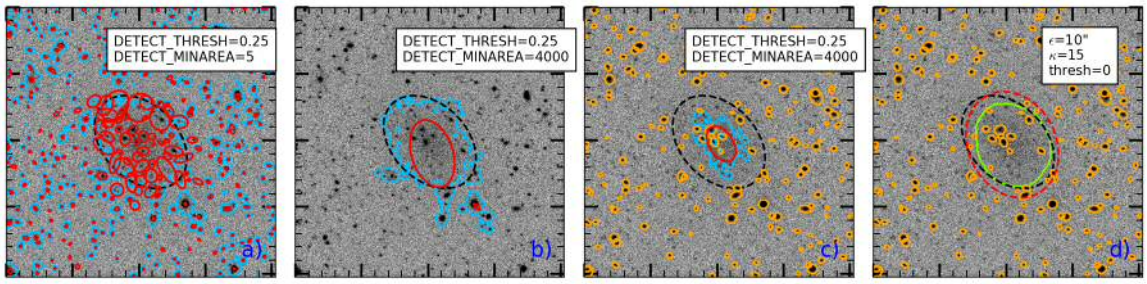

Figure 2.5: This is a real 810×810 pixel NGVS data cut-out with a synthetic LSB galaxy with a convolved central surface brightness of $\sim 26.9 g\mu$ and an exponential profile (n=1) of effective radius 30″ (roughly equivalent to the lower size limit of a UDG at Virgo). The black dotted ellipse represents this source out to one effective radius. In frames a) to c), the light blue contour traces objects on the SExtractor `SEGMENTATION` check plot. The red ellipses represent SExtractor estimates of the effective radius for each detection, measured in the same way as in §2.3.5. In figures c) and d), the orange ellipses bound the sources that are masked. In figure d), the green contour traces the DeepScan `segmap`, whereas the dashed red ellipse represents the effective radius as measured by `DeepScan`. See text.

in figure 2.5a). A much better job is done of identifying the LSB source as a single object, but I note that the detection suffers two problems: the shape of the segment corresponding to the LSB source is significantly perturbed by background objects; and spurious detections still exist around groups of background objects despite very high values of `DETECT_MINAREA`. Activating the de-blending here exacerbates the situation and the LSB source is missed entirely, with its flux being solely attributed to some of the background objects rather than any central object.

Herein lies a downside of using SExtractor to detect LSB sources. The areas of the segmentation image corresponding to LSB objects detected on smoothed frames is made significantly unstable because of the presence of background objects that have either not been properly de-blended or have been erroneously assigned to the source in the cleaning stage. This is made clear by the morphology of the detection in figures 2.5b) and c). As this significantly effects the number of pixels an object contains, the usage of `DETECT_MINAREA` becomes an inherently unreliable tool to identify genuine LSB objects, yet is required to discriminate against background objects.

The problem is partially alleviated by applying a source mask to the image before its input to SExtractor as is also done by Greco et al. (2018b) (note that there

is no source mask handling within `SExtractor`) in that now the only detection that appears in the output catalogue seems to be associated with the LSB object itself rather than brighter objects in its vicinity. However, the problems associated with the LSB segment still exist - the segment is irregular and contains several unmasked background objects that significantly perturb it's shape. It is also notable that the elliptical fit (determined by `SExtractor`'s half-light `FLUX_RADIUS` and elliptical parameters) does not do a reasonable job at measuring the object. This size underestimation can lead to the source being missed, as it is typical for authors to perform a cut on the minimum size of objects.

In contrast, the DeepScan detection (figure 2.5d) is much smoother and does a better job of tracing the shape of the LSB structure. It is arguable that this is because I used a large value of ϵ ($10''$=50 pixels), but I note that DeepScan is designed to work with such large kernels. The work of Greco et al. (2018b) have also shown that using much larger kernels than $1''$ is not feasible in SExtractor because of blending with unmasked faint/background galaxies (Koda et al., 2015; Sifón et al., 2018), although this could be improved with a better masking strategy. I also note in passing that SExtractor does not allow kernels larger than 31×31 pixels, so it is impossible to use such a kernel from `SExtractor`. A second point of consideration is that the core points of the DeepScan detection define its shape, and these are identified as pixels with relatively high SNR on the original (i.e. non-smoothed) frame. Further, unmasked sources that cause the perturbations in the SExtractor segmentation map have less of an effect on the shape of the DeepScan detection because the pixels aren't flux weighted in the `DBSCAN` algorithm.

To summarise, while it might be preferable to use DeepScan to trace extended LSB light it is also possible to use SExtractor to detect very low surface brightness objects, provided certain criteria are met:

(i) A ready-made source mask is provided.
(ii) Large smoothing kernels are used (we note that the largest kernel size acceptable is 31×31 pixels).
(iii) Large values of `DETECT_MINAREA` are used.
(iv) De-blending is deactivated.
(v) One treats parameters derived from the `SEGMENTAION` check plot with some scepticism, particularly the `FLUX_RADIUS` as this seems to be systematically underestimated for large, diffuse objects.
(vi) Cleaning is on (`CLEAN=Y`) to avoid spurious detections.

2.5 Application to the NGVS

To demonstrate the DBSCAN algorithm I applied it to a subset of the publicly available NGVS data that I acquired from the Canadian Astronomy Data Centre[10]. This data was taken with the square-degree MegaCam instrument on the Canada France Hawaii Telescope, and covers ∼100 square degrees of the Virgo Cluster in the u, g, r, i, z bands. The NGVS was chosen because it offers deep imaging of the Virgo cluster at high resolution (0.186″ pixels), and was the same data used by Davies et al. (2016) with which I wish to compare. I use the g-band data as it has the best coverage, a low maximum seeing FWHM (1″) and an extended-source limit of 29μ (Ferrarese et al., 2012). The subset covers a five-degree^2 area projected radially eastwards from the centre of the Virgo cluster (i.e. M87). The subset is made from five overlapping frames, each covering an area of 1 degree^2 with corresponding sizes of 21000×20000 pixels. The frames are each 1.74Gb in size. This area overlaps with part of the region explored by Sabatini et al. (2005) so comparisons can also be made with their work. Objects detected in this region likely belong to sub-cluster A, the largest sub-cluster in Virgo (Mei et al., 2007).

My general strategy is to use SExtractor to identify sources for masking, before using DeepScan to search the remaining area for LSB objects. The following processes were performed using four 2.5 GHz Intel Core i5 processors with 8Gb of RAM. The overall pipeline I used is as follows:

(i) SExtractor source masking.
(ii) Sky modelling.
(iii) Source detection.
(iv) Source selection.
(v) Human validation.
(vi) Sérsic fits with ProFit.

2.5.1 Source masking

The first stage in the mask generation was to run DBSCAN over the raw data in order to identify saturated stars and their associated LSB halos. This is done because the SExtractor masks generated for such objects were not sufficient to cover the sources. The parameters I used were ϵ=10″ (∼50 pixels), thresh=0, κ=20. These are similar to those that I used for the actual LSB detection, but were modified based on trial and error masking of large saturated stars. Any detection that contained a saturated

[10]http://www.cadc-ccda.hia-iha.nrc-cnrc.gc.ca/en/

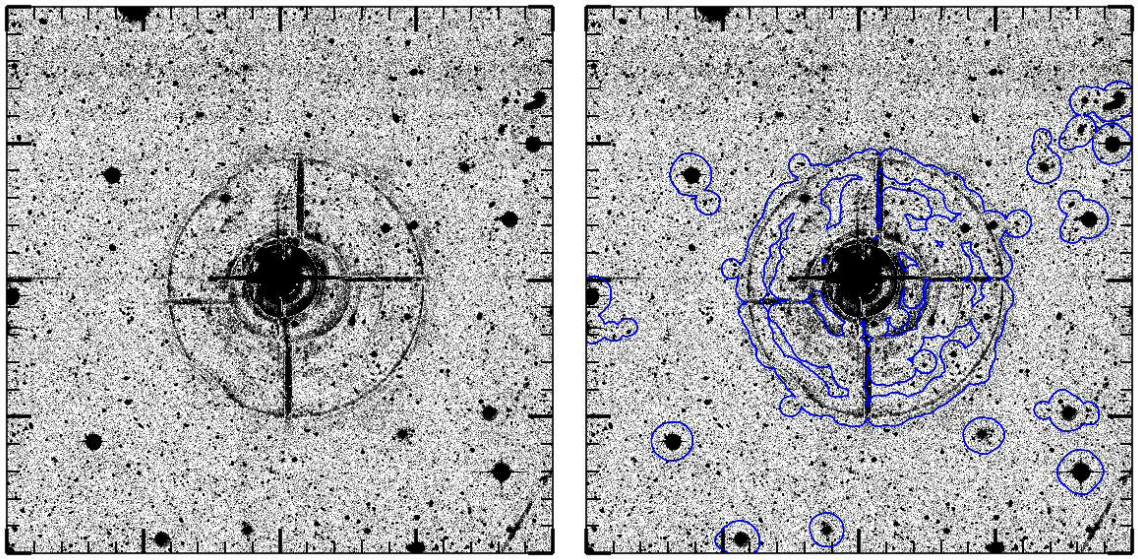

Figure 2.6: Saturated star masking. The masked regions are contoured in blue against a smoothed cut-out of NGVS g-band data. The image size is 4000×4000 pixels, or ∼12.5×12.5 arcminutes. The large mask in the centre successfully masks out the LSB ring around the saturated star and I find no detections in its vicinity.

pixel within its `segmap_dilated` was masked within it. An example of the result of this saturated star masking is shown in figure 2.6.

The masked regions were set to zero and this data was used as the input to `SExtractor`. For this I used a `DETECT_THRESH` of 6 (see §2.5.3) and convolved the image with the default filter (a 5×5 pixel Gaussian filter of FWHM 2 pixels). This makes me sensitive to sources with surface brightnesses $\geq\sim 25 g\mu$ for the final `DBSCAN` run. I disabled de-blending in order to prevent the fragmentation of sources close to the detection threshold as this produced poor masks in their vicinities. I allowed SExtractor to perform its own background and RMS estimates in small meshes of size 64×64 pixels to better detect smaller sources against their local background. All other parameters were left to their defaults. The isophotal radii were then calculated as described in §2.3.5 for the $29 g\mu$ isophote.

In a minority of cases, the initial SExtractor mask did not cover the full source. This was particularly true for bright galaxies with extended LSB halos and bright unsaturated point sources where the approximate Sérsic fits were inadequate. Requiring a fainter masking isophote did not solve the issue, so I was forced to enlarge the apertures by a factor of 1.5 for these sources (this was determined empirically).

Each mask took approximately 30 minutes to generate. On average, 28 percent of

each frame was masked excluding the border regions.

2.5.2 Sky modelling

We use the source mask as an input to DeepScan's sky modelling routine. I add to the map by masking sources detected by `DBSCAN` with the default parameters described in §2.3.4; this allows sources to be masked to well below the sky RMS level. I only do one iteration of the DeepScan detection/masking process as the mask is already quite complete. The sky itself was measured in 500×500 pixels and is median filtered over 3×3 pixels. The large sky-mesh along with the median filtering increases my robustness against bias in the sky measurement due to unmasked LSB haloes. A second advantage is that 500 pixels (∼90″) is large compared to the sources I am searching for so they should not be significantly subtracted with the sky.

2.5.3 Source detection

The relevant parameters for the segmentation map generation using `DBSCAN` are the detection threshold, the search radius ϵ and κ. My approach for setting the parameters was to perform empirical tests on a field image (this is an NGVS frame of an area of sky displaced from Virgo where we expect a low density of LSB objects). The detection threshold was set to 0.5 times the RMS ($\sim 27.7 g\mu$) because lowering it much further made it sensitive to image artefacts such as background defects. An ϵ value of 10" (∼50 pixels) was used as this is the smallest aperture that can critically sample an UDG at Virgo distance (UDGs have minimum size of 1.5 kpc van Dokkum et al. (2015) so at 16.5Mpc this is ∼20″). With these settings I can expect to detect UDGs with average SB within their effective radii of $\sim 28 g_\mu$. I ran the field image through the overall pipeline for several values of κ and SExtractor detection thresholds (for the mask generation), with results shown in figure 2.7.

From the figure, it is clear that the results converge for high values of κ. I select a value of κ=32.5 based on this plot by requiring less than 5 detections on the field image. With regards to the SExtractor detection threshold, I was interested in a value that was low enough in order to have a reasonably low number of contaminant objects while being high enough not to partially mask out LSB sources with shredded detections. I adopt a `DETECT_THREHSH` of 6 to lower the number of contaminant sources. Note that I did not probe lower values than 6 for the masking because this encroaches too far into the LSB regime, with surface brightnesses fainter than ∼25 magnitudes per square arcsecond.

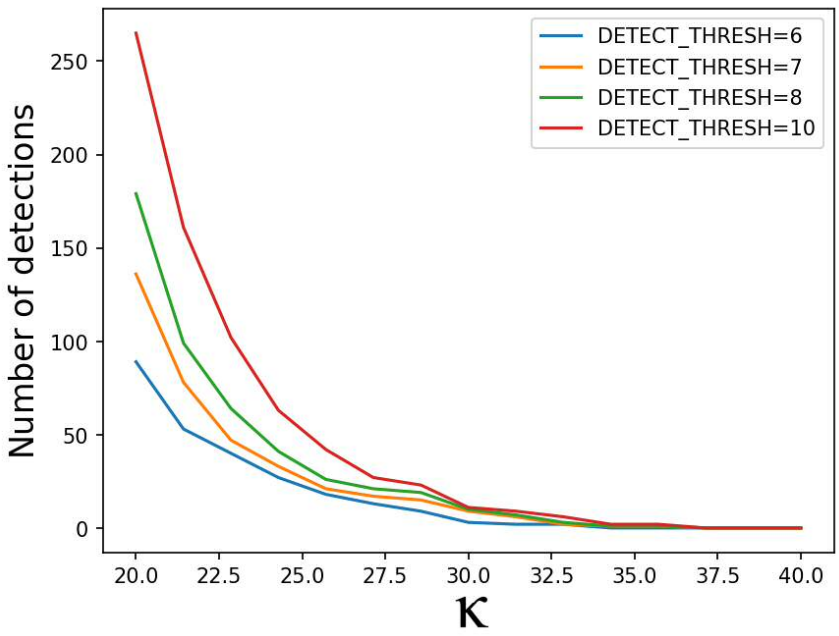

Figure 2.7: DeepScan parameter tuning. Here I show the effect of varying the SExtractor detection threshold `DETECT_THRESH` and κ parameter on the number of detections on a background NGVS frame. The number of sources detected begins to level off for $\kappa \geq 30$ for each value of `DETECT_THRESH`.

There is actually a significant difference in the background RMS level between the frames which makes it sensitive to different surface brightnesses from frame to frame. When I ran all the frames with the same settings as above, I found there was much more contamination of the output sample from spurious sources on some frames compared to others. I therefore normalised the threshold for each frame to the absolute surface brightness corresponding to that which was used on the field image, with settings as above. This is because lowering the threshold (in SNR) for frames with relatively high background RMS increases η sufficiently to protect against the spurious detections.

After setting the parameters, `DBSCAN` was run and took approximately 12 minutes per frame. In total, 67 objects were detected.

2.5.4 Detection analysis

Each source was assessed visually in order to determine whether it was an astrophysical LSB source or miss-detection. I define miss-detections to encompass data artefacts (such as stellar diffraction rings and satellite trails) and the real LSB component of

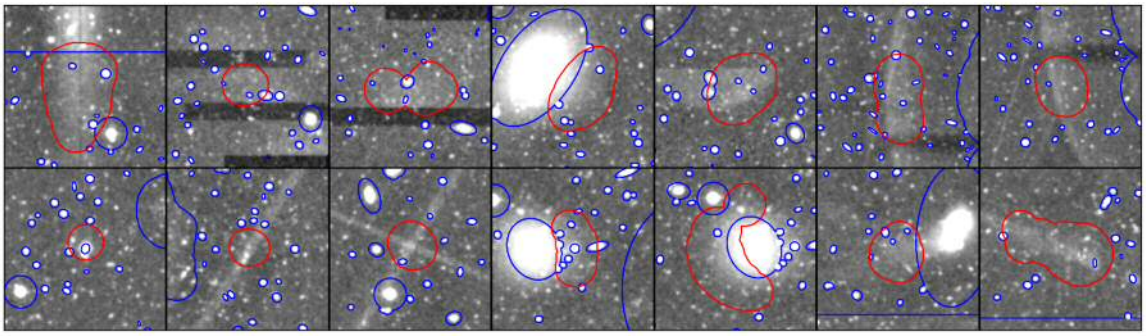

Figure 2.8: Smoothed g-band cut-outs of the rejected detections. The maximum brightness of the dynamic range is approximately 25.5 magnitudes per square arcsecond. The red contour show the dilated segmentation maps produced with `DBSCAN` and the blue contour traces the mask.

bright sources that have been inadequately masked. 14 of the raw detections were deemed to be miss-detections, leaving me with a sample of 53 objects. Of the rejected sources, 5 were associated with bright objects, 3 were unmasked stellar halos, 3 were satellite trails, 2 were caused by artefacts from the data stacking procedure and one was an extended LSB bloom caused by a bright source outside the FOV. The rejected sources are shown for clarity in figure 2.8.

The remaining sources were cross-matched with the VCC (Binggeli et al., 1985), LSBVCC (Davies et al., 2016) and Sabatini et al. (2005) catalogues, using a search radius of 20″ (chosen so large to account for positional uncertainty in other surveys). 23 of the sources had matches, leaving a sample of 30 new LSB galaxies.

I used DeepScan to fit 1D Sérsic profiles to the detections ignoring masked regions. I used the `segmap_dilated` to estimate elliptical parameters and the centroid positions. 500×500 pixel cut-outs were obtained from the original data, the sky and RMS maps as well as the source mask and dilated segmentation map, centred on these centroids.

The cut-outs were then used as inputs to the 2D Bayesian profile fitting packageProFit, with initial parameter guesses given by the 1D fits. I follow the methodology suggested by the ProFit team[11], which consists of a three stage fitting process. First, a BFGS gradient decent fit is obtained. The results from this are then used as the initial parameters for a Laplace approximation using the method of Levenberg-Marquardt (LM). Finally, the results are used as the initial guesses for a more robust MCMC fit using the component-wise hit-and-run metropolis (CHARM) algorithm,

[11]http://rpubs.com/asgr/274695

with 1000 iterations. In the fitting I used a simple Gaussian PSF of FWHM 1″.

The residuals for each fit were judged by-eye to ensure they were reasonable. In general they were, but for 8 of the sources I had to slightly modify the mask in order to get a good fit. Taking the standard deviation of the posterior distributions for each parameter results in uncertainty underestimates, likely because the high-quality of the initial parameter estimates and limited number of iterations mean only a narrow region of parameter space can be explored. To get a more realistic error, the upper and lower range of the posterior distributions were used as the parameter uncertainties. While a better estimate of the uncertainty could be obtained by running the MCMC for longer, the fits already took approximately one hour for each galaxy and the difference would be of little consequence for the present analysis.

In figure 2.9 I show the cut-outs for my final sample that contains new detections and those that had a match only in the catalogue of Sabatini et al. (2005). This is done because this catalogue does not cover the whole of Virgo in the same way that the VCC and LSBVCC do. These matches are of genuinely very low surface brightness and are obtained from a different dataset using a matched filter approach, making their re-detection a good coincidence test. I denote the names of galaxies in my final sample that had matches within the Sabatini catalogue with an asterisk after their name. On the figure I also show the elliptical annuli corresponding to the ProFit fits out to the effective radius. Note that the galaxy VLSB23 does not have any measurements due to its highly unusual morphology, an odd over-density of point sources superimposed on a LSB fuzz, and may be worth investigating further.

2.5.5 Results

I plot the effective radius (r_e) vs the mean surface brightness within the effective radius ($\langle\mu_e\rangle$) based on my ProFit models in figure 2.10 for the final sample (includes matching Sabatini et al. (2005) sources). On the plot I also show the LSBVCC sample. However, the measurements presented in Davies et al. (2016) are in central surface brightness and exponential scale size units, and all assume a Sérsic index of 1. To try and estimate the scatter this introduces on the r_e - $\bar{\mu}_e$ plane, I take their initial results and calculate the relevant parameters using Sérsic indices randomly generated based on my sample ($\langle n\rangle = 1.0 \pm 0.4$). I also plot the theoretical `DBSCAN` upper detection boundary assuming $n = 1.4$ (i.e. 1σ above the mean) which is consistent with my findings.

For context, I also show the selection criteria used by van der Burg et al. (2017)

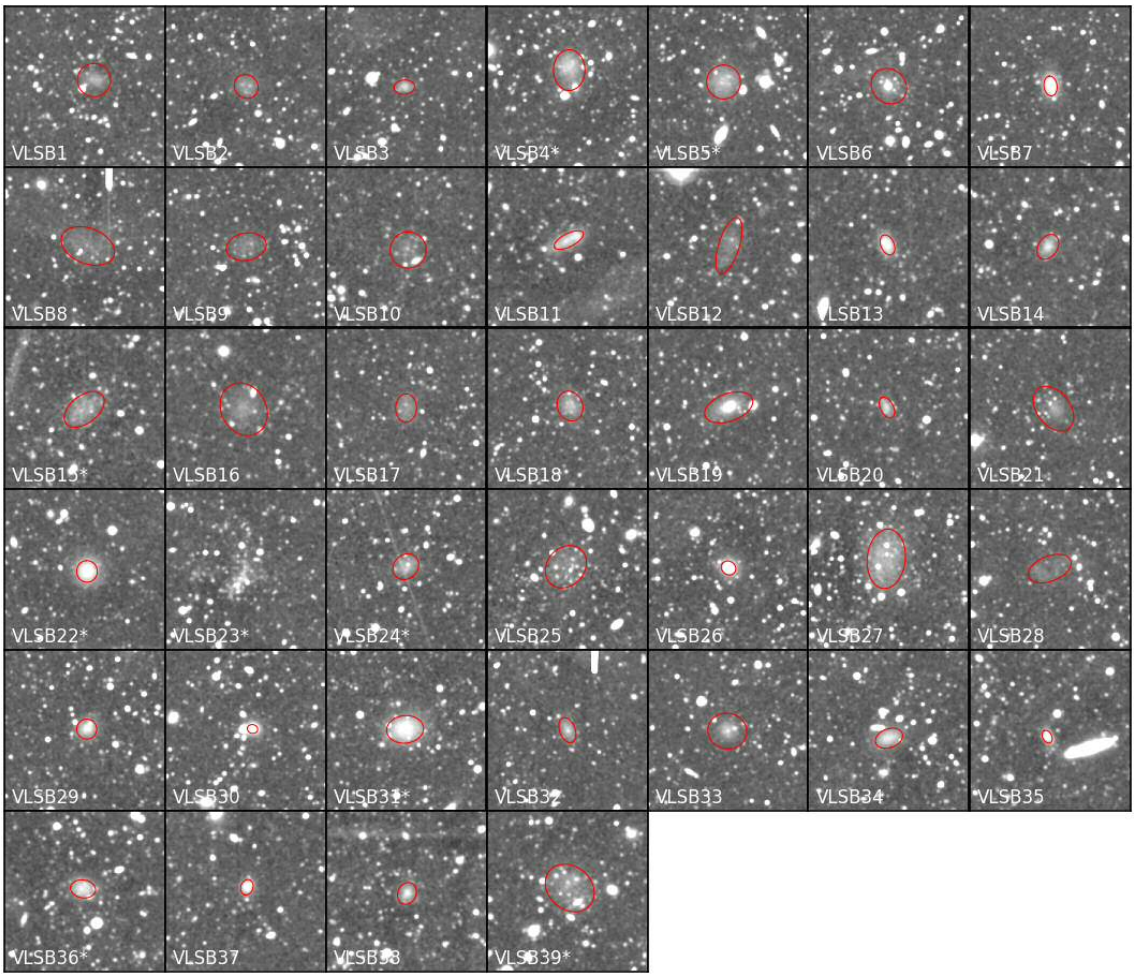

Figure 2.9: Smoothed g-band cut-outs of my new detections. While some of the sample appear reasonably bright, the maximum brightness of the dynamic range is approximately 25.5 magnitudes per square arc-second, which is close to the faint end limit of detection for the LSBVCC. The red ellipses show the ProFit models out to the effective radii. Sources with an asterisk following their name are also present in the catalogue of Sabatini et al. (2005).

in the figure, who used MegaCam imaging in the search for UDGs around groups and clusters. Further, I plot the sample of Yagi et al. (2016), who obtained a catalogue of LSB galaxies in Coma with deep Subaru-R Suprime-Cam imaging, using their single Sérsic GALFIT (Peng et al., 2002) fits. These results have been mapped to Virgo g-band data by assuming a Virgo distance of 16.5Mpc and a Coma distance of 99Mpc. The Subaru-R to g conversion was done using a fiducial $(g-r)$ value of 0.45 (Roediger et al., 2017).

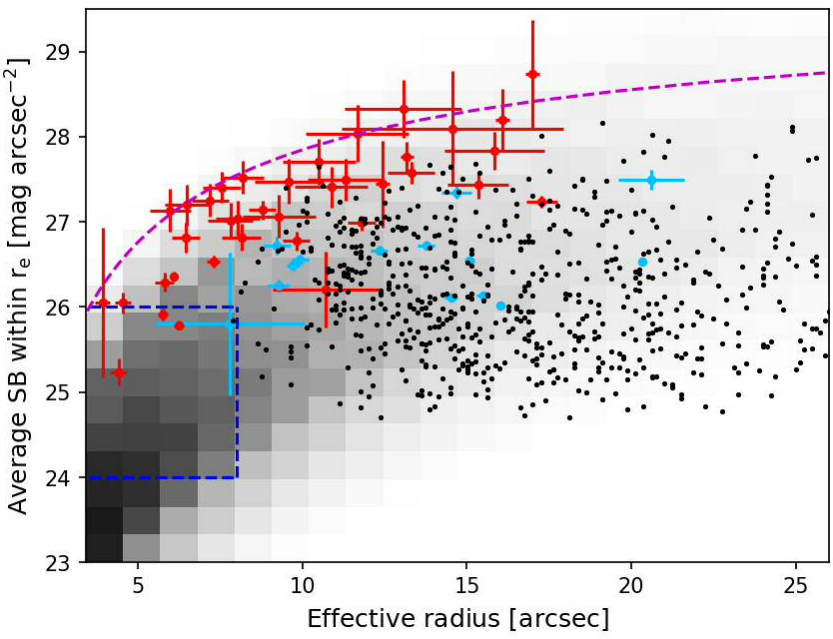

Figure 2.10: Effective radius vs mean surface brightness within the effective radius for the new sample (red) and the extrapolated LSBVCC data (greyscale heatmap, see text). The light blue points represent the complementary sample. The blue box is the selection criteria used in van der Burg et al. (2017), while the purple dashed line is the theoretical upper limit of detection given my `DBSCAN` settings. The black dots represent the catalogue of Yagi et al. (2016) projected at Virgo distance.

It is clear that the sample in this work represents an extension of the parameter space explored by other surveys towards the very low surface brightness regime, with $\bar{\mu}_e >\sim 26.5$. It is perhaps surprising that no larger LSB objects were found and this may be in part due to the initial background subtraction performed on the public NGVS data, which is done over scales of 20″. I am hesitant to draw conclusions from this until I have a more complete sample, and completeness estimates, which I intend to acquire in a follow-up study.

On the figure I also show the complementary sample, that is the sources that I detected but had matches in the VCC or LSBVCC. Two of these, VCC1331 and VCC1882, have measured effective radii larger than 20" and therefore may warrant reclassification to UDGs from their original classification of dwarf ellipticals.

I measured the $(g-i)$ colours of the sample in elliptical apertures out to the effective radii measured in the g-band with `ProFit`. For these measurements, I ignored the masked pixels, the results of which are shown in figure 2.11. For the i-band data I again used the publicly available NGVS data (Ferrarese et al., 2012). I also show measurements from the VCC and LSBVCC obtained by Keenan (2017). The general trend follows the Virgo red-sequence with fainter sources having redder colours and a flattening of the colour towards the faint end, as observed by Roediger et al. (2017). Many of my final sample are consistent with this picture, but there are exceptions. Most noticeably, a collection of sources seems to depart the red sequence at the faint end, in favour of lower $(g-i)$ values. Note that this trend was still observed when recalculating the colours taking into account masked pixels. Two of the sources have unusually high values of $(g-i)$. The source with the largest value (VLSB30) may have a biased colour due to its proximity to a star and the second, (VLSB19), seems to be associated with a large nucleated source.

Despite the red colours of the sources, they are generally better-detected in the NGVS g-band because fiducially the RMS level of g is ~ 1.2 magnitudes per square arc-second fainter than that in i.

I briefly note that 10 of the 14 rejected detections have $(g-i)$ colours below the minimum measured from my final sample, as can be seen in the figure. It may therefore be possible to increase the purity of the output automatically by applying a colour selection for future surveys.

Figure 2.12 shows the effective radii, both in units of arc-seconds and kpc (at the Virgo distance of 16.5±1.1Mpc, Mei et al., 2007) against the stellar mass calculated using the empirical relation derived by Taylor et al. (2011). The galaxies have a mean (logged) stellar mass of $10^{6.3\pm0.5}M_{\odot}$, making them fairly less massive than the sample of UDGs presented in van Dokkum et al. (2015), which have a median stellar mass of $6\times10^{7}M_{\odot}$. Note that if the colours are measured without their source masks in place, the average stellar mass rises only slightly to $10^{6.4\pm1.0}M_{\odot}$. There is an outlier in the plot that corresponds to VLSB30 which is in proximity to a star. It is likely that the colour has been considerably effected by the star such that the stellar mass estimate may be erroneous.

On the figure I have also plotted estimates of the stellar masses of the Yagi et al. (2016) sample projected at Virgo. Clearly there are several uncertainties in this procedure, but to attempt to get a representative picture I have randomly generated a set of data in which uncertain parameters have been perturbed within their errors, including the original error estimates from the `GALFIT` models as well as uncertainties

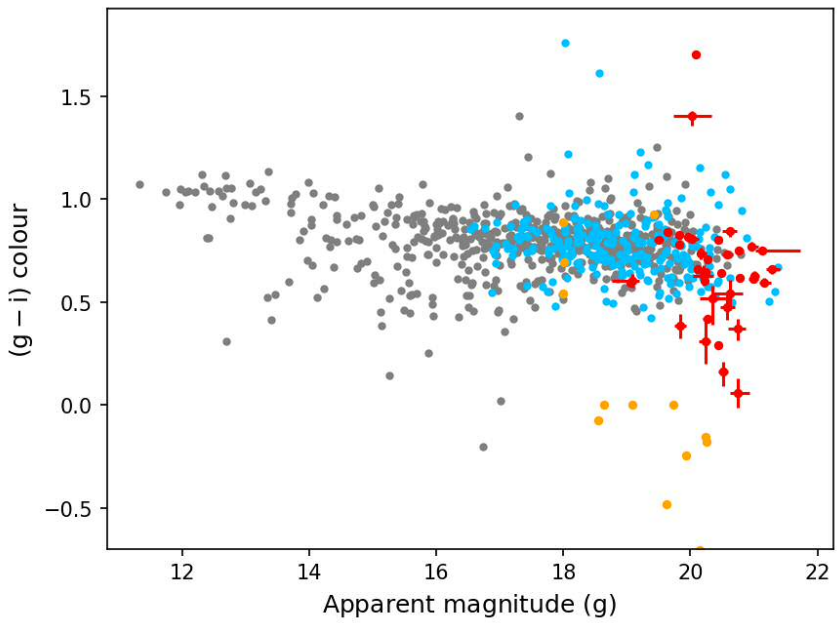

Figure 2.11: Colour-magnitude diagram for galaxies in the VCC (grey), LSBVCC (blue) and the new sample (red). Many of the new sample are consistent with the Virgo red sequence, but there is a non-negligible sample occupying bluer colour space. The rejected detections are shown in orange.

in distance and colour (we used $(g - i) = 0.7 \pm 0.2$ based on figure 2.11). My final sample seems to be both smaller in terms of size and also stellar content compared to their sample. It is interesting that some of my sources that matched with the VCC/LSBVCC agree well with the projected distribution from the Yagi sample, as it suggests that a re-inspection of the VCC/LSBVCC may result in the reclassification of some objects to UDGs.

None of the final sample are larger than the 1.5 Kpc lower limit required for UDG classification; there is a notable dearth of large LSBs. Given their sizes and low stellar content, I classify them as ultra-faint dwarfs (UFDs). It could also be that the UDG population is already present in the catalogues of the VCC and LSBVCC but has not been explicitly identified as such; an idea supported by the fact that two of the galaxies in the complementary sample likely meet the UDG criteria. I note that the original NGVS background subtraction over 20" may have the effect of causing my measurements of the sizes of galaxies to be underestimates.

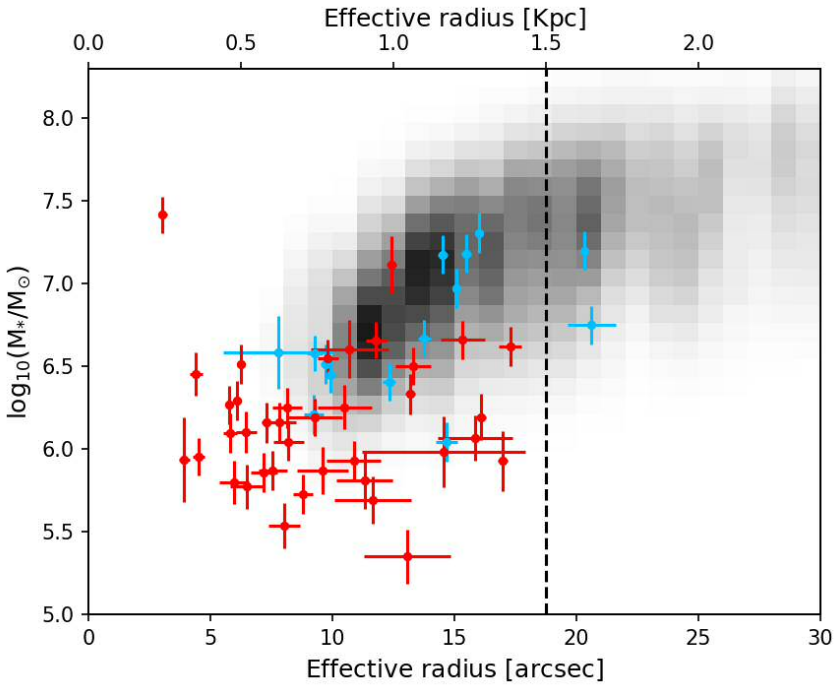

Figure 2.12: Stellar masses vs effective radius for the final sample (red) and complementary sample (blue) assuming a Virgo distance of 16.5 Mpc. The heat map is the extrapolated data of Yagi et al. (2016) projected at Virgo distance (see text). The black dotted line represents the size of the background estimation kernel in the original data reduction.

2.6 Summary

In this chapter I have introduced a new software package that I have used to detect low surface brightness features in wide area survey data. The software is capable of measuring the background distribution and producing source masks, currently based on SExtractor catalogues. The major novelty of DeepScan is that it uses a highly efficient implementation of the `DBSCAN` algorithm to detect LSB features using much larger search radii than has been done before, allowing for the detection of extremely faint extended sources.

As with any detection process, there is a trade-off between completeness and the purity of the output sample. In `DeepScan`, this is controlled by setting the `DBSCAN` input parameters: the clustering radius ϵ, the confidence parameter κ and the detection threshold `thresh`. In general, larger values of ϵ allow for fainter objects to be detected,

but using excessively large values may result in source confusion. κ must be chosen high enough to protect against spurious detections of e.g. groups of background point sources, but setting it too high may result in unacceptably low completeness.

The purity of the output sample is dependent on the quality of the source mask. Creating a good mask is fairly difficult because of the problems involved in masking out LSB components associated with bright sources that do not adhere to elliptical profiles. My current approach is to use SExtractor to detect bright sources and mask out to an isophotal radius derived based on fitting a Sérsic profile using outputs from the SExtractor catalogue. This technique is successful in the majority of cases, but is not perfect. In the example application, I masked the LSB components associated with saturated regions using the `segmap_dilated` produced using DBSCAN. A disadvantage of this is possible source confusion, but is favourable because of its ability to mask large LSB features of arbitrary shape. One promising future approach to creating the source mask could be to use the dilation until convergence approach used by `ProFound`, which can trace objects of arbitrary shape and thus provide non-parametric source masks.

In the application to the NGVS data, the κ value was chosen by measuring the number of objects detected on the frame as a function of κ and choosing a value which had a low number of detections. Using such a high value of 32.5 means I have been limited in my capability to fully exploit `DBSCAN` because of the need to mitigate against contaminant sources in the output sample. Even with such a high value, I still reject 14 out of 67 detections, which consist mainly of satellite trails and unmasked regions associated with bright objects such as saturated stars. It is conceivable that some of these objects may be removed automatically using a colour analysis in future surveys on a larger scale.

Of the remaining 53 sources, 30 do not have matches in either the VCC, LSBVCC or Sabatini catalogues. Keeping the Sabatini sources, I am left with a sample of 39. These measurements have ranges between $26.0 \leq \bar{\mu}_e \leq 28.5$ and $19 \leq m_g \leq 21$ following fitting of Sérsic profiles with `ProFit`. Of this sample, none are large enough to be classified as UDGs and I classify them as UFDs (assuming cluster membership). My current evidence for cluster membership is that they are reasonably consistent with the colours of other Virgo galaxies, and have angular sizes larger than the optimal selection criterion of >3" for Virgo galaxies. Assuming cluster membership, the galaxies have very low stellar masses, with an average of $10^{6.3\pm0.5} M_{\odot}$.

Comparing my final sample with those from other surveys, I find that I have probed a different region of parameter space, characterised by very low stellar mass estimates

and surface brightness. I hypothesise that the dearth of larger detections stems from the initial background subtraction performed on the publicly available NGVS data. Following measurements of two galaxies VCC1331 and VCC1882, it is further hypothesised that some of the UDG population of Virgo may be contained within existing catalogues such as the VCC and LSBVCC.

I have not made any efforts to estimate the completeness of my new sample. In future work[12], I plan to perform a similar analysis on the whole of the NGVS. One method to quantify the completeness is to inject artificial sources into the data and measure which ones I am able to recover in a similar way to what has been done by van der Burg et al. (2017); only then would I be able to draw robust astrophysical conclusions based on my findings. The main conclusion from the experiment I have performed here is that I am able to detect new LSB features in areas that have specifically been searched for them before, which are some of the most diffuse detected in Virgo and reside in a different region in parameter space compared with those in the VCC and LSBVCC catalogues.

[12] see also chapter 3.

Chapter 3

Ultra-diffuse Galaxies in the Field

3.1 Introduction

An obvious question arising from the detection of the faint LSB Virgo galaxies discussed in chapter 2 is how many similar objects exist in the Universe. If their number is high enough, they could make up a significant proportion of the missing baryons in the cosmological baryon budget. The related discovery of a surprising abundance of large LSB galaxies (i.e. UDGs) in the Coma Cluster by van Dokkum et al. (2015) has prompted systematic searches from them in group and cluster environments (e.g. van der Burg et al., 2016; Román & Trujillo, 2017b; van der Burg et al., 2017; Mancera Piña et al., 2018). The number density of UDGs in groups and clusters are presently not thought to be high enough to warrant a significant steepening of the stellar mass function at the low mass end (Jones et al., 2018). However, one issue that has been brought up by these studies is whether UDGs are preferentially formed inside, as found by van der Burg et al. (2017), or outside (cf. Román & Trujillo, 2017b; Mancera Piña et al., 2018) of galaxy clusters. If it is true that UDGs form preferentially (or are not destroyed so efficiently) outside of clusters, then it may be that there is a relatively significant population of UDGs in the field that are not yet accounted for in the baryon budget.

While much is known of the abundance and properties of UDGs in dense environments like galaxy groups and clusters, relatively little is known about the field population[1], expected to form preferentially through secular mechanisms[2] (see, for

[1]We note that my working definition of the field is a representative piece of the Universe in which galaxy groups and clusters are included, but massive haloes naturally make up a relatively small fraction by mass.

[2]Although for the present study I cannot rule out all external processes such as accretion from gas

example, the work of Graham et al., 2017; Janz et al., 2017, regarding the evolution of early type dwarf galaxies). This is mainly because of the difficulties involved in measuring distances to large samples of LSB galaxies without prior information such as cluster association. Observational studies in groups and clusters alone have been unable to disentangle the relative importance of in-situ vs. environment-driven formation because of the need to perform a statistical background subtraction of interloping (i.e. non-group or cluster) UDG candidates. Some studies (Das, 2013; Leisman et al., 2017; Papastergis et al., 2017; Greco et al., 2018a; Zaritsky et al., 2019) have shown that a field population of UDGs does indeed exist, yet the global properties of these galaxies are poorly understood.

One particularity relevant piece of work is that of Leisman et al. (2017) and Jones et al. (2018), who have shown that not only do HI-rich UDGs exist in the field (as theoretically predicted by Di Cintio et al., 2017), but also that their number density is too high to be explained by an extrapolation of the empirical relation between the number of UDGs and M_{halo} measured by van der Burg et al. (2017). Further, their sample appear systematically bluer than anticipated for UDGs in the field when compared with semi-analytic models (Rong et al., 2017; Jones et al., 2018). However, this analysis was limited to HI-rich field UDGs and it is unknown how this population relates to the overall field UDG population.

In this chapter, I use deep, wide-field imaging combined with an empirical UDG model to statistically constrain the global properties of UDGs in the field without knowing the distances to any of my sources. This includes an analysis of their colours, number density and mass-formation efficiency. The chapter is structured as follows: I describe my data in §3.2. I describe my sample of UDG candidates in §3.3 and quantify my recovery efficiency. In §3.4 I describe my empirical model of UDGs and potential interlopers in my UDG candidate sample. My results are presented and discussed in sections 3.5 and 3.6 respectively. I conclude in §3.7. All magnitudes are quoted in the AB magnitude system. Cosmological calculations are performed assuming ΛCDM cosmology with Ω_m=0.3, Ω_Λ=0.7, H_0=70 kms^{-1}Mpc^{-1}.

3.2 Data

For source detection and structural parameter estimation, I use a 180 deg^2 subset of data from the Astrowise (McFarland et al., 2011) reduction of the Kilo-Degree Survey (KiDS; de Jong et al., 2013; Kuijken et al., 2019) that overlaps with the GAMA

clouds.

spectroscopic survey (Driver et al., 2011) equatorial fields. I use the r-band for source detection because it is the deepest and has the best image quality. This is the same data[3] used by van der Burg et al. (2017) in their study of the UDG populations of galaxy groups and so I can make direct comparisons with their findings. Despite the GAMA overlap, redshift measurements are not available for most of my sources because they are generally much fainter than the limiting depth of GAMA at 19.8m$_r$.

The pixel size of KiDS is 0.2″, small enough to properly sample the point spread function (PSF) that has a typical full-width at half-maximum FWHM<1″. The sky-background is estimated in meshes of 20″ that are median filtered over in 3×3 meshes.

While the KiDS r-band is sufficient to reach a limiting surface brightness of $\bar{\mu}_{e,r}\sim$26.5, I additionally use the first data release of the overlapping Hyper-Suprime-Cam Subaru Strategic Program (Aihara et al., 2018) to measure colours. The HSC-SSP data is around 0.5m$_r$ deeper than KiDS (more so in the g-band), but has a smaller overlapping footprint by about a quarter compared to the KiDS area I consider. This leaves us with ∼39 deg^2 of unmasked data from which I can measure HSC-SSP colours.

Compared to the 180 deg^2 KiDS-GAMA overlap I use here, the remaining footprint that I have HSC-SSP data for is fairly limited. This may make us partially sensitive to cosmic variance. However, I note that my footprint is spread uniformly over three GAMA regions (G09, G12 and G15), each separated by at least 26 degrees. Additionally, I can account for local galaxy groups and clusters using the GAMA group catalogue (Robotham et al., 2011, see §3.6.6). In the future, my analysis can be easily upscaled to larger footprints.

I note that I do not use the HSC-SSP for detection because of its limited footprint and because its background subtraction is more aggressive compared to KiDS (mesh grid of ∼20″ but with no median filtering over meshes), meaning that it could restrict the maximum angular size of sources I could measure accurately. For the present analysis, I restrict myselves to the g and r bands but note that this can be expanded in future studies.

3.3 Measurements

Since I do not know the distance to any of my sources (apart from a small subset, cf. section 3.6.3), I must rely on selection criteria defined in observable parameter space (i.e. that which I measure as projected on the 2D surface of the sky). Specifically,

[3]Although they use the THELI (Erben et al., 2013) KiDS reduction, the depth is essentially equivalent.

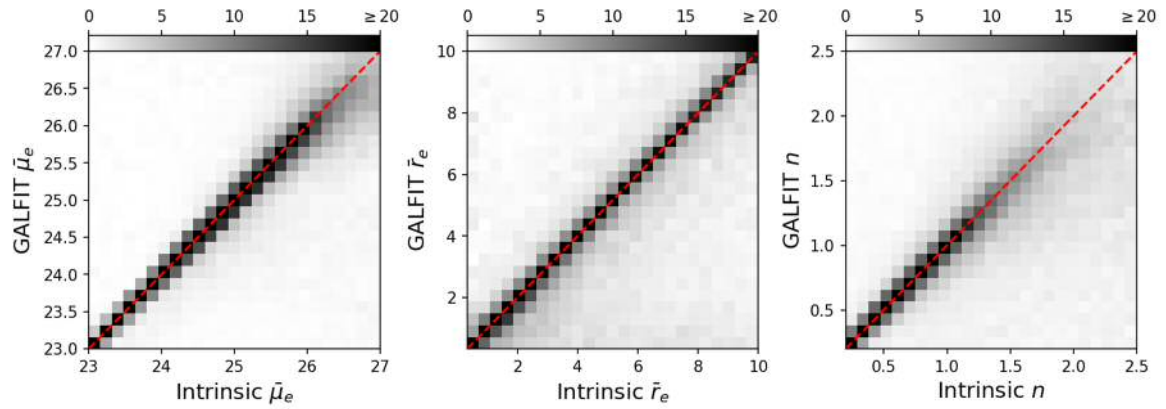

Figure 3.1: This figure shows how the measured GALFIT parameters compare with the intrinsic Sérsic parameters of the synthetic galaxies that I inject into the data. The black 2D-histograms show this data for my sample of recovered injections. The red dashed line is the one-to-one relation and is not a fit. Clearly I am able to recover the intrinsic parameters with good accuracy and precision over the range of parameter space that I am interested in. However, my precision diminishes slightly when recovering high n (Sérsic index) sources, but I retain accuracy such that there is little bias in the recovered parameter estimates. The units of the histogram are the percentage of sources with intrinsic parameters (given by the column) that occupy a particular bin in measured parameter space. The colour map is capped at 20% in order to increase the contrast for bins of lower completeness.

I target the LSB regime $24.0 \leq \bar{\mu}_{e,r} < 26.5$, where the lower limit is chosen for consistency with the literature and the upper limit is defined by the depth of KiDS. The upper selection limit on $\bar{r}_e$ is chosen to be much smaller than the spatial scale of the KiDS background subtraction and I set it as 8.0" in line with van der Burg et al. (2017). The lower limit on $\bar{r}_e$ is more difficult to set; while technically I am limited by the size of the PSF (FWHM ∼0.6"), it is also worth considering that the number of contaminant sources (i.e. non-UDGs) that satisfy my selection criteria quickly increases as this limit is lowered because of massive galaxies in the background. This issue is compounded by the fact that I do not have the advantage of a directly-measurable background surface density compared to similar studies in groups and clusters. Here, I use a lower limit of $\bar{r}_e \geq 3''$ for my selection (e.g. Sabatini et al., 2003; Davies et al., 2016). At a redshift of z=0.2, $\bar{r}_e$=3″ corresponds to ∼10 kpc. My upper limit of $\bar{r}_e$=8″ corresponds to 1.6 kpc at z=0.01.

3.3.1 Source detection and measurement

I choose to improve upon the catalogue used in van der Burg et al. (2017), who used SExtractor for source extraction, by using software optimised for the detection of LSB sources; this enables us to probe slightly deeper than their catalogue. I have experimented with several different detection and segmentation algorithms, including MTObjects (Teeninga et al., 2016), ProFound (Robotham et al., 2018)[4] and DeepScan (chapter 2). After some consideration, I selected MTObjects as the most suitable for my analysis because it seemed to produce less spurious detections around large, bright galaxies in my pipeline. I note that during this work ProFound has been updated with an alternative segmentation algorithm that improves its reliability around such objects, but I have not tested this. I used default parameters from MTObjects: α=10^{-6} and `move_factor`=0.5, where α sets the statistical significance level for the de-blending, and `move_factor` determines the spread of large objects.

I used the KiDS weight images to mask out regions in the data which have less than three exposures contributing to the imaging prior to the MTObjects run. This was done to ensure uniform sensitivity over the full data set, as MTObjects relies on a global estimate of the background distribution.

Point spread function measurement

I took advantage of my decision to split the KiDS frames into 3×3 subframes by making one PSF model per subframe (i.e. nine PSF models per square degree). This was accomplished by targetting point sources in the `R_e` and `mag` plane (MTObjects estimates of the effective radius and total magnitude respectively) based on my MTObjects catalogues from each subframe. Point sources are required to have an axis ratio as estimated by MTObjects greater than 0.9. I then fit Moffat profiles to the individual point source candidates using GALFIT (Peng et al., 2002). My final PSF model for each subframe was taken as the model corresponding to the mean Moffat FWHM after a sigma-clipping algorithm was applied to remove outliers. I measure a mean FWHM of 0.6″ with a standard deviation of 0.1″ over the full KiDS area that I use. I measure a median value of the Moffat β parameter of 2.2, with a standard deviation of 0.1.

[4]https://github.com/asgr/ProFound

Source extraction pipeline

Following measurement of the PSF, my overall detection and measurement consists of the following steps:

(i) Use MTOBJECTS to produce a segmentation image and preliminary source catalogue for each subframe.

(ii) Apply a pre-selection to the preliminary catalogue to identify candidates suitable for input to GALFIT. This is necessary to ensure a practically feasible number of fits. Specifically, I required `R90`$>1.5''$ and `mu_mean`>23.5 mag arcsec^{-2}, where `R90` and `mu_mean` are a proxy for the radius containing 90% of the galaxy light and the average segment surface brightness, respectively.

(iii) Use GALFIT to fit a combined Sérsic plus inclined sky plane model to each pre-selected source, ignoring masked pixels and additionally masking other segments from the MTOBJECTS segmentation images. Parameter estimates from MTOBJECTS are used as the initial guesses for GALFIT. The sky RMS is estimated directly from pixels in the cut-out region that are unmarked in the segmentation image. The Sérsic profile is defined as,

$$I(r) = I_e \exp\left\{-b_n\left[\left(\frac{r}{r_e}\right)^{\frac{1}{n}} - 1\right]\right\} \tag{3.1}$$

where $I(r)$ is the galaxy's intensity as a function of radius, I_e is the intensity at the effective (half-light) radius r_e, and b_n is a constant determined only by the index n, which in turn governs the profile slope, defined by:

$$\Gamma(2n) = 2\gamma(2n, b_n) \tag{3.2}$$

where Γ is the complete gamma function and γ is the incomplete gamma function. I note that all conversions between Sérsic parameters are performed using the prescriptions of Graham & Driver (2005).

(iv) Apply the selection criteria to the resulting GALFIT models in order to produce a final catalogue of UDG candidates.

The PSF models are used as an input to GALFIT such that the measurements correspond to deconvolved parameters. The GALFIT cut-out size was 400 pixels, large enough to recover the intrinsic Sérsic parameters properly over the full range of parameter space I explore here, and was tested with the synthetic source injections (see §3.3.2). I note that it was important to include the inclined sky plane component in the GALFIT modelling in order to retrieve unbiased measurements in the cases of high $\bar{r}_e$ or high n profiles.

Following the full set of pre-selections, I am left with $\sim 2\times10^6$ sources that are input to GALFIT over the full 180 degree squared KiDS area. Following the selection using GALFIT parameters ($24.0 \leq \bar{\mu}_{e,r} \leq 26.5$, $3.0" \leq \bar{r}_{e,r} \leq 8.0"$, $n \leq 2.5$), I am left with 829 UDG candidates. After selecting sources also within the HSC-SSP footprint, my final catalogue of UDG candidates consists of 212 sources; some examples are shown in figure 3.2. I note that contrary to UDGs in clusters and groups, my sample comprises of sources that appear far more irregular, with features suggestive of active star formation.

3.3.2 Recovery Efficiency

I define the recovery efficiency ϵ as the fraction of sources with intrinsic observable parameters (i.e. without the effects of measurement uncertainty) that have measurements that meet my selection criteria. As such, sources that do not meet my selection criteria in terms of their intrinsic observable parameters may be selected ($\epsilon > 0$) because of measurement uncertainty. Anti-correlated with ϵ is therefore the selection purity, defined as the fraction of detections with intrinsic observable parameters that do not meet the selection criteria, but have measured properties that do and thus make it into the final catalogue of UDG candidates. This is different compared to the purity of the UDG candidates, which is defined as the fraction of sources in the UDG candidate catalogue that are intrinsically UDGs as defined by their physical properties.

I have used synthetic source injections to quantify ϵ as a function of intrinsic Sérsic parameters, $\epsilon(\bar{\mu}_{e,\mathrm{int}}, \bar{r}_{e,\mathrm{int}})$. To do this, I create mock images by inserting artificial galaxy profiles (PSF-convolved, one-component Sérsic) into each frame of the real KiDS data and run them through the full detection and measurement pipeline described in §3.3.1. The Sérsic parameters are drawn uniformly from the ranges presented in table 3.1, where q is the observed axis ratio. As noted by van der Burg et al. (2016), $\bar{r}_e$ (the circularised half-light radius) is robust against the intrinsic distribution of axis ratios and $\bar{\mu}_e$ is a better indicator of a sources detectability than other parameters such as the central surface brightness. Note that I do not include the Sérsic index n as a free parameter in $\epsilon(\bar{\mu}_{e,\mathrm{int}}, \bar{r}_{e,\mathrm{int}})$ in order to simplify the analysis. This does not severely impact my results because the intrinsic range in n for UDGs is narrow (e.g. Koda et al. (2015) find a mean of $n \sim 1$ with a standard deviation of 0.34).

The sources are injected at a surface density low enough to ensure that injected profiles are separated on average by 6.5 times the maximum value of $\bar{r}_e$ given in table 3.1. I repeated the process several times in order to increase the number statistics for

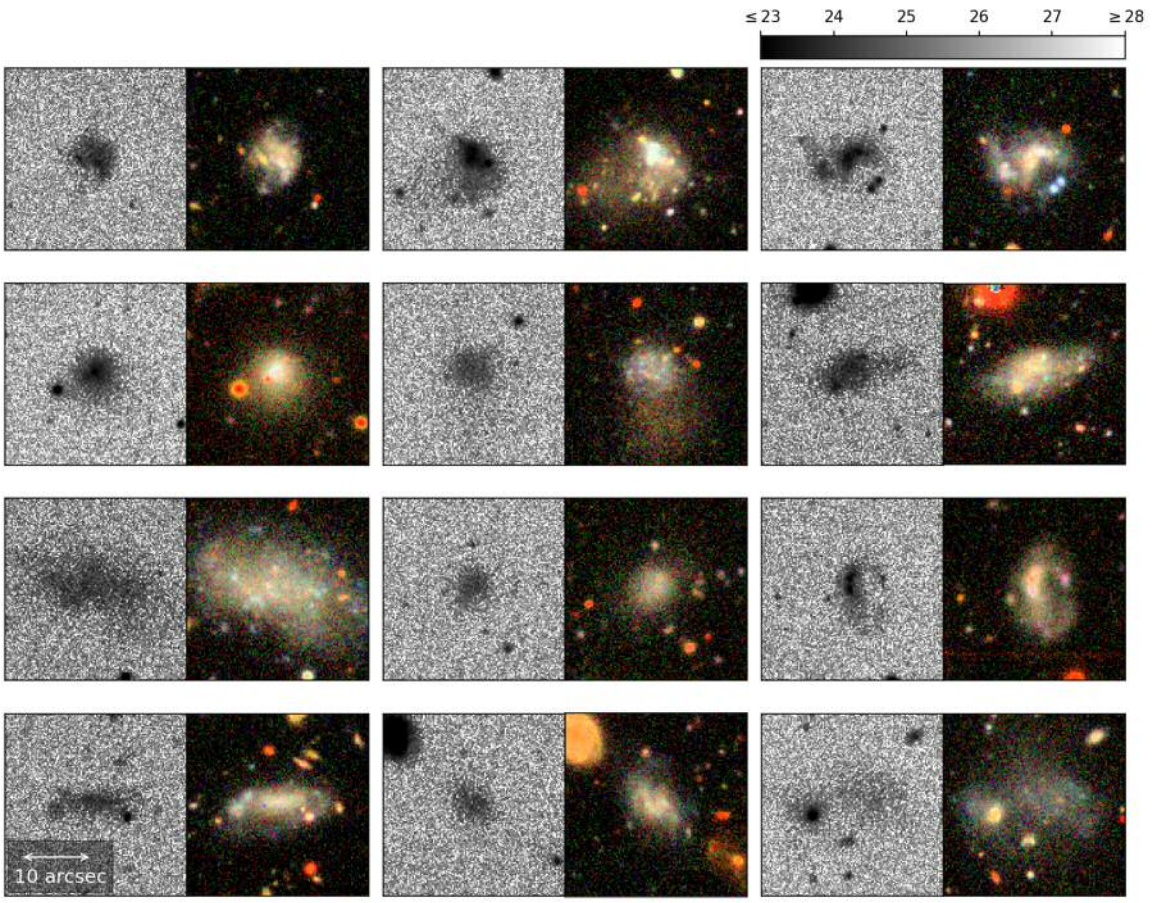

Figure 3.2: Examples of large LSB sources in my UDG candidate sample. Each source in shown in KiDS r-band (left panels) vs. a colour image made from the g, r & i HSC-SSP bands according to Lupton et al. (2004) (right panels). The cut-out size is 25″, much less than the 80″ regions that I use to fit the sources. The colour bar shows the surface brightness in units of magnitudes per square-arcsecond for the KiDS data. In comparison to the regular morphologies of UDGs in groups and clusters, many of my sources are rather amorphous, with possible signs of discrete sites of active star formation.

the $\epsilon(\bar{\mu}_{e,\mathrm{int}}, \bar{r}_{e,\mathrm{int}})$ measurement, simulating ∼735,000 sources overall.

We only considered unmasked sources for the estimate of $\epsilon(\bar{\mu}_{e,\mathrm{int}}, \bar{r}_{e,\mathrm{int}})$, which was measured with ∼550,000 artificial galaxy injections spread evenly over my full KiDS subset. I show my ability to precisely measure the intrinsic parameters of my injected sources in figure 3.1. My fiducial measurement of $\epsilon(\bar{\mu}_{e,\mathrm{int}}, \bar{r}_{e,\mathrm{int}})$is shown in figure 3.3.

We note that I additionally tested Imfit (Erwin, 2015) in place of GALFIT in the pipeline described in §3.3.1 and found that it made no significant difference to my

Parameter	Lower limit	Upper limit
$\bar{\mu}_e$ [mag arcsec^2]	22	30
$\bar{r}_e$ [arcsec]	1	25
q	0.1	1
n	0.2	2.5

Table 3.1: Parameter ranges for the artificial galaxy injections. The parameter realisations are drawn uniformly within the ranges, which are much wider than my selection criteria for UDG candidates.

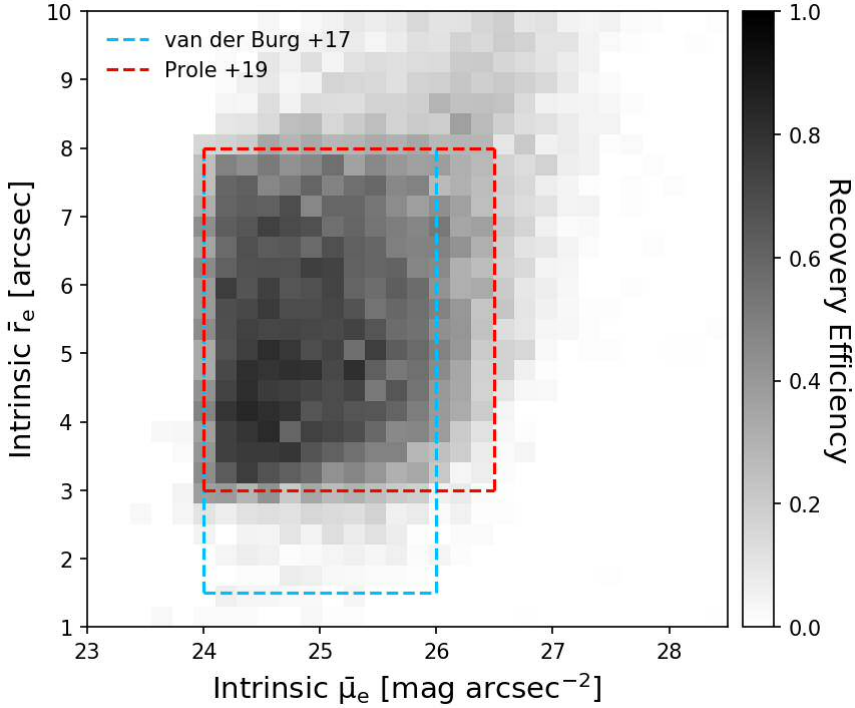

Figure 3.3: The recovery efficiency of synthetic sources injected into the data as a function of circularised effective radius and mean surface brightness within the effective radius given my selection criteria. The red box indicates my selection criteria, while the blue box is that used in van der Burg et al. (2017).

measurement of $\epsilon(\bar{\mu}_{e,\mathrm{int}}, \bar{r}_{e,\mathrm{int}})$.

A criticism of the above method is that it relies on field galaxies being well fit by Sérsic profiles. While this is certainly justified in dense environments where UDGs show little evidence of tidal features (Mowla et al., 2017), it is not clear whether this is justified for the field population. Since field galaxies are expected to be relatively

isolated, they are likely to show little evidence of tidal disruption. This does not however rule out irregular morphologies caused by bright star-forming regions and secular processes like stellar feedback from e.g. supernovae.

In addition to the injections described above and in equal measure, I inject nucleated profiles (Sérsic + Moffat PSF model) to the data assuming that approximately 1% of the galaxy light is contained within the nucleus. This allows us to quantify any systematic differences in my recovery efficiency that might be caused by the presence of nuclei and adapt my selection criteria accordingly. I note that I do not attempt to fit nucleated profiles for my measurements because experiments with artificial galaxy profiles showed that the fits are not reliable. I find that the presence of a nucleus is sufficient to positively bias recovered values of the Sérsic index by approximately 30% at n=1. Since almost all recorded UDGs have Sérsic indices less than around 1.5 (there is both observational and theoretical evidence for this, e.g. Koda et al., 2015; Román & Trujillo, 2017b; Di Cintio et al., 2017; Venhola et al., 2017), my upper selection limit of n=2.5 is justified. The effect of the selection in n on my results is discussed further in §3.5.2.

In comparison to van der Burg et al. (2017), my selection criteria allow us to probe ∼0.5 magnitudes deeper in surface brightness. The increased depth comes in-part from the proficiency of MTOBJECTS over SEXTRACTOR for detecting LSB sources. I also use a higher cut in $\bar{r}_e$; since I do not have the advantage of a measurable background level, imposing a higher minimum cut in $\bar{r}_e$ allows us to mitigate against an excessively contaminated sample of UDG candidates. In figure 3.3, it can be seen that I expect some contamination from (apparently) large, faint sources that do not intrinsically meet my selection criteria (top right of the figure). However, since the number of apparently large sources is very small compared to the number of smaller ones, the decrease in purity from such sources is negligible.

3.3.3 Colours

We exploit the overlap of the KiDS survey with the HSC-SSP footprint in order to measure $(g-r)$ colours for my UDG candidates. I remind the reader that while HSC-SPP is ∼0.5 magnitudes deeper than KiDS, the background subtraction in HSC-SPP is slightly more aggressive.

I use an aperture-based strategy to measure $(g-r)$ colours. Specifically, all colours are measured within the $1\bar{r}_{e,r}$ ellipses from my GALFIT measurements. I estimate the sky level along with its uncertainty using a random aperture approach, whereby I

place 100 equally shaped apertures in close vicinity to (but not touching) the source. Before measuring the median background level and its uncertainty, fluxes are sigma-clipped at 2σ[5] in order to lower the potential for over-estimating the background level because of nearby sources. I do not perform additional aperture corrections because my sources are much more extended than the PSF.

Due to the increased depth offered by HSC-SSP, I am able to measure positive fluxes in the g and r-bands for close to 100% of my sources. The typical measurement error in $(g-r)$ due to the background fluctuations is approximately 0.04 magnitudes; this is discussed further in §3.4.3. I show comparisons between KiDS and HSC-SSP imaging in figure 3.2.

3.4 The empirical model

Without knowing the distances to any of my sources, it is difficult to tell how many are intrinsically UDGs and how many are cosmologically dimmed background galaxies. In this section, I describe an empirical model that can be used to generate a synthetic population of UDGs in order to compare with observation. This is supported by an additional model for massive galaxies that allows us to estimate the number of non-UDG contaminants in my observational sample.

3.4.1 Empirical UDG model

Empirical properties of UDGs

One of the simplest models can be created by assuming field UDGs share similar empirical properties with UDGs in clusters. Of importance for my analysis are prescriptions for $\bar{\mu}_{e,\mathrm{int}}$ (the surface brightness corrected for cosmological projection effects), $\bar{r}_{e,\mathrm{phys}}$ (physical size) and $(g-r)_{\mathrm{int}}$, the rest-frame colour. As discussed in §3.3.2, I assume that all UDGs occupy the range of 0.2 to 2.5 in Sérsic index.

van der Burg et al. (2016) recorded that the distribution of average surface brightness $\bar{\mu}_e$ is approximately uniform in group environments for UDGs. This has been complimented by the findings of Danieli & van Dokkum (2018), who found that the distribution of absolute magnitude at fixed size is approximately uniform for large, red galaxies in the Coma Cluster after accounting for the newly discovered UDGs. These two observations are equivalent, since at a given size the mean surface brightness is

[5]The bias in the recovered standard deviation when sigma-clipping normally distributed data at 2σ is approximately 25%; I therefore correct my estimates by this factor.

uniquely defined by its magnitude (i.e. there is no dependence on Sérsic index). However, the work of Danieli & van Dokkum (2018) showed that this relation extends from the low surface brightness regime and to much brighter galaxies. I therefore adopt a uniform distribution $U_{x_{\rm min}}^{x_{\rm max}}(x)$ for $\bar{\mu}_{e,\rm int}$:

$$\bar{\mu}_{e,\rm int} \sim U_{24.0}^{26.5}(\bar{\mu}_{e,\rm int}) \tag{3.3}$$

The subsequent observational study of van der Burg et al. (2017) (supported theoretically by Carleton et al., 2018) has shown that the size distribution of UDGs in groups and clusters is well described by a power law of slope -2.71 ± 0.33 in logarithmic size bins, such that smaller UDGs are much more common than larger ones. The intrinsic distribution of physical sizes in kpc is therefore taken as:

$$\bar{r}_{e,\rm phys}[\rm dex] \sim \bar{r}_{e,\rm phys}^{-2.71} \tag{3.4}$$

where I assume the range of $\bar{r}_{e,\rm phys}$ lies between 1.5 and 7.0 kpc, consistent with van der Burg et al. (2017). I probe the effect of varying the power-law slope on my result in section 3.6.1.

It has been noticed by several authors that UDGs in clusters tend to lie on the red-sequence (e.g. Koda et al., 2015; van der Burg et al., 2016) and this is also expected theoretically. There have been hints that UDGs may tend to be much bluer in less-dense environments (Román & Trujillo, 2017b; Jones et al., 2018; Jiang et al., 2018), although this is not always clear from an observational perspective (Román & Trujillo, 2017a). As such, colours of field UDGs remain relatively poorly understood. I therefore leave the distribution of $(g-r)_{\rm int}$ as a variable of my model, and discuss it further in §3.4.1.

Estimated number density of UDGs

I use the (almost linear) empirical relation between the number of UDGs and the mass of their parent halo measured by van der Burg et al. (2017) to estimate the formation efficiency of UDGs per unit mass in clusters and groups. From this, I can calculate the total number of UDGs that should exist out to redshift $z_{\rm max}$, using my cosmological model to estimate the total enclosed mass. I calculate the total mass $M_{\rm tot}$ contained within the volume V probed by solid angle ω out to $z_{\rm max}$ using equation 3.5,

$$M_{\rm tot} = \int_0^{z_{\rm max}} \Omega_m(z)\rho_{\rm crit}(z)\frac{dV}{dz\,d\omega}dz\,d\omega \tag{3.5}$$

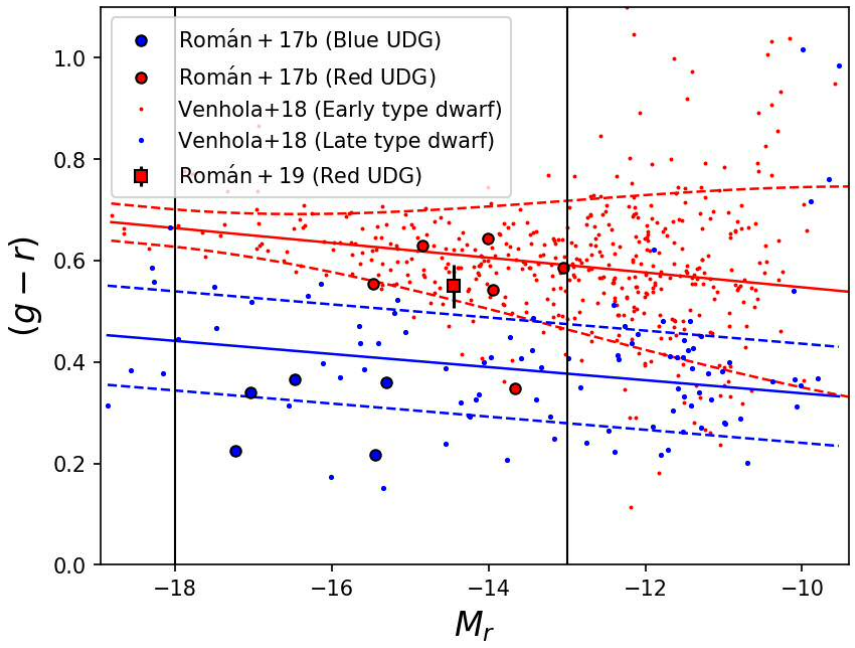

Figure 3.4: Our model colour distributions for the mock UDG catalogues. The data I fit are from Venhola et al. (2018) and correspond to late-type (small blue points) and early-type (small red points) dwarf galaxies in the Fornax cluster. I fit simple linear models (uninterrupted coloured lines) and their 1σ uncertainties (dashed lines) after clipping outliers at 3σ. Under the assumption that UDGs have similar colours to dwarf galaxies, I use the late/early type fits for my star-forming/quiescent mock UDG catalogues. The vertical black lines span the approximate range of absolute magnitudes occupied by UDGs (van der Burg et al., 2016). The bold points show the red/blue UDGs of Román & Trujillo (2017b) and Román et al. (2019a).

where Ω_m is the fractional contribution of matter to $\rho_{\rm crit}$, the critical density of the Universe. I additionally assume that the UDGs are spatially distributed smoothly according to the integrand of equation 3.5, such that the redshift distribution of field UDGs follows the mass.

Note that the adopted value of $z_{\rm max}$ does not impact the result, provided that the number of sources I predict to observe out to z (given my recovery efficiency) has converged, i.e. has stopped increasing, before $z_{\rm max}$. For my modelling I use $z_{\rm max}$=1, which meets this criterion (see §3.5.1).

If I assume that UDGs form with an average efficiency equivalent to a $10^{15}M_{\odot}$ cluster according to equation 1 of van der Burg et al. (2017), I derive a volume density of $\sim 9 \times 10^{-3}$cMpc^{-3}. This is a factor of six higher than the total number density of HI-bearing UDGs measured by Jones et al. (2018), who measured a value of $1.5\pm0.6\times10^{-3}$cMpc^{-3}. Using a different value than $10^{15}M_{\odot}$ for the halo mass would not strongly modify the initial number density estimate since the slope of the relation between M_{halo} and the number of UDGs hosted by the halo is approximately one (at least down to $M_{halo}\sim10^{12}M_{\odot}$, Román & Trujillo, 2017b; van der Burg et al., 2017; Mancera Piña et al., 2018). However, by selecting a halo mass of $10^{15}M_{\odot}$, I am essentially comparing the field abundance with that in a $10^{15}M_{\odot}$ cluster in my later analysis. I estimate the impact that the uncertainty in the van der Burg et al. (2017) relation has on my result in section 3.6.1.

Accounting for cosmological effects

We account for the cosmological distance modulus, angular diameter distance d_a (describing how physical sizes map to angular sizes as a function of the redshift, z) and k-corrections (the filter and spectral energy distribution (SED) dependent effect that modifies a source's apparent brightness with z, independently from the distance modulus, Hogg et al., 2002). In combination, these quantities allow us to project the surface brightnesses and angular sizes of my mock sources out to a certain redshift.

While d_a is simple to account for, the exact k-correction depends on the assumed SED for the UDGs. Quiescent UDGs are thought to be old, metal poor galaxies (e.g. Ruiz-Lara et al., 2018; Ferré-Mateu et al., 2018; Fensch et al., 2018). I adopt the average UDG properties from Ferré-Mateu et al. (2018) to estimate the k-corrections for such galaxies, namely an age of 6.7 Gyr and [Z/H]=-0.66. In the case of star-forming UDGs, I assume the same age and metallically as in the quiescent model, but introduce star-formation at a uniform rate until the time of observation. While this is an idealised scenario, I probe the significance on the assumed model for k-corrections in section 3.6.1. All stellar population models and k-correction estimates are calculated using the Flexible Stellar Population Synthesis (FSPS, Conroy et al., 2009; Conroy & Gunn, 2010) code. For the KiDS r-band, I assume SDSS-like filters for the k-correction estimates. For the HSC-SSP colours, I use Subaru Suprime Cam filters.

I construct mock catalogues by sampling intrinsic parameters from the appropriate distributions. Following this, I convert the units into apparent, observed quantities through the cosmological distance modulus, angular diameter distance (for the angular sizes) and band-specific k-corrections.

UDG colour models

The inclusion of colour into my analysis is critical because the colour of a galaxy contains some information about its distance thanks to the cosmological redshifting of spectroscopic features. While I have assumed a stellar population model for the UDGs in order to estimate the k-corrections, I cannot use these models to assign colours to my mock catalogue because of the need to include some intrinsic scatter. One alternative approach is to model the colours using measurements from the literature.

Several authors have shown that UDGs occupy the red sequence in clusters (e.g. Koda et al., 2015; van der Burg et al., 2016) and this is also supported theoretically (Rong et al., 2017). However, modelling the colour distribution of star-forming UDGs is slightly harder because there is not as much available data for them. Since UDGs have stellar populations similar to dwarf galaxies (e.g. Fensch et al., 2018), one viable method is to assume that star-forming UDGs have colours similar to late-type dwarf galaxies.

Venhola et al. (2018) have measured the $(g-r)$ colours as a function of absolute magnitude for such galaxies in the Fornax cluster. Using these measurements, it is possible to fit the relationship between colour and absolute magnitude with a simple linear model separately to each of their early and late type samples. For the late type galaxies, I use a constant scatter term, while I interpolate the standard deviation of the colours in bins of absolute magnitude for the early type galaxies. I show the corresponding fits in figure 3.4, where I have clipped outliers at 3σ. I note that my fit to the early-type dwarf galaxies is consistent with the approximate fit to the red-sequence UDGs in clusters from van der Burg et al. (2016).

In figure 3.4 I also compare with the UDGs discovered by Román & Trujillo (2017b)[6], which have been decomposed into red and blue populations based on their $(g-i)$ colour. While their red population is fairly consistent with my red-sequence model, the blue UDGs seem to be systematically bluer than my colour model for blue galaxies. This may be explained by the fact that my model is based on measurements from the Fornax galaxy cluster where environmental processes, for example ram-pressure stripping, may cause reddening of the galaxies. In comparison, the UDGs of Román & Trujillo (2017b) are found in isolated galaxy groups where such effects are less prolific.

[6]where the $(g-r)$ colours have been kindly provided by Javier Román.

3.4.2 Empirical model for massive galaxies

Not all of the UDG candidates in my observational sample are intrinsically UDGs. As large, bright galaxies are shifted towards higher redshift, they become both fainter and smaller in terms of their angular size and may eventually satisfy my selection criteria. Equally, small foreground dwarf galaxies not meeting the UDG criteria have the potential to contaminate the sample. Since I am not able to directly measure the number of these interlopers (as can easily be done when considering a group or cluster environment), I am forced to use empirical relations from the literature to estimate the level of contamination.

It is standard practice to broadly categorise galaxies as either late-type or early-type based on their morphology and/or colour (e.g. Bell et al., 2003; Baldry et al., 2004; Driver et al., 2006; Taylor et al., 2015). Massive early-type galaxies (ETGs) are typically quiescent and therefore redder than late-types. Additionally, massive ETGs generally have higher Sérsic indices compared to late types. For ETGs, the Sérsic index increases with total stellar mass (e.g. Caon et al., 1993; Graham et al., 1996; Danieli & van Dokkum, 2018). While early-type dwarf galaxies exist with Sérsic indices around one (chapter 4), I am probing the field population and therefore expect that the main contribution from ETG interlopers will be from higher mass galaxies with correspondingly higher Sérsic indices. I discriminate against recovering such objects in my UDG candidate sample through the upper-limit cut in Sérsic index at n=2.5. As such, I expect the dominant source of contamination in terms of massive galaxies ($M_* > 10^9 M_\odot$) to be mainly constituted of massive late-type galaxies. By contrast to massive ETGs, late-type galaxies are systematically bluer, with Sérsic indices n <2.5 (Vulcani et al., 2014).

This population of interlopers we model here is by no means complete. For example, low n early type galaxies may still appear in the observed catalogue which have not been accounted for. I have limited the interloper model to massive late type galaxies because they are expected to be the dominant source of interloper and they can be subtracted cleanly from our observational sample given my selection criteria. Including additional types of contaminant objects would of course lower the number of UDGs I estimate to exist in the field, meaning that the present analysis can only produce an upper-limit estimate of the number of field UDGs.

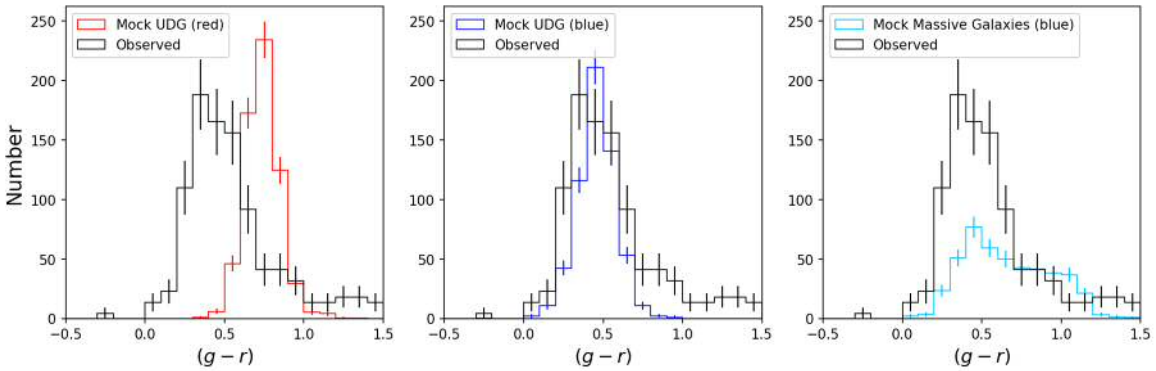

Figure 3.5: Synthetic distributions of $(g-r)$ colour for my mock red UDG, blue UDG and massive blue galaxy catalogues, weighted by the probability of observation, compared to the actual observed histogram. The absolute numbers are normalised to an area of 180 square degrees. The error-bars show the Poisson uncertainties in each bin. Colours are in the observed reference frame. I note that I include the effect of measurement error in my mock colours. It is clear that the red UDG model is not consistent with the observations, being much more consistent with the blue model.

Canonical empirical distributions

While relatively little is known about the population of field galaxies with stellar masses lower than around $10^9 M_\odot$, much is known about objects at higher stellar masses. I can therefore produce mock observational catalogues of high stellar-mass galaxies and use them to estimate the contamination level in my actual observed catalogue. Such an estimate would be naturally conservative owing to the truncation at $10^9 M_\odot$ that essentially excludes all dwarf galaxies including UDGs (chapter 4) The ingredients of my model catalogues are:

- The stellar mass function (SMF). I have used measurements from GAMA (Baldry et al., 2012) and COSMOS/UltraVISTA (Muzzin et al., 2013) to model the galaxy SMF of field galaxies, including its redshift dependence between z=0 and z=1. These measurements are additionally decomposed into red and blue galaxy populations and I have incorporated this into my mock catalogues.
- The stellar-mass to size relation. I used the measurements of van der Wel et al. (2014) to assign effective circularised radii to each of my random sample of stellar masses, as a function of redshift. Again, I make use of their early/late-type colour decomposition. I also include scatter in the sampling based on their measurements.

- I assign $(g-r)$ colours to my sampled galaxies as a function of their stellar mass by using empirical data gathered by the GAMA survey[7] (Taylor et al., 2011). Specifically, I bin their $(g-r)$ measurements in stellar mass and assign intrinsic (i.e. observed at z=0) colours to my mock observations in each bin by randomly sampling from the corresponding GAMA $(g-r)$ measurements.
- I calculate k-corrections in the same way as described in §3.4.1, this time assuming an onset of star formation 8.9 Gyr ago and [Z/H]=-0.38, with a uniform star formation rate. These values are based on the high surface brightness, late type sample of Ferré-Mateu et al. (2018). The effects of modifying this are discussed in section 3.6.1.
- As before, I assume the redshift distribution of my catalogue to be smoothly distributed with the mass in the Universe, according to equation 3.5.

Clearly such an exercise is approximate in nature, and is designed only to get a first-order estimate of the number of contaminants in my UDG sample. A discussion regarding the contribution to my observations from LSB galaxies smaller than the fiducial UDG limit of $\bar{r}_e$=1.5 kpc can be found in section 3.6.2.

Many of the mock massive galaxies are brighter than m_r=19.8. This means that it is possible to compare the redshift distribution of my mock catalogue with that of the GAMA spectroscopic survey. I show in section 3.6.3 that my mock catalogues are consistent with that observed by GAMA.

3.4.3 Measurement Errors

Before the catalogues can be directly compared with my observations, it is important to consider the effect of measurement errors on the predicted distributions of observed parameters. Of particular importance is the uncertainty in $\bar{r}_e$, which increases for larger and fainter galaxies. This is significant because there is typically a steep gradient in the distribution of $\bar{r}_e$, whereby there are far less large objects than small ones, both in terms of physical and angular size. Thus, including the measurement error in the mock catalogues causes an increase in the predicted number of galaxies observed with large angular sizes.

The measurement uncertainty on the Sérsic parameters is estimated directly from the synthetic source injections described in §3.3.2. Measurements of the injected sources are used to estimate the recovery efficiency, defined in intrinsic observable

[7]Specifically, I use the public `StellarMasses v19` catalogue available from http://www.gama-survey.org/dr3/.

parameter space. Consequently, the effects of measurement error (including any bias) as a function of intrinsic size and surface brightness are already contained in the recovery efficiency estimate. I can therefore account for the effect of the measurement uncertainty in my mock catalogues by using the recovery efficiency to assign probabilities of detection (see §3.5.1). The limitation of this approach is that I cannot directly compare structural parameters in my mock catalogues with the observations.

Also of importance is the measurement error in $(g-r)$ colour. Starting from my estimates of fluxes and their errors described in §3.3.3, I perform monte-carlo realisations of flux ratios in order to estimate the distribution of uncertainties in the magnitude. I fit a log-normal distribution to the result, and use it to randomly sample uncertainties in colour; I then "jiggle" (randomly perturb within error) the colours in the mock catalogues according to the result. The mean uncertainty in colour is $\sim$0.04 mag.

3.5 Results

3.5.1 Observations vs. Model Predictions

We are now in a position to compare my mock catalogues with the observations. I note that for this analysis, all absolute numbers are normalised to an area of 180 square degrees. Each source in my mock catalogues is assigned a probability of recovery using the recovery efficiency discussed in §3.3.2, which are used as weights in the analysis. I note that after using such weights, the number of UDGs I predict to observe converges (i.e. does not increase further) by $z\sim$0.2 (see figure 3.9). Similarly, the mock massive galaxy catalogue converges by $z\sim$0.5. This is mainly an effect of the lower limit angular size cut at $\bar{r}_e \geq 3''$. I probe the accuracy of my modelling with reference to the redshift distribution of GAMA spectroscopic sources in section 3.6.3.

I compare the $(g-r)$ histogram of my observed UDG candidate catalogue with each of my mock catalogues (red UDGs, blue UDGs, massive blue galaxies) in figure 3.5; the results of which are fairly striking. Clearly either the assumption that all UDGs are on the red sequence as they are in clusters is not correct (as made clear by the significant offset between the peaks of the observed and predicted distributions), or UDGs in general do not form in the field with a mass-efficiency anywhere near what they do in clusters. However, since I already know that blue UDGs do exist in abundance in the field (e.g. Leisman et al., 2017; Jones et al., 2018), it is clear that I can rule the latter hypothesis out completely. From this result I would expect isolated

red UDGs, like the ones found by Martínez-Delgado et al. (2016) and Román et al. (2019a), to be relatively rare.

The discrepancy is further compounded if one considers the estimates for the massive blue galaxy interlopers. I argue that since the massive blue galaxies are the dominant source of contamination in my UDG catalogue (see §3.4.2 and section 3.6.2), I can obtain an observational sample that is representative of the UDG population by statistically subtracting the massive blue galaxy population from the observed catalogue of UDG candidates. The result is displayed in the left panel of figure 3.10, along with my mock UDG catalogues.

It is clear that the mock blue UDG catalogue is in much better agreement with the observed colour distribution than the red UDG catalogue. However, the observations are ∼0.05 magnitudes bluer than what my empirical models predict. This means that the colours may be more consistent with the blue UDGs of Román & Trujillo (2017b) (see figure 3.4). This is not particularly surprising; late-type galaxies in clusters are typically redder than those in the field because of environmental quenching from e.g. ram-pressure stripping.

By comparing my mock catalogues to the interloper-corrected observations, it is possible to estimate the total field density of UDGs, along with a corresponding mass formation efficiency. This is accomplished by comparing the predicted number of UDGs from my empirical model with the estimated number of observed UDGs. From the appearance of figure 3.10 (left panel), it is clear that I have overestimated the number density (and therefore mass formation efficiency) of UDGs in my model.

For the empirical model to predict the correct number of UDGs, I would require a mass-formation efficiency ∼0.8±0.2 times that what it is in clusters, taking into account uncertainties discussed in section 3.6.1. This translates into a field density of $8\pm3\times10^{-3}$cMpc^{-3}. This is an upper limit on the true field abundance of UDGs because the estimated number of observed UDGs is likely an overestimate; I have only considered massive blue galaxies as contaminant sources. This number density is still approximately two orders of magnitude lower than that of the overall dwarf galaxy ($M_*\leq10^9$M$_\odot$) population and is consistent with the value found by Dalcanton et al. (1997) for field LSB galaxies.

I note that these estimates apply only to the range of physical parameters that I have probed here, i.e. sizes in the range $1.5\leq\bar{r}_{e,r}$[kpc]≤7.0 and intrinsic (i.e. not cosmologically dimmed) surface brightnesses spanning $24.0\leq\bar{\mu}_{e,r}\leq26.5$. If I were to consider even fainter sources, this number density would likely increase.

3.5.2 Dependence on Sérsic index

UDGs typically have Sérsic indices n<1.5. The justification for my cut at n=2.5 is as follows:

(i) My measurements are conflated by measurement error, which gets worse as a function of surface brightness. In figure 3.1 I show that I am nevertheless able to recover essentially all non-nucleated profiles with n<1.5 by imposing a cut at n∼2.

(ii) Some UDGs are nucleated (e.g. Venhola et al., 2017). In this analysis, I have only fit single Sérsic profiles. The presence of a nucleus can bias my recovery of n by +30%, so a higher cut than n=2 is justified to preserve completeness.

(iii) I do not explicitly include the Sérsic index distribution in my empirical model for background interlopers[8], which is statistically subtracted from the observational sample in the analysis. It is typical in the literature to take n=2.5 as the dividing line between "early" and "late" type samples (e.g. van der Wel, 2008; Vulcani et al., 2014; Vika et al., 2015). Since I do not want to over-subtract the interloper population, it is important to use a consistent cut for the sample selection.

(iv) In this analysis, I am striving to place upper-limits on quantities like the UDG field number density; this is motivated by the fact that my interloper subtraction is likely incomplete. Lowering my Sérsic index selection cut would reduce the size of the observational sample and thus lower the inferred number density. In the interests of upper-limits, it is therefore prudent to keep a relatively high cut in n.

(v) My results are closely compared with the work of van der Burg et al. (2016), who used an even higher cut at n=4.

In summary, while the cut at n=2.5 might be relatively high compared to the observed values of n for UDGs, I account for the resulting contamination in my observational sample using the empirical model. However, it is important to discuss the effects of varying the index cut on my results. Recall that with the cut at n=2.5, the upper-limit mass formation efficiency is estimated to be ∼0.8±0.2 times that in clusters. If I instead take the cut at n=2.0, this drops to ∼50% of the value for clusters. Here I have likely increased the purity of UDGs in my sample, but for the reasons given above it is difficult to quantify the effect on the completeness. If instead I drop to n=1.5, the formation efficiency estimate drops to ∼30% of its value in clusters. However, it is likely that this value suffers from significant completeness effects and is

[8] although this could potentially be implemented in future studies

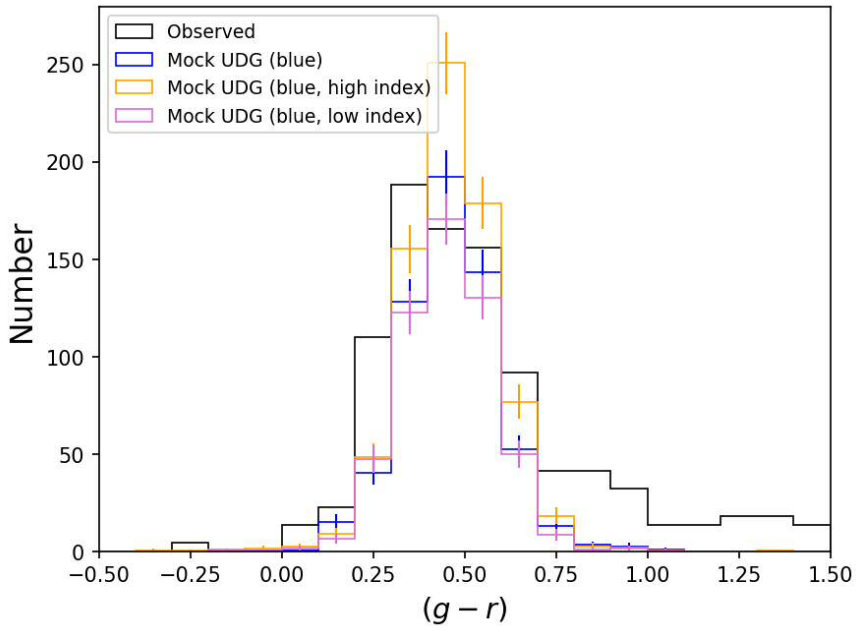

Figure 3.6: Impact of changing model parameters within their errors. The observed colour histogram is shown in back. The blue UDG model is shown in dark blue. The mock catalogue with the size power law index lowered/raised by 1σ is shown in purple/orange. All error-bars are Poisson uncertainties.

an underestimate.

3.6 Discussion

3.6.1 Model uncertainties

While my analysis here is approximate in nature, it is still important to quantify how uncertainties in the model ingredients may impact my result. In particular, I have not discussed how the uncertainties in the assumed UDG size distribution propagates. van der Burg et al. (2017) have measured a power-law index of -2.71±0.33 for the distribution of circularised radii (equation 3.4). The result of varying the slope by 1σ are shown in figure 3.6. It is clear that lowering the index (more small UDGs) lowers my estimate of the number of UDGs I expect to observe by around 15%. Conversely, increasing the index (more large sources) causes the predicted number to increase by around 25%.

An additional source of uncertainty is that which arises from my estimate of the

mass formation efficiency. From the empirical relation of van der Burg et al. (2017), I estimate a ∼20% error. This uncertainty also propagates to my estimate of the field density of UDGs.

A separate issue is how the assumed stellar population (i.e. that which defines the k-corrections) affects my analysis. For the UDGs, I have explored a red and blue colour model, using quiescent and star-forming populations for the k-corrections respectively. One of the uncertainties for the star-forming population model is the star formation history to assume; for this analysis I have assumed a uniform star formation rate. As a means to test whether this assumption impacts my result, I can also model the blue UDG population using the quiescent model for k-corrections. The results of this are shown in figure 3.7. From this figure, it appears that the change is small, with a slight shift towards redder colours. The impact on my analysis is negligible; this is not surprising as most of my observed UDGs are expected to be at low redshift where k-corrections are small.

I have repeated a similar process for my late type interloper model, replacing the stellar population model with that used for star-forming UDGs. I also find that this makes no significant difference to my results.

3.6.2 Extending the model

One shortcoming of my analysis is that I do not account for galaxies smaller than the fiducial UDG size limit of 1.5 kpc, a fairly arbitrary cut-off. One way to probe how the inclusion of such galaxies may alter the results presented in figure 3.5 is to extrapolate the empirical size distribution that I use in my UDG model (van der Burg et al., 2017) to lower size limits. This exercise is approximate in nature because it is not clear whether an extrapolation of this relation is valid for smaller galaxies. There are two competing effects: while smaller galaxies are more numerous because of the steep power law (equation 3.4), their smaller size means that they are much less likely to be observed given my selection criteria and the corresponding recovery efficiency.

I display the result of reducing the lower physical size limit from 1.5 kpc to 0.5 kpc in figure 3.8. The total number of sources I generate are increased proportionally to equation 3.4. Despite the increase in the number of sources (by a factor of ∼20), there is only a minor difference between the number of sources I would expect to observe.

As a further extension to the model, I can consider what happens when I decrease the lower limit in surface brightness at $\bar{\mu}_e$=24 to allow brighter sources into the selection. Since Danieli & van Dokkum (2018) have shown that the distribution of

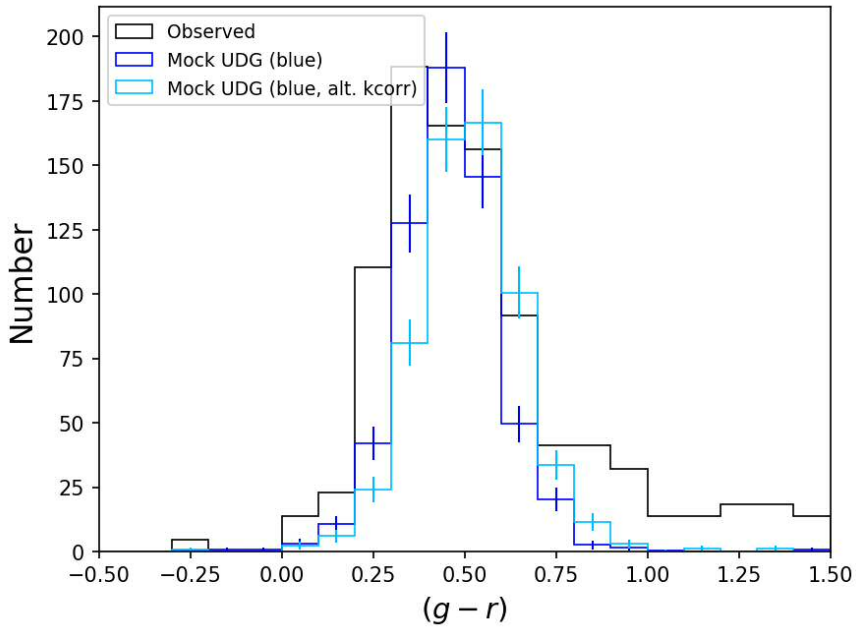

Figure 3.7: Comparison of the model colour distributions for blue UDGs assuming a star-forming stellar population for the k-corrections (dark blue), and a quiescent one (light blue). Overall, I find the difference to be negligible for my analysis.

intrinsic size of large red galaxies is approximately uniform with absolute magnitude (and therefore surface brightness), this extension can be interpreted as including large red galaxies with Sérsic indices meeting my selection criteria. As an example, I show in the figure the effect of using a bright-end surface brightness cut of $\bar{\mu}_e$=22, increasing the number of sources by 80% according to equation 3.3. As in the previous test, the difference with the result in figure 3.5 is fairly insignificant. I am left to conclude that my observational sample is indeed likely made up of large low surface brightness galaxies.

3.6.3 Comparison with GAMA redshifts

We can compare my measurements and mock catalogues against measurements from the GAMA spectroscopic survey in order to test how well my mock catalogues represent reality. For this test, I use my best model: the combination of blue UDGs with massive blue galaxy interlopers. I assume that UDGs form with a mass efficiency as calculated in §3.5.1.

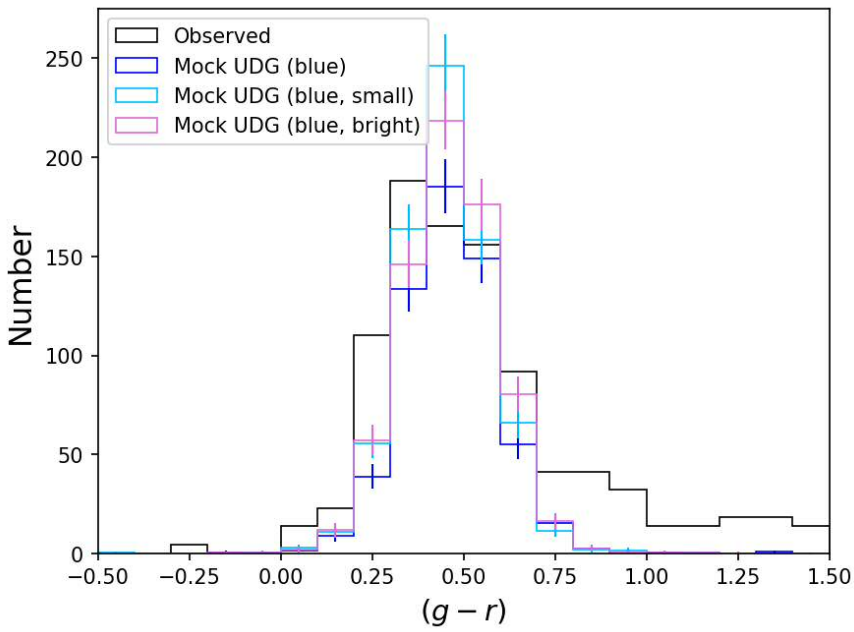

Figure 3.8: Results of extending my empirical model. The observed colour histogram is shown in back. The blue UDG model is shown in dark blue. The mock catalogue with the reduced lower limit of $\bar{r}_e \geq 0.5$ kpc is shown in light blue. The mock catalogue that includes bright galaxies ($\bar{\mu}_e \geq 22$) is shown in purple. All error-bars are Poisson uncertainties.

Using the public data release 3 data obtained from the GAMA website[9], I cross-matched the `SpecObj v27` catalogue (containing spectroscopic redshifts) with the `SersicCatSDSS v09` table (containing Sérsic profile fits to GAMA targets in SDSS data from Kelvin et al., 2012). I imposed my selection criteria on the Sérsic parameters and additionally required `SURVEY_CLASS`≥ 4 in order to select legitimate sources with $m_r < 19.8$, leaving us with 209 GAMA sources. I also applied the $m_r < 19.8$ criterion to my mock catalogue, retrieving 379 sources. The results of the comparison are shown in the top panel of figure 3.9.

Despite the surface brightness limits of GAMA (e.g. Wright et al., 2017), I find that 45 of my UDG candidates (over the full unmasked KiDS area) have matches in the GAMA `SpecObj` catalogue within $3''$. This allows us to make the same comparison as above, using my measurements in place of those of Kelvin et al. (2012); the result is also shown in the top panel of figure 3.9. Clearly the overall distribution of my mock

[9] http://www.gama-survey.org/

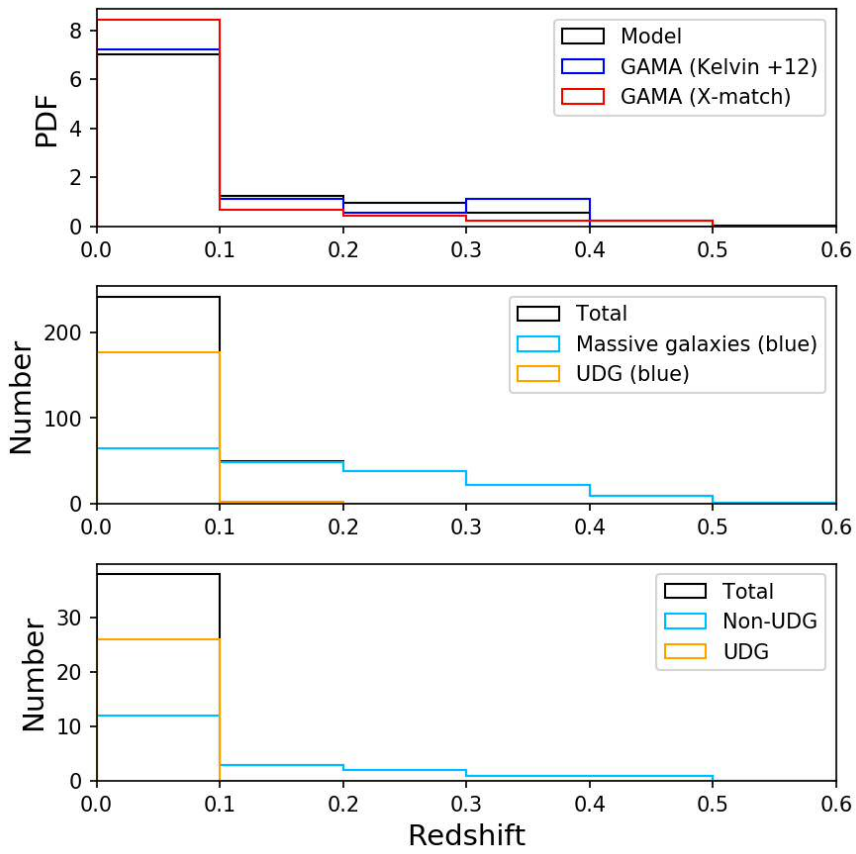

Figure 3.9: *Top:* Comparison of the overall redshift distribution from my mock catalogue (black histogram) with that of Kelvin et al. (2012) (blue histogram) and a crossmatch between my observed UDG candidate catalogue and the GAMA spectroscopic survey (red histogram). *Middle:* Decomposition of my mock catalogue into UDGs (orange) and interlopers (blue) as a function of redshift. *Bottom:* The same as the middle panel, but for my crossmatch with GAMA.

catalogues is in good agreement with the observations.

However, I can go one step further; I can now estimate the intrinsic parameters of the matching sources because I know their redshifts and apparent structural parameters. In doing so, I can test whether the individual distributions for UDGs and interlopers are approximately correct. For this test, I define a UDG as having $\bar{r}_e \geq 1.5$ kpc and $M_* \leq 10^9 \mathrm{M}_\odot$. The stellar mass was estimated assuming my blue UDG stellar population model from 3.4.1 together with the GALFIT m_r measurements. I can then decompose the catalogues into UDG and non-UDG populations.

The decomposition of the mock catalogue is shown in the middle panel of figure 3.9. This compares with the decomposed observed catalogue in the lower panel. Clearly the distributions are similar; at low redshifts UDGs dominate my sample, while at higher redshifts, massive interlopers dominate.

3.6.4 Comparison with HI-bearing UDGs

In this section, I compare my observed, contaminant-corrected $(g-r)$ histogram to other measurements from the literature. I do not consider values of $(g-r) > 1$ because they are almost certainly not part of the UDG population. One catalogue that I can directly compare with is that of Leisman et al. (2017)[10], who measured the colours of isolated HI-bearing UDGs using SDSS data. While these measurements are conflated with measurement error because of the limited depth of SDSS, I can perform a qualitative comparison between the reported results (figure 3.10).

In figure 3.10, I show how the colour distribution of my corrected observational sample of UDG candidates compares with that of Leisman et al. (2017). The two PDFs are very similar, providing an indication that UDGs in the field are predominantly blue independently of the colour models I assumed in §3.4.1. There are some differences between the two distributions: I observe slightly more UDG candidates on the red side of the peak. There are several possible explanations: they could be sources that Leisman et al. (2017) were not sensitive to thanks to low HI content, or they are contaminant objects that I have not properly accounted for in my UDG sample such as massive early-type galaxies. Alternatively, since I am sensitive to UDGs out to $z \sim 0.2$ comparing to their maximum distance of 120 Mpc ($z \sim 0.03$), it may be that k-corrections play a role. The Leisman et al. (2017) catalogue also seems to have an excess of blue UDGs compared to what I observe. This could be explained either by the measurement error arising from the limited SDSS depth, or perhaps because blue UDGs are intrinsically brighter and I miss them in my selection (Leisman et al. (2017) use a slightly brighter bright-end selection cut).

Jones et al. (2018) used the Leisman et al. (2017) catalogue to estimate the field density of HI-bearing UDGs, obtaining a value of $1.5\pm0.6\times10^{-3}\mathrm{cMpc}^{-3}$. This is approximately one fifth of my upper-limit estimate of the overall UDG field density. However, comparing such field densities is difficult because the limited depth of the SDSS imaging used by Leisman et al. (2017) to identify UDGs creates significant

[10]We note that I use the HUD-B sample, which contains 115 sources and was selected using selection criteria consistent with that of van der Burg et al. (2016).

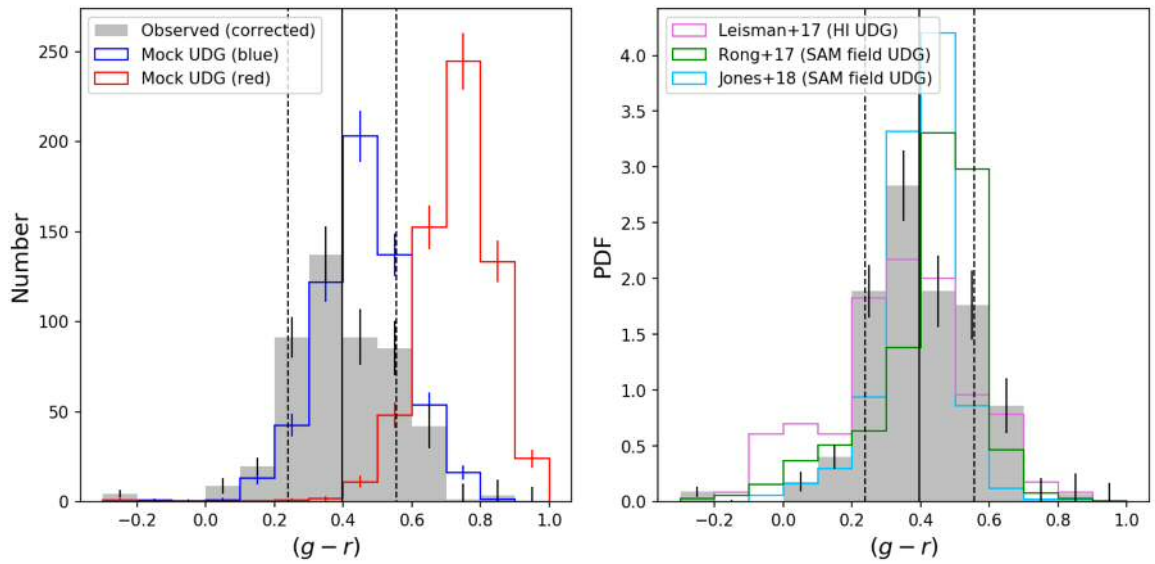

Figure 3.10: The observed distribution of colour after subtracting the estimated contribution from massive blue galaxies (grey histogram). *Left:* Comparison with the empirical red and blue UDG models from this study. I show the mean and 1σ dispersion of the observations with the vertical lines. *Right:* Normalised comparison with observations of HI-bearing field UDGs in the literature (Leisman et al., 2017). and predictions from the semi-analytical models (SAM) of Rong et al. (2017) and Jones et al. (2018). Poisson error-bars are shown. See text for discussion.

measurement uncertainty, blurring their selection boundaries and leading to an uncertainty of ∼25% in their sample size. However, using the fact that my estimate of the UDG density is an upper-limit, I can estimate that HI-bearing UDGs comprise at least one-fifth of the overall population.

3.6.5 Comparison with Semi-Analytic Models

I also compare my results with the work of Jones et al. (2018), who used the Santa Cruz semi-analytic model (SAM, Somerville et al., 2015) to generate a UDG sample in order to compare with the observations of Leisman et al. (2017). Their results are also displayed in figure 3.10. I note that I jiggle their $(g-r)$ colours to match my measurement error for the comparison. The peak of their $(g-r)$ distribution is in reasonable agreement with my observations, yet it is slightly shifted towards the red and narrower. This may be because my observed catalogue is not entirely made from UDGs but also contains some contaminant sources, or perhaps because the SAM does

not reproduce the correct amount of scatter for UDG colours. Alternatively, it may be a projection effect; my observed colours are in the observed frame and therefore are subject to k-corrections.

A similar comparison can be made with the work of Rong et al. (2017), who used the Guo et al. (2013) SAM to obtain a catalogue of simulated UDGs. I again jiggle their colours using my measurement error for the comparison. The colour distribution of their field UDGs is shifted towards the red compared to my observations, as also noted by Jones et al. (2018). If I were to include additional reddening of their colours because of k-corrections (i.e. to make a fair comparison with my observations), then this discrepancy would be exaggerated.

I calculate the total field density of UDGs in the Santa-Cruz SAM by integrating the stellar-mass function for UDGs (figure 4 of Jones et al., 2018). I obtain a value of $2{\times}10^{-2}\mathrm{cMpc}^{-3}$, approximately twice the upper-limit estimate from my measurements.

3.6.6 Impact of Nearby Galaxy Groups

While my observed catalogue of UDG candidates is dominated by field sources, it is important to consider the effects of nearby galaxy groups on my result. After all, if such sources are predominantly quiescent and exist in similar number to my field sample, I should expect to find a population of red UDG candidates.

Fortunately, my KiDS/HSC-SSP footprint overlaps with the GAMA spectroscopic survey (Driver et al., 2011) and thus the GAMA group catalogue (Robotham et al., 2011). I can therefore make an estimate of the contribution of group/cluster UDGs to my field sample using a similar method to van der Burg et al. (2017). Working in my favour is the fact that I have applied a relatively high cut in angular size ($\bar{r}_e{\geq}3''$) and there are no massive groups that are close enough to dominate my selection.

I select all GAMA groups from the group catalogue that satisfy $\texttt{Nfof}{\geq}5$ and $0.01{<}\texttt{Zfof}{<}0.2$ (where `Nfof` and `Zfof` are respectively the number of friends-of-friends sources and an estimate of the group redshift). For each group, I count the number of sources that are compatible with being UDGs at that redshift, using my selection criteria and a physical radius range of $1.5{\leq}\bar{r}_e$ [kpc]$\leq$7.0. I also subtract a statistical background correction based on the total number of sources across the entire footprint that meet the same criteria. I estimate that up to 8% of my UDG candidate catalogue is associated with relatively massive groups (i.e. the ones that have at least five friends-of-friends members), with the uncertainty coming from the background count estimate. The colour histogram of these sources is statistically indistinguishable from

that of the whole catalogue, and I conclude that their inclusion does not significantly impact my results.

Using the empirical scaling relation between the group mass and total r-band luminosity from Viola et al. (2015), I estimate that $\sim 6\%$ of the available mass out to $z = 0.2$ is taken up by the groups I consider here. This is very similar to the fraction of observed UDG candidates associated with groups. Taken with the fact that I expect all observed UDGs to be at $z < 0.2$, this provides an independent indication that the mass formation efficiency of UDGs in the field is comparable to that in groups and clusters. It also shows that the presence of the massive groups does not severely impact my result.

3.7 Summary

In this chapter I have used deep, wide-area optical imaging from the KiDS survey to detect sources with low surface brightness ($24.0 \leq \bar{\mu}_{e,r} \leq 26.5$) and large angular sizes ($3.0" \leq \bar{r}_{e,r} \leq 8.0"$). Following the detection and measurement of these sources with MTOBJECTS and GALFIT, I measured colours using the HSC-SSP survey data. My catalogue of UDG candidates consists of 212 sources over $\sim$39 square degrees. Compared to UDGs in groups and clusters, my sample consists of sources that appear to have much more irregular morphologies and show hints of active star formation.

These observations are compared to mock observations of UDGs created by sampling empirical distributions of UDG properties from the literature. My key assumptions are intrinsic size, surface brightness and colour distributions for the UDGs. All the assumptions I made are justified based on the current understanding of UDGs.

By comparing my mock catalogues with the observations, I have shown that it is very unlikely for a significant population of UDGs that are as red in colour as they are in clusters to exist in the field. It is much more likely that almost all UDGs in the field are instead much bluer, with colours similar to late-type dwarf galaxies in clusters. An immediate conclusion based on the predominantly blue colours is that secular evolutionary processes are not producing large numbers of cluster-like quenched (red) UDGs.

This finding means that isolated red UDGs, like the ones found by Martínez-Delgado et al. (2016) and Román et al. (2019a), should be quite rare. At first glance this contrasts with the work of (Singh et al., 2019), who find a population of UDGs with low specific star formation rates (compared to the star forming main sequence) in the field. Taking the appearance of my detections in figure 3.2 into account, it is

likely that UDGs in the field are forming stars only in a few localised sites; this locality may result in relatively low specific star formation rates that nevertheless make their integrated colours bluer. As noted by (Singh et al., 2019) and (Zaritsky et al., 2018), their sample of field UDGs is still systematically bluer than that observed for cluster UDGs. Assuming that UDGs across different environments share similar metallicities, this is good evidence that star formation in field UDGs can be quite tentative and easily quenched in cluster environments.

I also created mock observations of massive blue galaxies, thought to be the primary source of contamination in my UDG candidate sample, using canonical empirical relations. I statistically subtracted these from my observations to acquire a contaminant-corrected catalogue of UDGs. The normalised distribution of $(g-r)$ colour is very similar to that estimated for HI bearing field UDGs measured by Leisman et al. (2017). The observed distribution is also similar to that predicted for UDGs in SAMs (Rong et al., 2017; Jones et al., 2018), but slightly bluer. While my colour distribution appears to have greater dispersion, this is likely due to systematic shortcomings in comparing simulations with observations.

Using my mock catalogues as a reference, I estimate an upper limit on the field density of UDGs as $8\pm3\times10^{-3}$cMpc^{-3}, equivalent to a mass formation efficiency $\sim0.8\pm0.2$ times that in clusters. Perhaps surprisingly, this density actually implies that UDGs form with a mass efficiency in the field that is quite close to what they do in cluster environments. The field density applies for UDGs with physical sizes $1.5\leq\bar{r}_{e,r}[\mathrm{kpc}]\leq7.0$ and intrinsic (i.e. not cosmologically dimmed) surface brightnesses $24.0\leq\bar{\mu}_{e,r}\leq26.5$. This number density also suggests that current SAMs over-predict the number of UDGs by at least a factor of two. However, I note that if UDGs exist in abundance at lower surface brightnesses than what I have probed here, the total number density of large LSB objects could be much higher. Based on the field density measured by Jones et al. (2018), HI-bearing UDGs comprise at least one fifth of the overall UDG population in the field. This is consistent with what is predicted from the Santa-Cruz SAM.

I note that the analysis I have performed in this work has been approximate in nature because of the absence of any distance measurements. Acquiring large samples of spectroscopic redshifts for LSB galaxies in the field is not currently feasible. In the near-term the second data-release of the HSC-SSP will provide an opportunity to follow-up the present work thanks to its expanded footprint; this analysis can easily be expanded to larger areas. In the longer term, deep all-sky imaging (perhaps combined with photometric redshifts) from LSST may provide the ultimate data set for providing

statistical constraints on LSB galaxies in the field.

Chapter 4

Halo Mass Estimates

4.1 Introduction

The intrinsic properties and formation histories of UDGs are still not fully understood and it is not clear whether they represent a distinct population from smaller LSB dwarf galaxies, which can form naturally in high spin halos expected from hierarchical galaxy formation models (e.g. Dalcanton et al., 1995; Jimenez et al., 1998) and from harassment of normal dwarf galaxies (Moore et al., 1998; Mastropietro et al., 2005).

In chapter 3, I showed that blue UDGs are prominent in the field, with quiescent UDGs occupying a relatively small fraction of the overall number. This suggests that UDGs are quickly quenched in the cluster environment (e.g. from ram pressure stripping) as is supported by zoom-in cosmological simulations. I also showed that I cannot exclude the possibility that field UDGs are forming with mass formation efficiencies comparable to those in galaxy groups and clusters. These results could be interpreted to suggest that most UDGs are not quenched massive galaxies because a correspondingly dense population of blue UDGs seem to exist in the field, where quenching mechanisms are expected to be less significant. This would make UDGs with L$*$-like halo masses, e.g. DF44 (van Dokkum et al., 2016), a minority. The most direct property of UDGs that can be measured (or at least estimated) to address this question is the total mass M_{halo}; if UDGs are massive galaxies that were quenched early (perhaps during cluster in-fall), then they should have retained their dark matter content, which is not sensitive to the stripping processes that affect the gas.

There are several possible mechanisms to explain the existence of UDGs other than the failed L$\star$ scenario. The halo mass is a key parameter in distinguishing between formation models of UDGs. Typically, current models favour dwarf-sized halos with truncated star formation histories (e.g. Rong et al., 2017; Amorisco & Loeb, 2016),

making them similar to normal LSB galaxies but larger. UDGs are abundant in high density environments such as in the centres of clusters (e.g. Mihos et al., 2015; Koda et al., 2015; Venhola et al., 2017) where they require a relatively high dark matter fraction in order to survive. However, it is not clear whether UDGs can form with lower mass-to-light ratios (M/L) in less dense environments such as the field (van Dokkum et al., 2018; Trujillo et al., 2018).

Globular clusters offer an interesting insight into the formation mechanisms of LSB galaxies. They are thought to form mainly in the early epochs of star formation within massive, dense giant molecular clouds that are able to survive feedback processes that might otherwise shut off star formation in their host galaxy (Hudson et al., 2014; Harris et al., 2017). The halo mass of galaxies has been shown to correlate well with both the number of associated GCs ($N_{\rm GC}$) and the total mass of their GC systems ($M_{\rm GC}$; e.g. Spitler & Forbes, 2009; Harris et al., 2013, 2017), which means measurements of either $N_{\rm GC}$ or $M_{\rm GC}$ can be used to constrain M_{halo}. However, Forbes et al. (2018) show that the traditional relation between $N_{\rm GC}$ and M_{halo} may lose accuracy in the low M_{halo} regime, perhaps because lower mass galaxies tend to have lower mass GCs without a common mean GC mass. Additionally, it has been shown that there is a correlation between the GC half-count radius and M_{halo} (Forbes, 2017; Hudson & Robison, 2018).

The majority of studies of the GC populations of UDGs have up until now focussed on the Coma galaxy cluster, the most massive ($M_{\rm tot} \sim 6 \times 10^{14} {\rm M}_{\odot}$, Hughes, 1998) galaxy cluster within 100 Mpc. In this chapter I analyse exclusively galaxies in the core of the Fornax cluster. In comparison to Coma, it is around five times closer ($d \sim$20 Mpc, Blakeslee et al., 2009) but less massive ($M_{\rm tot} \sim 7 \times 10^{13} M_{\odot}$, Drinkwater et al., 2001). Using the empirical relation of van der Burg et al. (2017), there are approximately 10 times less UDGs expected in Fornax than in Coma, many of which have been catalogued already (Muñoz et al., 2015; Venhola et al., 2017).

While overall I have a relatively small sample of UDGs for the present study, an advantage of working with the Fornax cluster is that cluster members have much larger projected sizes compared to the background galaxy population, so I can analyse the population of smaller LSB galaxies at the same time as the UDG population without contamination from interlopers. Indeed, much of the new literature surrounding LSB galaxies focusses on UDGs and this may be in-part due to the relative ease of distinguishing larger galaxies from background objects in group or cluster environments. A second advantage of Fornax over Coma is that GCs are brighter in apparent magnitude by $\sim$3.5 mag due to their relative proximity, meaning that I can probe further into the GC luminosity function.

I note that the relatively large number of galaxies I analyse in this study is important for at least partially overcoming systematic uncertainties involved in measuring halo masses with low numbers of tracers as made clear by Laporte et al. (2019) and the possible stochastic nature of the $M_* - M_{halo}$ relation at low mass (Brook et al., 2014; Errani et al., 2018).

In this work I provide constraints on the halo masses for a selection of LSB galaxies first identified by Venhola et al. (2017) using the optical Fornax Deep Survey (FDS, Iodice et al., 2016). The structure of the chapter is as follows: I describe the data in §4.2. In §4.3 I describe the method to detect globular cluster candidates (GCCs) and infer the total number of GCs associated with my target galaxies. I provide my results in §4.5, where I estimate the halo masses from the inferences on $N_{\rm GC}$ and $M_{\rm GC}$ using the empirical scaling relations of Harris et al. (2017). I discuss my results and provide conclusive remarks in §4.8. I use the AB magnitude system throughout the chapter, and adopt a distance of 20Mpc to the Fornax cluster.

4.2 Data

We use the four central 1×1 degree^2 frames of the FDS (FDS IDs 10, 11, 12 & 16), i.e. the same region used by Venhola et al. (2017) in their by-eye classification of low surface brightness sources in the Fornax galaxy cluster. These data were obtained using the OmegaCAM (Kuijken, 2011) instrument on the 2.6m ESO VLT Survey Telescope (VST, Capaccioli et al., 2012) in the u', g', r' & i' bands. I note that Fornax GCs are unresolved in my data such that I consider them as point sources throughout the chapter.

I specifically used the VSTtube-reduced FDS data (Grado et al., 2012; Capaccioli et al., 2015), which is optimised for point-source photometry but is not as deep as the data used by Venhola et al. (2017), which is reduced using a combination of the OmegaCAM pipeline and AstroWISE (McFarland et al., 2011), but with a slightly wider PSF than the VST-tube reduction because images with poor seeing were included in the stacks.

I performed additional photometric corrections to bring my photometry into the AB magnitude system. Starting from the VSTtube-reduced data, I used ProFound to detect and select point sources (see §4.4.1). I additionally measured fixed-aperture magnitudes for each source with an estimate of that magnitude. These aperture magnitudes had to be corrected for both the limited size of the aperture with respect to the PSF in each band, but also the absolute calibration to AB magnitudes.

While there is no ideal set of standard stars in my footprint with which to calibrate the photometry, Cantiello et al. (in prep) have used a set of existing, overlapping calibrated catalogues (ACSFCS (Jordán et al., 2007); APASS (Henden et al., 2012); SkyMapper (Wolf et al., 2018)) to calibrate their photometry in the same data. I have calibrated my own aperture magnitudes by matching my catalogue with theirs, selecting point sources as in §4.4.1 and applying a multiplicative correction to my measurements to nullify the mean offset between the measurement pairs. The RMS between my corrected aperture magnitudes with theirs is ∼0.05 mag in g, r, i and $\sim$ 0.2 mag in the shallower u-band, for all matching point sources with corrected g magnitudes brighter than 23 mag.

During the calibration it was noticed that the reference catalogue of Cantiello et al. (in prep.) contained minor systematic offsets in the stellar locus between individual FDS frames, suggesting a systematic error in the absolute calibration. I have dealt with this by shifting each locus to a common position in colour-colour space. The net result of this is a maximum systematic uncertainty of $\sim \pm 0.05$ mag in each colour plane. Finally, I note that since there is currently no available reference catalogue for FDS frame 12 I calibrated the photometry for that frame in accordance with FDS frame 11. This calibration is accurate enough to have negligible effects on my results.

4.3 Methodology

In this work I target the GC populations of galaxies identified by-eye in the (Venhola et al., 2017) (hereafter V17) catalogue. I split the sample into two groups: Low surface brightness galaxies (LSBGs), defined as those with r-band effective radii $r_{e,r} < 1.5$ kpc and UDGs, defined as those with $r_{e,r} \geq 1.5$ kpc. The sources are defined as LSB because they were measured to have central surface brightness $\mu_0^r \geq 23$ by V17. I omit two UDGs (FDS11_LSB1 and FDS11_LSB17) from the sample because they are in significantly crowded locations and measuring their properties accurately would require a more sophisticated analysis.

Before running my detection algorithm, I subtract model galaxy profiles in each band using Imfit (Erwin, 2015, see §4.4). I was unable to get a stable Imfit model for three sources (FDS11_LSB16, FDS12_42, FDS12_47) because they were too faint, so I adopt the measurements of V17 (made from deeper stacks) for these sources and rely on a separate background subtraction procedure to remove the galaxy light (see §4.4.2). I select only galaxies with measured r-band effective radii greater than 3″ ($\sim$ 0.3kpc at Fornax distance) so that I target cluster members with confidence (Sabatini

Parameter	Constraint
`mag` (g)	14 to 19 [mag]
`axrat`	≥ 0.95
`Nobject`	0
`Nmask`	0

Table 4.1: ProFound measurement constraints for point source selection in the PSF modelling. `axrat` is the axis-ratio. `Nobject` is the number of pixels belonging to the segment that are touching another source. `Nmask` is the number of pixels belonging to the segment that are touching a masked region. See the ProFound documentation for more details of these parameters. Further criteria are discussed in the text.

et al., 2003; Davies et al., 2016).

I used the ProFound[1] package (Robotham et al., 2018) for the source detection and photometry, with the following settings: `box=100` pixels (∼20"), `sigma=2` pixels, `threshold=1.03`, `tolerance=1`, `skycut=1`. All other settings were defaults. My detection was performed exclusively on the g-band (the deepest) so that I could easily measure and account for my detection efficiency without considering the colours of individual sources. I note that I split the four FDS frames into 9×9 subframes to ease the memory requirements for ProFound.

4.3.1 PSF models

I obtained point spread function (PSF) models for each band and subframe using my ProFound measurements as follows. Bright, unsaturated point sources were selected in the ProFound `mag - R50` (approx. half-light radius) plane, using the selection criteria listed in table 4.1. Additionally I sigma-clipped the measurements in `R50` (approximately flat over the magnitude range for point sources) and offset the relation by 4σ with respect to the median to measure an upper-limit on `R50` for the selection.

I used Imfit to fit a model Moffat profile (keeping the axis-ratio as a free parameter) to each point source following a local sky subtraction. I did not stack individual point source cut-outs to avoid artificial widening of the PSF caused by misalignment of the images. The resulting distribution of model Moffat fits was then sigma-clipped at 3σ in the FWHM-concentration index plane to remove outliers caused by bad fits. I finally selected a fiducial model PSF for each band and subframe by adopting the fit with the average FWHM.

[1] https://github.com/asgr/ProFound

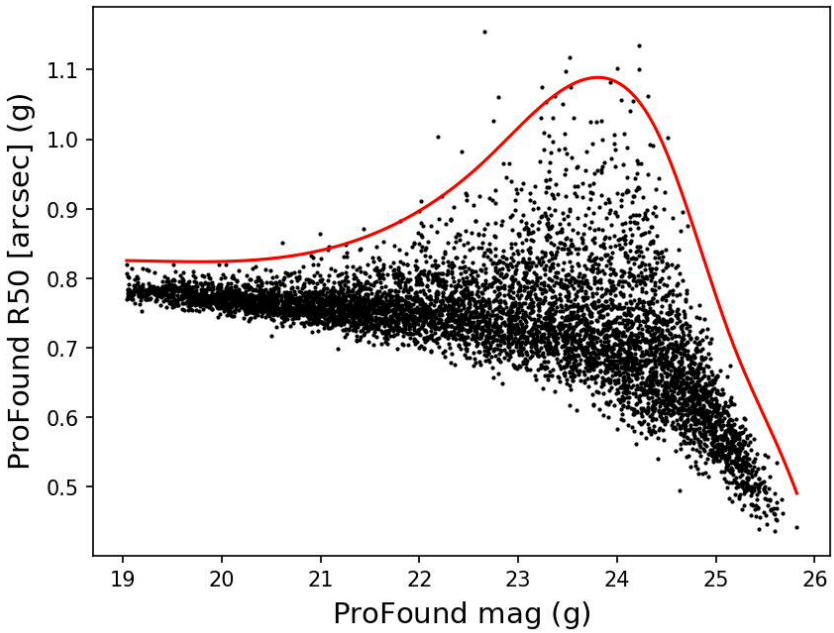

Figure 4.1: Point source selection function for a subframe (red line) obtained by fitting a cubic spline to measurements from the synthetic source injections (black points) after offsetting `R50` by 4σ in bins of magnitude with respect to the median. All points lower than the red line are selected. The increase in scatter towards the faint end means that a simple cut in `R50` would either result in a low recovery efficiency or high level of contamination from my point source selection.

Of primary importance for my analysis are the g-band PSF models. While for a specific FDS frame I found little variation of the PSF over its subframes, on a frame-by-frame basis the IMFIT FWHM ranges between approximately 0.7 and 1.2″.

4.4 Galaxy Modelling

I used IMFIT to fit single Sérsic profiles to each target galaxy. Fortunately, Venhola et al. (2017) (hereafter V17) already provide such fits in the r-band. While here I choose to remeasure the profiles for consistency with the other bands, I do make use of these data as initial guesses in the fitting. My approach was to iteratively fit the galaxy in the r-band, each time improving the mask of pixels to ignore in the fit. The general procedure for a galaxy is as follows:

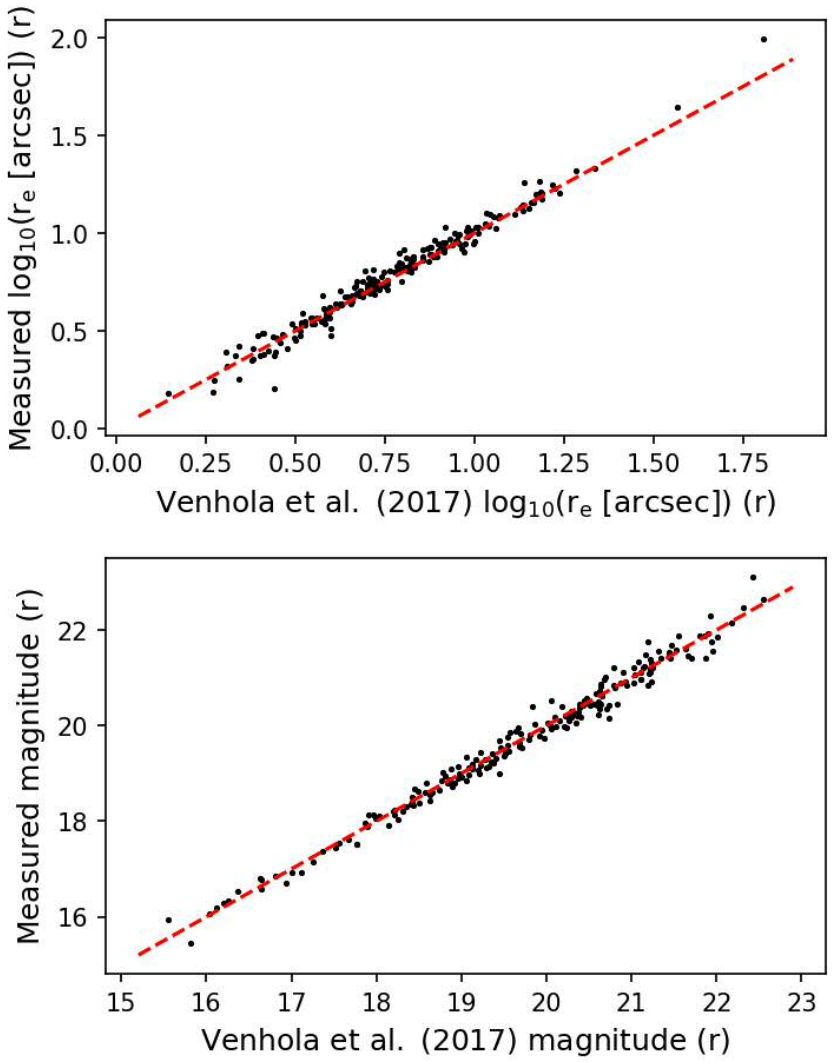

Figure 4.2: Our galaxy measurements (IMFIT) vs. those of V17 (`GALFIT`). Red dashed line: The one-to-one relation. The RMS of the residuals are 0.18 mag and 0.04 dex in r_e.

(i) Obtain an $8 \times 8r^r_{e,V17}$ r-band cut-out.

(ii) Subtract the V17 model from the result (include nuclear PSF if indicated by V17).

(iii) Use DEEPSCAN[2] (chapter 2) to get sky and RMS estimates from the result.

(iv) Create a smoothed image by applying a Gaussian filter with RMS=2 pixels.

(v) Mask all pixels with $\geq 6\sigma$ significance on the smoothed image.

(vi) Use IMFIT to fit a Sérsic model to the original data with the sky subtracted,

[2]https://github.com/danjampro/DeepScan

ignoring pixels in the mask.

(vii) Repeat steps 2 to 6 three times, each time updating the model image and mask.

(viii) Repeat steps 1, 3 & 6 for the other bands, using the same r-band mask in each.

For the DeepScan sky estimates I used a mesh size equal to the image size and performed three masking iterations. If the galaxy was indicated as nucleated by V17, I also fit a Moffat profile simultaneously with the Sérsic model. I found that in a minority of cases the residuals from the V17 fits were quite large, such that I had to modify the masks manually.

In the case of FDS11_LSB2, the largest galaxy in my sample (with $r^r_{e,V17} = 76''$), I re-binned the data by a factor of 5 to make the fit easier (the original fitting region was $\sim 3000 \times 3000$ pixels). Over this region the sky background level varies significantly, so I modified the DeepScan sky modelling to use mesh sizes of $\sim 200''$ and median filtered in 3×3 meshes. I note that an image of FDS11_LSB2 is displayed in figure 20 of V17.

Overall my results are consistent with V17 (figure 4.2), with a few exceptions. These include FDS11_LSB2, which I measure to be 1.5 times larger than originally reported. This result was robust against changes in the size of the background mesh. I also note that I measure a slightly lower Sérsic index n for this object, and n is generally anti-correlated with r_e. This discrepancy likely arises from the difficulty involved in measuring such a large, diffuse galaxy in a reasonably crowded field with a varying sky; I use `DeepScan` whereas V17 fit a 2D sky plane in their GALFIT (Peng et al., 2002) modelling. I also note that V17 did not leave the central coordinate of their model profiles as a free parameter.

Finally, I note that I was not able to obtain stable Imfit models for several sources because they were too faint: FDS12_LSB42, FDS12_LSB47, FDS11_LSB16 & FDS12_LSB34. I therefore adopted the fits of V17 for these sources[3].

4.4.1 Point source selection

I used synthetic source injections based on my Moffat PSF models from §4.3.1 to produce my point source selection function and quantify my recovery efficiency (RE). I injected ∼25000 synthetic profiles per subframe into the real data at random locations in the vicinities ($420 \times 420''$ cut-outs) of my target galaxies after subtracting the galaxy models from the data. This was done in the g-band, with apparent magnitudes ranging

[3]Since V17 did not measure $(g - i)$ colours, I have omitted them from stellar mass calculations and from figure 4.9.

between 19 and 26. My matching criteria for the synthetic sources was that the central coordinate of the injected source had to lie on top of a segment in the ProFound segmentation map. Additionally, I only considered sources that did not match with segments from the result of running ProFound over the original frames (i.e. without the injected sources). I note that the measurements of the synthetic point sources are in good agreement with measurements of real sources when plotted on the `mag-R50` plane.

Once I had acquired the ProFound measurements of the synthetic sources, I fitted a smooth cubic spline to the data in the ProFound `mag` and `R50` plane. Specifically, the spline was fit to the data binned in `mag`, and positively offset by 4σ of the `R50` values within the bin. See figure 4.1. The rationale behind this was that as sources become fainter, the scatter in `R50` increases such that a simple cut at a specific value would either be too high for bright objects or conversely too low for some of the fainter objects with large values of `R50`. I obtained a different point source selection function for each subframe.

4.4.2 Colour measurement

I obtained aperture magnitudes of the ProFound sources in fixed apertures of diameter 5 pixels in all the bands. The sky level and its uncertainty were calculated for each detected source by placing many identical apertures in 51×51 pixel cut-outs (the fiducial FWHM of $1''$ is ~ 5 pixels) and recording the median and standard-deviation of the contained flux values after sigma clipping these at 2σ to remove contamination from other sources. Additionally the sky apertures were placed at radii greater than 20 pixels from the centre of the source.

These magnitudes were then corrected for the PSF size in each band through calibration against an existing catalogue of PSF-corrected point sources made using the same data (Cantiello et al.; in prep.). I note that I have not used this catalogue for this work because of the need to subtract galaxy profiles from the data and the need to quantify the RE.

4.4.3 Recovery Efficiency

I quantified the RE separately for each subframe using the point source selection functions with the synthetic source measurements. I imposed a faint-end limit on the corrected g-band aperture magnitude of 25 mag because measuring accurate colours at fainter magnitudes is more difficult and because the degeneracy between point sources

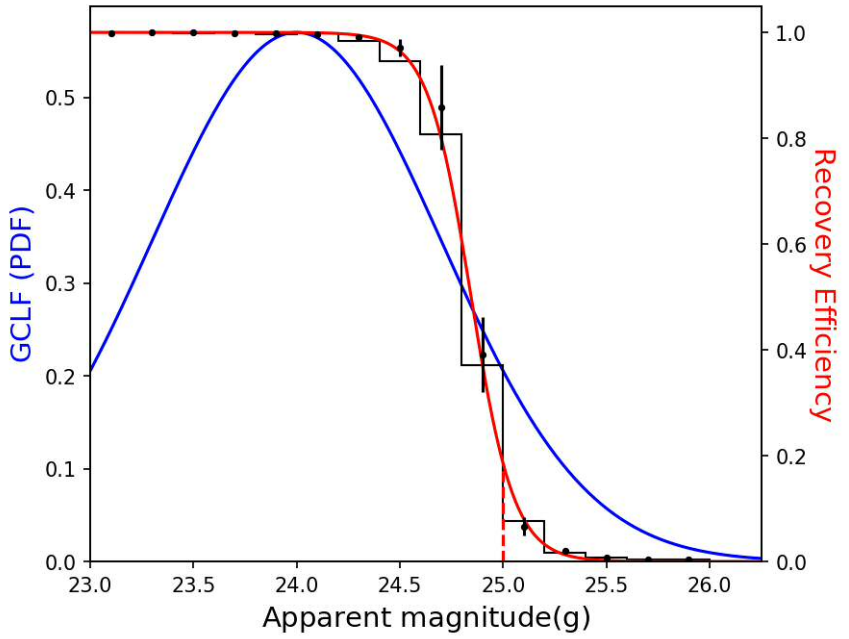

Figure 4.3: g-band recovery efficiency for a subframe of injected point sources (black histogram) after applying my selection criteria, along with the RE (red line) and a fiducial GCLF (blue line). The dashed red line shows my additional cut at m_g=25. The GCLF turnover magnitude is clearly reached by my detection pipeline. Integrating the recovery efficiency over the GCLF (including the magnitude cuts) yields an overall recovery efficiency of $\sim 80\%$. Also shown is the mean and standard deviation of the RE across all the subframes for FDS frame 10 (black errorbars).

and other faint sources in the `mag-R50` plane is exacerbated in this region. Additionally, I apply a lower-bound cut in the corrected g-band aperture magnitude of 21 mag to reduce possible contamination from bright stars, ultra-compact dwarf galaxies (UCDs) and nuclear star clusters (NSCs).

The RE itself was measured by taking the ratio of detected and selected point source injections to the total number of injected sources in bins of intrinsic magnitude. A sigmoid function,

$$\epsilon(m) = [1 + \exp(-k_1(k_2 - m))]^{-1} \tag{4.1}$$

was fit to the result (see figure 4.3). The recovery efficiency $\epsilon(m)$ is sufficient to reach the turnover magnitude of the g-band GC luminosity function (GCLF), which

is approximated by a Gaussian function centred at $24m_g$ at the distance of Fornax (Villegas et al., 2010). I adopt a value of 0.7 for the GCLF standard deviation, which is a reasonable estimate for low surface brightness galaxies (Trujillo et al., 2018). Under these assumptions, my estimated GC completeness ranges between ∼60% and 90% depending on the subframe. The mean completeness is estimated to be 82% across all the subframes. Of course, this number depends on the exact form of the adopted GCLF. While its peak at $24m_g$ is fairly well known (The peak of the GCLF can sometimes be used as a standard candle, see Rejkuba, 2012), a degree of uncertainty is attributed to its width. I discuss the effects of varying the GCLF on my results in §4.5.6.

4.4.4 Colour Selection

I have applied a colour selection to my point sources to produce a catalogue of globular cluster candidates (GCCs) for each target galaxy, using as few assumptions about the underlying GC colours as possible. The full colour space in u, g, r, i was used for the selection. This is important because of the need to remove interloping point sources from my final GCC catalogue, which include foreground stars and unresolved background galaxies. However, I point out that both interloping populations are partially degenerate in colour space with the actual GCs (see also Pota et al., 2018) and these sources must be accounted for using spatial information (see §4.4.5).

This colour selection was accomplished by first cross-matching my point sources from the four FDS frames with a compilation of spectroscopically confirmed Fornax compact objects (Schuberth et al., 2010; Wittmann et al., 2016; Pota et al., 2018). This resulted in a catalogue of 992 matching sources. I note that I partially account for bright UCDs with my bright-end magnitude cut-off. The external catalogue has a magnitude distribution that drops off quickly at magnitudes fainter than $\sim 23m_g$ and so is not complete for my purposes and this limited depth has to be accounted for.

I used the density-based clustering algorithm `DBSCAN` (Ester et al., 1996) to define regions in the $(u-g)-(g-r)$ and $(g-r)-(g-i)$ planes separately for my GCC selection. I used a clustering radius of 0.1 mag and required at least 5 spectroscopic GCs within this radius for clusters to form. After acquiring the DBSCAN clusters, I fitted a convex hull to all the clustered points and used this as the boundary of the selection box; the results of this are shown in figure 4.4. Approximately 93% of the spectroscopic GCs occupy the selection region and I correct for this factor in my later inferences on $N_{\rm GC}$.

The fraction of GCs that occupy the colour selection box decreases as a function of magnitude because of measurement error. Thus, I have used a probabilistic approach to identify all sources that *could* occupy the box, given their uncertainty. Specifically, I selected all sources that were consistent within 2σ of their measurement uncertainty of the box, separately in each colour-colour plane. While the colour selection box was measured in magnitude units, I actually converted it into linear flux-ratio units (accounting for the photometric calibration) to select GCCs. This was done primarily to overcome the effects of the shallow u-band which would otherwise impact my estimate of the RE. A visual example of my combined point source selection criteria with colour selection for one of my target galaxies is shown in figure 4.5.

I also performed a separate analysis using a much wider colour-selection box, also shown in figure 4.4. I measured the minimum bounding rectangles in each colour-colour plane from the matching sources, forming a 3D colour selection box. The box is bounded by -0.18< $(g-r)$ <1.23, 0.32< $(g-i)$ <2.00, 0.37< $(u-g)$ <5.07; the high upper-limit on $(u-g)$ is likely due to scatter caused by the shallow u-band. While conservative in nature, the box is sufficient to contain all the matching spectroscopic GCs down to $m_g \sim 23$. I note here that my overall results are not significantly impacted by this change. I refer to the results obtained using the `DBSCAN` colour box for the remainder of the chapter.

I note that I do not fit for the intrinsic colour distribution of GCs and interlopers as was done in Amorisco et al. (2018). The reason for this is that simple statistical representations (e.g. Gaussian) are inappropriate to describe my data in the multi-dimensional colour space. This can be gathered from the appearance of figure 4.4. It may be possible to include extra colour-terms in the mixture models described in §4.4.5, but I leave this for future work. This would allow the colours of the GCs to be constrained at the same time as their other properties and increase the quality of the GC selection.

4.4.5 Bayesian Mixture Models

We adopt a simplified version of the Bayesian mixture modelling of Amorisco et al. (2018) to measure the properties of the GC systems of my target galaxies. I am similarly motivated to rescale the spatial coordinates of the GCCs into units of the $1r_{e,r}$ (half-light radius of the galaxy) ellipse. My model consists of two surface density components: A central Plummer profile to represent the GCs associated with the target galaxy,

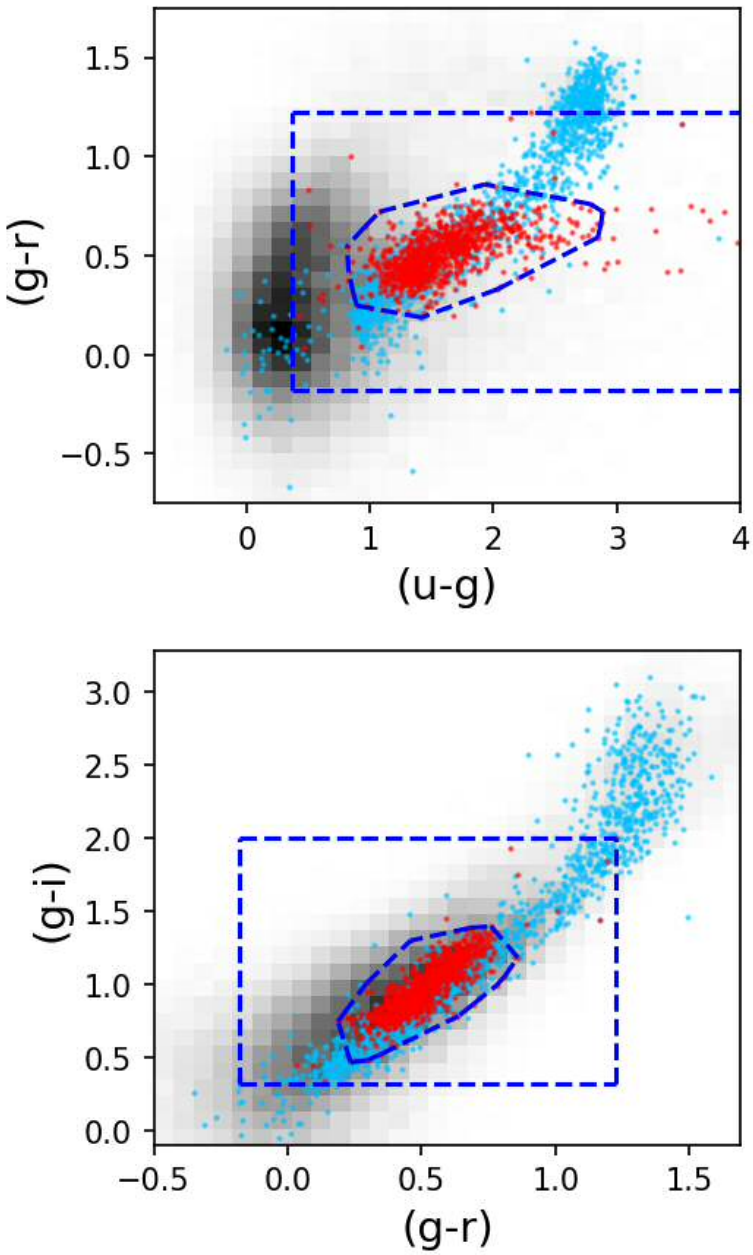

Figure 4.4: Colour-colour measurements of spectroscopic Fornax compact objects (red points), galactic stars (blue points) against the empirical distribution of all detected point sources over a subframe (greyscale histogram). Also shown are two colour selection boxes (blue dashed lines). In each panel, the large box corresponds to the minimum-bounding rectangle of compact object measurements in this colour-colour plane, and the smaller box is produced using the `DBSCAN` algorithm.

$$\Sigma(r, r_h) = \frac{1}{\pi} \frac{1}{r_h^2(1 + r^2/r_h^2)^2} \tag{4.2}$$

where r_h is the half number radius, in units of $r_{e,r}$, and a uniform distribution to represent the background, which mainly consists of stars, background galaxies and intra-cluster GCs. The presence of NGC1399 in the centre of Fornax means that its GC system may contribute to a non-uniform background in its vicinity. However, it can be shown that for my galaxies the gradient in the surface density of GCs belonging to NGC1399 is negligible, with a maximal gradient value of $\sim 10^{-4}$ objects arcmin^{-3} in the vicinity of my sources. For this calculation, I have used the de Vaucouleurs' fit to the GC system of NGC1399 from Bassino et al. (2006). The total model likelihood takes the form:

$$\mathcal{L} = \prod_{i=1}^{N_{\rm GCC}} \left[\frac{f\Sigma(r_i, r_h)}{\int_0^{r_{\rm max}} S(r')\Sigma(r', r_h)r'dr'} + \frac{1-f}{\int_0^{r_{\rm max}} S(r')r'dr'} \right] \tag{4.3}$$

where i runs over all detected (and not masked) GCCs within the transformed radius $r_{\rm max}$ of the galaxy, which I fix as $15r_{e,r}$; large enough to include all the galaxies' GCs and a large number of background GCCs. I do not consider larger regions because of the increased potential of contamination from steep GCC gradients in the Fornax core. The spatial completeness function $S(r)$ encodes the fractional unmasked area as a function of radius. There are two free parameters: f, the mixing fraction (i.e. the fraction of all sources that are GCs belonging to the target galaxy) and the ratio r_h/r_e. I do not explicitly include morphological or colour terms in the model likelihood, but account for this in the GCC selection described in sections 4.4.1 and 4.4.4.

I impose a Gaussian prior on the ratio r_h/r_e based on the results of Amorisco et al. (2018). The prior is centred at $\sim 1.5r_e$ with a standard deviation of 0.8 and truncated at zero. The choice of prior is very influential, particularly in the low f regime in which most of my sources are anticipated to lie. However, since Amorisco et al. (2018) probe a similar sample of sources in a similar environment (the Coma cluster) to my sample and that the $r_h/r_e \simeq 1.5r_e$ relationship appears elsewhere in the literature (van Dokkum et al., 2017; Lim et al., 2018) it is a reasonable estimate. I probe the effects of modifying the prior on r_h in §4.7. The prior width is much greater than the RMS of the median values quoted by Amorisco et al. (2018), so that if there is any significant deviation it should be recognised in my analysis.

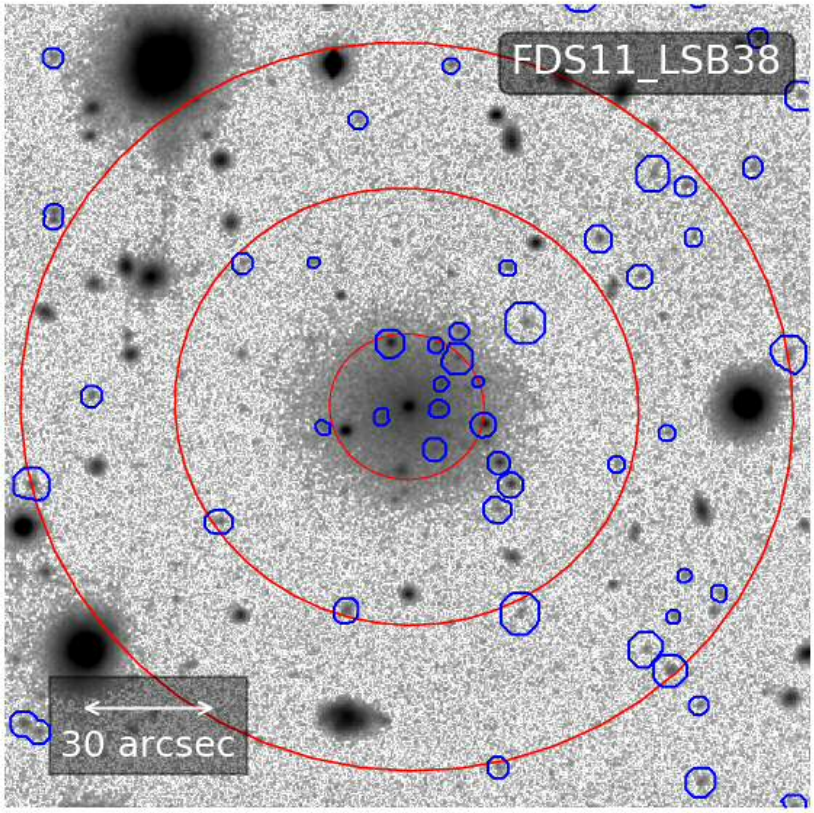

Figure 4.5: Result of full GCC identification and selection for one of my target galaxies. The selected sources are shown using their ProFound segments in blue. I also show the 1, 3 and 5 $r_{e,r}$ contours from the Imfit modelling. A clear over-density of sources can be seen close to the centre of the galaxy.

4.5 Results

4.5.1 Inference on Globular Cluster numbers

I made data cut-outs in each band for each source that were 15×15 $r_{e,r}$ in size. I chose this size because tests with mock datasets (with realistic numbers of interlopers derived from the data) revealed that the measured number of GCs was negatively biased for much smaller values, and this particularly affected systems with less than 10 intrinsic GCs. At 15×15 $r_{e,r}$, I was able to recover unbiased measurements of $N_{\rm GC}$ even for systems with no GCs.

The GCCs were selected according to the criteria described in §4.4.1 and by their colour, described in §4.4.4. All non-selected sources were masked using their Pro-

FOUND segments. I additionally automatically masked the areas around PROFOUND sources with g-band magnitudes brighter than 19 in an effort to remove interloping GCs belonging to other systems. This was accomplished by placing elliptical masks scaled to 2 times the PROFOUND `R100` radius.

All sources in the GCC catalogues that had central coordinates overlapping with the masks were removed. The spatial completeness function could then be measured by measuring the masked fraction in concentric annuli centred on the galaxy, spaced by $0.01r_{e,r}$ and linearly interpolating the result. I note that two sources[4] were omitted from the analysis because they were almost completely masked.

I then ran the Monte-Carlo Markov chain (MCMC) code EMCEE[5] to obtain the posterior distributions of f and r_h for each individual target galaxy. The final inference on the number of GCs associated with each galaxy, $N_{\rm GC}$, was calculated as

$$N^j_{\rm GC} = f^j \frac{\int_0^{r_{\rm max}} \Sigma(r, r^j_h) r dr}{\int_0^{r_{\rm max}} S(r)\Sigma(r, r^j_h) r dr} \frac{\int_{m_1}^{m_2} g(m) dm}{\int_{m_1}^{m_2} \epsilon(m) g(m) dm} N_{\rm GCC} \tag{4.4}$$

taking into account the masked fraction and magnitude incompleteness. Here, j indicates the posterior index and $g(m)$ is the Gaussian g-band GCLF. The results of this are shown in figure 4.6, where I convert my galaxy photometry to V-band magnitudes using the prescriptions of Jester et al. (2005). Note that I also plot specific frequency contours, defined as the number of globular clusters per unit galaxy luminosity. As a means of comparison, I show in §4.6 that my inferences on $N_{\rm GC}$ are consistent with the measurements of Miller & Lotz (2007) for a small sample of overlapping galaxies using a chi-squared test.

I record the following information from the $N_{\rm GC}$ posterior: The 10th, 50th, and 90th percentiles, the 15.9 & 84.1 percentiles (i.e. the 1σ limits centred on the median). The numbers I quote for $N_{\rm GC}$ are the median values and the uncertainties span the range of the 1σ limits centred on the median; these are the error-bars shown in figure 4.6. Note that these estimates are corrected for the colour incompleteness from §4.4.4. Trials with mock datasets showed that the median value is not significantly biased despite the marginal posterior in f being naturally truncated at zero by my model. I find that 0 out of 12 UDGs have median values of $N_{\rm GC}$ below one, compared to 12 across the whole sample. However, 106 of the whole sample of target galaxies are consistent with having no GCs within 1σ.

Overall, my results show a general increase of $N_{\rm GC}$ with M_V that is qualitatively

[4] FDS10_LSB33, FDS11_LSB32

[5] http://dfm.io/emcee/current/

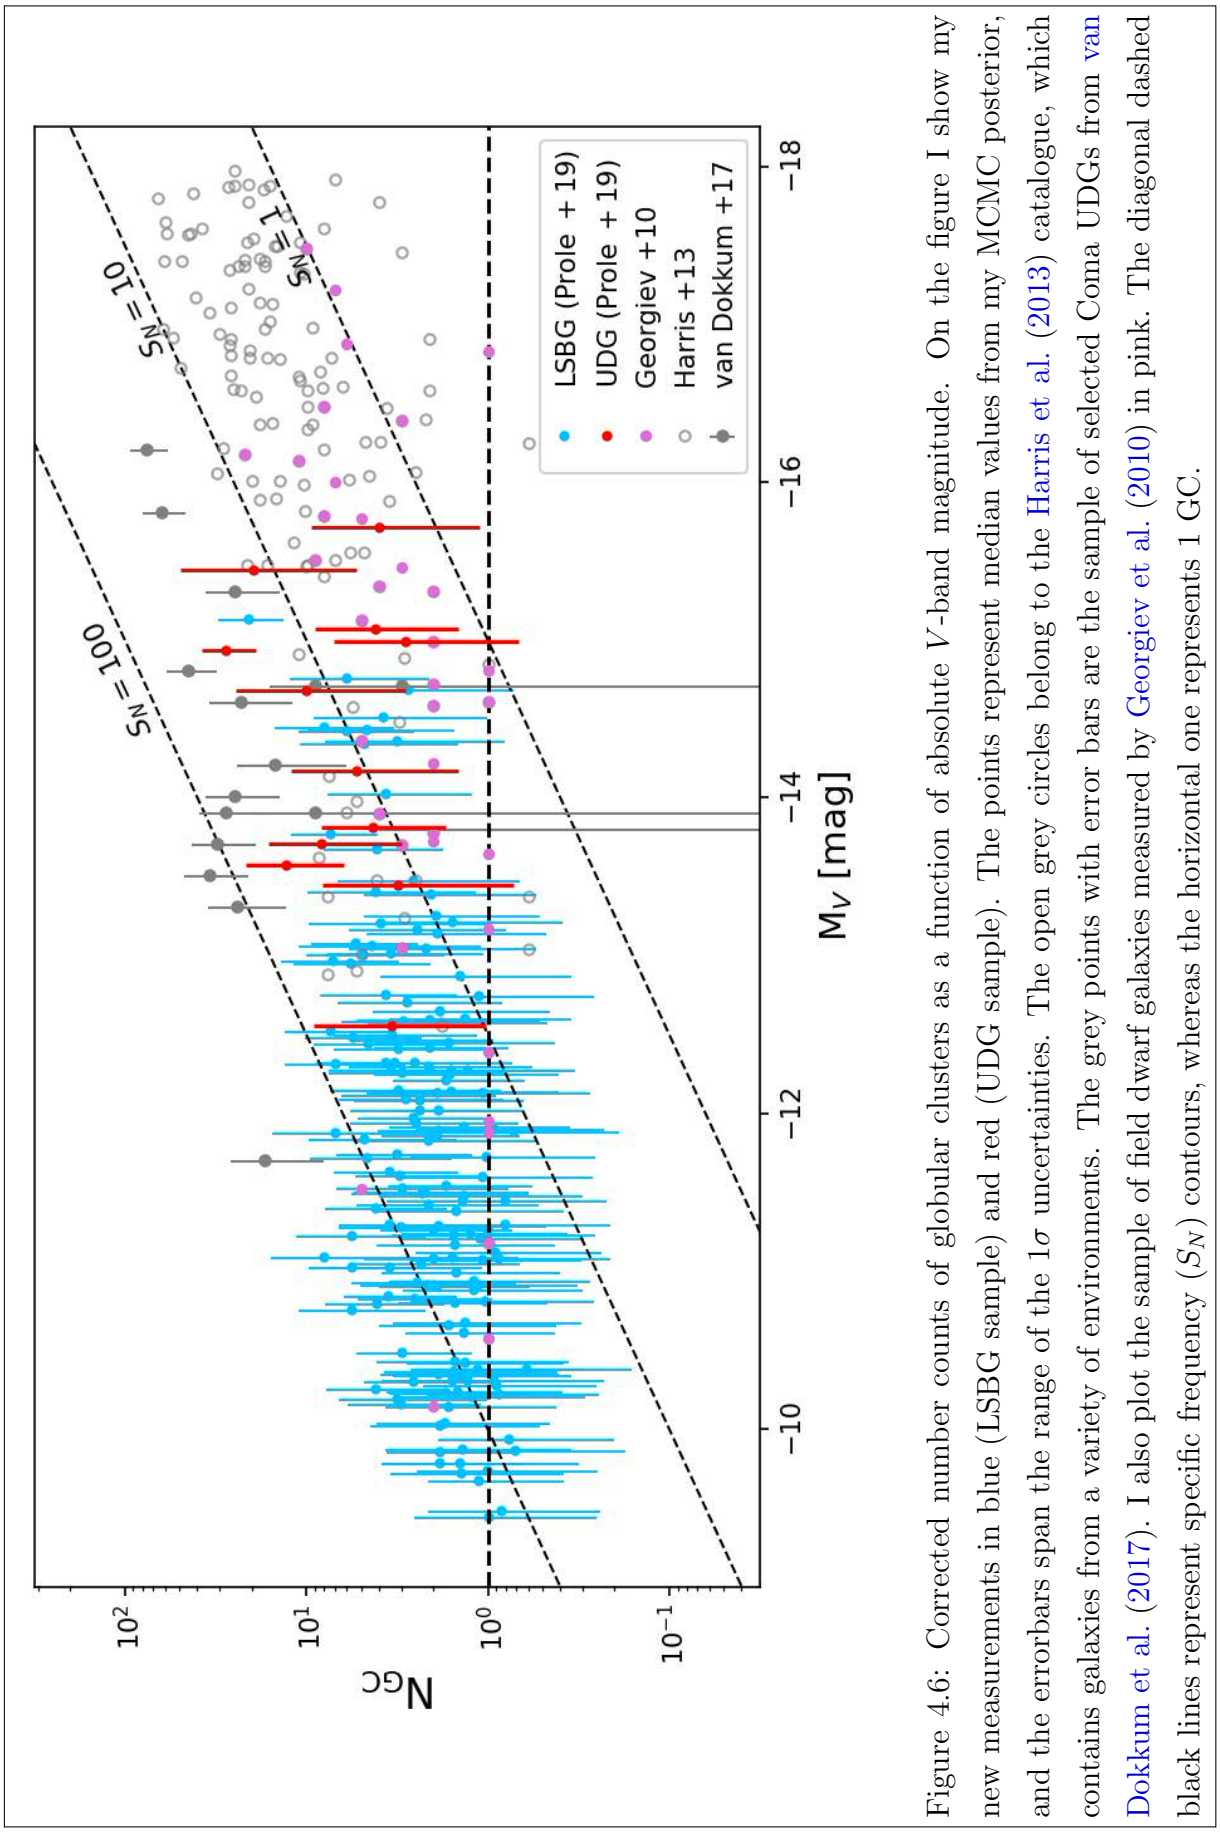

Figure 4.6: Corrected number counts of globular clusters as a function of absolute V-band magnitude. On the figure I show my new measurements in blue (LSBG sample) and red (UDG sample). The points represent median values from my MCMC posterior, and the errorbars span the range of the 1σ uncertainties. The open grey circles belong to the Harris et al. (2013) catalogue, which contains galaxies from a variety of environments. The grey points with error bars are the sample of selected Coma UDGs from van Dokkum et al. (2017). I also plot the sample of field dwarf galaxies measured by Georgiev et al. (2010) in pink. The diagonal dashed black lines represent specific frequency (S_N) contours, whereas the horizontal one represents 1 GC.

consistent with normal dwarf galaxies. While some UDGs are comparable with those of van Dokkum et al. (2017), most of their objects are quite remarkable when compared to my measurements in terms of having much higher $N_{\rm GC}$ for a given luminosity. It remains to be seen whether these sources are comparatively rare among LSB galaxies and because Fornax contains less galaxies I see fewer UDGs with GC excess, or that perhaps the increase in environmental density in the Coma cluster plays a positive role in producing such galaxies; this is discussed further in §4.8.

4.5.2 Colours

Despite already imposing a conservative colour selection criterion in §4.4.4, I can use my results to assess the distribution of colour within the selection box. For each posterior sample, one can assign a probability of belonging to the Plummer profile (i.e. the galaxy) to each GCC given by

$$P_{\rm GC}(r)^j = \left[1 + \frac{2(1-f^j)\int_0^{r_{\rm max}} \Sigma(r', r_h^j) r' dr')}{f^j \Sigma(r, r_h^j) r^2}\right]^{-1} \tag{4.5}$$

where j loops over the posterior sample. The result of selecting high-probability GCCs is shown for a selection of galaxies in figure 4.7. I display the full colour distributions for all my GCCs weighted by their probabilities of cluster membership in figure 4.8. It is clear from these distributions that one-or-two component Gaussian fits are inappropriate, so I limit myself to a qualitative discussion based on the weighted histograms.

Comparing the weighted $(g-i)$ histogram with the un-weighted version, it is clear that a narrow peak emerges that is coincident with the blue component measured by D'Abrusco et al. (2016) at $< g-i >$=0.74. I conclude that the GC population of my sample is mainly blue. This is consistent with the results of Peng et al. (2006), who have shown that low luminosity galaxies tend to have predominantly blue GC systems. The blue nature of the GCs is suggestive of young and/or low-metallicity stellar populations.

In figure 4.8 I also show the $\pm 1\sigma$ span of the colours of the target galaxies. Clearly the blue peaks I observe in $(g-i)$ and $(g-r)$ are consistent with these colours. In $(g-r)$, the blue peak of the GCs appears shifted to the blue compared to the galaxy colours. However, since this effect is within the 1σ, it is not a significant result.

4.5.3 Stellar mass vs Halo mass

Using my estimates of $N_{\rm GC}$ together with the empirical trend of Harris et al. (2017) (accounting for the intrinsic scatter in the relation), I am able to estimate the halo mass M_{halo} of the sample of galaxies. For the estimate to be valid, one must assume that $N_{\rm GC}$ is indeed a reasonable indicator of M_{halo} in the LSB regime. There is limited evidence to support this (Beasley et al., 2016; van Dokkum et al., 2017) based on comparisons between M_{halo} measurements inferred from $N_{\rm GC}$ and those inferred from kinematic measurements. However, it is possible that mechanisms such as ram pressure / tidal stripping in the Fornax core can move galaxies off of the assumed empirical relation by stripping either stars, GCs or even dark matter from the galaxies.

I estimate the stellar mass using the empirical relation of Taylor et al. (2011) (their equation 8), who used the GAMA survey (Driver et al., 2011) to calibrate stellar mass as a function of g and i magnitudes with an intrinsic scatter of $\sim$0.1 dex.

I plot my estimates of M_{halo} vs. M_* in figure 4.9. I also display other measurements from the literature, including the sample of Coma UDGs from van Dokkum et al. (2017) and the median values measured by Amorisco et al. (2018) (it is worth noting that only three of their sources have $M_{halo} > 10^{11}$M$_\odot$ at 90% confidence), along with measurements of other dwarf galaxies, including dwarf ellipticals in clusters (Miller & Lotz, 2007) as well as late-type dwarfs from a variety of environments including the field (Georgiev et al., 2010). I also show the 2σ credibility upper limit on the average mass of UDGs derived from weak lensing by Sifón et al. (2018), with which my results are consistent. Also I show the extrapolated theoretical predictions from abundance matching of Moster et al. (2010), Behroozi et al. (2013) and Brook et al. (2014), which were calibrated using observed stellar masses greater than approximately 10^8, 10^7 and 10^7M$_\odot$ respectively.

Forbes et al. (2018) show that the $N_{\rm GC}$ to M_{halo} relation may lose accuracy for $M_{halo} \leq 10^{10}$M$_\odot$, giving systematically higher values of M_{halo} than measured for their sample. According to their study, a better estimator of M_{halo} is the total mass associated with the GC system; however, they note that the assumption of a common mean GC mass is not valid at the low-mass end such that individual GC masses should be measured to get an unbiased estimate of M_{halo}, using the empirical relation of Spitler & Forbes (2009). While I have not measured the individual GC masses in this work, I note that my estimates of M_{halo} should be considered as upper-limits in light of their result.

Every UDG in my sample is consistent with inhabiting a dwarf sized halo (M_{halo}<10^{11}M$_\odot$)

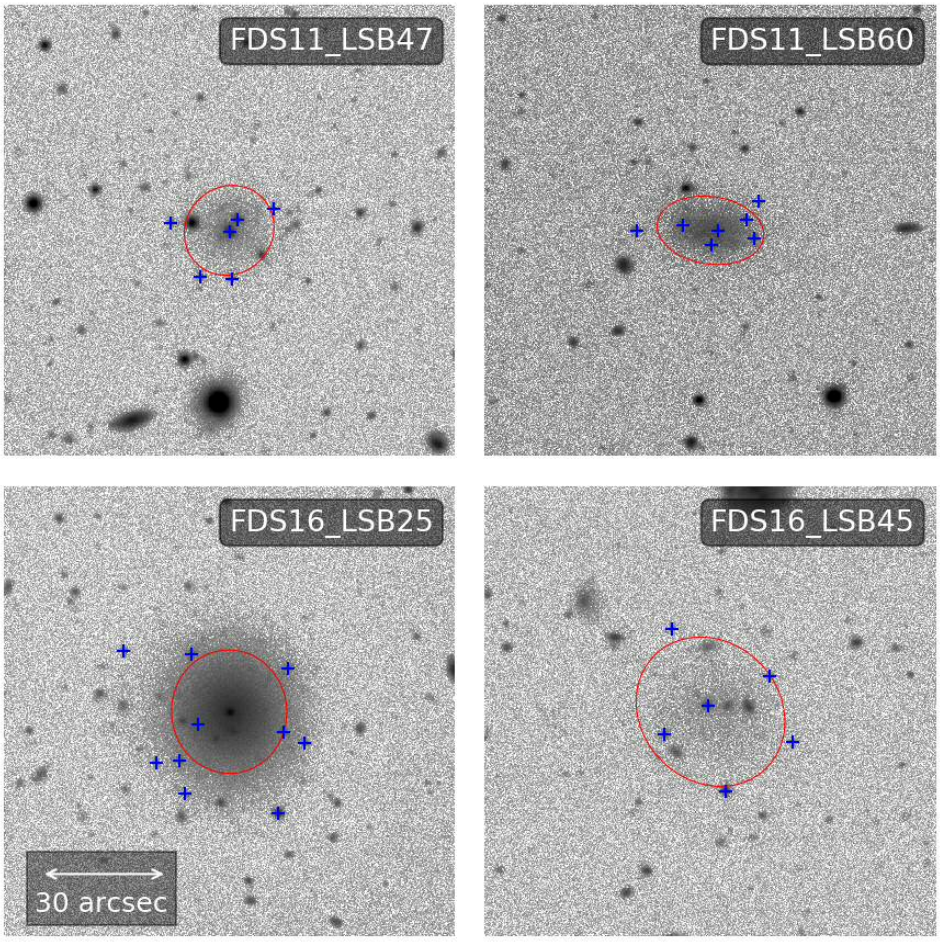

Figure 4.7: r-band cut-outs of four of my target galaxies, selected from galaxies with $N_{\rm GC}$ ≥5. The red ellipses represent $1r_{e,r}$ contours. The blue points identify GCCs with $P_{\rm GC} \geq 0.5$. Of the four galaxies shown, FDS16_LSB45 is the only one I measure to be large enough to be classified as a UDG. 30″ is ∼2.9kpc at Fornax distance.

to within 1σ. There appears to be no significant tendency for UDGs to have enhanced GC populations and therefore enhanced halo mass for their stellar mass. Indeed, there is a qualitatively continuous trend from the LSBGs towards the UDGs.

The overall population is most consistent with an extrapolation of the Brook et al. (2014) relation (calibrated with local group dwarf galaxies), but I cannot rule out consistency with that of Moster et al. (2010) or Behroozi et al. (2013) because of the potential for my estimates of M_{halo} to be overestimates. I emphasise however that all models require extrapolation, below stellar masses of $10^8 M_\odot$ for the Moster et al. (2010) relation, and $10^7 M_\odot$ for that of Behroozi et al. (2013) and Brook et al. (2014).

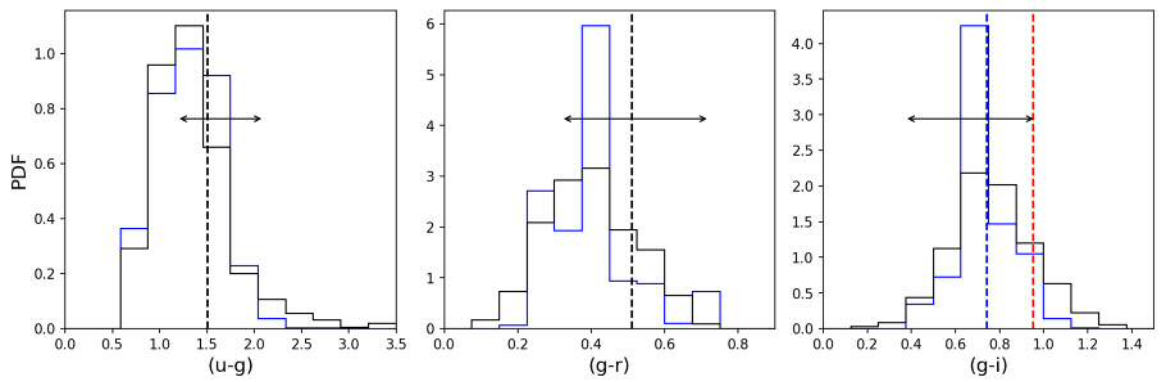

Figure 4.8: Normalised distribution of colours for the GCC sample (black histograms) vs. those of the same sample after being weighted by their probability of belonging to a target galaxy's GC system according to equation 4.5 (blue histograms). I note that I select only sources with $m_g \leq 23$ to overcome the measurement error and ease comparison with D'Abrusco et al. (2016), who used a similar limit. For $(u-g)$ and $(g-r)$, the vertical black dashed lines correspond to the median colours of the sample of spectroscopic sources described in §4.4.4. The dashed blue and red lines in the $(g-i)$ panel correspond to the means of the blue and red components measured by D'Abrusco et al. (2016). The results a consistent with a predominantly blue GC population. The horizontal arrows indicate the $\pm 1\sigma$ range of the galaxy colours.

While no UDGs have estimates of M_{halo} above what might be expected for enriched GC systems (according to the empirical relation of Amorisco et al., 2018), several of the LSBG sample do show evidence for excess. This might suggest a continuation of GC-enriched systems down to very low stellar mass.

Another point of interest is that my overall sample of LSB galaxies (including UDGs) appears offset from the mean trend of dwarf galaxies, having higher M_{halo} for a given M_*. While my estimates of M_* for the objects from the literature require assumptions about their colours, this may hint that LSB galaxies have systematically higher M/L ratios than normal dwarfs. However this might be an environmental bias; perhaps only LSB galaxies of high M/L ratio are able to survive in the Fornax core.

4.5.4 GC system sizes

Despite imposing a prior on r_h (the GC half number radius) with a mean of $1.5r_e$, I find that my GC systems are typically slightly larger. The median value of r_h/r_e recovered from the full sample of galaxies is 1.73, with a standard deviation of ~ 0.27

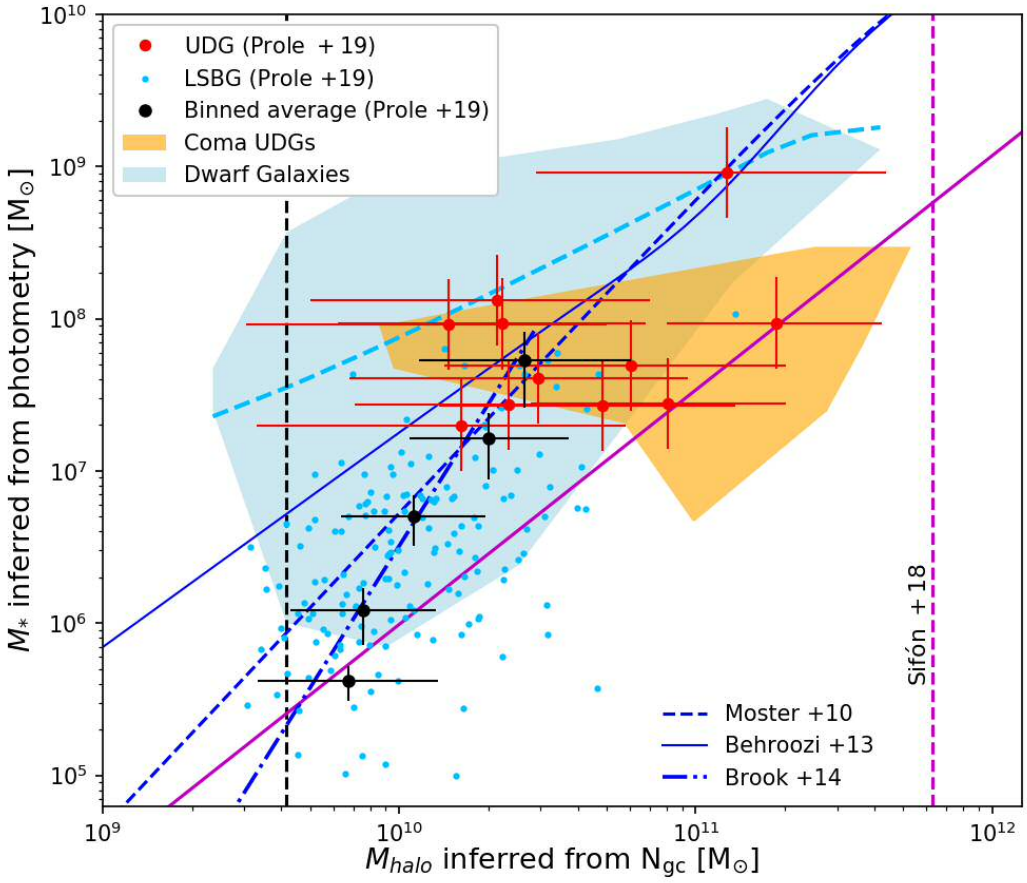

Figure 4.9: Halo mass (derived using $N_{\rm GC}$) vs. stellar mass. In the figure I show my new measurements, the light blue points correspond to LSBGs and the red to UDGs. The stellar masses for my points were derived from g and i band photometry using the scaling relation of Taylor et al. (2011). The black error-bars show illustrative binned averages of all my new measurements (calculated in logarithmic bins). The orange region bounds UDG measurements from the literature (van Dokkum et al., 2017; Amorisco et al., 2018; Lim et al., 2018), I also plot measurements of other "normal" dwarf galaxies (Miller & Lotz, 2007; Georgiev et al., 2010, blue region). The dashed light blue line represents the mean of these sources. The darker blue lines show extrapolated theoretical predictions of Moster et al. (2010), Behroozi et al. (2013) & Brook et al. (2014). The diagonal purple line corresponds to "enhanced" GC systems Amorisco et al. (2018), using the Harris et al. (2017) conversion between $N_{\rm GC}$ and M_{halo}. The dashed purple line indicates the 2σ credibility upper-limit on the average mass of UDGs derived from weak lensing Sifón et al. (2018). Finally, the vertical black dashed line corresponds to $N_{\rm GC} = 1$.

and range between 0.4 and 2.8 (in units of $r_{e,r}$). I note that the median value of r_h for the UDGs is consistent with that of the full sample.

If I use the relation between r_h and M_{halo} presented in Hudson & Robison (2018), the resulting M_{halo} estimate is much larger than previously estimated using $N_{\rm GC}$. For example, for a UDG with r_h=1.5 kpc should have M_{halo} of around $10^{11.6}$M$_\odot$, much higher than many of the estimates presented in figure 4.9 and generally inconsistent with UDGs with halo mass measurements in the literature (e.g. van Dokkum et al., 2017). While I note that Hudson & Robison (2018) make clear that the relation is calibrated only for $M_{halo} \geq 10^{12}$M$_\odot$, I advocate a relation more in-line with that of Forbes (2017) in this regime.

4.5.5 LSBGs vs UDGs

Now that I have estimates of $N_{\rm GC}$ for each of my target galaxies, I am in a position to directly compare the LSBG population with the UDGs. The two questions I want to answer are: Does the UDG population show any statistical excess of GCs when compared with the LSBGs in the same luminosity range?; and, is the observed distribution of $N_{\rm GC}$ for the UDGs discontinuous from that of the LSBGs?

I note that from the appearance of figure 4.6, it seems that $N_{\rm GC}$ vs. M$_V$ can be modelled approximately as a power law. I omit all UDGs from the sample and fit such a relation to my LSBG sample (see also figure 4.12):

$$\bar{N}_{\rm GC} = (0.04 \pm 0.02) \times 10^{(-0.15 \pm 0.02) \times {\rm M}_V} \tag{4.6}$$

I note that the scatter in the relation is approximately 0.2 dex across the full magnitude range. Using the fit, I can ask whether my sample of UDGs (ignoring other UDGs from the literature) are consistent with this description. I perform a chi-squared test with the null-hypothesis that the UDGs are drawn from equation 4.6. This results in a p-value of 0.30, which means I cannot reject the null hypothesis with an acceptable level of confidence. I therefore conclude that my UDG sample is quantitatively consistent with a continuation of the LSBG sample in this parameter space. I also note that since there is no UDG that has a $N_{\rm GC}$ measurement convincingly more than 3σ above the power-law predicted value, there is no compelling evidence that my UDGs have excessive GC populations.

As as means of comparison, I also do the same test for the population of GC-enriched UDGs from van Dokkum et al. (2017). While the two tests are not directly comparable since the sample of van Dokkum et al. (2017) was at least partially biased

to select extreme objects (as in the cases of galaxies DF44 and DFX1), I find that their sample is not consistent with equation 4.6, with a p-value much less than 1%.

Aside from DF44 and DFX1, the galaxies measured by van Dokkum et al. (2017) also include a list of 12 UDGs selected from the Yagi et al. (2016) catalogue of LSB galaxies that are also present in the Coma Cluster Treasury Program[6] footprint. Importantly, this should represent a small but unbiased sample of Coma UDGs. After selecting only these sources and repeating the test, I find that the Coma sample is still inconsistent with equation 4.6. This may indicate that UDGs in Coma have more GCs than galaxies in Fornax in the same luminosity range. I find that the choice in prior for the GC half-number radius does not impact this result; for a detailed discussion see §4.7.

There could be a number of explanations for this discrepancy. It could be that the coma sample have lower luminosity at a given halo mass and thus have excessive globular clusters for their luminosity - perhaps the coma sample were quenched earlier (in terms of star formation history) than the Fornax sample. Alternatively, the Fornax sample of UDGs may have had their GCs stripped more efficiently in the denser core of the Fornax cluster.

4.5.6 Effect of the GCLF

As stated in §4.4.3, I have adopted a Gaussian GCLF with a mean of $24m_g$ and standard deviation of $0.7m_g$. However, dwarf galaxies can have varied GCLFs and it is important to show that my results are robust against this. Villegas et al. (2010) have measured the g-band GCLFs for 43 early-type galaxies in the Fornax cluster, down to galaxies with absolute B-band magnitudes of around -16. I use this catalogue as a means to test what would happen to my measurements if the GCLF was wider and has turnover magnitude fainter than my adopted value, i.e. to get an upper-limit on the inferences on $N_{\rm GC}$.

From the Villegas et al. (2010) catalogue, I measure a mean GCLF with mean $24\pm0.1m_g$ and a standard deviation of $0.84\pm0.21m_g$ after clipping outliers at 2σ. I note that I selected from their catalogue only galaxies with absolute B magnitudes fainter than -18 to target dwarf galaxies for this calculation. This suggests that the GCLF might be wider than what I have assumed previously. Integrating the RE over the 1σ deeper and wider GCLF and comparing to my previous estimates of the observed GC fraction, I find that the maximum correction in my $N_{\rm GC}$ is an increase

[6] https://archive.stsci.edu/prepds/coma/

of ∼20%. I find that this is not sufficient to impact or change the overall results of my work (a 20% increase in $N_{\rm GC}$ is sufficient to increase a M_{halo} estimate by ∼0.1dex).

4.5.7 Effect of Nuclear Star Clusters

We do not treat potential NSCs any differently from GCs in my analysis; GCCs are defined by their magnitude and colour. While I have imposed a bright-end magnitude cut on my sample of GCCs, there is still potential for faint NSCs to contaminate my sample and therefore increase the number of GCCs for a target galaxy by one. For the galaxies with low estimates of $N_{\rm GC}$, this can amount to a significant source of error. However, most of my target galaxies have $M_* \leq 10^8 {\rm M}_\odot$ and are thus expected to have a low nucleation fraction (between 0.7 at $10^8 {\rm M}_\odot$ to 0.0 at $10^5 {\rm M}_\odot$, as shown in figure 8 of Sánchez-Janssen et al., 2018).

Removing GCCs close to the centres of galaxies introduces a subjective bias. However, I note that all the galaxies in my sample have already been visually classified as either nucleated or non-nucleated by Venhola et al. (2017). This number amounts to 10% of the catalogue. After applying my bright-end magnitude cut, this leaves us with 13 galaxies that are potentially contaminated by a NSC. To quantify the effect this may have on my estimates of M_{halo}, I simply drop these sources from the sample and repeat the analysis. I find that the results do not change; the new binned-average estimates of M_{halo} are consistent within much less than 1σ with those displayed in figure 4.9.

4.6 Comparison with Miller and Lotz (2007)

As several of the sources in the Venhola et al. (2017) catalogue were also identified in the Fornax cluster catalogue (Ferguson, 1989), I was encouraged to search for matches in the catalogue of dwarf ellipticals studied by Miller & Lotz (2007), who used the HST WFPC2 Dwarf Elliptical Galaxy Snapshot Survey to measure the GC populations for a sample of 69 galaxies. They measured $N_{\rm GC}$ using apertures of 5 times the exponential scale size of the galaxies, which roughly equates to $3r_e$ for a Sérsic index n=1.
I find seven matches: FDS16_LSB33 (FCC0146), FDS12_LSB10 (FCC0238), FDS12_LSB4 (FCC0246), FDS11_LSB62 (FCC0254), FDS16_LSB58 (FCC0110), FDS16_LSB32 (FCC0144) and FDS12_LSB30 (FCC0212). Overall my results are reasonably consistent (albeit with large error-bars), as is shown in figure 4.10. Note that in the figure one of the sources is not visible because it overlaps with another.

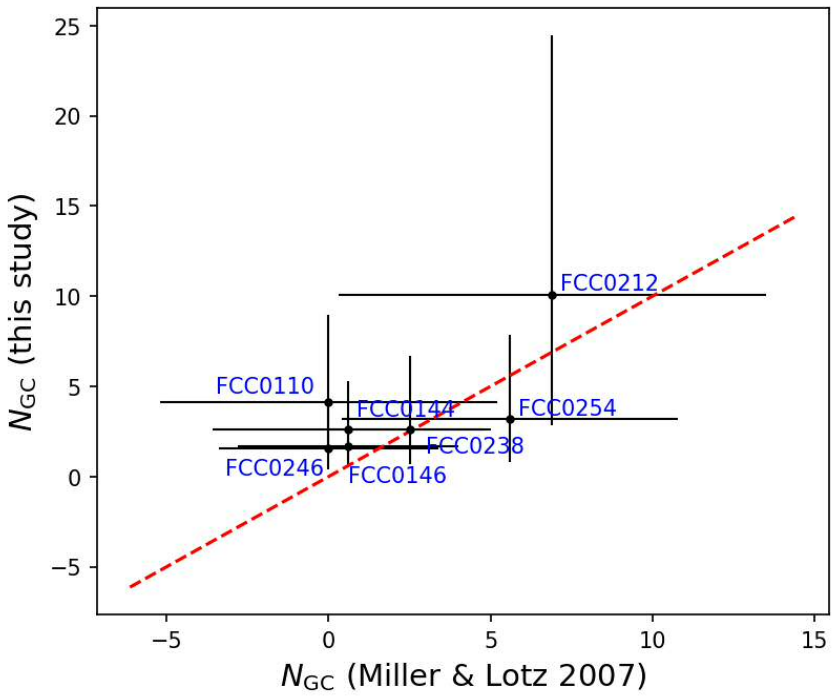

Figure 4.10: Comparison of GC number counts. The black points are median values from the MCMC posterior. The error-bars span the range of 1σ. The red dashed line is the one-to-one relation. A chi-square test (accounting for the errors) shows that the null hypothesis that the measurements are consistent with the one-to-one relation cannot be rejected with confidence greater than around 4%.

4.7 Choice of prior

While the choice for the prior on the GC half-number radius r_h is justified from previous literature measurements (van Dokkum et al., 2017; Amorisco et al., 2018; Lim et al., 2018), it is important to show how different choices may affect the results. This is particularly relevant because the spatial distributions of GCs for UDGs is not well known. By running the MCMC using different priors, I show in figure 4.11 that, despite the choice of prior in r_h strongly influencing the r_h posterior, the estimates of $N_{\rm GC}$ are robust.

A small increase in the mean of the prior on the GC half-number radius, $\bar{r}_{h,prior}$, is not sufficient to significantly impact my results. However, more dramatic modifications may produce a more pronounced change. In general, lowering $\bar{r}_{h,prior}$ increases the

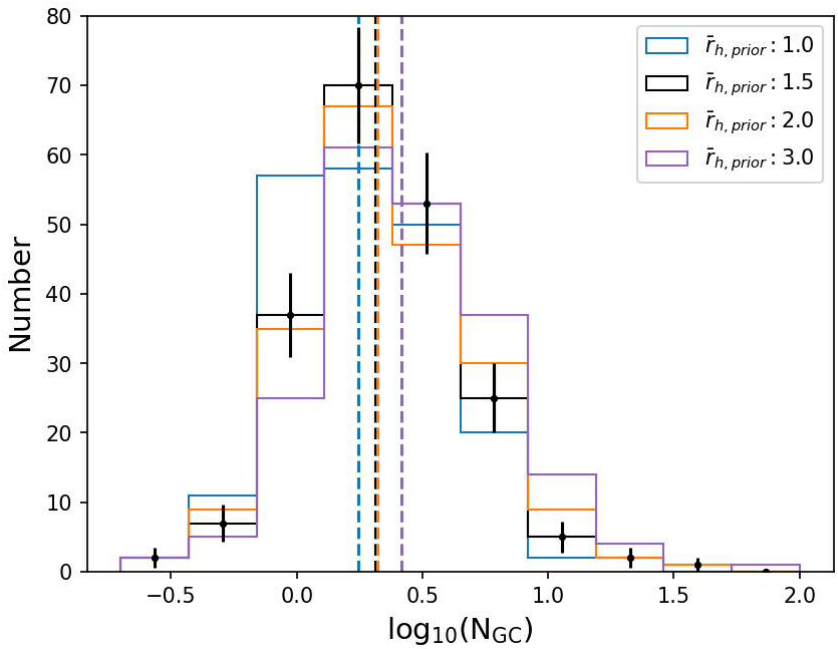

Figure 4.11: The distribution of $N_{\rm GC}$ estimates for my sample of galaxies as a function on the mean of the prior on the GC half-number radius, $\bar{r}_{h,prior}$, in units of galaxy half-light radius. Also shown are the median values (dashed lines). My value of $\bar{r}_{h,prior}$=1.5r_e is adopted from the literature and Poisson error-bars are shown for this value.

number of GC-poor systems, while increasing it results in more GC rich systems. However, the median value for the overall population is not significantly altered by using different priors. I finally note that repeating the analysis from §4.5.5 with $\bar{r}_{h,prior}$=3r_e leads us to the same conclusions; the overall result is robust against changes in the prior.

4.8 Discussion and Summary

In this chapter I have estimated the halo masses of a sample of 175 LSB galaxies in the Venhola et al. (2017) catalogue using the sizes of their GC populations, including a sub-sample of 12 UDGs. This constitutes the largest sample of low mass galaxies so-far analysed for this kind of study. Candidate globular clusters were identified in the g-band using measurements from the ProFound photometry package. I also applied a colour selection based on photometric measurements of a set of spectroscopically

confirmed Fornax cluster GCs, using PSF-corrected aperture magnitudes measured in the u, g, r, i bands. Following this, I used a Bayesian Mixture model approach (influenced by the work of Amorisco et al., 2018) to infer the total number of GCs associated with each target galaxy, assuming a GCLF appropriate for my sample.

My estimates of $N_{\rm GC}$ for the overall population are qualitatively consistent with more compact dwarf galaxies when plotted against M_V. I find that the sample of UDGs are statistically consistent with a power-law fit to the $N_{\rm GC}$ measurements for LSBGs, indicating that there is no discontinuity between the two populations; my sample of UDGs does not have a statistically significant excess of GCs compared to smaller LSB galaxies in the same luminosity range.

I converted the inferences on $N_{\rm GC}$ to M_{halo} using the empirical relation of Harris et al. (2017). I additionally derived stellar masses for the galaxies from the empirical relation of Taylor et al. (2011), using IMFIT galaxy models. Overall, the M_* estimates are consistent with dwarf galaxies and the M_{halo} estimates are consistent with dwarf sized halos. The LSBG galaxy population appears consistent with the extrapolated Brook et al. (2014) abundance-matching relation between M_* and M_{halo} and as an extension of measurements from typical dwarf galaxies, but perhaps with slightly larger M_{halo} for the average dwarf at a given M_*. I suggest that this might be a systematic effect due to the environment; it is possible that only LSB galaxies with high M/L ratios are able to survive in the Fornax core. However, as Forbes et al. (2018) have shown, the M_{halo} estimates may be too large because of a breakdown in accuracy of the $N_{\rm GC}$-M_{halo} relation in the low mass regime, and it is not yet clear how this affects my estimates.

None of my UDGs have median values of $N_{\rm GC}$ above the empirical boundary marking GC-rich systems measured by Amorisco et al. (2018). However, 5 are consistent within their 1σ uncertainties. Several LSBGs also have potential for GC-richness, and 13 are at least 1σ above the required threshold. Such objects could make interesting sources for a follow-up study, given that they could represent a continuation of GC-rich objects down to very low stellar mass. If genuine, they could mean that enhanced GC systems are not unique to UDGs and the mechanisms by which UDGs are produced are separate from those by which LSB galaxies gain enriched GC systems, something also observed by Amorisco et al. (2018).

Using a weighted histogram approach, I have shown that the GC population of my target galaxies is predominantly blue compared to the overall GC population in Fornax. My result is consistent with the blue peak in $(g-i)$ recorded by D'Abrusco et al. (2016), with a relative depletion of red GCs. Further still, the blue peak of my

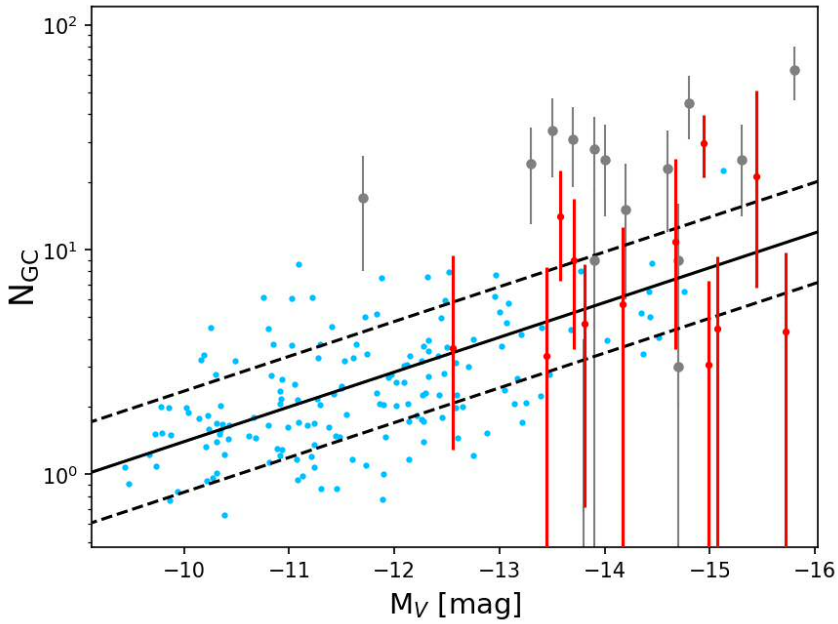

Figure 4.12: Power-law fit to $N_{\rm GC}$ vs. M_V for my LSBG sample (solid black line) along with an estimate of the 1σ scatter (black dashed lines). Blue points: My LSBG sample. Red errorbars: My UDG sample. Grey errorbars: The sample of GC-enriched UDGs from van Dokkum et al. (2017).

GC coincides with the $\pm1\sigma$ range of the galaxy colours. There is tentative evidence in $(g-r)$ that the galaxies may be slightly redder than the GCs, but since this is not a significant effect I do not comment on this further.

The Coma cluster UDGs measured by van Dokkum et al. (2017) seem to have significantly more GCs than what I see in the Fornax cluster. It is notable that my sample is confined to the core of the Fornax cluster. While Lim et al. (2018) show that there is no particular trend of specific frequency S_N with cluster-centric radius for bright UDGs in Coma, they also show that S_N decreases with cluster-centric radius for fainter galaxies; if anything this could mean that I probe a population with systematically higher $N_{\rm GC}$ at a given M_* than in the cluster outskirts. Two possibilities are that GC-enriched UDGs are comparatively rare objects and I simply do not observe them because Fornax is much less massive than Coma, or the denser environment of the Fornax cluster core causes GCs to be stripped more efficiently than those in Coma. I suggest that future studies could provide complete measurements of $N_{\rm GC}$ for UDGs in other clusters (e.g. Virgo) to address this question.

My measurements are sufficient to rule out the failed L_* formation theory for UDGs because the halo mass estimates indicate that they reside in dwarf sized halos. I find a continuation in properties between UDGs and smaller LSBGs such that it does not seem that UDGs have a unique or special formation mechanism. Since few of my UDGs are convincingly GC-rich compared to those in Coma, I speculate that this property may be related to environmental density. Perhaps the Coma objects are more efficiently stripped of gas in the Coma core, thus forming fewer stars relative to their halo mass, resulting in systems that appear GC-rich for their stellar mass. A consequence of this effect is that the fraction of GC-rich UDGs should decline with cluster-centric radius, and this may be a valuable way to estimate the relative strengths of secular vs. environmentally-driven formation mechanisms.

Chapter 5

Discussion

During the timespan of my PhD, the field of LSB research has grown rapidly, sparked by the advent of deep wide area surveys as well as the discovery of large numbers of UDGs in the Coma galaxy cluster by van Dokkum et al. (2015). This interest is set to continue over the coming years, thanks to the advent of next-generation deep wide-field galaxy surveys like Euclid and LSST.

The work presented in this thesis has contributed to science in two ways: The first has involved the design and development of a new source detection package, DeepScan, intended as an alternative to SExtractor that offers improved results in the LSB regime. The second has involved forwarding our understanding of ultra-diffuse galaxies both in clusters and in the field; much of this work has been presented in the context of UDG formation and evolution. In the following section, I will review the main conclusions from each chapter and put them in context with the wider literature.

5.1 Automated Extraction of LSB sources

Limitations in source extraction software are a primary concern for the detection of LSB sources in modern and upcoming datasets like those produced by Euclid and LSST, which, due to their enormous sizes have to be analysed automatically, robustly and efficiently. In particular, one of the most popular source extraction packages, SExtractor, is known to perform quite poorly when detecting and de-blending LSB sources, often leading to severe fragmentation known as shredding (e.g. Davies et al., 2016).

As part of this thesis I have written the DeepScan software, which can offer improvements over SExtractor for detecting extending LSB structure. The DeepScan package uses the `DBSCAN` algorithm to identify sources as over-densities of sig-

nificant pixels compared to what is expected given the noise level. Since the algorithm clusters significant pixels that are not necessarily contiguous using the `DBSCAN` algorithms, the method can be thought of as a generalisation of contiguous pixel clustering approaches used by e.g. SExtractor. This gives DeepScan an advantage over such algorithms for detecting extended LSB structure, which can be detected at a higher statistical significance level over broader areas. In essence, this can lead to a lower false (noise) detection rate for a given pixel detection threshold.

Besides a novel implementation of the `DBSCAN` algorithm that is particularly efficient for the analysis of large images, DeepScan also contains procedures to measure the sky background level, its RMS level, as well as routines to perform automated masking of the data. The full code is available online and is in a state of continual improvement and modification.

I have demonstrated DeepScan by identifying 30 new (uncatalogued) LSB galaxies in the Virgo Cluster using the public NGVS data, despite the same region having been searched before using a variety of different methods. The sample contains galaxies with $26.0 < \bar{\mu}_e < 28.5$ and $19 \leq m_g \leq 21$. These galaxies are likely members of the Virgo Cluster based on their angular sizes and colours, which are consistent with the red sequence of Virgo galaxies (Roediger et al., 2017). If they are indeed cluster members as seems likely, they are some of the faintest dwarf galaxies known in the cluster. Assuming cluster membership, the sample has a mean stellar mass of $10^{6.3\pm0.5} M_{\odot}$. I used ProFit to fit Sérsic profiles to my detections and found that none of the new sources have effective radii larger than 1.5 kpc and so do not meet the criteria for ultra-diffuse galaxy (UDG) classification.

5.1.1 Future Work

While DeepScan does present a novel approach to pixel detection, its main limitation at present is the lack of a de-blending algorithm. This is where SExtractor suffers from many of its pathological errors and can produce spurious results. Much of the current and future work in the field of source extraction is and will be focussed on developing improved de-blending algorithms. At present there are two competing methods of de-blending: watershed (cf. ProFound) and hierarchical tree structures (cf. SExtractor, MTObjects).

In certain situations, the watershed approach is superior. This is particularly true for crowded fields, when the density of detectable sources is large compared to empty regions, and where two or more objects of comparable sizes are blended together.

This is because the watershed approach typically returns comparably-sized segments for comparably sized blended sources. However, it runs into problems when small sources are embedded in larger ones, for example a globular cluster embedded in a stellar halo of a galaxy. This is because watershed algorithms typically do not allow for nested sources, meaning that the smaller source (i.e. the globular cluster) gets a disproportionately large segment compared to how it would as an isolated source. This can lead to over estimates in size and brightness from simple segment statistics.

By contrast, hierarchical tree structures are able to identify nested components (like the globular cluster example) easily. Often however, as in the case of MTOBJECTS, the segments corresponding to nested sources are too small compared to how they would appear if isolated. This can lead to underestimates in size and brightness from segment statistics. An additional advantage of hierarchical tree de-blending is that it does not necessarily require the background to be subtracted; instead, the background can be treated as a separate source in the bottom of the hierarchy. This enables MTOBJECTS to out-perform other source extraction packages when there is a large range in the angular extent of sources, for example in galaxy clusters, where a background estimate over a single spatial scale may not be appropriate.

Despite the advances in these techniques, it does not appear as if there is currently a "one size fits all" source extraction package; each has its own set of advantages and disadvantages. In the era of LSST, where we expect fields to be rather crowded, it remains to be seen what the best de-blending strategy will be.

Aside from de-blending, there is much room for improvement in the way that LSB galaxies are measured. The current standard in the literature is to fit one or two component Sérsic profiles using software like GALFIT. Such routines have the advantage that they can account for the shape of the PSF, but often the simple parametric models do not sufficiently model the intricacies of the galaxy's morphology. Adding more components can increase the quality of the fit, but this comes at the expense of higher run times as well as an increase in the number of unstable or poor quality fits. This issue is particularly relevant for LSB galaxies in the field, which can often be irregular and so are poorly fit by such methods. Estimating the magnitude of objects is relatively easily done (e.g. by using a dilation approach similar to PROFOUND), but estimating their "size" is not at all trivial.

It is important to point out that even given a "perfect" source extraction pipeline, the quality of its output will be limited by the quality of its input, i.e. the data. This highlights the importance of continued improvements in the way that data is reduced. Some encouraging examples of this are the re-reduction of the SDSS Stripe-82 data

by Fliri & Trujillo (2016) & Román & Trujillo (2018), which greatly improve the way the background is subtracted, the second public data release of HSC-SSP, which also improves the sky subtraction over its first release. There is also work being undertaken in the community that will enable the automated subtraction of stellar PSFs over arc-minute scales, as well as efforts to distinguish galactic cirrus from extra-galactic sources based on colour (Román et al., 2019b). Such efforts will be key for advancing our abilities to detect the faintest possible sources and become truly limited by the astrophysical sky background.

5.2 The Nature of Ultra-Diffuse Galaxies

UDGs are puzzling objects. While for the most part they seem to display a continuum of properties with other types of galaxies (in mass, size, metallicity, surface brightness etc.), there are still some that could be genuine outliers in the stellar-mass halo-mass plane. Illustrative examples of this are DF44, which has a potential halo mass of $10^{12}M_{\odot}$ (van Dokkum et al., 2016), being almost entirely dark matter dominated, and DF2, which might have no dark matter at all (e.g. van Dokkum et al., 2018)[1]. Such differences are striking and suggest that UDG-like objects have several formation mechanisms that can lead to them having a wide variety of properties. For now, in my opinion, it is particularly important to find more independent examples of such extreme objects in order to demonstrate that they truly exist and are not simply manifestations of measurement error or some other systematic error. If confirmed, their abundances may provide valuable constraints on theoretical models.

However, for the most part it seems that UDGs are dark matter dominated, at least in rich cluster environments. This has been determined in part from kinematics of their stellar systems, kinematics of their globular clusters, apparent absence of tidal features close to the cluster centres, and richness of their globular cluster systems and even weak lensing. A great deal of this work has been focussed on estimating UDG masses in the Coma Cluster; this is likely because it is relatively easy to identify UDGs in such a rich, nearby cluster in which there are many UDGs.

The halo masses of UDGs are an important measurement for distinguishing between different formation scenarios of UDGs. Large, Milky Way-like halo masses would suggest that UDGs are formed from massive galaxies that are quenched before they can form many stars, perhaps through tidal stripping during cluster in-fall (cf.

[1] Although different authors have reached different conclusions for the same object, e.g. Trujillo et al. (2018).

van Dokkum et al., 2015; Yozin & Bekki, 2015). Alternatively, halo masses more akin to dwarf galaxies would suggest that UDGs are formed either through tidal heating of smaller dwarf galaxies or perhaps through secular mechanisms like stellar feedback. Reproducing the observed halo masses is an important test for cosmological and zoom-in simulations, as well as for our general understanding of galaxy formation and evolution in the LSB regime.

In chapter 4, I have contributed to this topic by estimating the halo masses of 175 LSB ($23 \leq \bar{\mu}_{e,r}[\text{mag arcsec}^{-2}] \leq 28$) galaxies in the Fornax galaxy cluster using using the Fornax Deep Survey (FDS) data; this is the largest sample of low mass galaxies so-far analysed for this kind of study. I used a Bayesian mixture model approach to measure the numbers of their associated globular clusters with each galaxy in the sample, and exploited the empirical scaling relation (Harris et al., 2017) to estimate the halo masses of the galaxies. The proximity of the Fornax cluster means that I was able to measure galaxies with much smaller physical sizes ($0.3 \leq r_{e,r}$ [kpc] ≤ 9.5) compared to previous studies of the GC systems of LSB galaxies in the Coma cluster, probing stellar masses down to $M_* \sim 10^5 \mathrm{M}_\odot$. For the time being, direct comparisons with the globular systems in simulated galaxies is not possible because of the limited mass resolution of the simulations in comparison to the typical globular cluster mass.

In general, we find that the UDG population is statistically consistent with a linear extrapolation of the smaller LSB galaxy population, suggesting there is nothing particularly unique about the way the Fornax UDGs have formed (or at least, in the way that their globular clusters have formed) in comparison to their smaller counterparts. In particular, the sample appear to be largely consistent with having dwarf galaxy sized halo masses ($M_{halo}<10^{11}\mathrm{M}_\odot$).

I found that the half-number radius of the globular cluster systems was consistent with other measurements in the literature at around $\sim$1.5 times the effective radius. Additionally, the colours of the globular clusters were consistent with those of the LSB galaxies themselves, suggesting similar ages and metallicities and that they had not been accreted from other galaxies or unbound globular clusters in the Fornax Cluster. In principle, accretion of such objects would cause our halo mass estimates to increase.

One interesting finding however was that at a fixed luminosity, the Coma UDGs seem to have systematically more globular clusters than in our Fornax sample, an effect that is statistically significant. While the reason for this is not yet clear, it is worth noting that many of the UDGs in our sample are likely located in the core of the Fornax cluster, where potential gravitational stripping of globular clusters may be significant. This may lead to underestimates in the halo masses for these objects

based on their globular cluster numbers.

A simple test to assess the significance of this effect would be to perform a similar analysis over a wider area of the Fornax Cluster. Differences in the numbers of globular clusters associated with LSB galaxies of fixed luminosity as a function of the cluster radius could be examined. A possible source of systematic error in such an analysis would be the fact that gas (and stars) are stripped more efficiently from galaxies deeper in the cluster potential, meaning that galaxies of fixed halo mass could have systematically lower luminosity towards the cluster centre.

The idea that UDGs in groups and clusters are generally dark matter dominated suggests that they could form when in-falling galaxies are quenched through ram pressure stripping. Originally, it was suggested that these would have to be massive galaxies (i.e. Milky Way-sized) in order to explain the large effective radii of the UDGs together with their ability to survive relatively unperturbed in dense cluster environments. However, based on their halo masses being more consistent with dwarf galaxies it now seems more likely that galaxies formed in such scenarios are in the minority. Many authors have argued that the large sizes can be explained through the secular evolution of dwarf galaxies in relatively high spin haloes (e.g. Amorisco & Loeb, 2016), or even from enlargement of the stellar component through gas outflows caused by stellar feedback. If these scenarios are dominant, then it could be expected that UDGs form with equal mass efficiencies in the field as they do in clusters. On the flip side of the argument, the large sizes of UDGs may be explained by the gravitational harassment of dwarf galaxies in the group or cluster environment, or from individual nearby galaxies.

While we have learned much about Ultra-Diffuse Galaxies (UDGs) in groups and clusters, relatively little is known about them in less-dense environments like the field field (i.e. outside of massive galaxy groups or clusters). This is the subject of chapter 3. More isolated UDGs are important for our understanding of UDG formation scenarios because they form via secular mechanisms, allowing us to determine the relative importance of environmentally-driven formation in groups and clusters. I have used the public Kilo-Degree Survey (KiDS) together with the Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP) to constrain the abundance and properties of UDGs in the field, targeting sources with low surface brightness ($24.0 \leq \bar{\mu}_{e,r} \leq 26.5$) and large apparent sizes ($3.0'' \leq \bar{r}_{e,r} \leq 8.0''$).

Accounting for several sources of interlopers in the UDG candidate selection based on canonical scaling relations and using an empirical UDG model based on measurements from the literature, I show that a scenario in which cluster-like red sequence

UDGs occupy a significant number of field galaxies is unlikely, with most field UDGs being significantly bluer and showing signs of localised star formation. An immediate conclusion is that UDGs are much more efficiently quenched in high-density environments. I estimated an upper-limit on the total field abundance of UDGs of $8\pm3\times10^{-3}$cMpc^{-3} within the selection range; this is the first observational constraint placed on the field density of the overall UDG population. The mass formation efficiency of UDGs implied by this upper-limit is similar to what is measured in groups and clusters. While it is not possible to attribute all of these field sources to pure secular formation (because we cannot guarantee that the observed sample is truly isolated), we cannot exclude the possibility that secular processes provide a significant contribution to UDG production overall.

Interestingly, the field sample of UDGs have systematically bluer colours than what would be expected from a quiescent population and are more similar to late-type dwarf galaxies in clusters; this implies that UDGs in clusters are efficiently quenched by the dense environment. While this has been suspected in the literature for some time, this is the first time it has been systematically proven observationally. Several of the sources show visual signs of isolated regions of star formation based on extremely blue colours that visually stand out against the background. However, red UDGs are also known to exist in the field. This is not surprising based on my analysis, which suggests that such sources do exist but constitute a minority of cases.

The study presented in chapter 3 represents the first systematic search and analysis of UDGs in the field that also attempts to constrain their *intrinsic* properties. The key point is that no distances were known for the observed sample, so I relied on an empirical model for both UDGs and interlopers in order to measure the UDG properties. This is important because even with future large scale surveys like Euclid and LSST, it will still be unlikely to have spectroscopic redshifts for large samples of LSB galaxies. However, photometric redshifts may provide a powerful tool to reject high redshift interlopers in a catalogue of LSB galaxy candidates. Perhaps in the future, a hybrid model, similar in essence to my empirical model, but one that also accounts for the information provided by photometric redshifts, may be used to statistically disentangle the properties of LSB galaxies in the field.

One of the key goals of modern astronomy has been to constrain the galaxy stellar mass function (GSMF), which can be used to test and calibrate the results from cosmological simulations. The authors of Jones et al. (2018) concluded from their measurements of HI-rich UDGs that the inclusion of such objects in the GSMF would make a negligible difference, with a maximal correction of approximately 1% compared

to that presented in Wright et al. (2017). In this work, I have found that HI-rich UDGs make up at-least around one-fifth of the overall UDG population. This would therefore suggest a maximal change in the GSMF of order ∼5%. The idea that there are significant amounts of baryonic material locked away in large low surface brightness galaxies is disfavoured, however, this can only be said for the relatively confined region of parameter space that has been explored in this work. Further improvements in both the quality of reduced data as well as in our ability to detect, measure and classify galaxies in deep imaging surveys will allow this parameter space to be expanded in the near future.

5.2.1 Future Work

Halo Mass Estimates

As discussed previously, there is an unexplained difference between the observed globular cluster numbers for UDGs in the Coma and Fornax clusters. One suggestion is that it is environmental effect; since the Fornax UDGs are all located in its core, it could be that the globular clusters have been stripped away from their UDG hosts. If this is the case, then there may well be an observable trend between the globular cluster specific frequency of UDGs as a function of the cluster radius. Perhaps the same study could also be extended to UDGs in the field, where the chances that their GCs have either been stripped or accreted are minimised. However, the problem with this is that it would be necessary to know the precise distances to the UDGs in order to perform the analysis.

One of the additional difficulties of the above is that it may be hard to observe such effects given the relatively large errors on the observed GC number counts for individual galaxies. One way to improve this would be to purify the GC selection process. Deeper u-band imaging in combination with k-band imaging, for example from the Next-Generation Fornax Survey could help select GCs and remove stellar interlopers.

Another alternative would be to perform the same analysis for UDGs in the Virgo cluster, which is slightly closer than the Fornax cluster. The increased depth of the Next Generation Virgo Survey would allow the GCLF to be probed to higher completeness levels and reduce the overall uncertainty. This may add to our understanding of why Fornax UDGs seem to be less enriched in terms of their globular cluster populations when compared to those UDGs in Coma.

Of course, one of the most reliable methods to measure the halo masses of galaxies

is through the kinematics of their stars or stellar systems (i.e. GCs). However such measurements are time consuming, so at present we must rely on other estimators to process large samples.

Ultra-Diffuse Galaxies in the Field

The study outline in chapter 3 can be improved in several ways. The first and most obvious is the repetition of the analysis for the larger footprint offered by the second public data release of the HSC-SSP. In fact, the whole analysis could be redone with detection in the HSC-SSP rather than KiDS, because it is much deeper and as of the second data release has a much better approach to background estimation. It may also be a worthwhile idea to estimate the number density of truly isolated UDGs in order to study their secular formation. The difficulty with this idea is that estimating the distances to large samples of LSB galaxies is not currently feasible, so it would be hard to say which galaxies are isolated.

There are other properties of UDGs asides from their intrinsic colours that could be tested using the empirical model approach. For example, the intrinsic size distribution of UDGs could provide clues into how UDGs have formed. If they are significantly larger inside cluster environments, then it would seem likely that environmental process play an important role in "puffing-up" dwarf galaxies as some theoretical models predict. Alternatively, if it is the other way around, then secular processes would appear as the dominant channel for UDG formation. One approach could be to use my empirical model approach to estimate the intrinsic size distribution of the UDGs. However, this may not be so easily compared to the size distribution measured by van der Burg et al. (2017) because their measurements are conflated with measurement error (this might be significant because of the steep power law distribution in size).

The empirical model could be further improved by including extra parameters like the Sérsic index. This would likely increase the purity of our UDG candidate sample because most UDGs are observed to have lower Sérsic indices than our cut at n=2.5. Including this in the model would enable us to account for the effects of lowering the cut in n. Alternatively, other colours could be probed in the analysis: HSC-SSP is not limited to the g and r bands, but also has the i, z & y bands. In combination with other bands, for example the Y, J, H and Ks bands from the VISTA VIKING survey (which also overlaps with GAMA), this would enable photometric redshifts to be estimated. These could be used to help constrain the properties of UDGs by improving our interloper rejection and thus improving the purity of the UDG sample. This kind of analysis would also be applicable to upcoming surveys like Euclid and

LSST.

It may also be possible to obtain relatively large samples of spectroscopic redshifts for a reasonably large sample of LSB galaxies from upcoming spectroscopic surveys. In the work presented in chapter 3, I found ∼30 matching sources in the GAMA spectroscopic catalogue, each of which having their own spectroscopic redshifts. It is important to have such measurements as for large samples it can be generally assumed that the distances to each of the sources is set by the Hubble flow rather than peculiar velocities, which can affect distance estimates from spectroscopic redshifts. Examples of such upcoming surveys are WAVES and DEVILS (Davies et al., 2018), which in combination attempt to improve on the GAMA survey in both area and depth.

Bibliography

Abraham R. G., van Dokkum P. G., 2014, Astronomical Society of the Pacific, Publications, 126, 55

Aihara H., et al., 2018, Publications of the Astronomical Society of Japan, 70, S8

Akhlaghi M., Ichikawa T., 2015, The Astrophysical Journal Supplement Series, 220, 1

Alpher R. A., Herman R. C., 1949, Physical Review, 75, 1089

Amorisco N. C., Loeb A., 2016, MNRAS, 459, L51

Amorisco N. C., Monachesi A., Agnello A., White S. D. M., 2018, MNRAS, 475, 4235

Astropy Collaboration et al., 2013, Astronomy and Astrophysics, 558, A33

Astropy Collaboration et al., 2018, The Astronomical Journal, 156, 123

Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Ž., Lupton R. H., Nichol R. C., Szalay A. S., 2004, The Astrophysical Journal, 600, 681

Baldry I. K., et al., 2012, MNRAS, 421, 621

Bassino L. P., Faifer F. R., Forte J. C., Dirsch B., Richtler T., Geisler D., Schuberth Y., 2006, Astronomy and Astrophysics, 451, 789

Beasley M. A., Trujillo I., 2016, The Astrophysical Journal, 830, 23

Beasley M. A., Romanowsky A. J., Pota V., Navarro I. M., Martinez Delgado D., Neyer F., Deich A. L., 2016, The Astrophysical Journal Letters, 819, L20

Behroozi P. S., Wechsler R. H., Conroy C., 2013, The Astrophysical Journal, 770, 57

Bell E. F., McIntosh D. H., Katz N., Weinberg M. D., 2003, The Astrophysical Journal Supplement Series, 149, 289

Bennet P., Sand D. J., Zaritsky D., Crnojević D., Spekkens K., Karunakaran A., 2018, The Astrophysical Journal Letters, 866, L11

Berry D. S., 2015, Astronomy and Computing, 10, 22

Bertin E., Arnouts S., 1996, Astronomy and Astrophysics Supplement, 117, 393

Bertone G., Hooper D., Silk J., 2005, Physics Reports, 405, 279

Bhattacharya S., Mahulkar V., Pandaokar S., Singh P. K., 2017, Astronomy and Computing, 18, 1

Binggeli B., Sandage A., Tammann G. A., 1985, The Astronomical Journal, 90, 1681

Blakeslee J. P., et al., 2009, The Astrophysical Journal, 694, 556

Bothun G. D., Impey C. D., Malin D. F., Mould J. R., 1987, The Astronomical Journal, 94, 23

Bothun G. D., Impey C. D., Malin D. F., 1991, The Astrophysical Journal, 376, 404

Brook C. B., Di Cintio A., Knebe A., Gottlöber S., Hoffman Y., Yepes G., Garrison-Kimmel S., 2014, The Astrophysical Journal Letters, 784, L14

Caon N., Capaccioli M., D'Onofrio M., 1993, MNRAS, 265, 1013

Capaccioli M., et al., 2012, in Science from the Next Generation Imaging and Spectroscopic Surveys. p. 1

Capaccioli M., et al., 2015, Astronomy and Astrophysics, 581, A10

Carleton T., Errani R., Cooper M., Kaplinghat M., Peñarrubia J., 2018, preprint, (arXiv:1805.06896)

Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, MNRAS, 319, 168

Coles P., Lucchin F., 1995, Cosmology. The origin and evolution of cosmic structure

Collins M. L. M., et al., 2013, The Astrophysical Journal, 768, 172

Conroy C., Gunn J. E., 2010, The Astrophysical Journal, 712, 833

Conroy C., Gunn J. E., White M., 2009, The Astrophysical Journal, 699, 486

Conselice C. J., 2018, Research Notes of the AAS, 2, 43

Contenta F., et al., 2018, MNRAS, 476, 3124

Cooper A. P., et al., 2010, MNRAS, 406, 744

Cowie L. L., Songaila A., Hu E. M., Cohen J. G., 1996, The Astronomical Journal, 112, 839

D'Abrusco R., et al., 2016, The Astrophysical Journal Letters, 819, L31

Dalcanton J. J., Spergel D. N., Summers F. J., 1995, arXiv e-prints, pp astro–ph/9503093

Dalcanton J. J., Spergel D. N., Gunn J. E., Schmidt M., Schneider D. P., 1997, The Astronomical Journal, 114, 635

Danieli S., van Dokkum P., 2018, preprint, (arXiv:1811.01962)

Das M., 2013, Journal of Astrophysics and Astronomy, 34, 19

Davies J. I., Phillipps S., Cawson M. G. M., Disney M. J., Kibblewhite E. J., 1988, MNRAS, 232, 239

Davies J. I., Phillipps S., Disney M. J., 1989, MNRAS, 239, 703

Davies J. I., Davies L. J. M., Keenan O. C., 2016, MNRAS, 456, 1607

Davies L. J. M., et al., 2018, MNRAS, 480, 768

Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, The Astrophysical Journal, 292, 371

Di Cintio A., Brook C. B., Dutton A. A., Macciò A. V., Obreja A., Dekel A., 2017, MNRAS, 466, L1

Di Cintio A., Brook C. B., Macciò A. V., Dutton A. A., Cardona-Barrero S., 2019, MNRAS, 486, 2535

Disney M. J., 1976, Nature, 263, 573

Disney M., Phillipps S., 1983, MNRAS, 205, 1253

Drinkwater M. J., Gregg M. D., Colless M., 2001, The Astrophysical Journal Letters, 548, L139

Driver S. P., et al., 2006, MNRAS, 368, 414

Driver S. P., et al., 2011, MNRAS, 413, 971

Erben T., et al., 2013, MNRAS, 433, 2545

Errani R., Peñarrubia J., Walker M. G., 2018, MNRAS, 481, 5073

Erwin P., 2015, The Astrophysical Journal, 799, 226

Ester M., Kriegel H.-P., Sander J., Xu X., 1996. AAAI Press, pp 226–231

Fensch J., et al., 2018, arXiv e-prints, p. arXiv:1812.07346

Ferguson H. C., 1989, Astrophysics and Space Science, 157, 227

Ferrarese L., et al., 2012, The Astrophysical Journal Supplement Series, 200, 4

Ferré-Mateu A., et al., 2018, MNRAS, 479, 4891

Fliri J., Trujillo I., 2016, MNRAS, 456, 1359

Forbes D. A., 2017, MNRAS, 472, L104

Forbes D. A., Read J. I., Gieles M., Collins M. L. M., 2018, preprint, (arXiv:1809.07831)

Freeman K. C., 1970, The Astrophysical Journal, 160, 811

Frenk C. S., White S. D. M., Davis M., 1983, The Astrophysical Journal, 271, 417

Friedmann A., 1922, Zeitschrift fur Physik, 10, 377

Geller M. J., Huchra J. P., 1989, Science, 246, 897

Genina A., et al., 2018, MNRAS, 474, 1398

Georgiev I. Y., Puzia T. H., Goudfrooij P., Hilker M., 2010, MNRAS, 406, 1967

Gerola H., Seiden P. E., Schulman L. S., 1980, The Astrophysical Journal, 242, 517

Grado A., Capaccioli M., Limatola L., Getman F., 2012, Memorie della Societa Astronomica Italiana Supplementi, 19, 362

Graham A. W., Driver S. P., 2005, Publications of the Astronomical Society of Australia, 22, 118

Graham A., Lauer T. R., Colless M., Postman M., 1996, The Astrophysical Journal, 465, 534

Graham A. W., Janz J., Penny S. J., Chilingarian I. V., Ciambur B. C., Forbes D. A., Davies R. L., 2017, The Astrophysical Journal, 840, 68

Greco J. P., et al., 2018a, The Astrophysical Journal, 857, 104

Greco J. P., et al., 2018b, The Astrophysical Journal, 857, 104

Guo Q., et al., 2013, MNRAS, 435, 897

Guth A. H., 1981, Physical Review D (Particles and Fields), 23, 347

Guttman A., 1984, SIGMOD Rec., 14, 47

Harris W. E., Harris G. L. H., Alessi M., 2013, The Astrophysical Journal, 772, 82

Harris W. E., Blakeslee J. P., Harris G. L. H., 2017, The Astrophysical Journal, 836, 67

Henden A. A., Levine S. E., Terrell D., Smith T. C., Welch D., 2012, Journal of the American Association of Variable Star Observers (JAAVSO), 40, 430

Herschel W., 1785, Philosophical Transactions of the Royal Society of London Series I, 75, 213

Hogg D. W., Baldry I. K., Blanton M. R., Eisenstein D. J., 2002, arXiv e-prints, pp astro–ph/0210394

Homma D., et al., 2019, arXiv e-prints, p. arXiv:1906.07332

Hubble E., 1929a, Proceedings of the National Academy of Science, 15, 168

Hubble E. P., 1929b, The Astrophysical Journal, 69

Hudson M. J., Robison B., 2018, MNRAS, 477, 3869

Hudson M. J., Harris G. L., Harris W. E., 2014, The Astrophysical Journal Letters, 787, L5

Hughes J. P., 1998, in Mazure A., Casoli F., Durret F., Gerbal D., eds, Untangling Coma Berenices: A New Vision of an Old Cluster. (arXiv:astro-ph/9709272)

Impey C., Bothun G., Malin D., 1988, The Astrophysical Journal, 330, 634

Iodice E., et al., 2016, The Astrophysical Journal, 820, 42

Irwin M. J., Davies J. I., Disney M. J., Phillipps S., 1990, MNRAS, 245, 289

Janssens S., Abraham R., Brodie J., Forbes D., Romanowsky A. J., van Dokkum P., 2017, The Astrophysical Journal, 839, L17

Janz J., Penny S. J., Graham A. W., Forbes D. A., Davies R. L., 2017, MNRAS, 468, 2850

Jester S., et al., 2005, The Astronomical Journal, 130, 873

Jiang F., Dekel A., Freundlich J., Romanowsky A. J., Dutton A., Maccio A., Di Cintio A., 2018, arXiv e-prints, p. arXiv:1811.10607

Jimenez R., Padoan P., Matteucci F., Heavens A. F., 1998, MNRAS, 299, 123

Jones M. G., Papastergis E., Pandya V., Leisman L., Romanowsky A. J., Yung L. Y. A., Somerville R. S., Adams E. A. K., 2018, Astronomy and Astrophysics, 614, A21

Jordán A., et al., 2007, The Astrophysical Journal Supplement Series, 169, 213

Kauffmann G., White S. D. M., Guiderdoni B., 1993, MNRAS, 264, 201

Keenan O. C., 2017, On the missing dwarf problem in clusters and around the nearby galaxy M33, `http://orca.cf.ac.uk/104065/`

Kelvin L. S., et al., 2012, MNRAS, 421, 1007

Kennedy R., Frenk C., Cole S., Benson A., 2014, MNRAS, 442, 2487

Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, The Astrophysical Journal, 522, 82

Kochoska A., Mowlavi N., Prša A., Lecoeur-Taïbi I., Holl B., Rimoldini L., Süveges M., Eyer L., 2017, Astronomy and Astrophysics, 602, A110

Koda J., Yagi M., Yamanoi H., Komiyama Y., 2015, The Astrophysical Journal, 807, L2

Kuijken K., 2011, The Messenger, 146, 8

Kuijken K., et al., 2019, Astronomy and Astrophysics, 625, A2

Laporte C. F. P., Agnello A., Navarro J. F., 2019, MNRAS, 484, 245

Leavitt H. S., Pickering E. C., 1912, Harvard College Observatory Circular, 173, 1

Leinert C., et al., 1998, Astronomy and Astrophysics Supplement, 127, 1

Leisman L., et al., 2017, The Astrophysical Journal, 842, 133

Lelli F., Fraternali F., Sancisi R., 2010, Astronomy and Astrophysics, 516, A11

Lemaître G., 1927, Annales de la Société Scientifique de Bruxelles, 47, 49

Lim S., Peng E. W., Côté P., Sales L. V., den Brok M., Blakeslee J. P., Guhathakurta P., 2018, The Astrophysical Journal, 862, 82

Lovell M. R., et al., 2012, MNRAS, 420, 2318

Lupton R., Blanton M. R., Fekete G., Hogg D. W., O'Mullane W., Szalay A., Wherry N., 2004, Astronomical Society of the Pacific, Publications, 116, 133

Mancera Piña P. E., Peletier R. F., Aguerri J. A. L., Venhola A., Trager S., Choque Challapa N., 2018, MNRAS, 481, 4381

Mancera Piña P. E., Aguerri J. A. L., Peletier R. F., Venhola A., Trager S., Choque Challapa N., 2019, MNRAS, p. 241

Martínez-Delgado D., et al., 2016, The Astronomical Journal, 151, 96

Mastropietro C., Moore B., Mayer L., Debattista V. P., Piffaretti R., Stadel J., 2005, MNRAS, 364, 607

McConnachie A. W., 2012, The Astronomical Journal, 144, 4

McFarland J. P., Verdoes-Kleijn G., Sikkema G., Helmich E. M., Boxhoorn D. R., Valentijn E. A., 2011, preprint, (arXiv:1110.2509)

McGaugh S. S., Bothun G. D., Schombert J. M., 1995, The Astronomical Journal, 110, 573

Mei S., et al., 2007, The Astrophysical Journal, 655, 144

Mihos J. C., et al., 2015, The Astrophysical Journal Letters, 809, L21

Mihos J. C., Harding P., Feldmeier J. J., Rudick C., Janowiecki S., Morrison H., Slater C., Watkins A., 2017, The Astrophysical Journal, 834, 16

Miller B. W., Lotz J. M., 2007, The Astrophysical Journal, 670, 1074

Moore B., Lake G., Katz N., 1998, The Astrophysical Journal, 495, 139

Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, The Astrophysical Journal, 524, L19

Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010, The Astrophysical Journal, 710, 903

Mowla L., van Dokkum P., Merritt A., Abraham R., Yagi M., Koda J., 2017, The Astrophysical Journal, 851, 27

Muñoz R. P., et al., 2015, The Astrophysical Journal Letters, 813, L15

Muzzin A., et al., 2013, The Astrophysical Journal, 777, 18

Papastergis E., Adams E. A. K., Romanowsky A. J., 2017, Astronomy and Astrophysics, 601, L10

Peng C. Y., Ho L. C., Impey C. D., Rix H.-W., 2002, The Astronomical Journal, 124, 266

Peng E. W., et al., 2006, The Astrophysical Journal, 639, 95

Penzias A. A., Wilson R. W., 1965, The Astrophysical Journal, 142, 419

Perlmutter S., et al., 1999, The Astrophysical Journal, 517, 565

Planck Collaboration et al., 2018, arXiv e-prints, p. arXiv:1807.06209

Pota V., et al., 2018, MNRAS, 481, 1744

Prialnik D., 2000, An Introduction to the Theory of Stellar Structure and Evolution

Prole D. J., Davies J. I., Keenan O. C., Davies L. J. M., 2018, MNRAS, 478, 667

Prole D. J., Burg R. F. J. v. d., Hilker M., Davies J. I., 2019a, MNRAS, p. 1792

Prole D. J., et al., 2019b, MNRAS, 484, 4865

RS P., Zaritsky D., Donnerstein R., Spekkens K., 2019, arXiv e-prints, p. arXiv:1904.06262

Rejkuba M., 2012, Astrophysics and Space Science, 341, 195

Riess A. G., et al., 1998, The Astronomical Journal, 116, 1009

Robotham A. S. G., et al., 2011, MNRAS, 416, 2640

Robotham A. S. G., Taranu D. S., Tobar R., Moffett A., Driver S. P., 2017, MNRAS, 466, 1513

Robotham A. S. G., Davies L. J. M., Driver S. P., Koushan S., Taranu D. S., Casura S., Liske J., 2018, MNRAS, 476, 3137

Roediger J. C., et al., 2017, The Astrophysical Journal, 836, 120

Román J., Trujillo I., 2017a, MNRAS, 468, 703

Román J., Trujillo I., 2017b, MNRAS, 468, 4039

Román J., Trujillo I., 2018, Research Notes of the American Astronomical Society, 2, 144

Román J., Beasley M. A., Ruiz-Lara T., Valls-Gabaud D., 2019a, arXiv e-prints, p. arXiv:1903.08168

Román J., Trujillo I., Montes M., 2019b, arXiv e-prints, p. arXiv:1907.00978

Rong Y., Guo Q., Gao L., Liao S., Xie L., Puzia T. H., Sun S., Pan J., 2017, MNRAS, 470, 4231

Rubin V. C., Ford Jr. W. K., 1970, The Astrophysical Journal, 159, 379

Rubin V. C., Ford Jr. W. K., Thonnard N., 1980, The Astrophysical Journal, 238, 471

Ruiz-Lara T., et al., 2018, MNRAS, 478, 2034

Sabatini S., Davies J., Scaramella R., Smith R., Baes M., Linder S. M., Roberts S., Testa V., 2003, MNRAS, 341, 981

Sabatini S., Davies J., van Driel W., Baes M., Roberts S., Smith R., Linder S., O'Neil K., 2005, MNRAS, 357, 819

Sánchez-Janssen R., et al., 2018, arXiv e-prints, p. arXiv:1812.01019

Sandage A., Binggeli B., 1984, The Astronomical Journal, 89, 919

Sandin C., 2014, Astronomy and Astrophysics, 567, A97

Sawala T., et al., 2016, MNRAS, 457, 1931

Schombert J. M., McGaugh S. S., Eder J. A., 2001, The Astronomical Journal, 121, 2420

Schuberth Y., Richtler T., Hilker M., Dirsch B., Bassino L. P., Romanowsky A. J., Infante L., 2010, Astronomy and Astrophysics, 513, A52

Shapley H., 1919, The Astrophysical Journal, 49

Sifón C., van der Burg R. F. J., Hoekstra H., Muzzin A., Herbonnet R., 2018, MNRAS, 473, 3747

Singh P. R., Zaritsky D., Donnerstein R., Spekkens K., 2019, The Astronomical Journal, 157, 212

Slater C. T., Harding P., Mihos J. C., 2009, Astronomical Society of the Pacific, Publications, 121, 1267

Somerville R. S., Popping G., Trager S. C., 2015, MNRAS, 453, 4337

Spitler L. R., Forbes D. A., 2009, MNRAS, 392, L1

Springel V., et al., 2005, Nature, 435, 629

Springel V., Frenk C. S., White S. D. M., 2006, Nature, 440, 1137

Springel V., et al., 2008, MNRAS, 391, 1685

Stetson P. B., 1987, Astronomical Society of the Pacific, Publications, 99, 191

Taylor E. N., et al., 2011, MNRAS, 418, 1587

Taylor E. N., et al., 2015, MNRAS, 446, 2144

Taylor R., Davies J. I., Jáchym P., Keenan O., Minchin R. F., Palouš J., Smith R., Wünsch R., 2017, MNRAS, 467, 3648

Teeninga P., Moschini U., Trager S. C., Wilkinson M. H., 2016, in International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing. pp 157–168

Teyssier R., Moore B., Martizzi D., Dubois Y., Mayer L., 2011, MNRAS, 414, 195

Toloba E., et al., 2018, The Astrophysical Journal Letters, 856, L31

Torrealba G., et al., 2019, MNRAS, p. 1548

Trujillo I., Fliri J., 2016, The Astrophysical Journal, 823, 123

Trujillo I., et al., 2018, preprint, (arXiv:1806.10141)

Valls-Gabaud D., 2016, Proceedings of the International Astronomical Union, 11, 199–201

Venhola A., et al., 2017, Astronomy and Astrophysics, 608, A142

Venhola A., et al., 2018, Astronomy and Astrophysics, 620, A165

Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, Physical Review D (Particles and Fields), 88, 043502

Vika M., Vulcani B., Bamford S. P., Häußler B., Rojas A. L., 2015, Astronomy and Astrophysics, 577, A97

Villegas D., et al., 2010, The Astrophysical Journal, 717, 603

Viola M., et al., 2015, MNRAS, 452, 3529

Vulcani B., et al., 2014, MNRAS, 441, 1340

Weigel A. K., et al., 2017, The Astrophysical Journal, 845, 145

Weinberg S., 2008, Cosmology. Oxford University Press

Wetzel A. R., Hopkins P. F., hoon Kim J., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, The Astrophysical Journal, 827, L23

White S. D. M., Rees M. J., 1978, MNRAS, 183, 341

Williams J. P., de Geus E. J., Blitz L., 1994, The Astrophysical Journal, 428, 693

Wittmann C., Lisker T., Pasquali A., Hilker M., Grebel E. K., 2016, MNRAS, 459, 4450

Wittmann C., et al., 2017, MNRAS, 470, 1512

Wolf C., et al., 2018, Publications of the Astronomical Society of Australia, 35, e010

Wright A. H., et al., 2017, MNRAS, 470, 283

Yagi M., Koda J., Komiyama Y., Yamanoi H., 2016, The Astrophysical Journal Supplement Series, 225, 11

Yozin C., Bekki K., 2015, MNRAS, 452, 937

Zaritsky D., et al., 2018, preprint, (arXiv:1811.04098)

Zaritsky D., et al., 2019, The Astrophysical Journal Supplement Series, 240, 1

Zwicky F., 1933, Helvetica Physica Acta, 6, 110

de Blok W. J. G., 2010, Advances in Astronomy, 2010, 789293

de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, Experimental Astronomy, 35, 25

van Dokkum P. G., Abraham R., Merritt A., Zhang J., Geha M., Conroy C., 2015, The Astrophysical Journal, 798, L45

van Dokkum P., et al., 2016, The Astrophysical Journal Letters, 828, L6

van Dokkum P., et al., 2017, The Astrophysical Journal Letters, 844, L11

van Dokkum P., et al., 2018, Nature, 555, 629

van Dokkum P., Danieli S., Abraham R., Conroy C., Romanowsky A. J., 2019, arXiv e-prints, p. arXiv:1901.05973

van der Burg R. F. J., Muzzin A., Hoekstra H., 2016, Astronomy and Astrophysics, 590, A20

van der Burg R. F. J., et al., 2017, Astronomy and Astrophysics, 607, A79

van der Wel A., 2008, The Astrophysical Journal, 675, L13

van der Wel A., et al., 2014, The Astrophysical Journal, 788, 28