

Advancing Techniques for Detecting Dwarf Satellite Populations Beyond the Local Group

by

Chengyu Xi

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2022

© Chengyu Xi 2022

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Marcin Sawicki
 Professor, Department of Astronomy and Physics
 Saint Mary's University

Supervisor: James E. Taylor
 Associate Professor, Department of Physics and Astronomy,
 University of Waterloo

Internal Member: Michael Balogh
 Professor, Department of Physics and Astronomy,
 University of Waterloo

Internal Member: Achim Kempf
 Professor, Dept. of Applied Mathematics,
 University of Waterloo

Internal-External Member: Ghazal Geshnizjani
 Associate Professor, Department of Applied Mathematics,
 University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis is in an integrated form. Chapters 2–4 correspond to three published or submitted papers, co-authored by me, my supervisor, and collaborators.

Chapter 2 was originally published as [Xi et al. 2018](#) (Xi C., Taylor J. E., Massey R. J., Rhodes J., Koekemoer A., Salvato M., 2018, MNRAS, 478, 5336). In this paper, I performed the data selection and analysis, including sections 2.2.4–2.2.5, figures 2.2 and 2.4; the clustering calculation in sections 2.4–2.5, figures 2.5–2.7; testing and optimization of the structural cuts in section 2.6, and figures 2.8–2.17. J. Taylor proposed the original idea and contributed most of the text in the introduction (2.1) and summary (2.8), as well as figure 2.1, the work related to the Local Volume Catalogue, including sections 2.2.2–2.2.3 and figure 2.3, the work related to serendipitous discoveries, including section 2.7 and figures 2.18–2.21, and appendix 2.9. R.J. Massey, J. Rhodes and A. Keokemoer provided input on COSMOS imaging and photometry, and M. Salvato provided the internal catalogue of COSMOS spectroscopy. All authors contributed to editing the text and discussing the science in the paper.

Chapter 3 was originally published as [Xi & Taylor \(2021\)](#) (Xi C., Taylor J. E., 2021, MNRAS). In this paper, I performed most of the analysis and calculations, produced most of the plots and wrote most of the text. J. Taylor proposed the initial idea, was involved in all the analysis, and edited the text throughout the paper. He also contributed most of the content of the conclusion (3.7) and the abstract; and wrote the explanation and discussion of the results from [Conroy, Wechsler & Kravtsov \(2006\)](#) in section 3.5.3.

Chapter 4 has been submitted to MNRAS for publication and is currently under review. (Note that some material in the thesis, notably Figures 4.13–4.16, and Appendix 4.7, have been revised relative to the submitted version, based on feedback from the anonymous referee and the thesis committee.) In this paper, I performed most of the analysis and calculations, produced most of the plots and wrote most of the text. J. Taylor proposed the specific analyses and edited the text throughout the paper. He also performed the comparison of primary and secondary mass range to other work from the literature, including a part of section 4.3.3 and figures 4.6–4.7; discussed the implications of figures 4.2–4.4 and compared them to [Zu & Mandelbaum \(2015\)](#) as a part of sections 4.3.1 and 4.3.2; calculated the theoretical predictions of quiescent fraction as a function to primary-to-satellite halo mass ratio, explained the calculation in section 4.5.3, and provided the produced figure 4.16.

Abstract

This thesis includes material from three papers that develop techniques for quantifying faint dwarf satellite populations outside the Local Group. Dwarf satellites are important for our understanding of small-scale structure formation, the history of galaxy formation, and the nature of dark matter. Identifying faint satellites is technically challenging, however, because accurate distance information for these objects is usually unavailable and very expensive to obtain.

In Chapter 2 we test a previously proposed method of estimating the average satellite population around nearby bright galaxies. The method uses structural cuts on size and magnitude to preferentially select low-redshift dwarf galaxies, and clustering to estimate the fraction of true satellites within the selected sample. Using the high-precision photometric redshifts of the COSMOS survey, we were able to test the effectiveness of different structural cuts and optimize them for several different redshift ranges. We also describe a set of very nearby dwarf galaxies (at distances $D < 200$ Mpc) identified morphologically in the COSMOS field.

In chapters 3 and 4, we introduce a new method for quantifying satellite abundance specifically designed for samples with high-quality photometric redshifts. The method allows us to measure satellite abundance around primaries in crowded fields, improving on previous methods that considered only isolated systems. The method also avoids the use of spectroscopic redshifts, which makes it much more expensive in observation time to reach a given depth in the satellite population.

Chapter 3 focuses on establishing and testing the method, and presents some initial science results. In Chapter 4, we measure various satellite properties, including the satellite stellar mass function, the relative stellar mass function, and the quiescent fraction, as well as their dependence on primary properties such as mass, colour and specific star formation rate.

Acknowledgements

I would like to thank all the people who made this thesis possible. In particular, I would like to thank Dr. James Taylor for years of supervision, and the mental support during the difficult times. I would like to thank Nicole Drakos, Saoussen Mbarek and Haoyang Wei for proofreading this thesis and providing valuable feedback.

Table of Contents

List of Figures	xii
List of Tables	xvi
1 Introduction	1
1.1 Galaxies	1
1.1.1 Galaxy classification	2
1.1.2 Galaxy masses; dwarf and giant galaxies	3
1.1.3 Evidence for a dark component on galaxy scales	4
1.2 Dark Matter and Structure Formation	5
1.2.1 Linear and nonlinear structure formation	6
1.2.2 Cosmological N -body simulations	6
1.2.3 Dark Matter Halos	7
1.2.4 Galaxy formation	9
1.2.5 The stellar-to-halo mass relation (SHMR)	11
1.3 Galaxy Groups and Clusters	12
1.3.1 Satellite and Central Galaxies	12
1.3.2 Groups and clusters	13
1.3.3 The Milky Way and The Local Group	13
1.3.4 Quenching and environmental effects	14

1.4	Previous Work and Motivation	15
1.4.1	Small-scale challenges in Λ CDM	15
1.4.2	Solutions to the SSCs	16
1.4.3	The need to go beyond the Local Group	19
1.4.4	Prior Work	20
1.4.5	Our research goals	21
1.5	Organization of the thesis	22
2	Quantifying the abundance of faint, low-redshift satellite galaxies in the COSMOS survey	23
2.1	Introduction	24
2.2	Data	26
2.2.1	The Local Group	27
2.2.2	The Local Volume	27
2.2.3	More Distant Objects	28
2.2.4	COSMOS	28
2.2.5	SDSS Photometry in the COSMOS Field	29
2.3	The Basis for Structural Selection	31
2.4	Clustering Measurement Method	34
2.4.1	Selecting Primaries	36
2.4.2	Selecting Secondaries	37
2.4.3	Masking Corrections	37
2.4.4	The Clustering Calculation	39
2.4.5	Figure of Merit for Clustering	40
2.5	The Clustering Signal	40
2.5.1	The Signal with no Additional Cuts on the Secondary Sample	40
2.5.2	The Signal with Photo- z Cuts on the Secondary Sample	42
2.6	Effect of Structural Cuts	45

2.6.1	Single-parameter Cuts	45
2.6.2	Purity of the Secondary Samples	47
2.6.3	Two-parameter Cuts	49
2.6.4	Optimal Structural Cuts	53
2.6.5	Completeness and Bias in Other Properties	58
2.7	Other Morphological Distance Indicators	61
2.7.1	Serendipitous Discoveries and their Redshifts	61
2.7.2	Notes on Individual Objects	67
2.8	Summary and Conclusions	69
3	A hierarchical clustering method for quantifying satellite abundance	72
3.1	Introduction	73
3.2	Data–COSMOS	76
3.2.1	The COSMOS photometric redshift catalogue	76
3.2.2	Additional spectroscopic redshifts	77
3.2.3	The base sample	77
3.2.4	Stellar Mass Completeness	78
3.3	Determining the Clustering Scale	80
3.3.1	Halo Mass Assignment	80
3.3.2	Defining a Region of Interest (ROI)	81
3.4	The iterative clustering method	84
3.4.1	Initial primary selection	84
3.4.2	Clustering of the First-run Primary and Secondary Samples	85
3.4.3	Iterating over the Fit	94
3.5	Estimating Satellite Abundance	95
3.5.1	Abundance Estimates Using Method A	95
3.5.2	Abundance Estimates Using Method B	97
3.5.3	Comparison to Previous Work	98

3.5.4	The Satellite Luminosity Function	102
3.6	Testing for Systematic Uncertainties in the Method	102
3.6.1	Null Test	104
3.6.2	Parameter and Systematic Tests	104
3.6.3	Discussion	105
3.7	Conclusions	107
3.8	Data Availability	109
4	The hierarchical clustering method: abundance and properties of local satellite populations	110
4.1	Introduction	111
4.2	Review of the method	113
4.2.1	The Hierarchical Clustering Method	113
4.2.2	Calculating Distributions of Satellite Properties	116
4.2.3	Spectroscopically Confirmed Subsample	117
4.3	Satellite Stellar Mass functions	118
4.3.1	Absolute Stellar Mass Functions	118
4.3.2	Relative Stellar Mass Functions	121
4.3.3	Comparison with Literature Results	125
4.4	Satellite-Central connection	130
4.4.1	Satellite Populations Split by Central Colour or sSFR	131
4.4.2	Conformity in the Spectroscopic Sample	133
4.5	Environmental Quenching	136
4.5.1	Quenching versus Secondary Stellar Mass	137
4.5.2	Quiescent Fraction versus Primary Morphology and Stellar Mass	139
4.5.3	Quiescent fraction vs relative mass	140
4.5.4	Quenching: Summary	144
4.6	Summary and Outlook	146

5 Conclusion	148
5.1 Discussion	150
5.2 Future Research Directions	153
References	154
APPENDICES	169
A Additional details	170
A.1 Additional references for the total mass of MW	170
B Appendix for chapter 2	172
B.1 The Serendipitous Catalogue	172
C Appendix for chapter 3	175
C.1 Bias in Halo Masses Derived from the SHMR	175
C.2 Measuring the Masking and field boundaries	178
C.3 Comparing Background Estimation Methods A and B	179
D Appendix for chapter 4	181
D.1 Catalogue completeness	181
D.2 Redshift distribution for the six primary mass bins	181
D.3 Comparison to SAGA II (Mao et al. 2021)	183
D.4 Primary Subsamples in Colour and sSFR	184
Glossary	187
Abbreviations	189
List of Symbols	191

List of Figures

2.1	The difference between the photometric and spectroscopic redshift (COSMOS 2015), converted to a distance error, as a function of their i^+ magnitude.	30
2.2	Absolute B -band magnitude versus distance D for nearby galaxies, from the catalogs of K13 and Tully16.	33
2.3	Magnitude, size and surface brightness for galaxies in the Local Volume Catalogue.	35
2.4	The cumulative and differential redshift distributions of the primary and secondary samples, over the range $z = 0-1$	38
2.5	Clustering between the full secondary sample and primaries in three redshift ranges.	41
2.6	The surface number density of primary-secondary pairs as a function of projected separation R_p and velocity offset $\Delta V = c\Delta z$	43
2.7	The clustering signal, as in Figure 2.5, but after applying photo- z cuts to select secondaries likely to be at the same redshift as their primary.	44
2.8	The clustering SNR as a function of bright and faint limits on the magnitude of the secondaries.	46
2.9	Purity of the secondary sample as a function of a faint magnitude limit and as a function of a lower size limit.	48
2.10	The magnitude-size distribution of the secondary sample, coloured by redshift.	50
2.11	The magnitude-surface brightness distribution of the secondary sample, coloured by redshift.	51
2.12	SNR of the clustering measurement as a function of size-magnitude cuts on the secondary sample, parameterized by an intercept r_0 and a slope m (Equation 2.4).	52
2.13	SNR of the clustering measurement as a function of surface brightness-magnitude cuts on the secondary sample, parameterized by an intercept r_0 and a slope m (Equation 2.5).	54

2.14	The clustering signal between primaries in the redshift range $z = 0.07\text{--}0.15$ and the secondary sample after an optimal size-magnitude cut has been applied. . . .	55
2.15	The redshift (mainly photo- z) distributions of the entire secondary sample, and secondary samples after the best cuts in magnitude, in size, or in both size and magnitude have been applied.	57
2.16	$g - r$ colour distribution of the entire secondary sample and the distributions after two optimized cuts.	59
2.17	Excess counts ΔN as a function of secondary colour and as a function of specific star formation rate.	60
2.18	Visual class versus distance, inferred from photometric (red squares) or spectroscopic (blue triangles) redshift, for the serendipitous sample.	62
2.19	Apparent magnitude versus distance for the serendipitous sample.	64
2.20	Cutouts from the COSMOS ACS mosaics, showing examples of the different visual classes from the serendipitous sample.	65
2.21	An enlarged ACS F814W mosaic image (Koekemoer et al., 2007) of one of the Class 1 objects (ID 549719 in the COSMOS 2015 catalogue), showing diffuse, low surface-brightness emission and multiple point sources.	66
3.1	The redshift versus i^+ band absolute magnitude distribution of our base catalogue, coloured by the i^+ band apparent magnitude.	79
3.2	Primary-secondary clustering signal in the base catalogue.	83
3.3	The redshift-halo mass distribution of the first-run primary sample.	87
3.4	The surface density of secondaries around each of the primary subsamples shown in Fig. 3.3.	89
3.5	Average S_{halo} versus halo mass ($\log_{10}(M_{\text{h}}/M_{\odot})$), with colour indicating the mean redshift of the primary sub-sample.	90
3.6	As Fig.3.5, but with sub-samples of similar mean halo mass combined into single bins for clarity.	90
3.7	The background surface density Σ_{bg} versus d_A^{-2}	92
3.8	Likelihood distributions for the clustering model parameter values (fitting the full model to all 20 primary sub-samples simultaneously).	93

3.9	Satellite abundance as a function of primary stellar mass, estimated using method A.	96
3.10	Satellite abundance as a function of primary stellar/halo mass, estimated using method B.	99
3.11	Comparison to the HOD-based results of Conroy, Wechsler & Kravtsov (2006).	100
3.12	Comparison to the observed and mock results of Besla et al. (2018).	101
3.13	Satellite luminosity functions, for three ranges of primary halo mass, estimated using method B.	103
3.14	Results of a null test where the halo component of the clustering signal is measured using galaxies outside the velocity limits of the ROI.	104
4.1	Number of pairs versus satellite probability, as well as the cumulative number over that probability, with and without the spectroscopically confirmed pairs.	115
4.2	The (cumulative) satellite stellar mass function, for 6 bins of primary stellar mass.	119
4.3	The (cumulative) satellite stellar mass function, for 6 bins of primary halo mass.	120
4.4	The SSMF per unit halo mass, or ‘efficiency’ of subhalo occupation, using the same same bins of primary halo mass as in Fig. 4.3.	122
4.5	Satellite abundance as a function of the primary-to-secondary stellar mass ratio μ_* (the RSMF).	124
4.6	Limiting sensitivity, in secondary stellar mass versus primary halo mass, of a number of previous studies of satellite abundance.	126
4.7	Amplitude of the cumulative stellar mass function at $M_* = 10^9 M_\odot$, for the current study, compared to previous results from Wang & White (2012) and Lan, Ménard & Mo (2016).	128
4.8	The RSMF for MW-mass hosts, compared to the SAGA II results of Mao et al. (2021).	129
4.9	The SSMF for red/blue primaries in three primary stellar mass ranges.	132
4.10	As Fig. 4.9, but for primary samples split by sSFR.	133
4.11	Colour index versus stellar mass for the spectroscopically confirmed satellites; points are coloured by the redshift.	134
4.12	The colour index of the spectroscopically confirmed satellites versus the colour index of their primaries.	135

4.13	Quiescent fraction versus stellar mass for the whole base catalogue, field galaxies, satellites and results from the SAGA II (Mao et al., 2021) and LG (Wetzel, Tollerud & Weisz, 2015).	138
4.14	Quiescent fraction as function of secondary stellar mass, for the red and blue primary subsamples, as well as the whole base catalogue.	141
4.15	Quiescent fraction as function of secondary stellar mass for three primary stellar mass bins, and for the whole base catalogue.	142
4.16	Satellite quiescent fraction versus the mass ratio of the primary halo mass to the satellite infall halo mass.	145
C.1	Stellar mass errors versus stellar mass in the mock sample (with the averaged relation).	176
C.2	Halo-to-stellar mass relation of the Monte-Carlo samples.	177
C.3	Cumulative satellite luminosity functions estimated using methods A (dashed) and B (solid), for primaries in three halo mass bins (as in Fig. 3.13).	180
D.1	Test of completeness for stellar mass, sSFR and colour.	182
D.2	Cumulative redshift distributions for the six primary halo mass bins used in the SSMF calculation.	183
D.3	The $B - i^+$ colour versus stellar mass distribution for the primaries.	185
D.4	sSFR versus $B - i^+$ colour for the primaries.	186

List of Tables

2.1	Optimal values for size-magnitude cuts and for surface brightness-magnitude cuts.	56
3.1	The model fitting results from the first run and first iteration	95
3.2	The fitting results of all tests	106
B.1	The Serendipitous Catalog	173
B.2	The Serendipitous Catalog	174

Chapter 1

Introduction

Dwarf galaxies are of particular importance in the study of galaxy formation and structure formation, as they provide important tests for different theories in these fields. However, due to their being low-mass and faint, identifying them, especially outside the Local Group, is observationally challenging. The goal of this thesis is to develop and test new methods for identifying extra-galactic dwarfs, or at least to measure their abundance statistically. In this introduction, we will first review background concepts in the study of galaxies, dark matter structures, and galaxy groups and clusters. (Specific technical terms are also explained in the glossary [D.4](#).) We will then introduce the research goals of this thesis, as well as some of the previous work on this problem.

1.1 Galaxies

In general, a **galaxy** refers to a gravitationally bound system of stars, stellar remnants, interstellar medium, and **dark matter**¹. Galaxies are the fundamental building blocks of most visible structure in the Universe, and understanding galaxy formation and evolution is one of the most important goals in astrophysics.

The Galaxy that hosts our Solar System was recognized very early on². The term “**Milky Way**” (**MW**) was used to describe the bright band that can be seen on a night sky. In 1610, Galileo Galilei introduced the usage of telescopes to astronomy and discovered that the MW was

¹See section [1.2](#) for explanations

²The following summary of the discovery and properties of galaxies was inspired by and closely follows that of [Evans \(1998\)](#), as cited in the page <https://en.wikipedia.org/wiki/Galaxy>

composed of a huge number of faint stars (Galileo, 1610, as cited in [Evans, 1998](#)). In 1750, English astronomer Thomas Wright correctly speculated that the MW might be a rotating body of a huge number of stars held together by gravitational forces, including our Solar System as a part of it (Wright, 1750, as cited in [Evans, 1998](#)). Thus, the basic nature of our Galaxy has been suspected for almost 300 years.

However, extra-galactic astronomy has a much shorter history. There are only a few large external galaxies that are visible to the naked eye on a dark night, such as the **Andromeda Galaxy**, the [Large Magellanic Cloud \(LMC\)](#) and the [Small Magellanic Cloud \(SMC\)](#). Early astronomers noticed their extended shapes, relative to point-like stars, and described them as “small clouds” or “nebulae” (Astronomers did not know they were galaxies yet!). Toward the end of the 18th century, the French astronomer Charles Messier built a well-known catalogue containing about 110 brightest celestial objects having nebulous appearances, among which the Andromeda Galaxy is listed as “**Messier 31**” and now commonly referred to as **M31**. Some early hints suggested that the nebulae may be separate objects outside our Galaxy, rather than being parts of the MW. In the early 1900s, Vesto Slipher discovered that the majority of the nebulae have high Doppler shifts compared to most of the stars he observed, indicating that they are moving away from us quickly ([Slipher, 1913, 1915](#)). In 1917, Heber Curtis found that nova events within the M31 appear to be much fainter than the usual nova events that occur in our Galaxy, hinting that M31 may be a remote object (Curtis, 1917, reprinted as [Curtis, 1988](#)). These hints led to a historical debate among astronomers regarding the existence of extra-galactic objects and the scale of the Universe. However, more direct evidence was still needed to settle the dispute. Finally, in the 1920s, individual [Cepheid variable](#) stars were resolved by Edwin Hubble ([Hubble, 1929](#)), providing robust distance estimates to external nebulae, and most astronomers became convinced that many of the “nebulous” objects in the sky were indeed external galaxies ([Mo, van den Bosch & White, 2010](#)).

1.1.1 Galaxy classification

Our knowledge of extra-galactic galaxies has advanced greatly since their first discovery, with hundreds of millions of galaxies having been identified. Galaxies can be classified into three main categories by their visual appearance: **spirals**, **ellipticals** and **irregular/peculiars**. The visual classification originated from the work by [Hubble \(1927\)](#), who developed a sequence of galaxy morphology that further divides galaxies into a set of sub-categories, an arrangement known as the **Hubble sequence**. A good review of modern galaxy classification can be found in [Buta \(2013\)](#). Large, bright galaxies are typically found in either spiral or elliptical forms. For spiral galaxies, most of their stars and other visible material are contained in a flattened, disk-shaped plane. Their appearance in the sky highly depends on the orientation of their disk relative to

our line of sight. In contrast, elliptical galaxies appear elliptical in shape, regardless of viewing angle (See [Conselice, 2014](#), for a detailed review of galaxy structure). Our MW is a typical spiral galaxy with a well-developed galactic disk of approximately 30 kpc (100,000 ly) in diameter and 0.6 kpc (2,000 ly) in thickness ([Mo, van den Bosch & White, 2010](#)). Our Solar System sits about 8.3 kpc (27,000 ly) away from the galactic center. (See [Bland-Hawthorn & Gerhard, 2016](#); [Gillessen et al., 2017](#), for reviews of the properties of the MW and its role as a benchmark in galactic studies.)

Conventionally, elliptical galaxies are often referred to as “**early-type**” galaxies while spiral galaxies are referred to as “**late-type**” galaxies. There is a common misconception that this naming convention implies that the morphology of galaxies evolves from ellipticals, which are pretty much featureless, to spiral galaxies with more structural complexity. However, this terminology was proposed simply to distinguish simple systems from complex ones, without any implication of an evolution connection. One property that is more closely related to galaxy age is the intensity of new stars being formed, or the star formation rate (SFR). Interestingly, the “early-type” galaxies typically consist mainly of old stellar populations, and have very low SFRs; while the “late type” galaxies have a much higher rate of star formation, and thus have younger stellar populations overall. The age difference in their stellar populations also leads to a bimodality of galaxy colours (dividing galaxies into “red” and “blue” galaxies), as young stars emit more blue light compared to old stars. I will discuss galaxy colour and star formation activity in section 1.3.4.

Besides ellipticals and spirals, peculiar galaxies are usually the result of strong interactions with other galaxies, during galaxy mergers. Irregular galaxies refer to galaxies that can not be readily classified into an elliptical or spiral morphology. Irregular/peculiar types are most often found in small, faint galaxy populations. I will discuss these “**dwarf**” galaxies next.

1.1.2 Galaxy masses; dwarf and giant galaxies

Here I will discuss the typical mass range of galaxies, and naming conventions in the field. It is worth noting that a galaxy is a complicated structure that includes multiple components such as stars, an interstellar medium, and a dark matter halo (discussed below); the definition of its mass can vary, depending on the definition of galactic boundary and the components that are included. Two descriptions are most commonly used to define the mass of a galaxy – **stellar mass** and **total mass**. Stellar mass (M_*) refers to the total mass of the stellar component of a galaxy, which is normally inferred from the observed luminosity of the galaxy. The total mass refers to the total mass of the visible part of the galaxy together with its host dark matter halo (which will be discussed in the second part of this introduction.) In addition, we also use **halo mass** (M_h) to

describe the mass of the host dark matter halo. Note that the stellar mass of a given galaxy is usually significantly smaller than the mass of its host dark matter halo, the halo mass is usually quantitatively close to the total mass.

The total mass of the MW system, including dark matter, has been estimated in many ways over the last few decades. Despite considerable effort, the uncertainties of these mass estimates remain about a factor of two, ranging from 0.5 to $2.5 \times 10^{12} M_{\odot}$ (Wang et al., 2015; Bland-Hawthorn & Gerhard, 2016; Callingham et al., 2018), while its stellar mass is less than $1/10^{\text{th}}$ of this. (See Appendix A.1 for more references of MW mass estimates.) This makes the MW a luminous, or “giant” galaxy. However, this thesis focuses mainly on “dwarf” galaxies. These are galaxies with much smaller masses. Although not as noticeable as giant galaxies, dwarf galaxies make up the majority of galaxies in the Universe in terms of abundance and have a huge implication on our understanding of galaxy formation. (I will discuss the particular research interest of dwarf galaxies in more detail in the next section.)

Observationally, there seems to be no limit to how small a galaxy can be. Prior to 2004, the smallest galaxy known was Draco³, which has a stellar mass of $5 \times 10^5 M_{\odot}$, that is 10^5 less massive than the MW. Today, we know of dwarf galaxies with stellar masses 1000 times smaller (Bullock & Boylan-Kolchin, 2017). See also Willman (2010) for a more detailed review of the search for faint dwarfs.

While the term “dwarf” generally applies to galaxies with stellar masses $M_* \leq 10^9 M_{\odot}$, dwarfs can be further divided into more sub-types depending on the mass ranges. Although there is not a clear definition for each type of dwarf, there are some naming conventions that we can adopt for this thesis: The “classical dwarfs”⁴ have a stellar mass range of 10^5 – $10^7 M_{\odot}$. Dwarfs with larger stellar mass are usually called bright or luminous dwarfs, while dwarfs with smaller stellar masses (around 10^2 – $10^5 M_{\odot}$) are called “ultra-faint dwarfs”.

1.1.3 Evidence for a dark component on galaxy scales

Studies of galaxy dynamics also played a key role in the discovery of dark matter (Frenk & White, 2012). Some of the earliest pieces of evidence for the existence of dark matter, based on the velocities of galaxies in the Coma Cluster⁵ (Zwicky, 1933), or the rotation curve of the Andromeda “nebula” (Babcock, 1939) did not receive much attention till the 1970s. The change

³Draco is one of the classic “satellite” galaxies of our Milky Way. See section 1.3.1 and 1.3.3 for the explanations of satellite galaxies and the Milky Way.

⁴Classical dwarfs usually refer to dwarf galaxies discovered early on within the Local Group (LG). See section 1.3.3 for the explanation of the LG.

⁵see section 1.3 for the explanation of galaxy cluster.

was brought through several influential papers that argued that a massive spherical halo is required to stabilize a spiral galaxy disk against instabilities (Ostriker & Peebles, 1973; Ostriker, Peebles & Yahil, 1974; Einasto et al., 1974), as well as several improved rotational curve measurements of edge-on spiral galaxies that suggest galaxies contain six times more mass than that being observed (Rubin & Ford, 1970; Rubin, Ford & Thonnard, 1980). Several years after, White & Rees (1978) proposed that a hierarchically merging population of dark matter halos (more discussion below) could provide the basis for galaxy formation, gas cooling and condensing within the gravitational potential well set by the dark matter halos.

1.2 Dark Matter and Structure Formation

There is now overwhelming evidence from many different tests that the majority of the matter in the Universe is **Dark Matter (DM)**, a mysterious substance that does not interact with light and thus cannot be detected through the whole **electromagnetic spectrum**. Although the exact particle nature of dark matter remains unknown, constraints on possible candidates have gradually grown tighter over the past few decades (see Frenk & White (2012) for a review of the origins of the **Λ CDM** model and recent progress on DM).

Observations of the **Cosmic Microwave Background (CMB)**, Planck Collaboration et al., 2016, 2018) indicate that dark matter comprises approximately 85% of the total matter density in the Universe, while dark and **baryon** matter together comprise roughly 31% of the energy density of the Universe. The remaining 69% of the energy density is the even more mysterious “Dark Energy” (DE). We also know dark matter needs to be “cold”⁶ to form the large scale structures we see in the current Universe (Blumenthal et al., 1984).

The Universe is both homogeneous and isotropic at very large scales, but it also has structures like the Solar System, **galaxies** and **galaxy clusters** at smaller scales. On the largest scales, galaxies and clusters are not distributed randomly but form web-like structures. One of the classical examples is the Sloan⁷ Great Wall, where hundreds of thousands of galaxies that are roughly one billion light-years away form a giant wall-like structure across the sky.

All these structures originate from the tiny density fluctuations in the very early Universe and grow through **gravitational instabilities**. As dark matter corresponds to the major part of the mass density in the Universe, it also dominates the structure formation process and sets the environment for the baryon mass to form visible structures like galaxies and galaxy clusters.

⁶“Cold” means that the motions of dark matter particles are non-relativistic in the early Universe; This differentiates from “hot” dark matter, which consists of particles moving with ultra-relativistic velocities.

⁷The Sloan Digital Sky Survey (or **SDSS**) is a multi-spectral imaging and spectroscopic **redshift** survey using a 2.5m wide-angle optical telescope at Apache Point Observatory in New Mexico, United States.

1.2.1 Linear and nonlinear structure formation

The early Universe was much more homogeneous than today, with only very small variations in density. The relative amplitude of these density fluctuations can be defined in dimensionless terms as $\delta \equiv \Delta\rho/\bar{\rho}$, where $\bar{\rho}$ is the mean mass density and $\Delta\rho$ is the excess density. These density fluctuations are seeded by [primordial fluctuations](#) in the early Universe, thought to be created by quantum fluctuations in the field driving inflation. From [CMB](#) observations we know that the primordial density fluctuations can be described by a power law in [wavenumber](#) k : $P(k) \propto k^n$ ([Mo, van den Bosch & White, 2010](#)), where $n \sim 1$ ⁸. The primordial fluctuations grew through gravitational instability, eventually forming all the visible structures in the Universe.

Initially, fluctuations grew slowly as long as their relative amplitude was small, but later the growth rate accelerated as their amplitude approached $\delta \sim 1$. At the early stage, the fluctuation growth over time can be described by linear perturbation theory, which predicts growth that is linear with the expansion of the Universe ($\delta \sim a$, where a is the scale factor of the Universe). This early-stage growth is called the “linear regime”, or linear evolution, of structure formation. The physics of structure formation within the linear regime is particularly simple, compared to later stages. Both the density map (matter distribution) and the velocity map (velocities of the dark matter and baryon particles) of the Universe during this epoch can be calculated using linear perturbation theory.

The linear growth model derived from perturbation theory starts to break down when the amplitude of the density fluctuation becomes comparable to unity, making a transition to a quasi-linear or non-linear regime. Individual particles start to orbit or oscillate around local potential centres, instead of accelerating continuously towards these centers. The physics of structure formation becomes increasingly complicated as fluctuations grow at this epoch. Analytic treatments are only available under certain restrictive assumptions, such as very simplified geometries. In general, non-linear structure formation is studied instead through N -body simulations.

1.2.2 Cosmological N -body simulations

An N -body simulation is a numerical simulation of the dynamics of N particles, usually under the influence of the gravitational forces of the system. In astronomy, the particles are typically used to represent the matter distribution of the Universe (cosmological simulations) or the matter distribution within smaller structures such as galaxy clusters or galaxies. As mentioned above, it is difficult to study structure formation analytically in dense regions that are beyond the linear growth

⁸Recent measurements show $n = 0.9665 \pm 0.0038$ ([Planck Collaboration et al., 2018](#))

regime. Instead, the cosmological *N*-body simulation has provided a key tool to understand the dynamics and evolution of the structures in the Universe.

Early efforts to use *N*-body simulations for astrophysics date back to the 1940s (e.g. [Holmberg, 1941](#)), where analog tools were used instead of modern computational tools (74 light bulbs were used to represent two “galaxies” and simulate their gravity). The first cosmological *N*-body simulation was conducted by [Press & Schechter \(1974\)](#), with only 1000 particles at the time. The particle number and corresponding resolution have improved dramatically over the last several decades, along with access to more powerful computational hardware and improvements in algorithms. For instance, the Millennium Simulation ([Springel et al., 2005](#)) used over 10 billion ($2,160^3$) dark matter particles, with each particle representing 10^9 solar masses of dark matter, while several subsequent cosmological simulations have reached 10 times this number of particles or more (e.g. [Angulo et al., 2012](#); [Heitmann et al., 2015](#)).

To simulate individual structures in a cosmological context, a technique called the “zoom-in” simulation was introduced to better reveal part of the simulation with significantly higher resolution. In a zoom-in simulation, several regions of interest, such as certain halos, are picked from the initial, low-resolution run and then the simulation is re-run with higher resolution within the region of interest, while the rest of the volume is modelled with the same resolution as in the initial run. The higher resolution in the region of interest helps to resolve finer structure, while lower resolution in the surrounding regions provides a robust model of the gravitational environment. Several parts of the Millennium run were reproduced with much higher resolution, in the “Aquarius” project ([Springel et al., 2008](#)), for instance, with each particle in the high-resolution volume representing $(0.64\text{--}1.4)\times 10^4$ solar mass of dark matter.

Modern *N*-body simulations have shown their strong capabilities of investigating the structural evolution of the Universe over cosmological spatial and time scales. The results of these simulations indicate that the current Universe is largely filled with “voids”, whose matter density is significantly lower than the mean density of the Universe. Most dark matter condenses and forms large, filamentary structures. The match between the simulation predictions and the observation evidence of these large-scale structures is considered one of the greatest successes of the CDM model. Among these filamentary structures, we can also find dark matter “halos”, where dark matter is highly collapsed, forming highly dense, roughly spherical structures.

1.2.3 Dark Matter Halos

The concept of a dark matter “halo” dates back to the 1970s, when several papers pointed out a massive spherical component was required to stabilize a spiral galaxy disk ([Ostriker & Peebles, 1973](#); [Ostriker, Peebles & Yahil, 1974](#); [Einasto et al., 1974](#)). At the time, it was widely considered

that a very faint “stellar halo” might be responsible for such a mass background. Now we know that it is dark matter halos, instead of stellar ones, that provide the gravitational potential wells that are needed for disk galaxies to form and retain their structure. Hereafter in this thesis, the term “halo” will always refer to a dark matter halo.

The dark matter halo is one of the basic non-linear structures predicted by cosmological N -body simulations. During the process of structure formation, self-gravity continually works against the cosmic expansion. Soon after the relative density enhancement in overdense regions reaches the non-linear regime, the local expansion will stop and start to turn around. The overdense dark matter cloud will then start to contract/collapse. During the collapse, the potential energy of the system will be converted into kinetic energy, until eventually the two energy components reach equilibrium. This process is also called the virialization of the system ($\langle P \rangle + 2\langle K \rangle = 0$ for an isolated self-gravitating system in equilibrium, where $\langle P \rangle$ refers to the averaged potential energy and $\langle K \rangle$ refers to the averaged kinetic energy). The net collapse will stop, as systems reach equilibrium. These stable systems are dark matter halos.

The shapes of dark matter halos are not necessarily symmetric. Many studies have demonstrated that dark matter halos are triaxial, and may even have different shapes in their inner and outer regions, showing more prolate shapes towards the centre and more oblate shapes in the outskirts (Frenk et al., 1988; Cole & Lacey, 1996; Jing & Suto, 2002; Thomas et al., 1998; Bailin & Steinmetz, 2005; Bett et al., 2007; Hayashi, Navarro & Springel, 2007; Kuhlen, Diemand & Madau, 2007; Vera-Ciro et al., 2011). Nevertheless, halos are usually analyzed in spherically averaged terms. A characteristic radius – the “virial radius” (r_{vir}) is commonly used to quantify the size of a halo. Within the virial radius, the matter can be considered virialized. Similarly, the “virial mass” (M_{vir} , the mass within the virial radius) is used to quantify the mass of dark matter halos.

Once higher-resolution simulations started to reveal the internal structure of the dark matter halos in 1990s, Navarro, Frenk & White (1996) found that a single, universal density profile (the ‘NFW’ profile hereafter, specifying density as function of radius r) provided a good description for all dark matter halos:

$$\rho(r) = \frac{\rho_0}{\frac{r}{r_s} \left(1 + \frac{r}{r_s}\right)^2}. \quad (1.1)$$

The scale radius (r_s) and the density (ρ_0) are two free parameters that vary from halo to halo. Subsequently, a slightly different model profile, the three-parameter Einasto profile (Navarro et al., 2004), was shown to be a more precise description to the halo density profile. However, the NFW profile works almost as well as the Einasto profile for most cases (especially for low-mass halos), and has the advantage of fewer free parameters. Thus, we will adopt the NFW profile as the default for the work in this thesis.

Halos are formed in a hierarchical manner, through mergers on progressively larger mass scales. As mentioned before, the density fluctuations in the early Universe have a power spectrum $P(k) \propto k^n$ with $n \approx 1$, which means the amplitude of density fluctuation is larger at smaller scales. During linear growth, fluctuations at different scales grow independently and linearly with the scale factor of the Universe. As a result, on average small structures are the first to reach non-linear growth regime and collapse to form small halos. These small halos subsequently grow into bigger halos by the accretion of the surrounding diffuse dark matter gas, and by merging with other pre-existing halos.

In the late 1990s, as the resolution of N -body simulations continued to improve, the substructure within halos was revealed (Ghigna et al., 1998; Klypin et al., 1999a). It became clear that the dense cores of the pre-existing small halos can survive the merging process, forming distinct, self-bound substructures or “subhalos” within the larger system. On large (galaxy cluster) scales, the prediction of halo substructure seemed a reasonable match to observations. However, the (very similar) predictions for MW-sized halos seemed much less consistent with observations of Local Group dwarfs, as discussed in detail in the next section.

Early cosmological simulations included only DM-like particles, and any non-gravitational interactions were neglected. While these simulations were successful in explaining large-scale structure, non-gravitational effects become more important at small scales and can no longer be ignored when attempting to explain the inconsistencies between observations and the substructure predictions mentioned above. Full hydrodynamical simulations including baryonic physics, despite being computationally expensive and much more complex and uncertain, are one of the effective directions of exploring the small-scale structure formation. Hydrodynamical simulations have made important progress in the past decades and have succeed in making predictions that generally match basic observations (Springel et al., 2018). Prominent examples include EAGLE⁹ (Schaye et al., 2015; Crain et al., 2015), FIRE¹⁰ (Hopkins et al., 2014), Illustris¹¹ (Vogelsberger et al., 2014; Genel et al., 2014), IllustrisTNG¹² (Springel et al., 2018; Marinacci et al., 2018), and Magneticum¹³ (Dolag, Komatsu & Sunyaev, 2016).

1.2.4 Galaxy formation

DM plays an important role during the process of galaxy formation. In the picture of the standard Λ CDM paradigm, structures such as DM halos form first and provide gravitational potentials into

⁹Evolution and Assembly of Galaxies and their Environments, <http://icc.dur.ac.uk/Eagle/index.php>

¹⁰Feedback In Realistic Environment project, <https://fire.northwestern.edu/>

¹¹<http://www.illustris-project.org/>

¹²<http://www.tng-project.org/>

¹³<http://www.magneticum.org/>

which nearby baryonic matter can fall (White & Rees, 1978; Blumenthal et al., 1984).

During this process, baryonic gas heats up, leading to increased thermal radiation, which helps the gas to cool and condense further into the center of the DM halos. This radiation is normally fairly isotropic, and thus has very low efficiency for carrying away angular momentum. The conservation of angular momentum during gas cooling should in principle lead to the formation of a galactic disk (though the exact details at very high redshift are unclear). The disk itself is not a stable structure, and the gas within it will further cool and condense. Eventually, gas condensation leads to the formation of molecular clouds, which in turn lead to star formation, and thus the formation of a stellar disk or the spiral galaxy.

Galaxies will continue to grow through the accretion of the surrounding gas and smaller galaxies, or even through mergers with other galaxies. Note that the stellar disks in spiral galaxies can be easily thickened or disrupted during the accretion of satellite galaxies or major merger events (Velazquez & White, 1999). It is generally accepted that mergers are one of the driving forces in the formation of elliptical galaxies.

Along with the initial gas condensation, there are also feedback mechanisms that regulate the star-forming activity during the process of galaxy formation:

- Supernova (SN) feedback: at the end of stellar evolution, some stars can explode as supernovae. These energetic events can significantly heat and redistribute nearby gas, and release a large amount of ‘metals’, or heavy elements¹⁴
- Active galactic nucleus (AGN) feedback: in the most massive galaxies, SN feedback becomes less effective. An AGN corresponds to a region at the center of a galaxy where there is a supermassive black hole (SMBH) actively feeding on its surrounding gas and also releasing powerful radiation. Radiation or winds from AGN can heat up or displace surrounding gas, also limiting the star formation rate.

Both SN and AGN feedback can regulate star formation, as they release energy that can heat up the gas and prevent it from cooling too quickly. SN feedback is more important for small galaxies where the gravitational potential wells are generally shallower. AGN feedback requires the existence of an SMBH inside the galaxy and is usually not relevant for small galaxies that lack large SMBH, but it serves as the most important feedback mechanism for massive galaxies. Several other environmental mechanisms can also limit star-forming in galaxies. Together with the internal feedback mechanisms mentioned above, these factors can sometimes cause star formation to shut down quickly, a process referred to as “quenching”. I will discuss the quenching process in more detail in section 1.3.4.

¹⁴In astronomy, elements that are heavier than hydrogen and helium are conventionally referred to as ‘metals’.

1.2.5 The stellar-to-halo mass relation (SHMR)

Understanding the co-evolution of galaxies and halos, and the physical processes involved, is one of the primary goals of studies of galaxy formation and evolution (Hudson et al., 2015). However, it remains difficult to directly associate our theoretical knowledge of dark matter halos with observations of individual galaxies. We can observe snapshots of the distribution of galaxy properties such as luminosity and stellar mass at different redshifts, whereas similar predictions of halo mass and abundance can be obtained from numerical simulations; the question is how to put the two together. Note that constraining the relationship between the total mass of a halo and the stellar mass within it (the “stellar-to-halo mass” relation – SHMR), is a complicated topic, which is hard to cover in the scope of this short introduction. Wechsler & Tinker (2018) gives a detailed review of the SHMR and other relations between dark matter halo and galaxies. Leauthaud et al. (2012, - in particular their Figure 10) also provides a good comparison for different techniques for constraining the SHMR.

One simple way to connect the two is by a technique called “abundance matching (AM)” (e.g. Marinoni & Hudson, 2002; Vale & Ostriker, 2004). In this approach, galaxies and halos are rank-ordered in each set and matched one-to-one from largest to smallest in mass. With this technique, Marinoni & Hudson (2002) found out that there is a “characteristic halo mass” $M_h^{\text{peak}} \sim 10^{12.5} M_\odot$ in the local Universe, at which the stellar-mass-to-halo-mass ratio is maximized/highest. A natural interpretation is that M_h^{peak} corresponds to the halo mass at which star formation, integrated over the entire assembly history of the galaxy, has been the most efficient (Silk, 2013; Hudson et al., 2015).

The AM method of Marinoni & Hudson (2002) was only considering all galaxies in one halo. It has since then been extended to consider both central galaxies and satellite galaxies separately¹⁵, in a variant known as subhalo abundance matching or SHAM (Yang, Mo & van den Bosch, 2003; Vale & Ostriker, 2004; Conroy, Wechsler & Kravtsov, 2006; Moster et al., 2010; Behroozi, Conroy & Wechsler, 2010; Hearin et al., 2013). Another popular approach to associating a galaxy population to its corresponding halos is called the Halo Occupation Distribution (HOD) model. This assigns a probability of each halo hosting one or more galaxies from a given sample, as a function of its halo mass. This can populate halos with galaxies in such a way as to reproduce the observed galaxy clustering pattern¹⁶ (Jing, Mo & Börner, 1998; Peacock & Smith, 2000; Cooray & Sheth, 2002; Seljak et al., 2005; Coupon et al., 2012).

Both AM and HOD are powerful tools, but they are also model-dependent methods. They both make critical assumptions about the statistical link between galaxies and their host halos. A more

¹⁵See section 1.3.1 for explanations.

¹⁶See section 1.3 for explanations.

direct way to connect galaxies to their halos is to probe dark matter halos gravitationally. There are several ways to do so. Traditionally, dynamical methods are commonly used to estimate the total mass surrounding each galaxy. For instance, the well-known Tully-Fisher relation (Tully, de Marseille & Fisher, 1975) shows an empirical correlation between the mass or intrinsic luminosity of spiral galaxies and their disk rotational speed or emission line width. Similar estimates can be pushed out further into the halo using satellite dynamics (van den Bosch et al., 2004; More et al., 2011). Note that all dynamical methods require some assumptions regarding the dynamical equilibrium of the galaxy and its surrounding halo.

Finally, the most powerful way to probe the mass distribution to large radii is using gravitational lensing. Light from distant sources, usually galaxies, will be bent by any intervening gravitational potential well before it reaches the observer. As a result, the images of the distant sources will appear distorted. ‘Strong’ lensing systems refer to cases where the lensing effects are so strong that the source images are split into multiple pieces. Strong lensing systems are rare, however, and thus less useful relative to weak lensing systems, where the source images retain only single, slightly distorted components. The signal from an individual weak lensing system is very small. Only by stacking thousands of such systems do they become a powerful tool for measuring the average mass distribution around galaxies (Brainerd, Blandford & Smail, 1996).

1.3 Galaxy Groups and Clusters

1.3.1 Satellite and Central Galaxies

During the hierarchical assembly of dark matter halos, small progenitor halos often survive accretion onto a larger system, creating a population of bound remnants, or “subhalos”. The galaxy formed by gas cooling within the central parent halo is called the “central” galaxy. Similarly, if the gas in the progenitor halos of the subhalos managed to cool and form galaxies before these were accreted onto the main system, it creates a population of “satellite” galaxies around the central galaxy. Satellite galaxies comprise a significant fraction of the entire galaxy population. Many studies have measured satellite abundance, using various methods, and generally find that about 30% of galaxies with stellar masses around $10^9 M_{\odot}$ are satellites. This fraction also decreases with increasing stellar mass (Mandelbaum et al., 2006; Tinker et al., 2007; van den Bosch et al., 2008; Yang, Mo & van den Bosch, 2008; Wetzell et al., 2013).

1.3.2 Groups and clusters

The observed structures that contain multiple galaxies (central and satellite galaxies) are referred to as galaxy groups or clusters. Structures that contain tens to hundreds of dwarf or giant galaxies are usually referred to as galaxy groups; while structures that contain hundreds or thousands of galaxies are referred to as galaxy clusters. (Multiple galaxy groups and clusters together can also constitute an even larger structure known as a supercluster, but this is not a single, relaxed, virialized object.)

Typically, galaxy groups only contain a few (optically) luminous galaxies and produce low levels of X-ray emission. In contrast, galaxy clusters contain large numbers of bright galaxies, and the space between the galaxies is filled with a hot intracluster medium that can be very bright in X-ray (e.g. [Giodini et al., 2009](#); [Dai et al., 2010](#); [Anderson et al., 2015](#)). Another interesting difference between groups and clusters lies in position on the SHMR. For low-mass galaxy groups, the stellar-to-halo mass ratio increases with mass; while for high mass groups and clusters, this ratio decreases with halo mass. This difference hints at different mechanisms that regulate the star formation in groups and clusters ([Bullock & Boylan-Kolchin, 2017](#)).

1.3.3 The Milky Way and The Local Group

The Milky Way and its immediate surroundings hold a special place in astrophysics, not only because they are the home of our planet Earth, but also they are the closest large structures to us. Thus, they will always be the places where the most detailed observational information is available across many branches of the field.

The Local Group (**LG**) is the group of galaxies consisting of the MW and nearby galaxies, including the Andromeda galaxy (**M31**) and the satellite galaxies of the MW and M31. M31 is the nearest bright spiral galaxy to the MW and is about 780 kpc from us ([Ribas et al., 2005](#); [McConnachie, 2012](#)). It has a similar (though probably slightly higher) luminosity and mass as the MW ([Karachentsev & Kashibadze, 2006](#)). The host halos of the MW and the M31 are both estimated to have virial radii of ~ 300 kpc ([Klypin, Zhao & Somerville, 2002](#)), so their halos do not quite overlap with each other.

Being the nearest galaxy group, the LG is the best-studied group environment. It has a large population of dwarf galaxies, with over 50 confirmed satellite members and more being discovered in recent years ([McConnachie, 2012](#); [Bechtol et al., 2015](#); [Drlica-Wagner et al., 2015, 2020](#)).

[McConnachie \(2012\)](#) reviewed the content of the LG and provided an updated online cata-

logue¹⁷ of the 100 nearest (within 3 Megaparsecs (Mpc) from the Sun) dwarf galaxies that were known by 2012, of which 76 galaxies are very likely to be members of the LG. Among these LG member candidates, there are 27 galaxies that are within 300 kpc (the suggested virial radius of the MW halo).

1.3.4 Quenching and environmental effects

Galaxies in the local Universe can be broadly classified into two populations by their relative rates of star formation: blue star-forming galaxies that are actively adding to their stellar mass, and red quiescent (or passive) galaxies that lack or have little active star-formation (Blanton et al., 2003; Kauffmann et al., 2003b, 2004; Balogh et al., 2004; Baldry et al., 2004; Brinchmann et al., 2004; Taylor et al., 2015). Based on our current understanding of galaxy evolution, galaxies are initially star-forming, and grow through continued star formation and mergers, but may eventually evolve into red, quiescent systems (Bell et al., 2004; Faber et al., 2007; Martin et al., 2007). The distribution of galaxy properties is clearly bimodal, for example in the colour versus stellar mass plane (Baldry et al., 2004; Taylor et al., 2015) or the star-formation rate (or specific star-formation rate) versus stellar mass plane (Balogh et al., 2004; Moustakas et al., 2013; Davies et al., 2016), with intermediate “green valley” objects being much rarer (Martin et al., 2007; Wyder et al., 2007; Salim, 2014). This suggests that the transition process from blue to red (i.e. quenching) is likely to be fast, across a broad range of stellar mass (Wetzel, Tinker & Conroy, 2012; Schawinski et al., 2014; Bremer et al., 2018). The cause of star-formation shut down is still a matter of vigorous debate and remains one of the most important goals in understanding galaxy evolution.

The mechanisms of galaxy quenching suggested by current observational evidence can be roughly divided into two main modes. The first is secular quenching (or mass quenching), which correlates with the internal properties of a galaxy, and can occur in a galaxy regardless of its environment (Kauffmann et al., 2003a; Driver et al., 2006; Wake, van Dokkum & Franx, 2012; Lang et al., 2014; Barro et al., 2017; Contini et al., 2020). The second quenching mode is driven by the local environment of a galaxy. The occurrence of quenching is also found to be correlated with the local environment density in galaxy groups/clusters (Peng et al., 2012), as well as the locations respective to the group/cluster-centre (Wolf et al., 2009; Wetzel, Tinker & Conroy, 2012; Woo et al., 2015).

Overdense environments such as galaxy clusters, groups or even close pairs can remove or inhibit the gas supply that sustains on-going star-formation (Patton et al., 2011; Robotham

¹⁷The catalogue of LG dwarfs from McConnachie (2012) is available [here](#); or [an easier-to-read version](#) from the SIMBAD database; or from the [NASA/IPAC Extragalactic Database search result](#), which also contains links to spectres and images of each dwarf.

et al., 2014; Davies et al., 2016), leading to quenching events (e.g. Peng et al., 2010; Schaefer et al., 2017). There are various physical processes that can drive this mode of quenching, including “starvation/strangulation”, where the fresh gas accretion is suppressed (Larson, Tinsley & Caldwell, 1980; Moore et al., 1999b; Peng, Maiolino & Cochrane, 2015; Trussler et al., 2019), tidal and ram-pressure stripping, where gas is actively removed from the satellites respectively by gravitational tidal force and the ram-pressure while moving through intra-cluster medium (Gunn & Gott, 1972; Moore et al., 1999b; Nichols & Bland-Hawthorn, 2011; Poggianti et al., 2017; Brown et al., 2017; Fattahi et al., 2018; Trussler et al., 2019), and “harassment”, where satellites get disturbed by the fly-by galaxy members (Moore et al., 1996). Galaxies with intermediate-to-low mass are more prone to be affected by this environmental quenching (Davies et al., 2019). This is likely due to the fact that it is more difficult for low-mass galaxies to retain or accrete gas while moving through dense environments.

Since environmental quenching should only apply to satellites, it implies that centrals and satellites will have different quenched fractions. At the low mass end, studies have shown that satellite galaxies are on average redder and more concentrated than central galaxies with the same stellar mass (van den Bosch et al., 2008), indicating quenching processes that are specific to satellites. Similarly, Geha et al. (2012) showed that while the quiescent fractions of small galaxies with $M_* = 10^8 - 10^9 M_\odot$ are consistent with zero for dwarfs that are sufficiently isolated (more than 1 Mpc away from their nearest MW-mass or larger hosts), this fraction rises sharply at smaller distances from massive systems. On the other hand, for halo masses larger than $10^{12} M_\odot$, central galaxies have a higher red fraction (Yang, Mo & van den Bosch, 2008), so in that case mass quenching mechanisms must dominate.

1.4 Previous Work and Motivation

As mentioned previously, the Λ CDM cosmological model has been a resounding success in matching observations of the large-scale structure of the Universe (Frenk & White, 2012). However, challenges have arisen on small scales, where the correspondence between theory and observed structure seems more complicated. In this section, I will review these challenges and some of the proposed solutions to them. Then I will briefly discuss previous efforts to quantify dwarf galaxy abundance in particular, as well as the work that still needs to be done in this area.

1.4.1 Small-scale challenges in Λ CDM

Cosmological N -body simulations with sufficiently high resolution have shown that MW-sized halos should contain thousands of subhalos massive enough ($> 10^7 M_\odot$) to support molecular

cooling (Moore et al., 1999b; Klypin et al., 1999a). We would expect such objects to form stars, and thus produce thousands of visible satellite galaxies. However, only about 50 satellite galaxies with stellar masses larger than $300 M_{\odot}$ are observed to orbit the MW within its estimated virial radius (Drlica-Wagner et al., 2015). Although the observational inventory is likely incomplete, and future surveys may bring up the total satellite number into the hundreds (Hargis, Willman & Peter, 2014), the large discrepancy between the theoretical prediction and the observations is unlikely to be resolved by improved observational data alone. The discrepancy in abundance between the predicted dark matter substructure and the observed population of local satellite galaxies is often referred to as the “Missing Satellite Problem” (MSP), or missing dwarf problem.

A complementary way to look at this discrepancy is to compare the mass function of dark matter halos predicted by CDM simulations and observed galaxies. It is a characteristic prediction of the Λ CDM cosmology that the dark matter halo mass function has a steep slope at the low-mass end. That slope is much steeper than the faint-end slope of the galaxy stellar mass function, which implies that low-mass CDM halos are significantly more abundant than faint dwarf galaxies (Kauffmann, White & Guiderdoni, 1993).

The missing satellite problem is usually resolved by assuming that dwarfs form preferentially in relatively massive halos because cosmic re-ionization and the energetic feedback from star formation remove baryons from the shallow gravitational potentials of low-mass systems (Bullock, Kravtsov & Weinberg, 2000; Benson et al., 2002; Somerville, 2002). However, Boylan-Kolchin, Bullock & Kaplinghat (2011) put this explanation to the test by comparing the central mass of MW satellites to the predictions from the Aquarius (Springel et al., 2008) and Via Lactea II (Diemand et al., 2008) simulations. They pointed out that the most massive subhalos from the simulations were systematically too centrally dense to host the bright satellites in the MW. While there were subhalos that were comparable in terms of central mass to the MW satellites, those were never among the top 10 most massive subhalos in the MW-size host halos. This raises the question, why would luminous satellites preferentially avoid the most massive subhalos, but occupy lower-mass ones? This question raised by Boylan-Kolchin, Bullock & Kaplinghat (2011) is another small-scale challenge for the Λ CDM model, and is known as the “Too-Big-to-Fail” (TBTf) problem. For a more comprehensive review of these small-scale challenges (SSCs), see the review by Bullock & Boylan-Kolchin (2017).

1.4.2 Solutions to the SSCs

From the theory side, there has been much work exploring possible solutions that can resolve the SSCs. These solutions mainly fall into two categories: fundamental modifications to the Λ CDM model, or solutions related to baryonic processes, within the Λ CDM model.

Alternatives to Λ CDM

The Λ CDM model can be modified in several possible ways, without disrupting its agreement with large-scale observations. In particular, while the usual version of the model assumes a power-law spectrum of initial fluctuations, subsequent processes will suppress fluctuations at high wavenumbers. The high- k is normally set by the free-streaming damping¹⁸, which is tightly related to the particle nature of DM. A larger free-streaming cut-off length means fewer seeds for structure formation at small scales and thus leads to a lower abundance of DM substructure. This in turn leads to a reduction in the expected number of satellites inside a MW-sized halo.

The effective free-streaming cut-off length scales inversely with the particle mass. The best-known alternative theories with smaller particle masses include hot dark matter (HDM) and warm dark matter (WDM). For HDM, the cut-off length scale set by the free-streaming damping is of the order of 1 pc in comoving scale (Bullock & Boylan-Kolchin, 2017), which is much smaller than the typical substructure scales, and therefore has a negligible effect on structure formation. “Hot” dark matter particles retain very high (relativistic) velocities until relatively late times, and their diffusion can damp out density perturbations even on super-cluster scales. This obviously conflicts with observations of large-scale structure and is ruled out as a result. On the other hand “warm” dark matter particles, whose masses are at the order of 1 keV, have smaller initial velocities and become non-relativistic earlier than hot dark matter particles. The corresponding damping will suppress perturbations on galactic or smaller scales (Bode, Ostriker & Turok, 2001). WDM is not ruled out by observation but the current observed number of satellites can set a lower limit on the WDM particle mass. (E.g. Lovell et al. (2014) claimed a constraint of $m > 2.3$ keV with 95% confidence. One can also refer to Kennedy et al. (2014) for a good review of constraining DM particle mass by comparing luminosity function predictions to MW satellite data.)

Another possible modification to the Λ CDM model is to assume DM physics that modifies non-linear evolution. As mentioned before, structures grow linearly in the early stage and then make a transition to the non-linear growth regime once the density starts to pass a critical value. One way to modify non-linear evolution is to assume non-negligible interactions between DM particles. The simplest form is called self-interacting dark matter (SIDM, also sometimes called collisional dark matter), which has a characteristic energy-exchange interaction cross-section σ (Spergel & Steinhardt, 2000). This gives a mean free path λ for DM particle motion $\lambda = (n\sigma)^{-1}$, where n is the local number density of DM particles. At large scales where the DM particle density is low, the self-interactions between DM particles are still negligible and the picture is consistent with CDM. It becomes important, however, in high-density regions, where the mean

¹⁸Free-streaming damping refers to the damping of density perturbation of collisionless particles due to their random velocities (Mo, van den Bosch & White, 2010).

free path is short and particles experience a high frequency of scattering events, making SIDM act like a fluid. The halos produced with SIDM usually have a relatively large, constant-density core in the center, compared to the sharp/cuspy density profile of conventional CDM halos.

Effects from Baryonic Physics

A less radical alternative to resolving the SSCs is to suppose that baryonic processes have a large effect on surviving dark matter structures, or on star formation within subhalos. Hydrodynamic simulations have shown that under the right conditions, stellar feedback can redistribute dark matter and baryonic matter in the central region of halos, forming a flat, core-like central density profile (Madau, Shen & Governato, 2014; Oñorbe et al., 2015; Read, Agertz & Collins, 2016). This effectively lowers the central density of sub-halos, possibly resolving the TBTF problem. However, we note that the effect of stellar feedback tends to be less significant for small dwarf galaxies, whose stellar-halo mass ratio drops as the mass decreases. Results from multiple simulations (Madau, Shen & Governato, 2014; Oñorbe et al., 2015; Read, Agertz & Collins, 2016; Fitts et al., 2017) suggests a threshold of $M_* \sim 10^6 M_\odot$ or $M_{\text{vir}} \sim 10^{10} M_\odot$, below which the effect of stellar feedback shaping the halo density profile should become negligible.

In addition to internal stellar feedback, several forms of external feedback from the host galaxy, such as tidal stripping, disk shocking and ram pressure stripping, can also reduce the central mass of satellites. Those external feedback processes are shown to be important in simulations (Wetzel et al., 2016) and may help to explain the TBTF problem as well. Disk shocking in particular may be more effective at disrupting substructure than usually imagined, reducing the number of surviving post-merger subhalos and partially explaining the MSP. If the external influences from the host galaxies are the key to explaining the small-scale challenges, there should be a significant difference between field dwarfs and satellite galaxies, and the properties of satellites may also have a systematic dependence on the relative position inside the host halo and their infall history. These expected differences need to be tested with observations.

Besides the theoretical explanations, there are also some observation complications that need to be considered:

- The completeness of LG satellites: The current catalogue of MW and LG satellites may be far from complete (Hargis, Willman & Peter, 2014). The galactic disk also adds an extra difficulty in identifying the fainter MW satellites.
- The mass of the MW: Wang et al. (2012) pointed out that the severity of the TBTF problem is very sensitive to the MW halo mass. The “problem” could be an indication of a smaller

MW halo mass than is usually considered. They claimed that for $M_{\text{halo}} = 10^{12} M_{\odot}$ halos in simulations, 40% of them show no TBTF problem.

- The conversion between HI line width and mass: The measured masses and density profiles of many dwarf galaxies rely on rotation curves derived from neutral hydrogen (HI) emission lines. Several factors complicate the interpretation of these results, however. There may be some contribution of non-rotational support from turbulent pressure, radiation pressure, etc., within these systems, and the neutral gas may not extend out to the regions where the dark matter halo potential would cause peak rotation speed. Finally, the uncertainty of inclination angles can also lead to systematics in the mass estimates (Bullock & Boylan-Kolchin, 2017). Please see Papastergis & Ponomareva (2017) for a full discussion of the issues related to HI line conversion.

1.4.3 The need to go beyond the Local Group

On one hand, the LG is currently the best-studied galaxy group, and it will likely continue to be so in the future. The nearby Universe will always be the easiest place to detect the faint end of the galaxy luminosity function and study resolved stellar populations (e.g. Weisz et al., 2011; McConnachie, 2012; Garrison-Kimmel et al., 2014). We can measure kinematics, stellar components and metallicities of the LG dwarf galaxies at a level of detail that would be very difficult to obtain for galaxies that are further away (Tolstoy, Hill & Tosi, 2009).

On the other hand, it remains an open question whether what we learned from the LG is representative of the rest of the Universe (van den Bergh, 2000; Weisz et al., 2011). Many recent studies suggest that the LG has several properties that may be atypical compared to the rest of the local Universe. For instance, observational studies have shown that it is relatively rare for galaxies like the MW to have two bright, star-forming satellites like the LCM and SMC (Guo et al., 2011; Liu et al., 2011; Strigari & Wechsler, 2012; Robotham et al., 2012; Speller & Taylor, 2014) and numerical simulations have reached similar conclusions (Boylan-Kolchin et al., 2010; Busha et al., 2011; Kang, Wang & Luo, 2016; Zhang, Luo & Kang, 2019; Evans et al., 2020)

We also know that the properties of satellite galaxies are strongly influenced by their environment. Various environmental effects could destroy or transform the satellites during and after their merging process, such as tidal heating or stripping (Gunn & Gott, 1972; Moore et al., 1999b; Peng, Maiolino & Cochrane, 2015; Trussler et al., 2019), the encounter with other galaxies (Moore et al., 1996), and external feedback (Patton et al., 2011; Robotham et al., 2012; Davies et al., 2016). These environmental effects provide important insights that can help us to get a better understanding of the complex picture of galaxy formation. The LG can only provide one

example of a group environment; we need a much greater number of examples to explore the effects of the environment on small-scale galaxy formation.

1.4.4 Prior Work

Given the motivation to study satellite populations beyond the LG, one of the basic observational challenges is to distinguish satellites from the foreground and background systems, by obtaining some form of distance information for them. There have been a few main approaches in the recent literature: (1) use existing spectroscopy to identify satellites around the nearest bright systems (e.g. [Merchán & Zandivarez, 2005](#); [Yang et al., 2007](#)); (2) conduct a dedicated spectroscopic survey of faint targets around a relatively small number of selected systems (e.g. the SAGA survey¹⁹: [Geha et al., 2017](#); [Mao et al., 2021](#)); (3) use photometric distance estimates such as the tip of the red giant branch (TRGB) (e.g. [Carlin et al., 2016](#); [Danieli et al., 2017](#); [Cohen et al., 2018](#); [Carlsten et al., 2019](#)), or surface brightness fluctuations (e.g. [van Dokkum et al., 2018](#); [Carlsten et al., 2019](#)); (4) measure statistical abundance with clustering-based methods (e.g. [Liu et al., 2011](#); [Guo et al., 2011, 2012](#); [Strigari & Wechsler, 2012](#); [Wang & White, 2012](#); [Wang et al., 2014](#); [Sales et al., 2013](#); [Speller & Taylor, 2014](#)).

The four approaches have different strengths and weaknesses. Method (1) requires only existing data, and spectroscopic distance estimates are reliable and accurate, but it is restricted to the brightest satellites in the nearest systems, and it may also suffer from incompleteness due to fibre positioning limitations in dense fields (e.g. [Guo, Zehavi & Zheng, 2012](#); [Smith et al., 2019](#)). Method (2) also benefits from accurate spectroscopy and has much higher completeness and depth compared to method (1). However, it is extremely expensive in terms of observing time and has a low efficiency of finding satellites given the very large background populations. The (SAGA) Survey (Stage I, [Geha et al., 2017](#)), for instance, found only 25 new satellites around 8 MW analog hosts, out of the 17,000 spectroscopic redshifts they obtained for satellite candidates. This low efficiency limits the method to small numbers of primary systems. Method (3) is very limited in target selection, requiring primary systems to be very nearby [<20 Mpc], at which distance virial radii subtend large angles on the sky. This selection criterion makes complete coverage difficult. Method (4) cannot confirm individual galaxies as satellites or centrals. However, it has been very successful in making measurements of the average satellite abundance and is the least resource-intensive method of the four.

Among the studies using method (4) mentioned above, [Speller & Taylor \(2014\)](#) proposed a new method of selecting nearby galaxies based on morphological cuts, which we will test and improve in the work in chapter 2. The method can effectively select dwarfs that are likely to be

¹⁹Satellites Around Galactic Analogs, <https://sagasurvey.org/>

nearby, without spectroscopic information or high-quality photometric redshifts. By focusing on the very nearby systems, they were able to reach fainter limits in terms of detection sensibility. However, this morphological selection method also introduced issues such as incompleteness for satellite detection. Besides, like most other studies with method (4), [Speller & Taylor \(2014\)](#) applied an isolation cut while selecting their primaries, which excluded about two-thirds of all potential systems. The isolation cut effectively limits the studies of dwarf galaxies to lower-density environments, which could potentially bias the results. The new method we established in chapters 3 and 4 will try to address this issue.

1.4.5 Our research goals

The main goal of the work included in this thesis is to advance existing techniques, as well as develop new ones, for quantifying dwarf satellite populations outside the LG.

The work consists of two main parts:

1. Testing the effectiveness and optimizing the morphological selection method in [Speller & Taylor \(2014\)](#). In particular, we verified the incompleteness and clustering signal boost introduced by the morphological cuts. We also evaluated the effectiveness of the method at a broader range of distances, exploring the limits where we can apply this method. Finally, we tested different models of the morphological cuts and optimized model parameters at different distances to achieve maximum effectiveness.
2. Developing and testing a new method of quantifying satellite abundance using clustering. This method requires higher-precision photo-zs, making it less general than the method in the first part, but still more cost-effective and applicable to fainter satellite candidates compared to using spectroscopy. This new method also features a hierarchical primary selection, which allows us to find systems in dense environments, and across a very large mass range.

Given the techniques we developed, we also measured various satellite and central galaxy properties such as their overall abundance, luminosity function, stellar mass function, and quiescent fraction, and determined how these properties are related to their environments. We compared the LG to a broader sample of similar groups and discussed whether the LG is representative of the rest of the local Universe. Through this work, we are hoping to improve our understanding of dwarf satellite populations and the effects of the environment on these objects, as well as to open new opportunities for future research.

1.5 Organization of the thesis

The main research chapters of this thesis, Chapter 2, 3 and 4, correspond to two published and one submitted paper; the first two are reproduced without modification, while the last has been modified slightly based on thesis committee and referee comments, as explained in the Statement of Contributions. Chapter 2 was published as [Xi et al. \(2018\)](#), “Quantifying the abundance of faint, low-redshift satellite galaxies in the COSMOS survey”. Chapter 3 was published as [Xi & Taylor \(2021\)](#), “A hierarchical clustering method for quantifying satellite abundance”. Chapter 4 is from the submitted paper: “The hierarchical clustering method: abundance and properties of local satellite populations”. Chapter 5 reviews the main findings of the thesis and discusses the limitations, potential improvements and future applications of this work.

Chapter 2

Quantifying the abundance of faint, low-redshift satellite galaxies in the COSMOS survey

Abstract

Faint dwarf satellite galaxies are important as tracers of small-scale structure, but remain poorly characterized outside the Local Group, due to the difficulty of identifying them consistently at larger distances. We review a recently proposed method for estimating the average satellite population around a given sample of nearby bright galaxies, using a combination of size and magnitude cuts (to select low-redshift dwarf galaxies preferentially) and clustering measurements (to estimate the fraction of true satellites in the cut sample). We test this method using the high-precision photometric redshift catalog of the COSMOS survey, exploring the effect of specific cuts on the clustering signal. The most effective of the size-magnitude cuts considered recover the clustering signal around low-redshift primaries ($z < 0.15$) with about two-thirds of the signal and 80% of the signal-to-noise ratio obtainable using the full COSMOS photometric redshifts. These cuts are also fairly efficient, with more than one third of the selected objects being clustered satellites. We conclude that structural selection represents a useful tool in characterizing dwarf populations to fainter magnitudes and/or over larger areas than are feasible with spectroscopic surveys. In reviewing the low-redshift content of the COSMOS field, we also note the existence of several dozen objects that appear resolved or partially resolved in the HST imaging, and are confirmed to be local (at distances of ~ 250 Mpc or less) by their photometric or spectroscopic redshifts. This underlines the potential for future space-based surveys to reveal local populations of intrinsically faint galaxies through imaging alone.

Key words: dark matter – galaxies: dwarf – galaxies: formation – galaxies: groups: general – galaxies: luminosity function, mass function – Local Group

2.1 Introduction

The Milky Way, M31, and other bright galaxies in the nearby universe are observed to have retinues of faint dwarf satellites. The ‘classical’ dwarfs of the Local Group, those identified decades ago, have magnitudes brighter than $M \sim -6$ in the B or V -band, while the more recently discovered ‘ultra-faints’ can be many magnitudes fainter (see [McConnachie, 2012](#), for a review). Given their high velocity dispersions and implied high mass-to-light ratios, dwarf satellites are inferred to trace the dense substructure seen in simulated dark matter halos. As such, they provide a very important test of models of structure formation. The relationship between dwarf satellites and halo substructure is complex, however, since the simplest models relating the two fail to match the number ([Klypin et al., 1999b](#); [Moore et al., 1999a](#)), spatial distribution ([Kravtsov et al., 2004](#)) and central densities ([Boylan-Kolchin, Bullock & Kaplinghat, 2011](#)) of the known dwarf galaxies of the Local Group. Detailed, careful modelling (e.g. [Brooks & Zolotov, 2014](#); [Sawala et al., 2015](#)) seems to be required to understand the properties of these objects.

Despite ongoing observational efforts (e.g. [Karachentsev, Makarov & Kaisina, 2013](#); [Chiboucas et al., 2013](#); [Sand et al., 2014](#); [Merritt, van Dokkum & Abraham, 2014](#); [Javanmardi et al., 2016](#); [Crnojević et al., 2016](#); [Müller et al., 2017](#); [Greco et al., 2018](#), and references therein), most of our information about faint satellites comes from the Local Group, and models of dwarf galaxy formation typically set out to reproduce its properties. Our view of the Local Group is limited, however, by obscuration and uneven (albeit gradually improving – [Laevens et al. 2015b,a](#); [Bechtol et al. 2015](#); [Drlica-Wagner et al. 2015](#)) sky coverage. Furthermore, studies of bright satellites around Milky Way analogues suggest that our Galaxy may be unusual in some respects; for instance, the presence of two bright, star forming satellites represents a 1 in 250 or rarer occurrence ([Robotham et al., 2012](#)). To reach robust conclusions about typical satellite populations, we really need to expand the inventory of host systems with well-sampled satellite distributions by a factor of 100 or more.

Identifying faint satellites around more distant primaries is, unfortunately, very challenging. Over a reasonably large volume, such objects should be bright enough to be detected in large-area surveys. But without some means of determining distances to faint galaxies, and thus of associating them with nearby bright ones, local satellites are swamped by the much larger number of faint background galaxies. Recent work by the SAGA survey ([Geha et al., 2017](#)) provides a good indication of the challenge; a massive spectroscopic campaign measuring more than 17,000 redshifts found only two dozen new dwarfs down to a magnitude of -12 , within a projected

virial radius around their nearby targets. Going fainter would decrease the efficiency further, at prohibitive cost in observing time.

There are alternatives to spectroscopic distance determinations. Photometric redshifts are one example; in cases where many bands are available, these can be quite effective for determining 3D structure, at least on large scales (Scoville et al., 2007a, 2013). Unfortunately, photometric redshifts of this quality are only available for a few small fields, notably the COSMOS field (Scoville et al., 2007b). A second possibility is to use clustering to estimate the average distance to a faint population of objects, by association with a brighter set of objects of known distance or redshift. Association inferred from proximity on the sky provides distance estimates for a number of the (relatively rare) dwarf galaxies in the ‘Local Volume’ out to 11 Mpc (Karachentsev, Makarov & Kaisina, 2013), while the related statistical technique of ‘clustering redshifts’ (Ménard et al., 2013; Rahman et al., 2015) has been used to determine mean redshifts for populations at greater distances (e.g. Rahman et al., 2016). Even here, without any further sample selection beyond a basic magnitude cut, the clustering signal from faint, nearby systems tends to be weak.

A third alternative for estimating distances (or at least selecting local galaxies preferentially) is structural (size, magnitude, and/or surface-brightness) selection. As we will show below, the intrinsically faint galaxies of the local universe occupy a distinct region of structural parameter space. Cuts in size, magnitude, and/or surface brightness are not enough to uniquely identify them at all redshifts, but can be quite effective for nearby objects. This method has been used implicitly several times, e.g. in Ferrarese et al. (2012) or Karachentsev, Makarov & Kaisina (2013), but without much systematic study. In (Speller & Taylor, 2014, ST14 hereafter), we proposed a specific structural selection criterion, based on size and magnitude, to identify satellites around primaries at distances of 10–40 Mpc. We demonstrated that our structural cuts preferentially select nearby dwarfs by measuring the clustering signal of the cut sample with respect to the primaries. Overall, our cuts increased the signal-to-noise ratio (SNR) of the clustering signal from undetectable levels up to a value ~ 9 , allowing us to measure a number of properties of the satellite population. One major limitation of this method, however, is the incompleteness of the resulting samples (down to a fixed magnitude or luminosity limit), which we estimated to be 50% or more. Furthermore, our selection was tuned to relatively nearby systems. It is unclear how well this selection method extends to fainter magnitudes, and more generally, how it depends on the detailed form of the structural cut.

The goal of the current work is to study and test the structural selection method in more detail. Ideally, we would do this with large ($\gtrsim 10^4$ object), complete samples of faint (21–22 magnitude) galaxies with measured spectroscopic redshifts. Unfortunately, no samples of sufficient depth and areal coverage are currently available. The closest equivalent is the photometric redshift catalog of the COSMOS survey (Mobasher et al., 2007; Ilbert et al., 2009; Laigle et al., 2016). This provides photo- z s with an accuracy of 1% or better down to magnitudes of $i^+ = 23$ (or even

deeper at low redshift), and thus gives a good indication of which faint galaxies are truly local, albeit over a very small field. Since the COSMOS field is so small, we will push the limits of the selection method developed in [ST14](#), extending the distance range considered by a factor of 25, in order to increase the size of the primary sample and allow a robust detection of the clustering signal.

We apply structural selection to galaxies in the COSMOS field, using various cuts based on structural properties measured at ground-based resolution by the Sloan Digital Sky Survey (SDSS hereafter – [York et al., 2000](#)), for consistency with [ST14](#). Defining samples of bright, nearby primaries with spectroscopic redshifts and fainter secondary samples selected structurally, we measure the clustering of secondaries with respect to primaries, and use this to estimate what fraction of the secondaries are true satellites. We study the effect of several different selection cuts on the purity and completeness of the resulting satellites samples.

The outline of the chapter is as follows. In Section 2.3, we first introduce the surveys and datasets used. In Section 2.3, we then present the basic argument behind structural selection, using known dwarf populations from the Local Group or the ‘Local Volume’ within 11 Mpc to estimate the intrinsic distribution of dwarf galaxy properties. In Section 2.4, we describe our selection of primary and secondary samples in the COSMOS field, and explain how the primary-secondary clustering amplitude is measured. To establish a baseline for the effectiveness of structural selection, in Section 2.5, we measure this clustering amplitude for a secondary sample with no structural cuts, as well as for a sample with photo- z cuts to isolate those objects most likely to lie at the distance of the primary. In Section 2.6, we apply cuts on secondary structure instead, and show how much of the clustering signal these can recover. Finally, in Section 2.7 we consider the very nearest systems in the COSMOS catalog, that appear to be resolved or partially resolved in the Hubble Space Telescope (HST) imaging available for the field ([Koekemoer et al., 2007](#)), and give an indication of the samples that future wide-field, space-based surveys will provide. In Section 2.8 we conclude by discussing the limitations of structural selection, and the future prospects for this technique. Throughout the chapter we calculate distances assuming a cosmological model with parameters $\Omega_{m,0} = 0.31$, $\Omega_{\Lambda,0} = 0.69$ and $h = 0.678$, consistent with recent Planck analyses ([Planck Collaboration et al., 2014, 2016](#)).

2.2 Data

The data considered in this chapter includes several local samples of nearby galaxies, and the galaxies of the COSMOS field. For the latter, we use information both from high-resolution and/or space-based imaging, and from lower-resolution SDSS imaging. Each of the catalogs or sets of measurements is described below.

2.2.1 The Local Group

Although the inventory of identified local galaxies is always expanding, the Nearby Galaxy Catalog of [McConnachie \(2012, M12 hereafter\)](#) provides a reasonably recent summary of all known objects, up to a few Mpc from the Milky Way. We use the version of the catalog available on the author’s web-site¹, which was last updated in 2013. We have verified that with a few exceptions (e.g. Canis Major), the objects in this version of the Nearby Galaxy Catalog also appear in the Local Volume Catalog described below. For internal consistency, we will use the distances, magnitudes and isophotal radii recorded in the latter, since it contains more objects overall. We will use the different size measurements given in the two catalogs to estimate a half-light radius for every object in the Local Volume Catalogue, as described below.

2.2.2 The Local Volume

As discussed in the Introduction, identifying distant dwarf galaxies is challenging, and current catalogs of nearby galaxies are probably very incomplete. The most extensive list of nearby systems beyond the Local Group is the ‘Local Volume Catalog’ (LVC), first described in ([Karachentsev et al., 2004](#)). This catalogue was updated in [Karachentsev, Makarov & Kaisina \(2013, K13 hereafter\)](#), and is available on-line². We will use the LVC as an indication of what more distant dwarfs might look like from a structural point of view. In particular, we will use the B_T magnitudes and a_{26} sizes given in the on-line database, and documented in [K13](#). These are, respectively, total magnitudes in the Johnson B -band, from various sources listed in the database, and diameters of the isophotal radius corresponding to $26.5 \text{ mag/arcsec}^2$ in the B band, estimated visually and calibrated using light profiles, as described in ([K13](#)). ([K13](#) also notes that for objects with a central surface brightness equal to or fainter than $26.5 \text{ mag/arcsec}^2$, the isophotal diameter definition no longer makes sense; in these cases the values listed in the LVC correspond instead to the exponential scale radius.)

For typical objects with exponential profiles, the radius $r_{26} = a_{26}/2$ should be roughly equal to the effective radius r_{eff} (or ‘half-light radius’ r_h in [M12](#)). In principle, we could assume a specific radial profile for each object and convert more carefully from a_{26} to the effective radius, but we will not need this level of precision for the general arguments presented here. Comparing the LVC r_{26} values to the effective radii r_{eff} for the same objects given in the Nearby Galaxies Catalog, we find that the median ratio of the two radii is 1.05, although with large scatter and a systematic dependence on morphological type. Thus, in what follows we will assume $r_{\text{eff}} = r_{26} = a_{26}/2$

¹https://www.astrosci.ca/users/alan/Nearby_Dwarfs_Database.html

²<https://www.sao.ru/lv/lvgdb>

for the LVC objects. When needed, we will calculate exponential scale radii r_{exp} assuming an exponential profile, such that $r_{\text{eff}} = 1.678r_{\text{exp}}$. We will also use the mean surface density interior to the effective radius, calculated as

$$\begin{aligned}\langle\mu\rangle_{\text{eff}} &= m_{1/2} + 2.5 \log(\pi r_{\text{eff}}^2) \\ &= m_{\text{tot}} + 1.995 + 5 \log(r_{\text{eff}})\end{aligned}\tag{2.1}$$

where $m_{1/2}$ and m_{tot} are the magnitudes corresponding to half the luminosity and the total luminosity, respectively.

2.2.3 More Distant Objects

At distances $D > 11$ Mpc and out to a few tens of Megaparsecs, the Extragalactic Distance Database (Tully et al., 2009)³ provides a summary of many of the objects with known distances. We use the ‘Cosmicflows-3’ sample from the database (Tully, Courtois & Sorce, 2016, T16 hereafter) as indicative of the state of knowledge about galaxy populations in the distance range 10–50 Mpc. The database version of this catalog includes total B -band magnitudes, as well as distances estimated as described in T16.

2.2.4 COSMOS

To test our structural selection methods, we need uniform imaging for a large sample of faint galaxies with reasonably accurate (e.g. $\sigma_D \lesssim 100$ Mpc) distance estimates. Given the difficulty of obtaining spectra for faint objects, photometric redshifts (photo- z s) provide the only realistic solution. While photo- z s derived from shallow, optical photometry in five or fewer bands are of little use at low redshift (e.g. Speller & Taylor, 2014; Geha et al., 2017), those derived from deeper imaging with large numbers of narrow- and intermediate-band filters across the ultraviolet, optical and infrared range can achieve accuracies of 1% or less (e.g. Ilbert et al., 2009). The largest sample of accurate photo- z s is from the COSMOS survey (Scoville et al., 2007b), a deep, multi-wavelength survey of a 2 deg² equatorial field.

COSMOS photo- z s were derived by template fitting, as described in Mobasher et al. (2007) and Ilbert et al. (2009). More recently, they have been updated with the addition of new, deeper NIR and IR data from the UltraVISTA (McCracken et al., 2012) and SPLASH (Spitzer Large Area Survey with Hyper-Suprime-Cam⁴) projects (Laigle et al., 2016). In this chapter, we will use

³<http://edd.ifa.hawaii.edu>

⁴<http://splash.caltech.edu>

this updated catalog (‘COSMOS 2015’ hereafter)⁵ for our analysis. The quality of the photo- z s has been verified by comparing to a large number (50,000 or more) of spectroscopic redshifts available in the COSMOS field, notably from the zCOSMOS-bright sample (Lilly et al., 2007). In the redshift range $z = 0-1.2$, photo- z s for objects of magnitude $i_{AB}^+ \leq 22.5$ have an r.m.s. scatter of $\sigma = 0.7\%$ with respect to the spectroscopic redshifts, while the fraction of ‘catastrophic failures’ with relative errors $|z_p - z_s|/(1 + z_s) > 0.15$ is 0.51% .

At very low redshift, these photometric redshift errors correspond to fairly small errors in distance. Figure 2.1 shows the absolute value of the difference between the estimated photo- z and the measured spectroscopic redshift, converted to a distance error using the approximation $\Delta D = c\Delta z/H_0$, for very local objects in the COSMOS 2015 catalog ($z_s < 0.06$), as a function of their i^+ magnitude. (Six objects have differences of zero to within roundoff errors, and have been placed at $\Delta D = 2$ Mpc for clarity.) We see that for most ($\sim 80\%$) objects brighter than $i^+ = 21$, the distance errors are less than 100 Mpc, and they are less than 40 Mpc for half the objects brighter than $i^+ = 22$. Thus, most bright objects from the catalog with very small photo- z s (e.g. $z_p < 0.05$) should be genuinely nearby. We will return to this point in section 2.7.

In the process of fitting templates and estimating redshifts, Laigle et al. (2016) also calculated stellar masses and star formation rates, which we will consider further below. Finally, high-resolution imaging with ACS and/or WFC3 is available over most of the catalog area (Koekemoer et al., 2007, 2011), via the IRSA cutout server⁶. Convenient visual browsers for the ACS mosaic⁷ and the multi-wavelength coverage⁸ of the field are also available.

While there is no single public redshift catalog for the COSMOS field, most measured redshifts for the field are now available via the NASA Extragalactic Database⁹. We have used these redshifts, and a few others available privately from the COSMOS collaboration, to correct the photo- z s when possible. Further work obtaining spectra in the COSMOS field is also ongoing, e.g. with the C3R2 survey (Masters et al., 2017).

2.2.5 SDSS Photometry in the COSMOS Field

The structural selection initially introduced in ST14 was based on photometry from SDSS. SDSS covers a large area, but is both shallow (with a typical limiting magnitude of 22.2 in r ¹⁰) and has

⁵ftp://ftp.iap.fr/pub/from_users/hjmcc/COSMOS2015

⁶http://irsa.ipac.caltech.edu/data/COSMOS/index_cutouts.html

⁷<https://www.mpia.de/COSMOS/skywalker>

⁸<http://www.cadc-ccda.hia-ihp.nrc-cnrc.gc.ca/en/megapipe/cfhtls/scrollID2.html>

⁹<https://ned.ipac.caltech.edu>

¹⁰<http://www.sdss.org/dr12/scope>

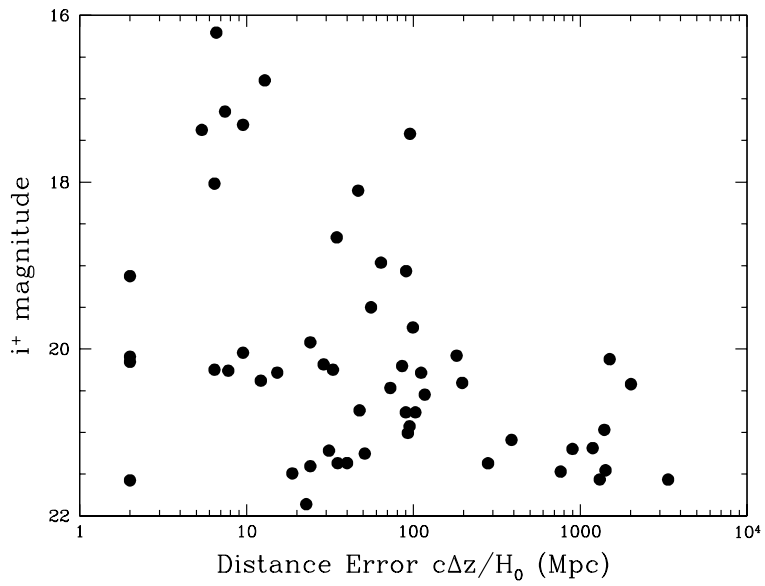


Figure 2.1: The absolute value of the difference between the photometric redshift and the spectroscopic redshift, converted to a distance error, for very local objects in the COSMOS 2015 catalog, as a function of their i^+ magnitude.

relatively poor seeing (a median value of $1.43''$ in r ¹¹). Thus it represents an image quality easily achievable by other large, ground-based surveys. For a fair comparison with the results of [ST14](#), we will also use SDSS photometry here, querying the spectroscopic and photometric galaxy catalogs from the latest SDSS Data Release 13 (DR13 – [Albareti et al., 2017](#)) and matching the results to the COSMOS 2015 catalog.

To match catalogs, we first selected a subsample of COSMOS 2015 objects likely to have detections in SDSS. From the original catalogue of half a million objects, we selected objects with $i^+ < 25.5$, $0 < z < 6.9$, $\sigma_z < 0.5$, and $z - 2\sigma_z < 0.3$, where i^+ , z and σ_z correspond to the catalogue quantities IP_MAG_AUTO, PHOTOZ, and $(ZPDF_H68 - ZPDF_L68)/2$ respectively. These cuts produced a subsample of roughly 22,000 objects. We also queried the DR13 SkyServer¹² to retrieve a photometric galaxy sample for the COSMOS region, with a ‘clean’ cut (as described at <http://skyserver.sdss.org/dr13/en/tools/search/sql.aspx>) to ensure photometric quality. We then associated objects from the reduced COSMOS and SDSS catalogs with positions identical to within $1''$ of each other. Galaxies without spectroscopic redshifts were assigned photometric redshifts and associated uncertainties from the COSMOS 2015 catalog, while those with spectroscopic redshifts were assigned the spectroscopic values, with an uncertainty of $\sigma_z = 0.0001$. The resulting matched catalogue contains 12,108 objects.

For each object in the matched SDSS-COSMOS catalog, we obtained and saved (r -band) magnitudes and sizes from SDSS. To be consistent with [ST14](#), we used the composite model (cmodel) magnitude, among the various magnitudes that SDSS provides. We did not apply a K -correction to these magnitudes, since our sample is relatively local. For galaxy sizes, we used the (r -band) exponential scale radius (expRad) provided by SDSS, as in [ST14](#). Where necessary, we convert from this scale radius to an effective radius using the relation appropriate for an exponential profile, $r_{\text{eff}} = 1.678r_{\text{exp}}$. The mean surface brightness within the effective radius is as calculated above,

$$\begin{aligned} \langle \mu \rangle_{\text{eff}} &= m_{\text{tot}} + 1.995 + 5 \log(r_{\text{eff}}) \\ \text{or } \langle \mu \rangle_{\text{eff}} &= m_{\text{tot}} + 3.1194 + 5 \log(r_{\text{exp}}), \end{aligned} \quad (2.2)$$

once again assuming an exponential profile.

2.3 The Basis for Structural Selection

Figure [2.2](#) shows a representative selection of nearby galaxies with distance estimates, including LocalGroup/LVC and more distant objects from [K13](#) and [T16](#) respectively. Galaxies are plotted

¹¹<http://classic.sdss.org/dr7/products/general/seeing.html>

¹²<http://skyserver.sdss.org/dr13/en/tools/search/sql.aspx>

in terms of their absolute B -band magnitude, estimated as described in Section 2. In general, both distances and magnitudes have considerable errors, particularly at faint magnitudes, but they give an indication of our knowledge of nearby galaxies. The upper and lower curves show the loci of objects with apparent magnitudes 17 and 22 respectively, roughly the completeness limits for current, wide-field spectroscopic and photometric surveys such as SDSS.

In the Local Group, at distances of less than 3 Mpc, approximately 120 galaxies are known, including ‘ultrafaints’ with absolute magnitudes $M_B > -6$. Within the Local Volume, the faintest identified objects generally correspond to the ‘classical’ dwarfs of the Local Group, with magnitudes $-15 < M_B < -6$. The total number of known objects in this volume is roughly 1000, though a comparison of the Local Group and LVC luminosity functions suggests the latter is incomplete by factor of up to 2 at $M_B = -10$, and a factor of 2–4 at the faintest magnitudes (see K13 for further discussion of the completeness of the LVC).

Beyond this there is a much larger volume, out to distances of 40–50 Mpc, where classical dwarfs should be easily detectable in the photometric catalogs of large-area surveys such as SDSS, given their photometric limits $M \sim 22$ (lower curve), but will lie below the typical spectroscopic limits of these surveys ($M \sim 17$ – upper curve). Given the number of objects identified in the Local Volume, for instance, we might expect $\sim (4-5)^3$ times as many, or $\sim 100,000$ galaxies, most of them dwarfs, out to $D = 50$ Mpc. On the other hand, at faint magnitudes background counts will overwhelm these local objects. The SAGA survey (Geha et al., 2017), for instance, counts roughly 3000 galaxies per square degree down to an extinction-corrected magnitude limit of $r_0 = 20.75$, versus the handful of nearby dwarf galaxies expected per square degree. Their spectroscopic follow-up around nearby bright galaxies obtained more than 17,000 spectra, but yielded only 25 new satellites, that is a detection rate of less than 1/500. This inefficiency raises the question of whether intrinsically faint, nearby galaxies could be preferentially selected by their photometric properties alone, and if so, over what range of distances.

One possible, albeit crude, alternative to complete spectroscopic surveys is to use the structural properties of dwarfs to separate them from background galaxies. Galaxies in the nearby universe show a clear trend in surface brightness with intrinsic luminosity. At fixed apparent magnitude, intrinsically faint galaxies have lower mean surface brightness on average, or equivalently, larger angular sizes on the sky. Thus it may be possible to select them preferentially using size or surface-brightness cuts. We can illustrate this by considering how the photometric properties of objects in the LVC sample of K13 would change if we saw them at progressively larger distances. The left and right panels of Figure 2.3 show how apparent magnitude, size and surface brightness change as we move the LV sample from their original distances ($D = 0-11$ Mpc; black points) to redshifts of 0.01–0.02 (blue), 0.05 (cyan), 0.1 (green), or 0.2–0.3 (yellow). Solid squares indicate intrinsically bright galaxies ($M_B < -16$), while open squares indicate intrinsically faint galaxies ($M_B \geq -16$).




Figure 2.2: Absolute B -band magnitude versus distance D for nearby galaxies, from the catalogs of [K13](#) and [T16](#). The colour scale shows the corresponding (total) apparent magnitude B_T . The upper and lower curves show the loci of objects with apparent magnitudes 17 and 22 respectively, roughly the limits for current, wide-field spectroscopic and photometric surveys such as SDSS. Note the lack of objects between the spectroscopic and photometric limits with distances $D > 11$ Mpc. (Objects in a few known clusters in this distance range, such as Virgo and Fornax, are not shown on this plot.)

In the left hand panel, objects get smaller and fainter, moving down and to the left, as their redshift increases. Intrinsically faint galaxies (open symbols) are more diffuse than intrinsically bright ones, however; as a result, at fixed apparent magnitude, dwarfs are typically $2\text{--}3\times$ larger than intrinsically bright galaxies on the sky. A cut in size and/or magnitude that selects the tail of the apparent size distribution will thus enhance the fraction of local, intrinsically faint galaxies in a sample.

The right hand panel shows a similar effect in magnitude versus surface brightness. With increasing redshift, objects move to fainter magnitudes, and then eventually shift to lower surface brightness as cosmological dimming becomes important. Intrinsically faint galaxies start at lower surface brightness, however, so the upper right hand side of the plot is dominated by low-redshift dwarf galaxies. Once again, cuts in surface brightness and/or magnitude may select these objects preferentially out of larger samples.

We note several caveats. First, the points in Figure 2.3 show the locus of typical galaxies at each distance, but do not account for changing abundance due to the increasing volume probed at larger distances. Second, we have assumed that the sample of K13 is representative of cosmological volumes in general, while in fact some galaxy types (e.g. those found in clusters) may be rare or missing entirely from the LVC sample. Finally, the region dominated by local dwarf galaxies in the right-hand panel lies at fairly low mean surface brightness. SDSS catalogs start to become significantly incomplete at central surface brightnesses of $\mu_0 \simeq 24\text{--}24.5$ (Blanton et al., 2005), although some objects can be recovered down to $\mu_0 \simeq 26\text{--}26.5$ (Kniazev et al., 2004). For the exponential surface-brightness profile typical of dwarf galaxies, these correspond to $\langle\mu\rangle_{\text{eff}} = 25.1\text{--}25.6$ or $\langle\mu\rangle_{\text{eff}} = 27.1\text{--}27.6$ respectively. Thus, objects in the most interesting region of parameter space may not be detectable in conventional, shallow surveys such as SDSS.

These complications motivate an empirical test of structural selection, using the COSMOS photometric redshift catalog, one of the only samples with accurate distance estimates for large numbers (tens or hundreds of thousands) of faint galaxies. In what follows, we will apply various structural cuts to this catalog and estimate their effect on the satellite population by measuring the resulting clustering signal.

2.4 Clustering Measurement Method

To confirm that our structural selection method works, we can measure the clustering of structurally-selected samples with respect to nearby bright galaxies that have spectroscopic redshifts, and thus reliable distance estimates. A positive clustering signal will indicate that at least part of the structurally-selected sample lies at the same distance as the primary sample, and thus




Figure 2.3: *Bottom Left Panel:* Apparent B -band magnitude versus apparent size, for LVC galaxies as seen at their original distances ($D = 0\text{--}11$ Mpc; black points), or at redshifts of 0.01–0.02 (blue), 0.05 (cyan), 0.1 (green), and 0.2–0.3 (yellow). Solid symbols indicate intrinsically bright galaxies ($M_B < -16$), while open symbols indicate intrinsically faint galaxies ($M_B > -16$). *Bottom Right Panel:* Apparent magnitude versus mean surface brightness $\langle\mu\rangle_{\text{eff}}$, for LVC galaxies seen at various distances. Symbols and colours are as in the left-hand panel. Top panels show the distributions of LVC galaxies at their original distances (i.e. the black points) only, for clarity. Note some regions of parameter space in either panel are dominated by low-redshift dwarf galaxies (open squares).

that we are preferentially selecting intrinsically faint, local galaxies. We describe the construction of the primary and secondary samples, the clustering measurement, and the corrections for masking below.

2.4.1 Selecting Primaries

Our goal in constructing the primary sample is to select bright galaxies similar to the Milky Way, at distances small enough that their satellites will be included in the COSMOS catalog, yet extending to high enough redshift that we have enough primaries to measure the clustering of their satellites with a reasonable SNR. We take as a starting point the photometry and photometric redshifts of the COSMOS 2015 photo- z catalog (Laigle et al., 2016), and proceed as follows:

1. First we select all galaxies with $M_{K_s} < -21.5$.
2. We then select those galaxies with photometric redshifts $z - 2\sigma_z < 0.3$, such that they have a reasonable chance of being low redshift objects (we choose a generous upper limit of $z = 0.3$ at this stage to make sure we do not exclude any primaries at the upper end of our highest redshift range.)
3. For this subsample, we then check for any available spectroscopy, and correct the redshift if necessary, adjusting the absolute magnitudes correspondingly. For objects with spectroscopic redshifts, the redshift errors are assumed to be $\sigma_z = 0.0001$, or $\sigma_v = 30 \text{ km s}^{-1}$.
4. Finally, we select only those galaxies with redshift errors $\sigma_z \leq 0.1$ (removing two galaxies with redshifts $z > 0.3$ from the primary sample.)

This selection process produces an initial sample of 735 primary galaxies. We estimate halo masses and virial radii for these objects from their stellar masses, assuming a standard stellar-to-halo mass relation (e.g. Leauthaud et al., 2012). The median stellar mass in this initial sample is $\langle M_* \rangle \sim 2.5 \times 10^{10} M_\odot$, corresponding to a halo of mass $M_h \sim 10^{12} M_\odot$, with a virial radius $R_{200c}^{13} \sim 200 \text{ kpc}$. We find that a few of the nearest and most massive systems (with $M_* = 2\text{--}3 \times 10^{11} M_\odot$) are predicted to have very large halo masses and projected virial radii $R_{200c} > 1 \text{ Mpc}$, complicating the clustering calculations. Thus, we make an additional cut, removing from the sample objects with absolute magnitudes brighter than -21.5 in the (SDSS) r -band. This cut reduces the final number of primaries to 527, and the median stellar mass to $\langle M_* \rangle \sim 2 \times 10^{10} M_\odot$.

¹³Where R_{200c} is the radius within which the halo has a mean density 200 times the critical density ρ_c .

The largest stellar masses in the final cut sample are $M_* \sim 7 \times 10^{10} M_\odot$, and have estimated virial radii $R_{200c} < 300$ kpc, such that our clustering calculations extend to more than three projected virial radii, even in the largest systems.

The primary sample is then divided into three redshift ranges:

- $z=0.07\text{--}0.15$, which contains 34 primaries;
- $z=0.15\text{--}0.20$, which contains 57 primaries;
- $z=0.20\text{--}0.25$, which contains 149 primaries.

The remaining 287 primaries have redshifts of 0.25 or more, which we will show is beyond the useful range for structural selection. The full redshift distribution of the primary sample is shown in Figure 2.4.

Finally, we note that ST14 also applied isolation cuts to their parent sample, to select primaries in the field or in poor groups (and thus close analogues of the Milky Way), as opposed to members of rich groups or clusters. Applying similar isolation cuts to our sample reduces the number of primaries considerably, so we will forego these cuts in the current chapter, since the focus here is on testing structure as a distance indicator, rather than on characterizing the satellite population of a given type of primary.

2.4.2 Selecting Secondaries

Our secondary source catalogue consists of those objects we were able to match between the COSMOS 2015 and SDSS catalogues, as described in Section 2.2.5. This sample contains 12,108 objects in total. We do not place any further cuts on this sample, since our initial goal is to test how much of the clustering signal we can recover without additional information. The photometric redshift distribution of the secondary sample is shown in Figure 2.4, over the range $z = 0\text{--}1$. (Note there are a few secondaries with redshifts beyond $z = 1$ not shown on the plot.)

2.4.3 Masking Corrections

The COSMOS field includes regions with poor photometry in one or more bands, due to contaminating halos from bright stars, ghosts from internal reflection, or other artefacts. While detailed mask files for these regions exist in each of the 30+ COSMOS bands, we have found it less computationally demanding to calculate clustering using a single, approximate mask image with




Figure 2.4: The cumulative and differential redshift distributions of the primary (lower curve & histogram) and secondary (upper curve and histogram) samples, over the range $z = 0-1$. Note that a few secondaries lie beyond the redshift range shown on the plot.

coarser sampling. We construct this mask empirically by making a 390×390 map of source counts in the COSMOS field, based on the entire (uncut) COSMOS 2015 catalog. Cells in this map with one or no counts are treated as potential masked regions. In a second round, any potential masked cell is determined to be masked if it has multiple neighbours with no counts. The resolution of our map file ($\sim 14''$) and threshold of one count were set such that the probability of masking a cell by chance due to Poisson fluctuations is extremely small (0.0026%). We have experimented with variants on this method, changing the source count map resolution from 200×200 up to 600×600 , and varying the threshold for counting cells as masked. We find that our clustering signals are stable to within $\sim 5\%$ with respect to these variations, but that the final mask looks most accurate for resolutions around the value (390×390) adopted here.

2.4.4 The Clustering Calculation

To measure clustering, we calculate the (cross-)correlation function of secondaries with respect to primaries, that is

$$\xi(R_p) \equiv \frac{\Delta N}{N_{\text{exp}}}(R_p) = \frac{N_{\text{obs}} - N_{\text{exp}}}{N_{\text{exp}}}(R_p)$$

where N_{obs} is the number of primary-secondary pairs observed at separation R_p , N_{exp} is the number of pairs expected for a uniform distribution, and R_p is the projected physical separation, assuming both members of the pair lie at the (spectroscopically determined) distance of the primary. We will also consider the ‘excess number’, which is simply $\Delta N(R_p) = \xi N_{\text{exp}}$.

Our method is essentially the same as that described in [ST14](#), with a few modifications in order to apply it at larger distances, so that we can obtain reasonable statistics given the relatively small field. First, we calculate the projected separations R_p of all the primary-secondary pairs, assuming the secondaries lie at the same distance as the primary. We then count the number of pairs as a function of separation, in linear bins of width 50 kpc, ranging from 50 to 1000 kpc (with the bins centered on separations of 75 kpc, 125 kpc, etc.). The innermost bin ($R_p = 0\text{--}50$ kpc, corresponding to $0\text{--}8.5''$ at $z = 0.25$) is excluded to avoid potential contamination from components (e.g. H II regions) of the primary detected independently in the catalog, and because it is comparable to the resolution of our mask for the highest redshift primaries.

To calculate the expected number of pairs, we use a local background density determined from the secondary counts between projected separations of $R_p = 600$ and 950 kpc (this range is also consistent with [ST14](#)). Given the stellar masses of our primaries, this range of separations should correspond to roughly 2–3 times the virial radius of their halos, and therefore measures the larger-scale local background (the ‘2-halo term’), rather than the overdensity associated with

the primary halo. The expected counts in the outer region are corrected for masking, and then scaled to the masked area of each inner bin to determine the expected number in that annulus.

The excess counts in bin i around primary j are thus:

$$\Delta N_{i,j} = \frac{A_{0,i}}{A_{i,j}} \left(N_{i,j} - \frac{A_{i,j}}{A_{\text{outer},j}} N_{\text{outer},j} \right) \quad (2.3)$$

where $A_{0,i} = 2\pi(R_i^2 - R_{i-1}^2)$ is the full geometric area of bin i in the absence of masking (and assuming small angles), $A_{i,j}$ is the area of bin i around primary j after masking, $N_{i,j}$ are the total counts in bin i around primary j , $A_{\text{outer},j}$ is the net area of the outer region used to calculate the background, after masking, and $N_{\text{outer},j}$ are the total counts in this region.

2.4.5 Figure of Merit for Clustering

To quantify the extent to which structural cuts can preferentially select local samples, it is convenient to define a single measurement of clustering that we can use as a figure of merit. In what follows, we will consider the SNR of the mean excess counts per primary ΔN (the ‘clustering signal’), summed over the range of separations $R_p = 50\text{--}450$ kpc, relative to a local background estimated from the secondary counts at separations $R_p = 600\text{--}950$ kpc. To calculate the error in ΔN , we assume the main uncertainty in the mean excess counts comes from the Poisson errors on the galaxy counts N_{inner}^b and N_{outer} , which are propagated into an error in the final value ΔN in the usual way.

2.5 The Clustering Signal

2.5.1 The Signal with no Additional Cuts on the Secondary Sample

To establish a baseline for subsequent measurements, we first calculate the clustering signal ΔN , by the method described in the previous section, using the entire secondary sample. Figure 2.5 shows the clustering signal of the full secondary sample with respect to primaries in the three redshift ranges, $z = 0.07\text{--}0.15$, $z = 0.15\text{--}0.20$ and $z = 0.20\text{--}0.25$ (left, middle, and right plots respectively). In each plot, the top panel shows the mean excess counts per primary in each annular bin; the middle panel shows the cumulative counts within R_p (excluding objects at $R_p < 50$ kpc), and the bottom panel shows the SNR of the cumulative excess, given the uncertainties in the excess counts in individual bins. (Note that since ΔN can be negative in any given bin, the cumulative counts and SNR do not necessarily increase monotonically with radius.)




Figure 2.5: Clustering between the full secondary sample and primaries in the redshift ranges $z = 0.07\text{--}0.15$, $z = 0.15\text{--}0.20$, and $z = 0.20\text{--}0.25$ (left, middle, and right plots respectively). In each plot, the three panels are, from top to bottom, the mean excess number of secondaries per primary in each radial bin, the cumulative excess number per primary as a function of radius (excluding objects with $R_p < 50$ kpc), and the total SNR of the cumulative excess detection.

In the lowest redshift range, we see that while there is some marginal evidence of clustering – the differential counts interior to 600 kpc are positive on average – the SNR of the cumulative excess is around 1 or less. We infer that more distance information is needed to determine which secondaries are associated with these nearby primaries, and to remove background galaxies from the secondary sample. The middle and upper redshift bins show stronger clustering, the SNR peaking at a value of 5.5–7, at projected separations $R_p = 450\text{--}500$ kpc. This scale corresponds to ~ 1.5 times the virial radius of our primaries, and matches the extent of the clustering signal seen in [ST14](#). In terms of our previously defined figure of merit, the SNR for ΔN cumulated over the range 50–450 kpc is 0.6, 4.5, and 6.4 for the three redshift ranges respectively.

2.5.2 The Signal with Photo- z Cuts on the Secondary Sample

Whereas photometric redshifts derived from a few broad bands are of limited use at low redshift (e.g. [Geha et al., 2017](#)), the COSMOS photo- z s claim percent-level accuracies, even for relatively faint galaxies at low redshift. In [Figure 2.6](#), we test this accuracy. The plot shows the surface density of secondaries around primaries, as a function of projected separation R_p and of velocity separation $\Delta V = c\Delta z$ as inferred from the photo- z s, the latter in units of the velocity/redshift error $e_{\Delta V} = c\sigma_z$ claimed in the catalog. We see a clear clustering signal at small projected separations, that is generally confined to the $\pm 2\sigma_z$ range around the primary velocity. Assuming this excess corresponds to physically associated satellites, the width of the velocity offset distribution indicates that the photo- z error estimates in the catalog are generally realistic.

Given the validity of the photo- z error estimates, we can select around each primary only those secondaries whose redshifts lie within $\pm 2\sigma_z$. (We note that secondaries should have real, physical velocity offsets with respect to the primary, but these will be negligible compared to the photo- z errors, which are typically several thousand km s^{-1} .) The resulting clustering signal for this cut sample is shown in [Figure 2.7](#). Comparing [Figures 2.5](#) and [2.7](#), we see that the photometric redshift cuts significantly improve the detection of the clustering signal, increasing the SNRs from less than 1 to 5.5, from 4.5 to 6.9, and from 6.4 to 9.8, in the three redshift ranges respectively.

If the photo- z selected sample were complete, these results and the results from [Section 2.5.1](#) would bracket the range of clustering amplitude and SNR we could expect from structural selection. If photo- z selection is relatively inefficient, we will measure clustering around the different primary samples with SNRs comparable to those in the lower panels of [Figure 2.5](#), while if it is extremely efficient, we may approach the SNRs shown in the lower panels of [Figure 2.7](#). (If the photo- z selection is incomplete, e.g. because of missing photo- z s or large redshift errors for certain objects, structural selection could actually produce a larger amplitude




Figure 2.6: The surface number density of primary-secondary pairs as a function of projected separation R_p and velocity offset $\Delta V = c\Delta z$, where the latter has been calculated from the photo- z s, and is expressed in units of the velocity uncertainty $e_{\Delta V} = c\sigma_z$.




Figure 2.7: The clustering signal, as in Figure 2.5, but after applying photo- z cuts to select only those secondaries likely to be at the same redshift as their primary.

signal than photo-z selection, albeit with lower a SNR.)

2.6 Effect of Structural Cuts

As shown in the previous section, the SNR of the clustering signal ΔN (the figure of merit defined in section 2.4.5) can be increased significantly by removing background galaxies from the secondary sample. We test the effect of five simple, single-parameter structural cuts, and two slightly more complicated two-parameter cuts, on the SNR of this measurement.

2.6.1 Single-parameter Cuts

The single-parameter cuts we test are:

- a cut on bright magnitudes, $r > r_{\text{bright}}$
- a cut on faint magnitudes, $r < r_{\text{faint}}$
- a cut on small sizes $r_{\text{exp}} > r_{\text{exp}}^{\text{low}}$
- a cut on high surface brightness, $\mu > \mu_{\text{bright}}$
- a cut on low surface brightness, $\mu < \mu_{\text{faint}}$

These are shown in the five panels of Figure 2.8, from top left to bottom right. In the latter two cases, the surface brightness is the mean value within the effective radius, $\langle \mu \rangle_{\text{eff}}$, as defined in Equation 2.2.

Reviewing the results of the first two cuts, in the top left and middle panels of Figure 2.8, we conclude that a bright magnitude limit on the secondary sample has little effect on the SNR, until this limit becomes faint enough that it starts reducing the size of the sample substantially (at which point the SNR drops correspondingly). A faint magnitude limit has more complicated effects. For the lowest redshift primaries, the maximum SNR is achieved by cutting out secondaries fainter than $r \sim 21$, while for the higher redshift primary samples, a faint magnitude cut has little effect, provided it is fainter than $r \sim 21$ – 21.5 . We note however that as we make the magnitude cut fainter, the SNR increases before the size of the secondary sample does. We conclude that objects brighter than $r \sim 21$ provide a large part of the clustering signal. (All of these results are of course contingent on the magnitude distribution of our secondary sample, which extends only to ~ 22.5 , since our secondaries are required to have SDSS photometry.)

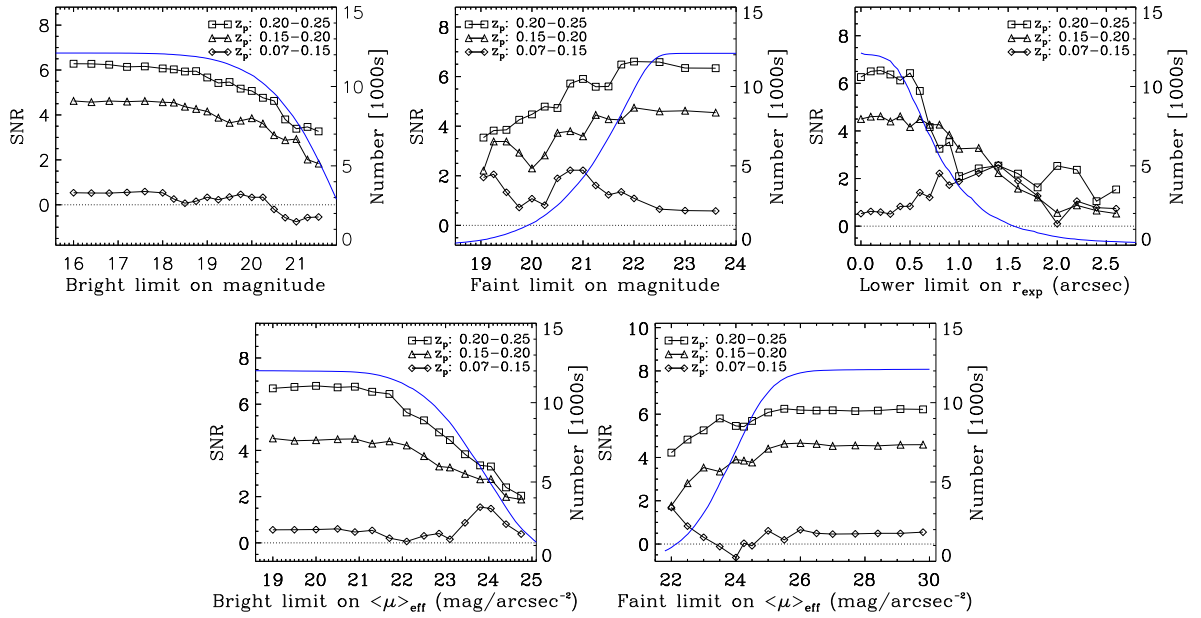


Figure 2.8: The clustering SNR as a function of bright and faint limits on the magnitude of the secondaries (top row, left and centre panels), a lower limit on secondary size r_{exp} (top row, right panel), and bright and faint limits on secondary surface brightness $\langle\mu\rangle_{\text{eff}}$ (bottom row, left and right panels). All quantities are measured in the SDSS r -band. In each panel, the three sets of points are for the three primary redshift ranges. The smooth (blue) curve in each panel indicates the total number of secondaries left in the sample after applying the magnitude cut (with the scale indicated on the right side of the plot).

The results of imposing a lower size limit depend on the redshift range of the primary sample (top right panel). At low redshifts ($z < 0.15$), the SNR starts out at ~ 0.5 , and increases to 2.5 as the value of the lower size limit increases to $1.5''$. The initial increase in SNR makes it clear that large objects are more often local, and that a lower size cut $r_{\text{exp}} \gtrsim 1\text{--}1.5''$ can enhance the fraction of local dwarfs in the sample. If the size limit is increased beyond this value, the SNR drops, probably due to the loss of objects from the secondary sample (as indicated by the smooth curve). Also the effectiveness of size cuts is restricted to low redshift; for the two upper redshift bins, a lower limit on secondary size reduces the SNR of the clustering signal overall.

Somewhat surprisingly, limits on surface brightness μ do not generally improve the SNR, except possibly at low redshift. In the lowest redshift bin, cutting out objects with surface brightness $\mu \lesssim 24$ increases the clustering SNR from ~ 0.5 to 1.5, but for the higher redshift bins, the highest SNR are achieved for no cuts at all. (A faint cut around $\mu \sim 22$ also appears to increase the SNR of the clustering measurement for the lowest redshift primaries, but in this case the size of the secondary sample is so small that we take this to be noise in the calculated SNR.)

Overall, we conclude that for low-redshift primaries, a lower limit on secondary size and/or a faint limit on magnitude can significantly increase the SNR of the clustering measurement. At higher redshift ($z = 0.2\text{--}0.25$), single parameter cuts generally have no effect, or reduce the SNR. Given the distribution of local dwarfs in magnitude-surface brightness or magnitude-size space (Figure 2.3), we expect that simultaneous cuts in two parameters may be more effective than single-parameter cuts. Before we consider these, however, we will briefly discuss the purity of the cut samples.

2.6.2 Purity of the Secondary Samples

While structural selection can enhance the SNR of the clustering signal significantly, the purity of the final cut sample, that is the fraction of the sample that is physically associated with the primaries, remains low. In terms of our previously defined quantities, the purity of a cut sample can be defined as the ratio $P = \Delta N / (\Delta N + N_{\text{exp}})$. The left panel of Figure 2.9 shows the purity of samples produced by a faint magnitude cut. When all but the brightest secondaries are cut out of the sample, the resulting purity is 20% or higher; on the other hand these cuts drastically reduce the size of the sample, and thus the SNR of the clustering measurements. Less severe cuts at $r = 21\text{--}22$ maximize the SNR, but reduce the purity to 5–15% or less. In particular, for the lowest redshift primaries, the magnitude cut with the highest clustering SNR produces a final sample with a purity $P \sim 8\%$.

The right panel of Figure 2.9 shows the purity of samples produced by a cut on small sizes. Here too, strict cuts on the secondary sample (removing all but the largest galaxies) produce higher

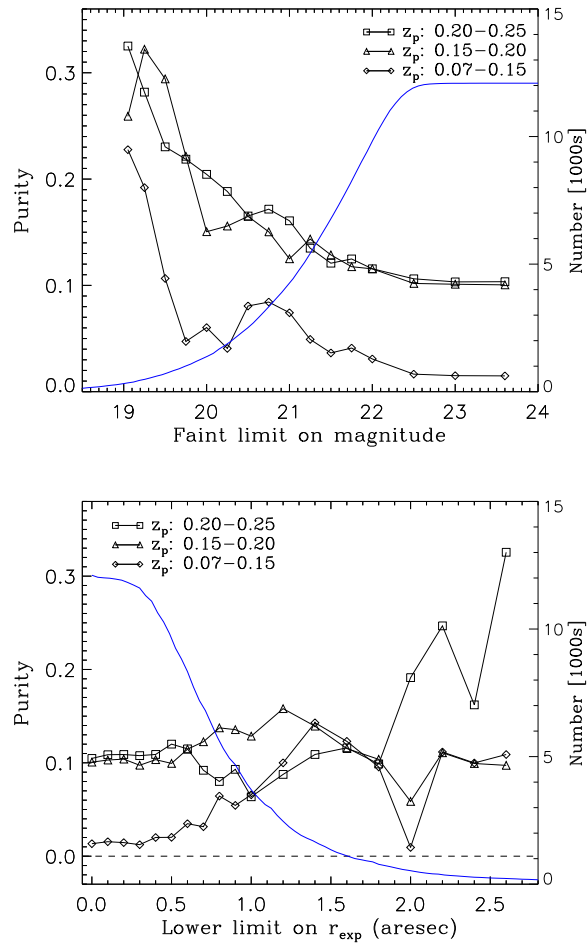


Figure 2.9: Top: Purity $P = \Delta N / (\Delta N + N_{exp})$ of the secondary sample, as a function of a faint magnitude limit. Bottom: Purity as a function of a lower size limit on the sample. The smooth (blue) curve in each panel indicates the total number of secondaries left in the sample after applying the magnitude cut (with the scale indicated on the right side of the plot).

purity (25–30%) but eliminate most of the sample, reducing the overall SNR of the clustering signal. Less strict cuts generally increase the SNR at the expense of purity. The exception is for the lowest redshift primaries, where size cuts around $1.4''$ maximize the SNR, while still retaining a purity of almost 15% .

The purity in these two examples, $P \sim 5\text{--}15\%$, is typical for all the single-parameter structural cuts we have considered in this chapter. Cuts on two parameters can produce slightly higher purity, as discussed below, but still have $P < 50\%$. Thus, while structurally-selected samples are useful for constraining overall satellite abundance, they should be used with caution when, e.g., targeting objects for spectroscopic follow-up. Even extreme magnitude or size cuts that eliminate most of the sample are relatively ineffective at conclusively identifying individual objects as low-redshift dwarf galaxies, in the absence of spectroscopic information.

2.6.3 Two-parameter Cuts

The distribution of apparent (SDSS r -band) magnitude versus size and versus surface brightness for the secondary sample is shown in Figures 2.10 and 2.11. The colour scale indicates the photometric redshift, while galaxies at $z < 0.1$ or $z > 0.7$ are denoted by larger squares/circles respectively, and shown separately in the side panels. We see that nearby galaxies ($z < 0.1$) are generally larger and lower surface brightness, but that the typical size and surface brightness depend on magnitude. Thus, two-parameter cuts in these planes seem promising for local galaxy selection.

Size-magnitude Cuts

First, we consider a size-magnitude cut. After experimenting initially with linear cuts, we found that cuts in $\log(\text{size})$ produced slightly higher SNRs. These cuts select objects with r -band magnitudes satisfying

$$r < r_0 + m \log[r_{\text{exp}}/1''] . \quad (2.4)$$

For positive/negative values of m , selected objects lie above a line sloping downwards/upwards (since magnitude increases downwards) in the magnitude-size space shown in Figure 2.10. The two free parameters are r_0 , the intercept of the line at $r_{\text{exp}} = 1''$, and m , the slope in $\log(r_{\text{exp}})$.

In Figure 2.12, we show the value of our figure of merit (the SNR of the mean excess counts, integrated between projected separations of 50–450 kpc), as a function of the parameters r_0 and m , for primaries in three redshift ranges. For the lowest redshift primaries, we find that a bright value for r_0 and a broad range of positive slopes (from $\sim 2\text{--}5$) can increase the SNR from ~ 0 to




Figure 2.10: The magnitude-size distribution of the secondary sample. Points are coloured by redshift (mainly photo-zs, corrected with spectroscopic redshifts where they are available). Larger squares and circles indicate the lowest and highest redshift objects, respectively, and are also shown separately in the side panels for clarity. A few data points with very large exponential scale radii are not shown on the plot. Lines indicate the optimal structural cuts in this space, for primary redshift ranges 0.07–0.15 (upper/rightmost line), 0.15–0.20 (middle line) and 0.20–0.25 (bottom/leftmost line). In each case, the structural cut selects galaxies *above* the line.




Figure 2.11: The magnitude-surface brightness distribution of the secondary sample. Points are coloured by redshift, with the larger symbols indicating the highest and lowest redshift objects, as in Figure 2.10. Lines indicate the optimal structural cuts in this space, for increasing primary redshift from top to bottom. In each case, the structural cut selects galaxies *above* the line.




Figure 2.12: SNR of the clustering measurement as a function of size-magnitude cuts on the secondary sample, parameterized by an intercept r_0 and a slope m (Equation 2.4). The three panels are for the three primary redshift ranges 0.07–0.15, 0.15–0.20, and 0.20–0.25, from left to right, top to bottom.

~4. This (fairly aggressive) cut removes small and/or faint objects, which are generally farther away. As we move to higher primary redshift, cuts with a fainter value of r_0 become optimal, including some with very large slopes m . For large values of m , these are close to pure size cuts. Finally at the highest redshift range, faint magnitude cuts produce the highest SNR. In particular, we need to include objects down to $r = 21$ – 22 or fainter to recover the maximum SNR. The lines on Figure 2.10 show the location of the best size-magnitude cuts for the three primary redshift ranges. Overall, comparing to our results from Sections 2.5.1 and 2.5.2, we find that size-magnitude cuts only improve the SNR significantly for low-redshift ($z < 0.15$) primaries.

Surface Brightness-magnitude Cuts

Next, we consider a cut in surface brightness and magnitude selecting objects with r -band magnitudes

$$r < r_0 + m(\mu - 25) \quad (2.5)$$

where r_0 , the intercept at $\mu = 25$, and m , the slope, are the two free parameters, and surface brightnesses are all $\langle \mu \rangle_{\text{eff}}$.

The results of this cut are shown in Figure 2.13. As in the previous figure, we see an initial pattern for low-redshift primaries (top left panel), that gradually changes as we move to higher primary redshift. At the lowest redshifts, this cut is relatively ineffective, except for one or two specific points in parameter space, which may simply reflect noise in the sampling or the clustering measurement. As the primary redshift limit increases, we find that cuts at fairly faint r_0 with slopes close to $m = 0$ (i.e. pure magnitude cuts) do best. Finally, for the highest redshift limit, any cut with a negative slope seems to work well. The lines on Figure 2.11 show the location of the best cuts for the three primary redshift ranges.

2.6.4 Optimal Structural Cuts

Table 2.1 lists the optimal parameter choices (i.e. those that maximize our figure of merit, the SNR of the clustering measurement) for the (log) size-magnitude cuts (first six columns) and the surface brightness-magnitude cuts (last six columns). For comparison, in the last two rows of each section of the table, we also list the corresponding SNRs for the secondary catalogue with no cuts (SNR_{nc}), or with photo- z cuts around each primary (SNR_{pz}). These SNRs were shown previously in the lower panels of Figures 2.5 and 2.7.

Figure 2.14 shows the clustering signal around the lowest-redshift primaries ($z < 0.15$), for the best of the structural cuts we have tested, a cut in (log) size and magnitude with the parameters




Figure 2.13: SNR of the clustering measurement as a function of surface brightness-magnitude cuts on the secondary sample, parameterized by an intercept r_0 and a slope m (Equation 2.5). Panels are as in figure 2.12.




Figure 2.14: The clustering signal between primaries in the redshift range $z = 0.07\text{--}0.15$ and the secondary sample after an optimal size-magnitude cut has been applied.

Table 2.1: Optimal values for cuts in size and magnitude (left six columns) and surface brightness and magnitude (right six columns). SNRs for clustering measurements with no cuts (SNR_{nc}) and with photo- z cuts (SNR_{pz}) are given for comparison.

Redshift Range	m	r_0	SNR	SNR_{nc}	SNR_{pz}	Redshift Range	m	r_0	SNR	SNR_{nc}	SNR_{pz}
0.07–0.15	3.3	18.7	4.3	0.6	5.5	0.07–0.15	-0.12	18.3	3.9	0.6	5.5
0.15–0.20	5.7	22.5	5.4	4.5	6.9	0.15–0.20	0.62	22.5	5.1	4.5	6.9
0.20–0.25	3.7	23.3	6.9	6.4	9.8	0.20–0.25	0.12	22.7	6.7	6.4	9.8

listed in Table 2.1. The SNR reaches a value of 4.3 at 450 kpc, compared to 5.5 for the photo- z -selected sample (Figure 2.7), or 0.6 for the uncut secondary sample (Figure 2.5). Thus, we recover about 80% of the maximum SNR obtainable with COSMOS-quality photo- z s. We can also calculate the purity of the cut sample, $P = \Delta N / (\Delta N + N_{\text{exp}})$. For the optimal size-magnitude cut this is relatively high, $P = 0.34$, so more than a third of selected objects are genuine satellites.

The net effect of the structural cuts on the redshift distribution of the secondaries can be seen by comparing the photo- z s of the uncut and cut samples. Figure 2.15 shows these distributions for the entire secondary sample, and after applying best single-parameter cuts in magnitude or size, or the best size-magnitude cut (our ‘optimal’ cut). We see that a size cut on its own is of limited use, as it reduces the size of the sample but not the shape of the redshift distribution, except perhaps at very low redshift. A cut in magnitude is more effective, reducing the number of objects at $z > 0.4$, and eliminating most objects beyond $z \sim 0.6$ – 0.8 . The optimal size-magnitude cut is most effective, however, eliminating most objects beyond $z = 0.4$, and shifting the peak of the redshift distribution from $z = 0.35$ to $z = 0.1$.

The completeness of our cut sample, relative to a photo- z selected one, is a little less clear. On the one hand, applying the optimal (size-magnitude) cut to the secondary catalogue reduces the number of objects with photo- z s below 0.15 to 18% of the uncut number, suggesting our completeness should be $\sim 20\%$ or less. The best magnitude or size cuts reduce the sample size by similar factors. On the other hand, comparing Figures 2.14 and 2.7, we see the excess counts reach a value of $\Delta N \sim 2$ at 450 kpc, or $\Delta N \sim 3$ at large radii in the structurally selected secondary sample with the optimal cut, versus $\Delta N = 3$ or $\Delta N = 4$, respectively, in the photo- z selected sample. This suggests that the cut sample contains 66–75% of the true satellites in the photo- z selected one (with an uncertainty of about 20% on that fraction). One possible resolution to this puzzle is if the photo- z selected sample is itself incomplete, due to missing photo- z s, catastrophic failures, or other problems. In this case, the amplitude of the clustering signal in Figure 2.7 would




Figure 2.15: The redshift (mainly photo- z) distributions of the entire secondary sample, and secondary samples after the best cuts in magnitude, in size, or in both size and magnitude have been applied.

be an underestimate of the true signal. At the moment, we will content ourselves with comparing the *relative* performance of structural selection and photo- z selection, and leave a discussion of absolute performance and completeness to future work.

Overall, we conclude that at low redshift ($z < 0.15$), structural selection can be very effective, recovering $\sim 80\%$ of the clustering signal obtainable with high quality photo- z s, with reasonable completeness and purity ($\sim 66\%$ and 33% respectively, albeit with some uncertainty in the absolute completeness). This result is particularly impressive, given that we have considered only simple cuts, defined either by a single limit in magnitude, size or surface brightness, or by a linear relation between magnitude and $\log(\text{size})$ or between magnitude and surface brightness. Furthermore, our cuts are based on relatively shallow SDSS photometry, whereas the COSMOS photo- z s are based on 30 bands of photometry, most of it from much deeper and/or higher-resolution imaging.

At higher redshift, it is worth noting that these simple structural cuts are *not* as effective. The highest SNRs we achieve, 5.4 for $z < 0.15$ – 0.20 and 6.9 for $z = 0.20$ – 0.25 , are only slightly higher than those obtained without any cuts on the secondary catalog (cf. Figure 2.5), indicating that we have not succeeded in separating foreground and background galaxies very effectively at these distances.

2.6.5 Completeness and Bias in Other Properties

While we have shown structural selection can be effective in preferentially selecting faint satellites around nearby galaxies, even out to redshift $z \sim 0.15$, one potential concern is the completeness of such samples, and any biases that structural selection may introduce in other satellite properties. In particular, since red and blue galaxies differ in structure, we might expect structural selection to bias the colour distribution of the final samples. To test this possibility, Figure 2.16 compares the (SDSS) $g - r$ colour distribution for the whole secondary sample, and the distributions after two of the optimal size-magnitude cuts are applied. The distributions look remarkably similar, modulo an overall scaling, although the cuts do shift the mean colour slightly to the blue (from $\langle g - r \rangle = 1.11$ for the whole sample to $\langle g - r \rangle = 1.10$ after the optimal size-magnitude cut for the redshift range 0.15 – 0.2 is applied, or $\langle g - r \rangle = 1.03$ after the optimal size-magnitude cut for the redshift range 0.07 – 0.15 is applied).

In fact, a significant part of the clustering signal comes from galaxies with blue or intermediate colours. The top panel of Figure 2.17 shows the excess counts integrated from 50 to 450 kpc around primaries in the redshift range $z = 0.07$ – 0.15 , as a function of satellite colour. As before, the optimal size-magnitude cut for this redshift range has been applied to the secondary sample. We see roughly equal signals from all three blue bins, but less signal for the reddest bin (albeit with large uncertainties). On the other hand, our satellites are not necessarily forming stars




Figure 2.16: From top to bottom, the (SDSS) $g - r$ colour distribution of the entire secondary sample (grey), the colour distribution of the sample after the optimal size-magnitude cut for the redshift range 0.15–0.2 is applied (orange), and the distribution for the sample after the optimal size-magnitude cut for the redshift range 0.07–0.15 is applied (blue).




Figure 2.17: *Top panel:* excess counts ΔN as a function of secondary colour, around primaries in the redshift range $z = 0.07$ – 0.15 . *Bottom Panel:* excess counts as a function of specific star formation rate, for the same redshift range. In both panels, the optimal size-magnitude cut for this redshift range has been applied to the secondary sample.

rapidly. The bottom panel of Figure 2.17 shows the clustering signal for secondaries binned by specific star formation rate (SSFR, as derived in the COSMOS 2015 catalog – cf. Laigle et al. 2016). Here we see that passive galaxies are generally more clustered than active ones. This suggests that the pattern of clustering with colour seen in the top panel may be a result of the redshift distribution of the secondary sample, rather than a dependence on rest-frame colour. In some applications, colour cuts might provide a useful addition to structural cuts in selecting satellites, albeit with significant implications for completeness.

2.7 Other Morphological Distance Indicators

Finally, while working with the COSMOS catalog, we have noted (and have had pointed out to us) many individual galaxies that appear to be nearby from their detailed morphology, showing features such as multiple point sources in the Hubble Space Telescope (HST) imaging. Although it is slightly tangential to our main argument, in this section we will briefly consider the use of these detailed morphological features to estimate distances to very nearby dwarfs.

2.7.1 Serendipitous Discoveries and their Redshifts

Over the years, close examination of COSMOS HST images has revealed a number of galaxies that appear to be resolved, partially resolved, or otherwise unusual. Through visual examination, we have divided these serendipitous discoveries into seven rough classes:

1. **Class 1** objects contain many clearly recognizable point sources, which together account for a significant fraction of their light. The implication is that they are close enough to be resolved into regions dominated by individual bright stars in the COSMOS ACS images (which have a resolution of approximately $0.095''$ in F814W – cf. Koekemoer et al. 2007).
2. **Class 2** may be resolved or partially resolved into point sources, but are less distinct than Class 1.
3. **Class 3** objects appear to be high surface-brightness galaxies at larger distances, but still close enough to have visible H II regions and the like.
4. **Class 4** objects are large and extremely LSB, suggesting some or all may be local LSB dwarfs.




Figure 2.18: Visual class versus distance, inferred from photometric (red squares) or spectroscopic (blue triangles) redshift, for the serendipitous sample.

5. **Class 5** objects appear to be distant galaxies whose light is significantly contaminated by a single bright galactic star superposed on the galaxy.
6. **Class 6** objects are LSB galaxies with a few superposed point sources that may or may not be foreground galactic stars.
7. **Class 7** includes all other strange or unusual objects that look like they might be nearby.

Class 1 appears to be complete, in the sense that an examination of bright, low-redshift objects in the photo- z catalog reveals no other similar objects that were not already discovered serendipitously. Class 2 appears to be fairly complete as well, although it may be missing some similar objects. The other classes are very incomplete, though enough objects are known in each to provide a representative sample.

Given the COSMOS photo- z s are accurate at the percent level, even at low redshift, as discussed in section 2.2, we can use them to verify the robustness of this visual classification. Figures 2.18 and 2.19 show the visual class and magnitude respectively, plotted versus distance inferred from the (photometric or spectroscopic) redshift.¹⁴ For Classes 1–3, we see that visual classification is surprisingly effective. All objects classified visually as being clearly nearby (Class 1) lie at distances less than $D = 130$ Mpc, and all but two are at $D < 80$ Mpc. The less certain Class 2 objects are also fairly local, but lie at distances up to 260 Mpc. Class 3 objects, which appear to be more distant visually, generally are further away, with minimum distances of 90 Mpc. The other classes consist of objects whose distances are harder to estimate, or may be incorrect due to foreground contamination; as expected, their photo- z s indicate that they lie at a wide range of distances (Classes 5–7 are possibly contaminated and/or confusing objects, so we do not include them in Figure 2.18).

The distribution of serendipitous identifications with distance and magnitude (Figure 2.19) also sheds some light on the net outcome of visual classification. Bright objects within 100 Mpc are generally assigned Class 1; bright objects at larger distances are generally assigned Class 3; Class 2 objects are generally fainter and lie at a range of distances, while the other classes, similarly, are faint and spread over a range of distances. We note that in some cases, the success of visual classification is circular; the objects in the serendipitous catalog come from many different sources, and some were flagged as having low photo- z s before they were examined visually. The majority of the serendipitous discoveries were identified directly in the HST imaging *before* their photo- z was checked, however, so overall we can confirm that visual classification works fairly well, even in the absence of other information.

From these figures, we conclude that visual classification of images with HST resolution can reliably identify bright ($i^+ < 19$) local galaxies out to distances of ~ 100 Mpc, and can identify some fainter ($i^+ = 19$ –21) galaxies out to ~ 250 Mpc. The COSMOS field alone has more than a dozen of each, in an area of less than 2 square degrees. This has interesting implications for future wide-field, space-based imaging surveys. Surveys such as those planned with Euclid¹⁵ and WFIRST¹⁶ can expect to discover tens of thousands of local, resolved galaxies, greatly enhancing our knowledge of faint, nearby galaxy populations.

¹⁴We note that in a few cases, objects in the serendipitous sample had neither a spectroscopic redshift, nor a single converged photo- z from template fitting. In these cases we took the midpoint between the 68% confidence limits as the estimated photo- z .

¹⁵<http://sci.esa.int/euclid>

¹⁶<https://wfirst.gsfc.nasa.gov>




Figure 2.19: Apparent magnitude versus distance for the serendipitous sample.




Figure 2.20: Cutouts from the COSMOS ACS mosaics (Koekemoer et al., 2007) showing examples of the different visual classes from the serendipitous sample. Each image is centred on the catalog coordinates and scaled to $6.6 r_{\text{eff}}$ on a side (with the exception of 709026, where the image is $15''$ on a side). Rows 1 and 2, from left to right, contain Class 1 (resolved) objects 213165, 260583, 331749, 401988, 458976, 561851, and 653748, 677414, 686606, 709026, 733922, 551648 respectively. (The last object on row 2, 551648 (ARK227), is Class 1 but may have the wrong spectroscopic distance.) Row 3 contains the Class 2 (marginally resolved) objects 259971, 279307, 589205, 627637, 642238, 997756. Row 4 contains the Class 3 (distant) objects 460674, 660791, 706494, 915194, 923647, 955856. Row 5 contains the Class 4 (LSB) objects 261496, 282078, 643833, 733610, 771819, 1038253. The final row contains Class 5 (contaminated) objects 377112, 484608, Class 6 (contaminated/LSB) objects 423926, 840592, and Class 7 (unclear) objects 518816, 731241.




Figure 2.21: An enlarged ACS F814W mosaic image ([Koekemoer et al., 2007](#)) of one of the Class 1 objects (ID 549719 in the COSMOS 2015 catalogue), showing diffuse, low surface-brightness emission and multiple point sources. The image is 15'' on a side.

2.7.2 Notes on Individual Objects

Table B.1 lists the IDs, coordinates, redshifts and magnitudes of the serendipitous discoveries, sorted by class. The IDs are from the COSMOS 2015 catalog, except where indicated. We note the following about individual objects:

260583 (LSBC L1-099) This is a bright Magellanic-type irregular, first catalogued by [Impey et al. \(1996\)](#), and detected in H I by [Taylor et al. \(1996\)](#). It has a spectroscopic redshift of 1816 km/s, and is part of a dwarf-dominated group in the COSMOS field at a distance of roughly 26 Mpc. This galaxy is highly fragmented in the COSMOS 2015 catalogue; as many as 18 separate catalog entries may correspond to star-forming regions or nebulosity associated with this galaxy.

279307 This irregular galaxy may be a superposition or merger between two or more objects. In the COSMOS 2015 catalogue, it is split into two separate components. It appears to have multiple faint/marginal point sources, so we have placed it in Class 2 (marginally resolved), although there is also a single, much brighter point source towards the edge of the object that could be a contaminating foreground star.

549719 This low surface-brightness object, close to a bright star, is resolved into several dozen faint point sources in HST images (see Figure 2.21). Unusually, it has imaging in three separate ACS filters, F814W and F475W from the COSMOS survey, and F606W (as well as F814W) from the CANDELS survey ([Koekemoer et al., 2011](#)). A comparison of the different HST images shows that the point sources have a broad range of colours, suggesting that they may be the brightest (supergiant) stars in an actively star-forming system. The object also appears bright in GALEX images of the COSMOS field. The photometric redshift puts this object at a rough distance of 21.5 ± 34 Mpc, but given many COSMOS galaxies in this distance range are part of the previously mentioned group at 26 Mpc, it seems likely that this object is another faint member of the group. If one or the other of these two distance estimates is correct, 549719 has an absolute magnitude of -12.4 or -12.7 in i^+ , making it one of the faintest resolved galaxies known at this distance.

On the other hand, another intriguing possibility is that 549719 could be a nearby analogue of the ‘ultra-diffuse galaxies’ (UDGs) recently discovered in the Coma cluster ([van Dokkum et al., 2015](#)). Slightly deeper *HST* imaging of these objects shows them to be diffuse, low surface-brightness, roughly spheroidal systems, with dozens of bright point sources corresponding to globular clusters ([van Dokkum et al., 2017](#)). While the point sources in 549719 show a broad

range of colours, and depending on its distance, may be too faint to be globular clusters, the possibility that this object is a field UDG warrants further investigation, as it does for several other objects in the serendipitous catalogue (e.g. 458976, 316142, the very faint 300323, and the objects in Class 4).

551648 (ARK227) This previously catalogued galaxy has a spectroscopic redshift of 1793 km/s, putting it at a distance of $D \sim 26$ Mpc, in the same group as 26058 and 677414, and giving it an absolute magnitude of -17.8 . It seems likely this redshift is incorrect, however, as the galaxy appears to be an intrinsically bright, regular elliptical with a large population of globular clusters. The brightest of these have magnitudes of $i^+ \sim 23-24$, suggesting a distance up to two times further away.

677414 (LSBC L1-100) This is another bright, Magellanic-type irregular, originally catalogued by [Impey et al. \(1996\)](#). It has a spectroscopic redshift of 1729 km/s, and is likely part of the same group as 260583. It is fragmented into at least four separate components in the COSMOS 2015 catalogue.

709026 The size of this object appears to be incorrect in the COSMOS 2015 catalogue, so we have included a $15''$ cutout in Figure 2.20. It has many distinct point sources, however, as well as extended diffuse emission, so it is clearly Class 1.

J100222.70+022520.3 This object is large and relatively diffuse, but is also located very close to a bright star. In the deeper ground-based COSMOS images it appears to have a central bar and twisting isophotes. It was masked out of the COSMOS 2015 catalogue, although it appears in earlier versions of the COSMOS photometric and photo- z catalogues ([Capak et al., 2007](#); [Mobasher et al., 2007](#)), where it has a photo- z of 0.09 (i.e. a distance of $D \sim 400$ Mpc). It is not clear whether it contains resolved point sources; the one or two in this area may be foreground stars seen in projection. Given its unusual size and surface brightness, we have included it in Class 7, although it is also another plausible candidate field UDG.

Finally, we note that two objects, 213165 and 259971, have multiple conflicting redshifts listed within $1''$ of each other. 213165 has redshifts 0.03 and 0.1529 listed, while 259971 has redshifts 0.01 and 0.8058 listed. Both objects appear to be local, however (particularly 213165, which appears to be resolved into multiple point sources), so the status of these objects is unclear.

2.8 Summary and Conclusions

From a theoretical point of view, dwarf galaxies are particularly important as tracers of small-scale dark matter structure, both in the field and within the halos of brighter galaxies. The faintest identified dwarfs are members of the Local Group, but this sample may not be representative of dwarf properties in general. In particular, if satellite populations depend on the properties – stellar mass, morphology, and/or detailed star formation history – of their central galaxy, many more examples of satellite populations will be needed to clarify these connections. Thus, identifying intrinsically faint galaxies in the nearby universe beyond the Local Group is of considerable importance.

Based on the local samples that exist, there should be a large population of objects just below the spectroscopic limits of current wide-field surveys, whose distinct structural properties can be used to separate them to some degree from the much larger number of faint background galaxies. In this chapter, we have experimented with structural selection as a tool for quantifying local dwarf populations, selecting samples with cuts in magnitude, size and surface brightness, and using their clustering with respect to bright galaxies with known redshifts to confirm that some fraction of the selected sample is indeed nearby.

We have tested this approach using the photometric redshift catalog of the COSMOS survey, since it is one of the only sources of accurate redshift estimates for large numbers of faint galaxies. In other ways, however, COSMOS is not the ideal survey for our purposes, as it covers only a small field. As a result, we have pushed our approach, originally introduced in ST14 to identify galaxies with ~ 40 Mpc, out to a redshift of $z = 0.15$ or more, that is roughly 15 times further away.

We find that structural selection does work surprisingly well even out to these distances, although it starts to fail beyond that. It produces samples enhanced in local dwarfs that are neither complete nor unbiased in magnitude or luminosity, but can nonetheless be useful in studying satellite abundance at a statistical level. Our best selection cuts recover two-thirds of the clustering signal measured using the extremely high quality COSMOS photo- z s, with 80% of the SNR, and a purity of $\sim 33\%$.

The structural selection methods tested here were deliberately based on SDSS photometry in the COSMOS field, which has both poor spatial resolution and fairly bright surface-brightness limits ($\langle \mu \rangle_{\text{eff}} \lesssim 25\text{--}27$). A new generation of low surface-brightness instruments (e.g. The Dragonfly Telephoto Array – Abraham & van Dokkum 2014) and/or surveys (e.g. MATLAS¹⁷, LSST¹⁸, HSC–SSP (Aihara et al., 2018)) will push detection limits much further into the region

¹⁷<http://irfu.cea.fr/Projets/matlas/MATLAS/MATLAS.html>

¹⁸<https://www.lsst.org>

of parameter space populated by the known local dwarfs. [Danieli, van Dokkum & Conroy \(2018\)](#), for instance, show that integrated light surveys with Dragonfly or similar instruments could detect typical local dwarfs in the magnitude range $M_V = -5$ to -10 out to distances of $D \sim 10$ Mpc. Thus, these surveys will fill in the gap between the ‘ultrafaints’, detected locally using star counts, and the brighter populations we are able to characterize at larger distances ($z \leq 0.15$, or $D \lesssim 600$ Mpc), using structural selection and clustering. We note, however, that spectroscopic follow-up may be challenging or impossible for very low surface-brightness objects, so even with these new samples, clustering analysis may still be required to determine the purity and true satellite fraction.

Our serendipitous discovery of dozens of local galaxies in the COSMOS field also augurs well for future space-based imaging surveys. While the COSMOS samples of very local galaxies are relatively small, surveys such as Euclid¹⁹ or WFIRST²⁰ should detect tens of thousands of similar objects. Here too, we expect structural selection to help significantly in separating nearby galaxies from distant ones, revealing the faintest galaxies of the local universe.

Acknowledgements

The authors acknowledge useful discussions with Simon Driver (on methods to identify nearby galaxies) and Alexandar Mechev (on the likely distance to ARK227). We also thank the referee, Bob Abraham, for a number of helpful comments and for pointing out the possible connection to UDGs. Finally, we thank our friends and collaborators from the COSMOS survey for many years of support and advice, and for first pointing out many of the objects in the serendipitous catalogue.

This chapter made use of the NASA Extragalactic Database (NED – <http://ned.ipac.caltech.edu>), the COSMOS cutout service at IRSA (<http://irsa.ipac.caltech.edu/data/COSMOS>), Knud Jahnke’s COSMOS Skywalker visual search engine (<https://www.mpia.de/COSMOS/skywalker>), and Stephen Gwyn’s interface to the multi-wavelength coverage in the COSMOS field (<http://www.cadc-ccda.hia-ihp.nrc-cnrc.gc.ca/en/megapipe/cfhtls/scrollD2.html>). We thank the creators of these resources for facilitating this work.

JET acknowledges support from the Natural Science and Engineering Research Council of Canada, through a Discovery Grant. JR was supported by JPL, which is run under a contract for NASA by Caltech. RM is supported by a Royal Society University Research Fellowship.

¹⁹<http://sci.esa.int/euclid>

²⁰<https://wfirst.gsfc.nasa.gov>

The COSMOS 2015 catalog is based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under ESO programme ID 179.A-2005 and on data products produced by TERAPIX and the Cambridge Astronomy Survey Unit on behalf of the UltraVISTA consortium.

Chapter 3

A hierarchical clustering method for quantifying satellite abundance

Abstract

We present a new method for quantifying the abundance of satellites around field galaxies and in groups. The method is designed to work with samples, such as local photometric redshift catalogues, that do not have full spectroscopic coverage, but for which some redshift or distance information is available. It consists of identifying the galaxies most likely to be centrals, and using the clustering signal around them as a template to iteratively decompose the full population into satellite and central populations. In that sense it is similar to performing crowded-field photometry, after having first used isolated stars to determine the point spread function of the image. The method does not identify individual satellites or centrals conclusively, but assigns a probability to each galaxy of being one or the other. Averaged over a large sample, it provides a statistical estimate of satellite abundance, even in crowded fields with large redshift uncertainties. We test the method using data from the COSMOS field, which includes a large set of local objects with accurate photometric redshifts. We measure satellite abundance as a function of central stellar or halo mass, as well as the satellite luminosity function, and find results consistent with previous studies, but extending over a broader range of central masses. We also consider a number of possible systematic uncertainties in the method, and show that they are generally smaller than our random errors. Having presented the method in this chapter, we will use it to study the properties of the satellite populations in a forthcoming one.

3.1 Introduction

In the current picture of hierarchical structure formation, cold dark matter (CDM) haloes merge together to form progressively larger systems as the Universe evolves. Smaller dark matter halos often survive accretion onto larger systems, leaving a population of distinct ‘subhalos’ within CDM halos. While the gas that cools and settles into the centre of the main halo will contribute to the growth of a central galaxy, smaller galaxies that formed within subhalos before they merged can survive as distinct satellite galaxies, subject to a broad range of environmental effects that may transform or destroy them, including tidal heating or stripping, encounters, and internal or external feedback. The observed abundance of satellites in the local Universe provides a detailed test of this complex picture and gives important insights into the overall effect of environment on galaxy formation.

The dominant galaxies of the Local Group (LG), the Milky Way (MW) and M31, have the best studied satellite populations in the Universe. Recent surveys have discovered many new, faint members of the LG (e.g. [Bechtol et al., 2015](#); [Drlica-Wagner et al., 2015](#); [Koposov et al., 2015](#)), such that the total abundance of LG satellites can be estimated with increasing confidence ([Newton et al., 2018](#)). Over the past two decades, however, several points of tension have arisen between the observed population of LG satellites and that expected from theory. The most famous is the “missing satellite problem”, which contrasts the small number of observed satellites with the large number of dark structures predicted by theory ([Klypin et al., 1999b](#); [Moore et al., 1999a](#)). A second, “too-big-to-fail” problem contrasts the low central densities estimated in the massive satellites of the MW with the much higher densities expected from theory ([Boylan-Kolchin, Bullock & Kaplinghat, 2011](#)). There may be other tensions as well, in the radial clustering (e.g. [Kravtsov, Gnedin & Klypin, 2004](#); [Taylor & Babul, 2004](#)) or 3-D spatial distribution (e.g. [Pawlowski et al., 2015](#)). We refer the reader to [Bullock & Boylan-Kolchin \(2017\)](#) for a detailed review of these challenges.

Many solutions have been proposed to resolve the tensions between theory and observations of the LG satellites, including internal feedback due to star formation (e.g. [Dekel & Silk, 1986](#); [Mashchenko, Wadsley & Couchman, 2008](#); [Governato et al., 2010](#); [Wetzel et al., 2016](#)), the effects of global (e.g. [Bullock, Kravtsov & Weinberg, 2000](#); [Gnedin & Kravtsov, 2006](#)) and/or inhomogeneous (e.g. [Lunnan et al., 2012](#)) reionization, tidal or other environmental effects (e.g. [Taylor & Babul, 2001](#); [Mayer et al., 2006](#); [Lokas, Kazantzidis & Mayer, 2012](#)), or modifications to the underlying dark matter model such as warm dark matter (e.g. [Macciò & Fontanot, 2010](#); [Anderhalden et al., 2013](#); [Kennedy et al., 2014](#); [Lovell et al., 2014](#)), self-interacting dark matter (e.g. [Spergel & Steinhardt, 2000](#); [Fry et al., 2015](#); [Elbert et al., 2015](#)), or fuzzy dark matter (e.g. [Nadler et al., 2019](#)).

There remains, however, the important question of whether the MW and/or LG satellites are

representative of all satellite populations. Many observational studies have shown that the MW is *not* typical in having two LMC/SMC-like satellites (Guo et al., 2011; Liu et al., 2011; Strigari & Wechsler, 2012; Robotham et al., 2012; Speller & Taylor, 2014), and similar conclusions have been suggested by numerical simulations (Boylan-Kolchin et al., 2010; Busha et al., 2011; Kang, Wang & Luo, 2016; Zhang, Luo & Kang, 2019). The ongoing Satellites Around Galactic Analogs (SAGA) survey (Geha et al., 2017) has also shown that there is a large variation in satellite populations from system to system. Theoretical models predict that the abundance of halo substructure should vary more than expected from Poisson statistics alone, and should be correlated with the formation redshift of the system (e.g. Jiang & van den Bosch, 2017; Chua et al., 2017). These complications caution us from relying too heavily on the properties of a single system to constrain models of galaxy formation. To determine whether the LG is representative, and to understand satellite properties across a broad range of environments, we should seek out satellites around as large a sample as possible of central galaxies.

Identifying satellites and distinguishing them from foreground or background systems requires some form of distance information. The main approaches in the literature include: (1) the use of existing complete spectroscopy to identify satellites around the nearest and brightest systems (e.g. Yang et al., 2007); (2) dedicated spectroscopic campaigns to obtain spectroscopy for fainter targets around a smaller number of selected systems (Geha et al., 2017); (3) the use of photometric distance estimates from techniques such as the tip of the red giant branch (TRGB) (e.g. Carlin et al., 2016; Danieli et al., 2017; Cohen et al., 2018; Danieli et al., 2019), or surface brightness fluctuations (e.g. van Dokkum et al., 2018; Carlsten et al., 2019); (4) statistical abundance measurements based on clustering (e.g. Liu et al., 2011; Guo et al., 2011, 2012; Strigari & Wechsler, 2012; Wang & White, 2012; Wang et al., 2014; Sales et al., 2013; Speller & Taylor, 2014; Xi et al., 2018). The four approaches have different strengths and weaknesses. Method (1) requires only existing data, but is restricted to the brightest satellites in the nearest systems, and may also suffer from incompleteness due to fibre positioning limitations in dense fields (e.g. Guo, Zehavi & Zheng, 2012; Smith et al., 2019). Method (2) is extremely expensive in terms of observing time, and thus limited to small numbers of systems. Method (3) is restricted to very nearby systems [<20 Mpc], whose virial radii subtend large angles on the sky, making complete coverage difficult. Method (4) cannot confirm individual galaxies as satellites or centrals; it has been very successful, however, in making measurements of the average satellite abundance, and is the least resource-intensive method of the four *a priori*.

Clustering-based methods have generally been applied to samples at redshifts ~ 0.05 – 0.2 , selected from the Sloan Digital Sky Survey (SDSS – York et al. (2000); e.g. Liu et al. (2011); Guo et al. (2011, 2012); Strigari & Wechsler (2012); Wang & White (2012); Wang et al. (2014); Sales et al. (2013)). A different strategy was adopted by Speller & Taylor (2014), who focussed on very nearby systems (out to 42 Mpc). This allowed them to estimate the abundance of intrinsically

faint satellites, at the expense of significant background contamination. They used a selection technique based on galaxy structural properties (mainly apparent size) to reduce the background contamination and boost the signal-to-noise ratio (SNR) of the clustering measurement. The technique was further developed and tested in [Xi et al. \(2018\)](#), using a broader range of morphological cuts. An optimized version was shown to be effective up to $z \sim 0.15$, far beyond the range considered in [Speller & Taylor \(2014\)](#).

These previous clustering-based studies have generally considered samples of primaries that are clearly isolated, in the sense that they have no brighter companion within fixed projected and line-of-sight separations. This approach works well for bright, massive primaries, but becomes inefficient for less luminous ones. By dropping isolation cuts, [Xi et al. \(2018\)](#) were able to detect a clear clustering signal and constrain satellite abundance using only observations from the fairly small COSMOS field, but this resulted in a broad selection of primaries, including many systems with overlapping virial regions. As a result, the interpretation of their results remains slightly unclear, relative to previous studies, as not all of their primaries are true central galaxies.

In this work, we introduce a new method to deal with the complications of overlapping systems and crowded fields. We start by identifying the subset of galaxies in a sample most likely to be true central galaxies, using a hierarchical search in which galaxies are checked for isolation in order of decreasing stellar mass, with isolation criteria that scale with the estimated virial radius of the system. The cross-correlation function of the sample with respect to this set of most likely primaries provides an initial template for the clustering signal. This template is used to estimate the probability that *any* member of the sample is a primary or a secondary. Finally, we can iterate through the last two steps, recalculating a probability-weighted cross-correlation function and the modified primary/secondary probabilities until convergence. The final primary/secondary probabilities for the whole sample then allow us to estimate satellite abundance, luminosity functions, and other distributions of secondary properties. Note that we have developed and optimized our method for low redshift samples. Some of our assumptions may need modification, in order to apply the method at higher redshifts.

In this chapter we present the method and give some simple estimates of satellite abundance; in the next chapter we will study the properties of the detected satellite populations in more detail. The chapter is structured as follows. In [Section 3.2](#) we describe our data selection, including the basic cuts that define our initial sample. In [Section 3.3](#) we measure the clustering signal and use it to define a “region of interest” around each primary likely to contain most genuine satellites. In [Section 3.4](#) we describe our iterative method for estimating primary and secondary probabilities for each galaxy. In [Section 3.5](#) we present our main results on satellite abundance. In [Section 3.6](#) we test the method for possible systematic uncertainties. Finally, in [Section 3.7](#) we summarize our results and discuss future prospects for this new method.

3.2 Data–COSMOS

The satellite galaxies we can hope to detect around a low-redshift primary (at most a few tens per system, based on abundances in the Local Group) will be seen in projection with a much larger number of foreground and background galaxies (on the order of thousands) that are not physically associated with the primary. Precise distance information is essential for separating true satellites from this foreground/background population. Spectroscopic redshifts are ideal for this purpose, but impractical for large samples. For instance, if we want to search for satellites brighter than -18 in absolute magnitude out to a redshift of 0.2, this requires distance information for galaxies down to an apparent magnitude of roughly 22. However, with a few exceptions (e.g. [Geha et al., 2017](#)), wide-field spectroscopic catalogues are usually only complete down to an apparent magnitude of 17 to 18, far from the depth required. Thus, using photometric redshifts (“photo-zs”) is the only realistic solution. The COSMOS field features high quality photo-zs generated from 30+ deep bands ([Scoville et al., 2007b](#); [Ilbert et al., 2013](#); [Laigle et al., 2016](#)), making it an ideal place to test our method.

3.2.1 The COSMOS photometric redshift catalogue

COSMOS is a deep ($AB \sim 25\text{--}26$), multi-wavelength ($0.25 \mu\text{m}\text{--}24 \mu\text{m}$) survey covering a 2 deg^2 equatorial field ([Scoville et al., 2007b](#)). The multi-wavelength imaging includes Hubble Space Telescope (HST) imaging with the Advanced Camera for Survey (ACS) and follow-up observations from many other facilities across a wide range of wavelengths – X-ray, UV, optical/IR, FIR/submillimeter and radio ([Scoville et al., 2007b](#)). In this chapter, we will use a recently updated photometric redshift catalogue ([Laigle et al., 2016](#), ‘COSMOS 2015’ hereafter) for our analysis. The main improvement of this catalogue compared to the previous releases is the addition of new, deeper NIR and IR data from the second data release (DR2) of the UltraVISTA and SPLASH (*Spitzer* Large Area Survey with Hyper-Suprime-Cam [Miyazaki et al., 2012](#)) projects. Compared to the first data release (DR1) of UltraVISTA, the exposure time of DR2 was significantly longer ([McCracken et al., 2012](#)), providing the deeper IR and NIR data as well as better SNRs ([Laigle et al., 2016](#)). On the other hand, the DR 2 data only covers a part (namely ‘ultra-deep stripes’, roughly 0.6 deg^2) of the COSMOS field. This causes a slight inconsistency in depth and SNR across the field, which we will address below by applying a magnitude cut.

The COSMOS photo-zs were derived using χ^2 template fitting, as described in [Mobasher et al. \(2007\)](#) and [Ilbert et al. \(2009\)](#). The Spectral Energy Distribution (SED) templates used in the COSMOS 2015 catalogue include a set of 31 spiral and elliptical galaxies from [Polletta](#)

et al. (2007) and a set of templates for young blue star-forming galaxies generated using Bruzual & Charlot (2003) models. Given the updated NIR and IR data and two additional star-forming galaxy templates, Laigle et al. (2016) further improved on photo-z quality relative to previous COSMOS catalogues (Capak et al., 2007; Ilbert et al., 2009, 2013). The accuracy of the photo-zs has been verified by comparing them to a large number of highly reliable (97% confidence) spectroscopic redshifts (Lilly et al., 2007) that are available in the COSMOS field. For the objects of magnitude $i_{AB}^+ < 22.5$ and redshift range of $z = 0-1.2$, the photo-zs have an r.m.s scatter of $\sigma = 0.7\%$ with respect to the spectroscopic redshifts, while the occurrence of “catastrophic failures” with relative errors $|z_p - z_s| / (1 + z_s) > 0.15$ is only 0.51%. For this work, we choose the median of photo-z likelihood distribution from the template fitting (“ZPDF” in the catalogue) as the base redshift. From this redshift we calculate angular-diameter and luminosity distances, and corresponding luminosities and projected separations, assuming all galaxies follow the Hubble flow. In the process of template fitting and photo-zs estimation, Laigle et al. (2016) also calculated stellar masses and star formation rates for the galaxy samples, which will be used in our analysis below. Specifically, we use “MASS_MED” and “SFR_MED”, the medians of the stellar mass and star-formation-rate probability distribution functions (PDFs).

3.2.2 Additional spectroscopic redshifts

We can further improve on our distance estimates by supplementing the COSMOS photo-zs with spectroscopic redshifts, where these are available. While there is no single public spectroscopic redshift catalogue for the whole COSMOS field, most of the measured redshifts in the region are now accessible through the NASA Extragalactic Database¹. In addition to these redshifts, we also obtained a few other unpublished redshifts from the COSMOS collaboration (M. Salvato, private communication). The redshifts used in this work will be mainly photo-zs from the COSMOS 2015 catalogue, but replaced with spectroscopic redshifts where possible. Given the numerous literature sources and slightly different qualities of the spectroscopic redshifts, a universal redshift uncertainty of 0.0001 is assigned to each galaxy whose photometric redshift is replaced with a spectroscopic value. Absolute magnitudes and stellar masses for those objects are also corrected, based on the resulting change in the distance modulus.

3.2.3 The base sample

As mentioned above, the depth of the COSMOS 2015 catalogue varies across the field, depending on whether the new “ultra-deep” (UltraVISTA DR 2) imaging is available or not. In general, the

¹<https://ned.ipac.caltech.edu>

catalogue appears to be relatively complete down to a magnitude of $i^+ < 25.5$ (MAG_AUTO), but becomes incomplete beyond this. [The 3σ depths in the i^+ band are 26.2 and 26.9 for $3''$ and $2''$ apertures respectively (Laigle et al., 2016).] To ensure reasonable completeness over the redshift range of interest, we apply the following initial cuts on the catalogue, which are the same cuts used in Xi et al. (2018):

- $i^+ < 25.5$
- $0 < z_{\text{pdf}} < 6.9$
- $z - 2\sigma_z < 0.3$
- $\sigma_z < 0.5$

where z_{pdf} refers to the median of photo-z likelihood distribution measured using galaxy template fitting, and σ_z refers to the photo-z error, estimated by using the 68% confidence level upper and lower limits of the photo-z likelihood distribution provided in the catalogue (i.e. $\sigma_z = (z_{\text{pdf}}^{U68} - z_{\text{pdf}}^{L68})/2$). Note that we include the broad redshift cut $0 < z_{\text{pdf}} < 6.9$ to exclude stars and X-ray sources in the catalogue, as well as objects in the masked regions, as these objects do not have robust z_{pdf} estimates; we include the upper limit cut of redshift $z - 2\sigma_z < 0.3$ to focus on the local volume while keeping a reasonable completeness over a target redshift range of 0–0.25; finally, we use a redshift error cut $\sigma_z < 0.5$ to exclude those galaxies with poor quality redshifts (mainly faint galaxies) from the further analysis. These cuts produce a base catalogue of 41,559 galaxies (37,578 after excluding galaxies with large redshift errors). Fig. 3.1 shows the redshift versus i^+ absolute magnitude distribution for our base catalogue after applying the cuts above. Given our cut in apparent magnitude, the sample galaxies have absolute magnitudes between -24 and -10 for the redshift range ($z=0-0.25$) we will consider below.

3.2.4 Stellar Mass Completeness

Laigle et al. (2016) estimated the stellar mass completeness of their catalogue; for redshift range $0 < z < 0.35$, they suggested a 90% completeness limit of $M_* = 10^{8.6} M_\odot$. We are considering systems at a slightly lower redshift range $z < 0.25$. Examining the stellar mass function and stellar mass errors for this redshift range, we conclude that we are complete down to at least $M_* < 10^{8.2} M_\odot$, where the differential mass function peaks, and cut the main sample at this value. For very low redshifts ($z < 0.07$) we appear to be complete down as low as $M_* < 10^{7.2}-10^{7.5} M_\odot$; we will discuss local satellite abundance at these lower stellar masses below and in subsequent work.




Figure 3.1: The redshift versus i^+ band absolute magnitude distribution of our base catalogue, coloured by the i^+ band apparent magnitude, as indicated in the right-hand colour scale.

3.3 Determining the Clustering Scale

We assume a model in which the galaxy with the largest stellar mass is the dominant galaxy within each halo, and resides at or close to, its geometric and dynamical centre. To tell whether a given galaxy is the dominant central galaxy (CG, or “primary” hereafter) within its own halo, or a potential satellite of another primary, we need to search its surroundings to see if there is another more massive galaxy close by. If there are more massive neighbours nearby, this suggests the galaxy may be satellite, whereas if all nearby galaxies are less massive, it suggests the galaxy is a primary. To quantify the characteristic scale on which satellites are associated with their primaries, we measured the clustering pattern of all the galaxy pairs in our base sample. Based on these clustering results, we will quantify “spatially nearby” and specify the region of interest (ROI) for the primary-secondary classification.

3.3.1 Halo Mass Assignment

We expect the characteristic extent of the satellite distribution in a halo to scale with its virial radius. We calculate a fiducial halo mass and virial radius for each galaxy, assuming that it is the CG of its host halo. These masses and virial radii will be used to characterize the clustering throughout this work.

To estimate halo masses, we could in principle use abundance matching, assuming a monotonic relation between stellar mass and halo mass that we derived empirically by comparing the observed stellar mass function and the predicted halo mass function within the observation volume. However, the effective area of the COSMOS field (after correcting for masking as discussed below) is only 1.46 deg^2 , giving an effective comoving volume of $1.67 \times 10^5 \text{ Mpc}^3$ up to $z = 0.25$. The cosmic variance in the mean density for a volume this size is a factor of approximately 0.4 (Somerville et al., 2004); if we consider ten independent, equal-volume redshift slices, the relative cosmic variance of each slice increases to 0.7. Thus, there is a large systematic uncertainty in the normalization of the halo mass function within this volume. Instead, we use the Stellar-to-Halo Mass Relation (SHMR) derived by Behroozi, Wechsler & Conroy (2013) (B13 hereafter). They provide a formula (Eqn. C.1) for the inverse Halo-to-Stellar Mass Ratio (HSMR), with parameters as listed in Appendix C.1.

We note, however, that a combination of observational errors in the stellar mass estimates, intrinsic scatter in the SHMR, and the non-linearity of the halo mass function will bias halo masses estimated directly from the B13 formula. As there are many more low mass halos than high mass ones, errors in stellar mass will more frequently scatter objects into a given mass range from below than from above. Thus, directly applying the HSMR to estimate the halo mass

corresponding to an observed stellar mass will lead to systematic overestimates, especially at the high-mass end. To quantify and correct this bias, we have performed a Monte-Carlo simulation of the effect of errors to obtain a bias-corrected HSMR based on [B13]. We explain this correction in more detail in Appendix C.1.

Given an estimate of the halo mass for each galaxy, we also assign a corresponding virial radius and virial velocity. The virial radius is taken to be $\left(\frac{3M_h}{4\pi\rho_c\Delta_c}\right)^{1/3}$, where ρ_c is the critical density of the universe and Δ_c is the mean overdensity of the halo within the virial radius, with respect to the critical density. We use the fitting formula from Bryan & Norman (1998) for Δ_c . Once the virial radius has been calculated, the virial velocity is given by $\sqrt{GM_h/R_{\text{vir}}}$.

The stellar mass completeness limit is estimated empirically, following the method described in Laigle et al. (2016). They calculated the K_s -band magnitude limits for the COSMOS catalogue to be 24.0 and 24.7 for the deep and ultra-deep fields respectively. In this work, we choose a limit of 24.0 in order to have uniform depth across the whole field. Given this magnitude limit, the limiting mass a galaxy would need to have to be observed at a given redshift is calculated as:

$$\log M_{\text{lim}} = \log M - 0.4 (K_{s\text{lim}} - K_s) \quad (3.1)$$

Next, a stellar mass limit is estimated for each redshift bin, within which 90% of the galaxies lie, given the stellar mass errors. We also calculate a corresponding halo mass limit in each redshift bin using the bias-corrected HSMR (although in this case we ignore the effect of scatter on the completeness threshold).

3.3.2 Defining a Region of Interest (ROI)

Around each CG, we define a ‘‘Region of Interest’’ (ROI) in which we will search for potential satellites. The size of ROI is determined by two considerations: first, the clustering signal should be consistent across systems with different CG masses, and second, the ROI should include most of the ‘‘one-halo’’ clustering signal associated with the main halo around the CG, while excluding the regions that are dominated by the ‘‘two-halo’’ or background terms. We have explored different possible choices of ROI boundaries by measuring the two-dimensional (line-of-sight and projected) clustering signal of all pairs in our base catalogue.

For each pair in the catalogue, the galaxy with the larger stellar mass is assumed to be the primary. We then count pairs as a function of separation in projected distance and line-of-sight velocity offset, scaled by the halo virial radius and the velocity error respectively. These scaling choices are found to give us relatively consistent clustering signals for primaries of different masses (as shown in Fig. 3.2 below). Note that some fraction of an annulus around a given

primary may be missing from the catalogue, as it overlaps with field boundaries and masked regions. We carefully measure the shape of the survey boundaries and masked regions to produce a single template for the whole field. Monte Carlo sampling of this template is then used to determine the area completeness η around each primary as a function of radius, as explained in Appendix C.2, and the counts are corrected by this factor.

Fig. 3.2 shows the density of pair counts in 2D phase space for all galaxies (top panel), and binned by primary halo mass (bottom panels), where the primary is defined to be the member of the pair with the larger stellar mass. A clear overdensity of pairs can be seen both in the projected separation and in the velocity separation directions. Overall, primaries with larger masses show a stronger clustering signal. At the same time, the clustering patterns in the different mass ranges have a similar dependence on separation scale: they all have the strongest clustering within $0.5 R_{\text{vir}}$. The signals all extend to fairly large radii, but start to drop significantly after $1.5\text{--}2 R_{\text{vir}}$. Along the velocity axis, which is scaled by redshift error, the signals in all three mass bins drop off at a similar rate, reaching the background level at $\Delta v \sim 1\text{--}1.5 \sigma_{\Delta v}$. It is worth noting the slight asymmetry of the clustering signal along the velocity separation axis, with slightly more negative velocity separations than positive ones. This is due to the incompleteness at the faint, high-redshift end of the survey volume (as shown in Fig. 3.1). We calculate the velocity offset with respect to the more massive (and thus more luminous) member of the pair, which as a result of Eddington bias due to incompleteness, trends to be further away on average. Thus, it produces a negative velocity offset more often than a positive one.

Based on these clustering patterns, around each primary galaxy we will define a ROI for potential satellites using the following cuts:

1. A cut in projected separation (assuming both galaxies are at the line-of-sight distance of the primary), scaled by the virial radius of the primary: $R_P/R_{\text{vir}} < A$
2. A cut in velocity difference, scaled by the circular velocity of the primary: $\Delta v/v_{\text{vir}} < B$
3. A cut in velocity difference relative to the uncertainty in velocity difference between the primary and secondary: $\Delta v/\sigma_{\Delta v} < C$

Note the final cut depends on the secondary properties as well as those of the primary, so this is determined for each galaxy pair individually. Any secondary galaxy in a pair that meets conditions (i) and ((ii) or (iii)) is considered a potential satellite of the more massive member. We will choose the values $(A, B, C) = (3.0, 2.0, 1.5)$ as our default, but test the effect of changing these definitions of the ROI in Section. 3.6 below.

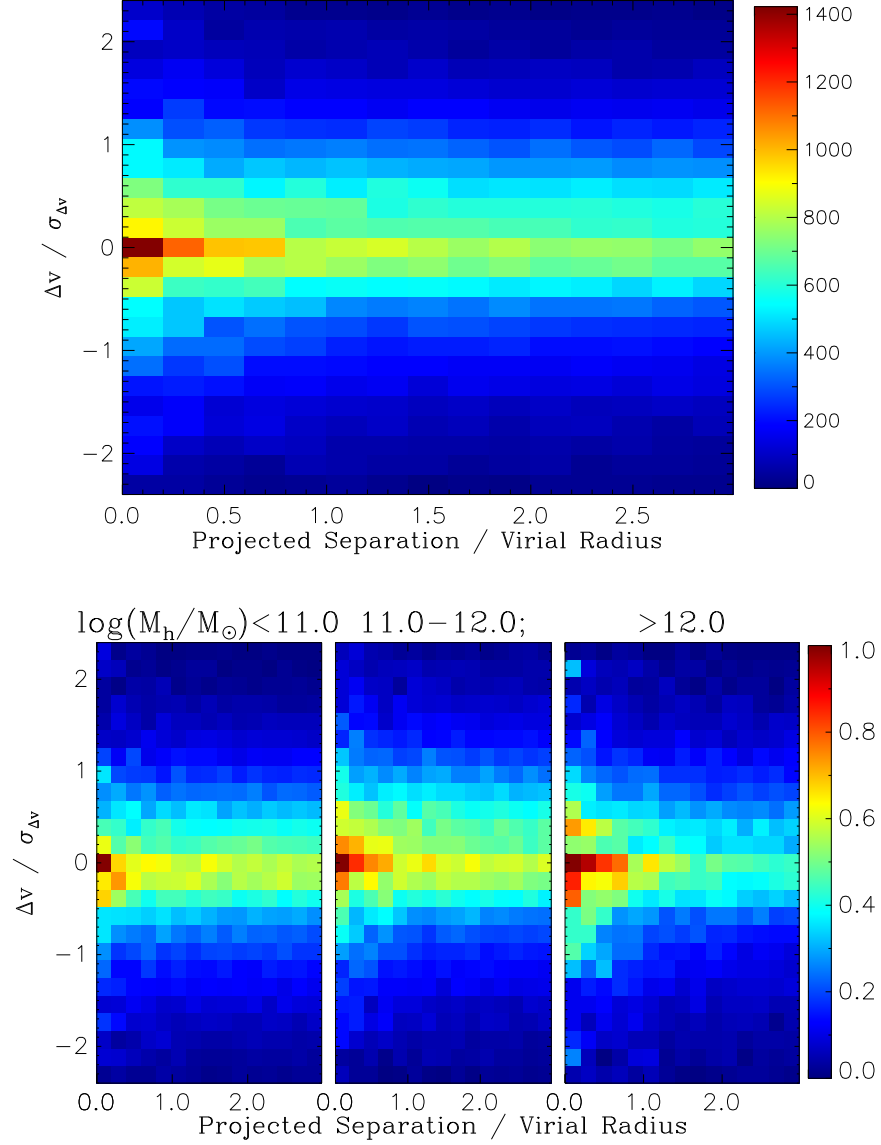


Figure 3.2: Primary-secondary clustering signal in the base catalogue. The colour scale shows the density of galaxy pairs $\Sigma = N/\Delta R\Delta v$ as a function of projected separation and velocity offset, where these have been scaled by the estimated virial radius and velocity error, respectively. Top panel: all primaries; bottom panels: same quantity normalized and binned by primary mass.

3.4 The iterative clustering method

Our goal is to quantify satellite abundance using the strength of the clustering signal. A obvious complication is that any given galaxy may be a central or a satellite; without further selection, the ‘raw’ clustering signal measured in the previous section consists of a complicated sum of central-satellite and offset-satellite terms (e.g. [Leauthaud et al., 2012](#)). Depending on the stellar mass and range of separations considered, the latter can significantly bias estimates of the true satellite abundance.

To avoid this complication, our strategy will be to identify the galaxies most likely to be centrals, and use the clustering signal around these objects as a template to separate out central and satellite contributions. In this section, we will first describe the initial, ‘first-run’ primary selection; then we will model the clustering of secondaries around these first-run primaries to determine primary and secondary probabilities for all galaxies in the sample, and finally we will test the results of iterating over this process, by adding new potential primaries to our initial sample, weighted by their primary probability, and remeasuring a weighted clustering signal around the enlarged primary sample.

3.4.1 Initial primary selection

We want to select primaries that dominate a ROI that scales with their halo mass, as described in the previous section. Since the ROI is larger for the more massive systems, a smaller system that has no more massive companions within its own ROI can still lie within the ROI of a larger system. This asymmetry naturally leads us to a hierarchical search algorithm, where we start searching around the most massive systems first. The detailed steps are as follows:

1. All galaxies in the catalogue are assigned a halo mass and virial velocity dispersion based on their stellar mass, as if they were the CG in their own host halo.
2. We then go through the catalogue in ranked order of stellar mass, selecting the most massive galaxy in the catalogue as the first primary.
3. All galaxies that lie in the ROI around the first primary (as defined in the previous section) are classified as its secondaries and removed from the list of potential primaries;
4. The next most massive unclassified galaxy is then selected as the next primary candidate.
5. We check the stellar masses of all galaxies within the ROI around this next candidate. If the candidate is the most massive galaxy in its ROI, then it is classified as a primary and the other galaxies in the ROI are classified as its secondaries.

6. We iterate over the last two steps until all the galaxies in the catalogue are classified as primaries or secondaries.

This produces our first run primary sample. To guarantee reasonable completeness, we make two additional cuts on this initial sample:

1. Primaries with ROIs that are heavily affected by survey boundaries or masking are removed from the initial primary list.
2. Primaries with redshifts higher than 0.25 are removed from the initial primary list.

These cuts are necessary to remove primaries close to the sample boundaries, either on the sky or in redshift. These galaxies may have more massive companions that lie just outside the field or beyond our redshift cut. Thus, there is a higher probability that they are not actual CGs, but are in fact satellites of another, more massive galaxy.

Overall, this process is very conservative in selecting primaries, producing a sample of 1,490 galaxies that is incomplete (in the sense of missing many genuine CGs), but relatively uncontaminated by satellites. Cutting out systems with redshifts exceeding 0.25 reduces the number of first-run primaries to 873, while excluding those with area completeness less than 0.65 (i.e. those with ROIs that are masked or cut off by field boundaries by more than 35%), reduces the number to 815.

3.4.2 Clustering of the First-run Primary and Secondary Samples

To study satellite abundance and its dependence on primary properties, we first need to model and separate the contributions to the clustering signal from the satellite population and the background galaxy population. Having classified potential primaries and secondaries, we measure the surface number density of secondaries within the ROI and in an extended region around it (with the same velocity offset limits, but extended out to $3.2 R_{\text{vir}}$ in order to have a better estimate of the background surface density).

Around each primary, we count secondaries in annuli spaced evenly in $\log[R_P/R_{\text{vir}}]$. The annuli range from $0.1-3.2 R_{\text{vir}}$, in steps of 0.25 dex. We exclude secondaries with projected separations of less than $0.1 R_{\text{vir}}$ (~ 25 kpc, for the Milky Way) to avoid outlying HII regions or other components of the central galaxy. We calculate the surface density of secondary galaxies around each primary in physical units (Mpc^{-2}), assuming all secondaries lie at the same distance

as the primary. Given counts $N^{i,j}$ in radial bin i of projected area A_i around primary j , the surface density is:

$$\Sigma_{i,j} = \frac{N^{i,j}}{\eta_i A_i} = \frac{N^{i,j}}{\eta_i 2\pi R^i \Delta R} \quad (3.2)$$

where η_i is the mean area completeness in radial bin i (estimated as described in Appendix C.2), R^i is the mean radius of the bin, and ΔR is the width of the bin. As we want to scale the surface density of secondary galaxies in units of R_{vir} , it is useful to define a separation variable $X_i \equiv R_P^{i,j} / R_{\text{vir}}^j$. Thus, Eqn. 3.2 can be written as:

$$\Sigma_{i,j} = \frac{N^{i,j}}{\eta_i 2\pi (X_i \cdot R_{\text{vir}}^j) \Delta X \cdot R_{\text{vir}}^j} = \frac{N^{i,j}}{\eta_i 2\pi X_i \Delta X \left(R_{\text{vir}}^j\right)^2}. \quad (3.3)$$

In what follows, we will fit the surface density in bins of primary redshift and mass. Where the primary sample contains more than 5 objects, we use the bootstrap method to estimate the uncertainties in the surface densities, by re-sampling the primary sample 120 times. As the bootstrap method does not work well when the sample size is smaller than 5, we have also calculated Poisson uncertainties for each bin. The final uncertainties for the bin are taken to be the larger of the two values.

The secondary surface density consists of two main parts: the contribution from clustered satellites and the contribution from background or foreground galaxies.

$$\Sigma(X|M_{\text{halo}}, z) = \Sigma_{\text{sat}}(X|M_{\text{halo}}) + \Sigma_{\text{bg}} \quad (3.4)$$

The first component Σ_{sat} should correlate with the halo mass of the primary, but should be roughly independent of redshift over the narrow redshift range considered here, while the second component should be roughly independent of halo mass, but should depend on redshift. There should also be a more extended clustered component due to large-scale structure (the "two-halo" term, e.g. Cooray & Sheth (2002) for a detailed review), but the characteristic scale of this component (cf. 4–8 Mpc) is much larger than the scales considered here. Thus, we treat it as a constant with respect to radius, and include it in the foreground/background term. To fit the two terms, we split the primaries into 5 fixed redshift slices and 2–5 halo mass bins per slice, with adaptive boundaries as shown in Fig. 3.3.

We assume that the satellite distribution roughly matches the subhalo distribution, which in turn approximately follows a Navarro-Frenk-White (NFW – Navarro, Frenk & White (1996)) density profile. Over the range of radii we are most sensitive to, a projected NFW profile scales as r^{-2} in the outer parts of the halo, and is somewhat shallower in the inner parts. We could fit the density profile of the satellite component with the exact form of a projected NFW profile,




Figure 3.3: The redshift-halo mass distribution of the first-run primary sample. The coloured boxes show the boundaries for the 20 subsets used to fit the mass and redshift dependence of the clustering signal.

assuming some mean concentration-mass relation. Given the low SNR of the satellite component, however, we choose to fit it with the simplified form:

$$\Sigma_{\text{sat}}(X) = S_{\text{halo}}F(X) = \frac{S_{\text{halo}}}{X^2 + \alpha X + \beta} \quad (3.5)$$

where S_{halo} is an overall normalization (in units of Mpc^{-2}) that is independent of X but depends on the primary halo mass, while α and β are parameters describing the radial dependence. We have compared this simplified form to a projected NFW profile over the range $X = 0.05\text{--}3$, for concentration parameters $c = 5\text{--}15$, and found that using fixed values $\alpha = 0.2$ and $\beta = 0$ gives a good fit in all cases. For this choice of parameters, $\int_0^1 2\pi F(X)X dX = \log[1+\alpha] - \log[\alpha] = 11.18$, so the total number of satellites within the virial radius is $N_{\text{sat}} = 11.18 S_{\text{halo}} R_{\text{vir}}^2$.

The resulting fits for the secondary surface density (fitting S_{halo} and Σ_{bg} jointly for each individual primary bin, with α, β set to fixed values (0.2, 0)) are shown in Fig. 3.4. Summing over all primary bins, a clustered excess in Σ is detected at a SNR of approximately 11. We can see that the overall background surface density (as measured in physical units) decreases with primary redshift, while the clustered satellite component increases with primary mass. In the next section, we will explore these correlations in more detail.

Satellite/halo component

We fit the halo mass dependence of S_{halo} over the 20 primary bins using a linear relation in log-log space:

$$\log_{10} \left(\frac{S_{\text{halo}}}{1\text{Mpc}^{-2}} \right) = a \cdot \log_{10} \left(\frac{M_{\text{halo}}}{10^{12}M_{\odot}} \right) + b \quad (3.6)$$

where a and b are free parameters. We choose $10^{12}M_{\odot}$ as the pivot mass in our fit, as this is roughly the median halo mass of our 20 primary bins. The projected area of a halo will scale as the virial radius squared, that is as $M^{2/3}$. If we assume systems have a fixed number of satellites per unit halo mass (as expected from subhalo abundance, e.g. Gao et al. 2004), then the projected surface density should go as $M^{1/3}$, so we expect $a \sim 0.33$. The value of b (the normalization at $M_{\text{halo}} = 10^{12}M_{\odot}$) will depend on the depth of the catalogue, as discussed below. From our fits, we find $a = 0.30^{+0.11}_{-0.10}$ and $b = 0.26^{+0.08}_{-0.11}$, so the scaling with halo mass seems fairly consistent with the expected value.

Fig. 3.5 shows our fit for S_{halo} as a function of mean halo mass, over all 20 bins in primary mass and redshift. Given its sensitivity to smaller radial bins with larger errors, the fitted value of the parameter has a SNR ~ 8 , significantly lower than the SNR for the whole clustering signal. To

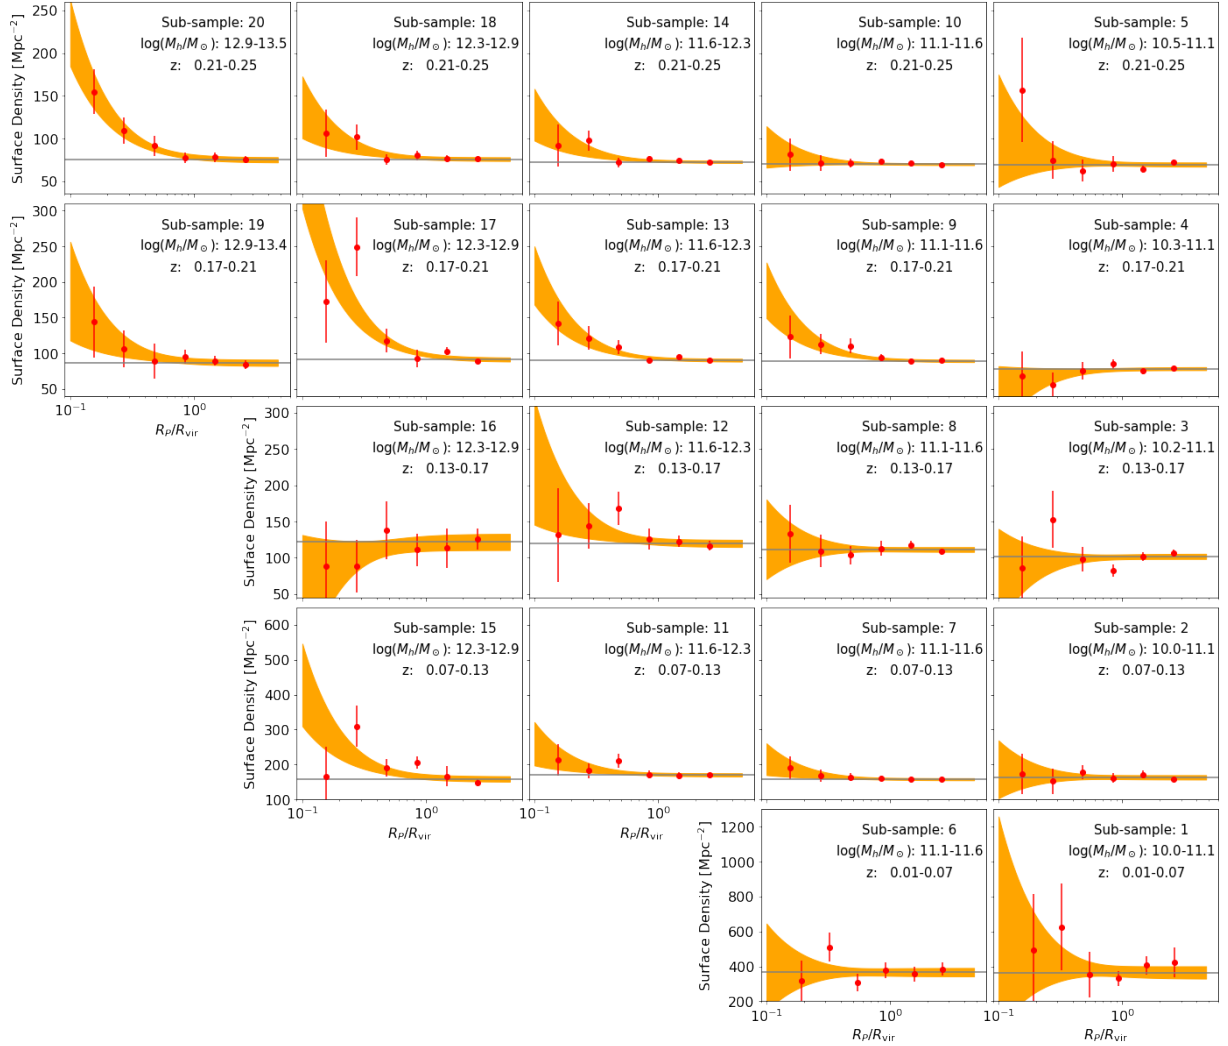


Figure 3.4: The surface density of secondaries around each of the primary subsamples shown in Fig. 3.3 (separate panels, with redshift increasing from bottom to top and halo mass increasing from right to left). Red points with error bars show the measured surface number density in radial bins, while the orange shading shows the $1-\sigma$ region around the best fit from Eq. 3.5. The horizontal grey lines indicate the background level in the best fit model.




Figure 3.5: Average S_{halo} versus halo mass ($\log_{10}(M_h/M_\odot)$), with colour indicating the mean redshift of the primary sub-sample. The blue curve and grey shaded area show the best-fit model of the mass dependence (Eqn. 3.6), together with the $1\text{-}\sigma$ uncertainty range.

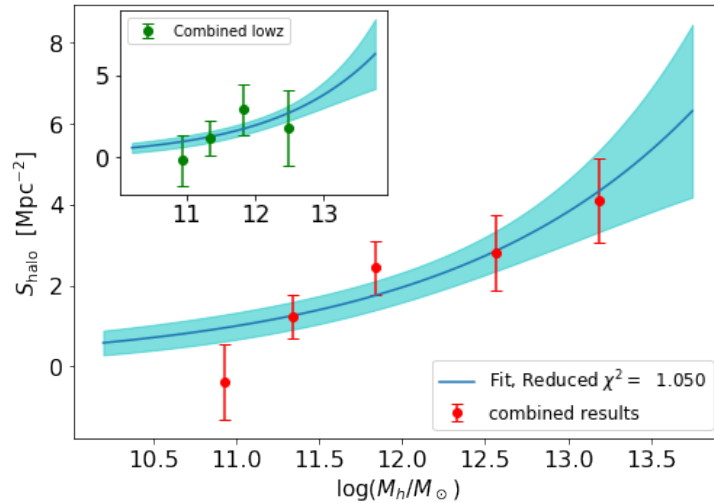


Figure 3.6: As Fig. 3.5, but with sub-samples of similar mean halo mass combined into single bins for clarity. The top-left subplot shows the results for primaries with redshift $z < 0.15$.

illustrate the dependence on halo mass more clearly, we also show in Fig. 3.6 a version combining bins with similar mean halo masses. We note that at the low-mass end ($\log(M_h/M_\odot) < 11$), the fitted value is actually negative (although consistent with zero, given the uncertainties). This may be partly due to completeness problems at low mass, which affect the high redshift bins in particular. Repeating the fitting process for low redshift primaries ($z < 0.15$) only, we obtain a less negative value, that is once again consistent with zero. (as shown in the sub-panel of Fig. 3.6).

Background component

Our surface densities are calculated in physical units (Mpc^{-2}) at the distance of the primary. Thus, if the foreground/background component consisted of a fixed field population with a broad redshift distribution and thus a fixed number per square degree, we would expect its inferred physical surface density to scale as d_A^{-2} , where d_A is the angular diameter distance of the primary. (It is worth noting that this assumption will not necessarily hold if extending the method to higher redshift.) Fig. 3.7 shows the fitted value of the background surface density Σ_{bg} in each bin, versus the average value of d_A^{-2} for that bin. We fit the trend with a simple linear model:

$$\Sigma_{\text{bg}} = c + d \cdot \left[\left(\frac{d_A}{10^3 \text{Mpc}} \right)^{-2} - 2.0 \right] \quad (3.7)$$

where c and d are free parameters.

Fitting the 20 bins gives tight constraints on the parameters: $c = 79.5_{-0.6}^{+0.6}$ and $d = 21.6_{-0.5}^{+0.5}$. Note that if the background scaled exactly as the inverse of the angular diameter distance, we would expect the constant term $c - 2d \sim 0$ to be small; in practice, various minor effects, notably the varying width of the redshift range Δz over which we measure the background, will cause the background density to deviate from the simple scaling. As it is, for our fitted values of the parameters we find $c - 2d = 36.3$, which is small relative to the typical values of Σ_{bg} .

Single-step Fit

While the two-step fitting procedure outlined above is useful to illustrate the features the model, it is more robust to fit the entire 4-parameter model for both terms in the surface density in a single step, given the potential covariance between the model parameters. We use the function `Minimizer.emcee`² from the Python module LMFIT (Least-Squares Minimization Fitting) to do MCMC sampling of the likelihood (Foreman-Mackey et al., 2013). The marginalized results of

²See [this page](#) for a detailed description of LMFIT and `Minimizer.emcee`




Figure 3.7: The background surface density Σ_{bg} versus d_A^{-2} . The grey area shows the $1\text{-}\sigma$ uncertainty range around the best-fit model (Eqn. 3.7).

this fit are shown in Fig. 3.8 and are also included in Table 3.4.3. We note that the parameters a and b are strongly (anti-)correlated; indeed, with higher SNR data we could imagine a more detailed, HOD-based fit to the halo-mass dependence of the satellite abundance. There is also some correlation between parameters c and b (or to a lesser degree c and a), indicating that satellite abundance estimates do require careful accounting for the background term.

Assigning Probability

Given our model fit to the clustering measurements, we can estimate the amplitude and radial distribution of the satellite component and background components around each primary. We define the probability of a secondary galaxy in the ROI being an actual satellite as:

$$P_{i,j}^{\text{sat}}(X, M_h, z^p) = \frac{\Sigma_{\text{sat}}}{\Sigma_{\text{tot}}} = \frac{\Sigma_{\text{sat}}(X, M_h)}{\Sigma_{\text{sat}}(X, M_h) + \Sigma_{\text{bg}}(z^p)}, \quad (3.8)$$

where M_h and z^p are the halo mass and redshift of the primary, respectively.

This equation can result in very small probabilities at large radii. Since real satellites (objects that have crossed the virial radius at least once) almost all lie within $3 R_{\text{vir}}$ (e.g. [Wetzel et al.](#),

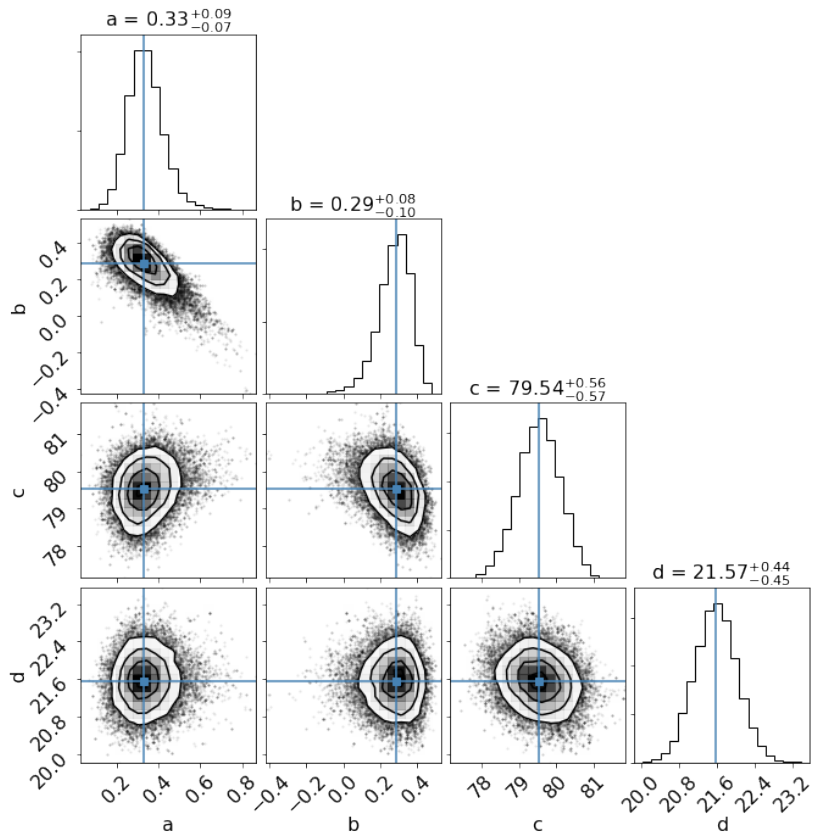


Figure 3.8: Likelihood distributions for the clustering model parameter values, derived by fitting the full model to all 20 primary sub-samples simultaneously. Panels show the distribution marginalized over the two (or, on the diagonal, all three) other model parameters.

2014), we truncate the probability around $X = 3$ as follows:

$$P'(X) = P(X) \cdot \frac{1}{1 + 1000^{X-3}}. \quad (3.9)$$

3.4.3 Iterating over the Fit

Our initial primary selection ignores many galaxies that may well be primaries, but appear close to more massive systems when seen in projection. To attempt to correct for this, we iterate over our clustering measurement, including a weighted contribution from all galaxies, proportional to their probability of being a primary.

For each iteration, we run the top-down selection again. During the new selection process, every galaxy starts with a 100% probability of being a primary. Starting with the most massive galaxy as the first primary, we assign probabilities of nearby galaxies being its satellites, using the method described in the last section. The probability of each of these galaxies being independent primaries is reduced accordingly. We then proceed through the catalogue in order of decreasing stellar mass. If a galaxy has a probability of being a primary between zero and 1, we estimate that nearby galaxies have a probability of being its satellites that is the *product* of its probability of being a primary and the satellite probability given in section 3.4.2. Proceeding through the entire catalogue in decreasing order of stellar mass ensures that the calculation is well-defined and that every galaxy is assigned a final probability of being a primary, equal to 100% minus the sum of all probabilities that it is a satellite of nearby systems.

Running through this process once, we find 8,920 primaries with probabilities greater than 0.99, of which 3,567 lie at redshifts of 0.25 or less. If we further remove systems with areal completeness of 0.65 or less, the number of high-confidence primaries is 3,246, versus 815 in our original sample; repeating this exercise for systems with primary probabilities of 0.999 or more, reduces the number to 1,478.

If we include primaries with probabilities greater than 0.99 (and weight all satellites by their CGs primary probability), the effect of iteration on the fit to the clustering signal is shown in Table 3.4.3. Overall, the parameter values after iteration show good consistency with our initial estimates, shifting by less than $2\text{-}\sigma$ in all cases. The uncertainties in the fitted parameter values drop, but only slightly. This suggests that in a dataset like the COSMOS catalogue that has extremely precise photometric redshifts, even the first-run sample of isolated primaries can provide a good estimate of satellite abundance. Since the use of lower probability primaries may dilute the clustering signal and introduce some bias, in what follows we will use our initial, first-run estimates of the fitted parameters to derive satellite abundance. We anticipate that in

Table 3.1: The model fitting results from the first run and first iteration

	a	b	c	d
F0	$0.33^{+0.09}_{-0.07}$	$0.29^{+0.08}_{-0.10}$	$79.5^{+0.6}_{-0.6}$	$21.6^{+0.4}_{-0.5}$
F1	$0.26^{+0.07}_{-0.07}$	$0.38^{+0.05}_{-0.06}$	$79.3^{+0.5}_{-0.5}$	$20.8^{+0.4}_{-0.4}$

datasets with less accurate redshifts, iteration will become more important in deriving accurate estimates of the clustering signal.

3.5 Estimating Satellite Abundance

In this section, we will make some basic estimates of the overall abundance of satellites, as well as their abundance as a function of properties such as stellar mass or luminosity. In each case, our estimate is based on the clustering signal, which is typically small compared to the background. The simplest way to estimate satellite abundance is to count every galaxy in the ROI, weighted by the satellite probability calculated in Section 3.4.2, so we will use this approach first in Section 3.5.1, referring to it as “method A”.

As explained below, method A assumes that the clustering amplitude is uncorrelated with stellar mass, luminosity, or any of the other secondary properties considered. More generally, we expect the fraction of galaxies in the ROI that are true satellites to depend on these other properties. In section 3.5.2 we develop a more sophisticated approach, “method B”, similar to the one introduced in Speller & Taylor (2014), that attempts to correct for possible correlations in the limit of a weak clustering signal. Future surveys with stronger detections of clustering should be able to bypass these complications by dividing the galaxy sample directly into bins of secondary property value before they measure the clustering amplitude, simplifying the analysis considerably; we call this “method C”.

3.5.1 Abundance Estimates Using Method A

In method A, to estimate satellite abundance N_{sat}^i around primary galaxy i we simply add up the probabilities P_j^i of each galaxy j in its ROI being a true satellite:

$$N_{\text{sat}}^i = \sum_j P_j^i. \quad (3.10)$$

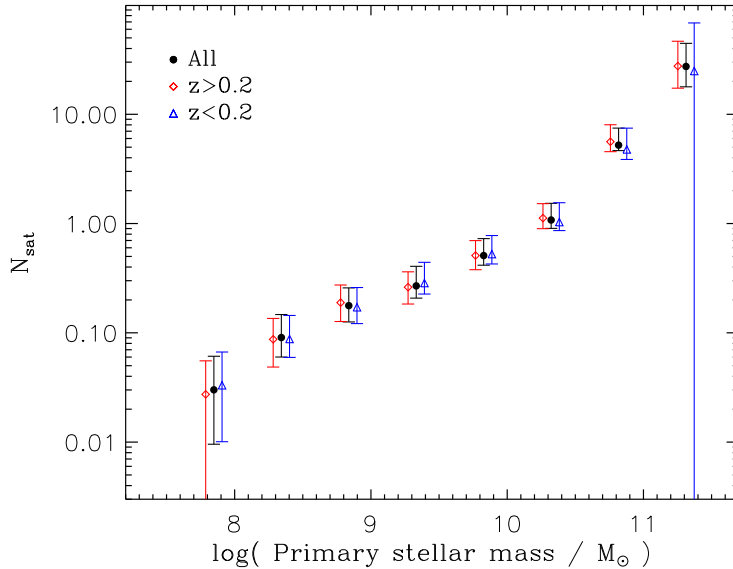


Figure 3.9: Satellite abundance as a function of primary stellar mass, estimated using method A. Red and blue points show the results for independent primary samples at higher redshifts ($z = 0.17\text{--}0.25$) and lower redshifts ($z = 0.07\text{--}0.17$), respectively.

Note that although P_j^i can remain non-zero at large distances from the primary, to compare to previous results from literature we set a radial limit of $1.5 R_{\text{vir}}$ by default, and only count towards the total satellite abundance secondaries that lie within this projected separation of the primary.

We can study the dependence of satellite abundance on primary mass by stacking systems with similar stellar or halo masses, as shown in Fig. 3.9 (black points with error bars). As expected, there is a strong trend in satellite abundance with halo mass. As a consistency check, we also split our primary sample in two by redshift, and calculate satellite abundance separately for each of the two sub-samples (red and blue points). The results for both sub-samples show good consistency with those for the whole sample. We note that N_{sat} increases faster with primary stellar mass at the high mass end of the range ($M_h > 10^{10.5} M_{\odot}$). This is consistent with the pattern seen in Halo Occupation Distribution (HOD) modelling (e.g. Seljak, 2000; Peacock & Smith, 2000; Berlind & Weinberg, 2002) and is a result of the changing slope of the SHMR; halo mass increases faster with stellar mass at large stellar masses, resulting in a faster increase in satellite numbers.

3.5.2 Abundance Estimates Using Method B

Method A provides the simplest estimate of satellite abundance, and is guaranteed to be correct when averaging over all galaxies in the secondary sample used to calculate the satellite probability (Eqns. 3.8 & 3.9). It is not necessarily correct, however, for subsamples of secondaries selected by luminosity, stellar mass, colour, or other properties, if these properties are correlated with the clustering amplitude. As a simple example, one can imagine a galaxy population that consisted of two distinct types, labelled “red” and “blue”. If the red galaxies were completely clustered, but the blue galaxies were completely unclustered, then we would measure some intermediate average clustering strength for the combined population, and give every galaxy a satellite probability based on this average value. If we weighted all galaxies by this average satellite probability, but then split them back into subsamples by colour, we would conclude that the satellite and background populations both had the same net colour distribution. In effect, the true colour distribution of the satellite population (100% red, in this example) would be contaminated by the colour distribution of the field population. More generally, whenever the background population is significantly different from the satellite population, the satellite properties inferred using method A will be biased towards those of the background population.

With sufficient SNR in the clustering signal, we could avoid this problem by selecting subsamples with a limited range of the desired satellite property (stellar mass, luminosity, colour, etc.) *before* measuring the clustering amplitude and calculating the satellite probability (we refer to this as “method C”). In the limit of low SNR, however, splitting the galaxies into narrow bins in a given property will increase the shot noise in the background estimate until it is unacceptably large. Instead, we have developed an intermediate solution, “method B”, based on the approach in [Speller & Taylor \(2014\)](#).

In method B, first the number of the background galaxies within a given radius is estimated for each primary, by summing up the non-satellite probabilities of each pair:

$$N_{\text{bg}}^i = \sum_j (1 - P_j^i). \quad (3.11)$$

We then measure the fraction of *all* background galaxies (i.e. summing over the ROIs of all primaries, or over the whole field) with a property of interest (e.g. stellar mass, luminosity, colour, etc.) in a given range, and scale the total number of background galaxies in the ROI by this fraction. This gives the expected background contribution to a particular subsample, that we then subtract to calculate satellite abundance. For instance, if we want to measure the luminosity function of satellites, Φ_{sat} , we first need to measure the total luminosity function within the ROI, Φ_{TTL} , and the total luminosity function for all background galaxies, Φ_{bg} . Then we remove the

background contribution from the total abundance in each bin k in luminosity, such that:

$$\Phi_{\text{sat},k} = \Phi_{\text{TTL},k} - \frac{N_{\text{bg}}}{N_{\text{TTL}}} \Phi_{\text{bg},k} . \quad (3.12)$$

Here N_{bg} and N_{TTL} refer to the number of background galaxies, and the total number of galaxies within the radial cut around each primary, respectively (before any cut in luminosity). Note that this approach can be used for individual primaries, except for the lowest-mass systems, where the galaxy counts are so small that Poisson fluctuations dominate. To correct for this, if the number of objects within the ROI is less than three times the number of luminosity bins, then we stack results for multiple primaries at similar redshifts, and use the average signal.

We used this method to calculate satellite abundance for different luminosity ranges (below), as well as full satellite luminosity functions (see Sec. 3.5.4). Fig. 3.10 shows abundance for various cuts in M_{i^+} . Overall, the dependence of abundance on primary stellar (top panel) or halo (bottom panel) mass has a similar form for different luminosity cuts, although there may be a truncation at lower stellar masses that depends on the luminosity cut. Here too, this pattern is very similar to those seen for brighter galaxies in HOD modelling. Plotted as a function of stellar mass, satellite abundance shows a change in slope around $10.5 M_{\odot}$ at all luminosities. Plotting as a function of halo mass, this feature disappears, confirming that it is a result of the non-linear SHMR. We also indicate on the plot the abundance of MW satellites brighter than $V = -14.5$ with a galactocentric distance greater than 20 kpc (brown diamond – three satellites meet these criteria), assuming a MW halo mass of $12.1 M_{\odot}$, as discussed below.

3.5.3 Comparison to Previous Work

For massive galaxies, a number of other estimates of satellite abundance exist in the literature. In Fig. 3.11, we compare our abundance estimates to the results of Conroy, Wechsler & Kravtsov (2006, C06). These are based on HOD modelling (e.g. Seljak, 2000; Peacock & Smith, 2000; Zehavi et al., 2002) of the luminosity functions and correlation functions of samples from the SDSS (York et al., 2000) and DEEP2 (Newman et al., 2013) surveys. HOD modelling provides an estimate of the average number $\langle N_{\text{gal}} \rangle$ of galaxies within a halo of a given mass, so in the limit where $\langle N_{\text{gal}} \rangle \gg 1$, $\langle N_{\text{gal}} \rangle - 1$ should match our measured satellite abundance. As $\langle N_{\text{gal}} \rangle$ decreases, some halos will contain no galaxies over a given magnitude limit, so we can only compare results in the large $\langle N_{\text{gal}} \rangle$ regime, i.e. for large halo masses. In addition, the results in C06 are given in bins of $M_r - 5 \log(h)$; we convert to our i^+ -band assuming a fixed mean colour $\langle M_r - M_{i^+} \rangle \sim 0.25$ (roughly the value measured for our sample in the COSMOS catalogue), which produces a shift of $(M_r - M_{i^+}) - 5 \log(h) \sim 1$ magnitude exactly in the C06 bin boundaries.




Figure 3.10: Satellite abundance as a function of primary stellar/halo mass (top/bottom panels respectively), estimated using method B. The individual shaded regions show results for different cuts in satellite luminosity, as indicated in the legend. The brown diamond shows the number of MW satellites brighter than $V = -14.5$ with a galactocentric distance greater than 20 kpc (three satellites meet these criteria).

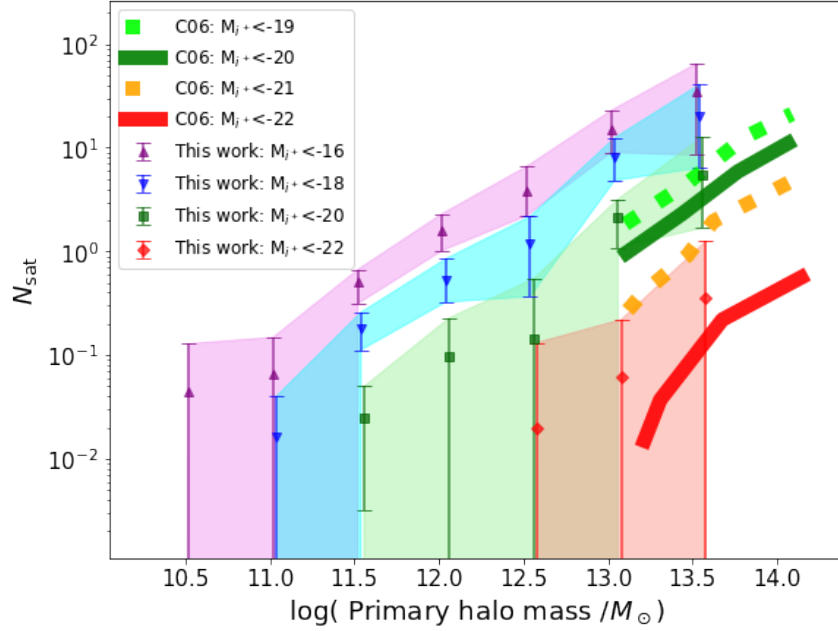


Figure 3.11: Comparison to the HOD-based results of [Conroy, Wechsler & Kravtsov \(2006, Fig. 5\)](#). Their luminosity bins have been converted to i^+ -band magnitudes, assuming a mean colour of $\langle M_r - M_{i^+} \rangle = 0.25$, and have had 1 subtracted from them to account for the central galaxy.

Given these conversions, examining [Fig. 3.11](#), we see that there is excellent agreement between our results and those of C06. Our estimated satellite abundance matches that measured by C06 to within half a standard deviation, for all four of the magnitude cuts in that study (two lie between our luminosity cuts, but are clearly consistent with our results.) The slope of the $N_{\text{sat}}-M_h$ relation is harder to judge given the limited baseline in C06, but generally it appears to be consistent with our inferred slope at halo masses of $\log(M/M_\odot) > 13$ or more. The agreement between these two sets of results is particularly striking, given that they employ completely different samples and methods, and that there is no parametric freedom in adjusting our results. Overall, SDSS provides a more robust estimate of satellite abundance for massive halos, but as a deeper survey with more accurate redshift estimates, COSMOS is better able to probe the low-halo-mass regime.

We have also compared our results to more recent work by [Besla et al. \(2018, B18\)](#), which is one of the few studies to estimate satellite abundance at lower halo mass. [Fig. 3.12](#) compares our results to theirs, for the primary stellar mass range $10^8 - 10^9 M_\odot$. The B18 results are based on a SDSS spectroscopic sample at redshift 0.013–0.0252, with r -band magnitudes between 14 and




Figure 3.12: Comparison to the observed and mock results of Besla et al. (2018, ; see their Fig. 11,12). (Note that the points are shifted to left or right by up to 0.02 dex for clarity).

17.77, and primary stellar masses in the range $0.2\text{--}5 \times 10^9 M_\odot$. (The apparent magnitude limit of 17.77 corresponds to absolute magnitude limits of -16.03 or -17.5 at $z=0.013$ or $z=0.0252$ respectively, equivalent to -16.3 and -17.8 in the i^+ -band.) They also compare these to a mock catalogue generated using the Illustris hydrodynamical simulations (Vogelsberger et al., 2014; Nelson et al., 2015). Four sets of results are shown. The "uncorrected SDSS" results are from raw counts of nearby companions; the "completeness corrected" version is after correcting for observational selection effects, using the mock catalogues. The "physical" simulation results show the abundance of real satellites, while the "projected" counts show the result including a background contribution introduced by projection effects. Completeness and projection corrections move the measured abundance up and down, respectively; overall, the best estimate of satellite abundance from B18 is the completeness corrected (orange) curve, decreased by a factor of $\sim 50\%$ to account for projection effects.

Over the mass range $0.2\text{--}5 \times 10^9 M_\odot$ covered by B18, our estimates for magnitude cuts $M_{i^+} < -16$ or $M_{i^+} < -18$ are consistent with their measured values. Given their effective magnitude limits range from -16.3 and -17.8 in the i^+ -band, there once again seems to be excellent consistency in the overall abundance estimated by the two methods. In contrast, B18

measure almost no trend in satellite abundance with primary stellar mass. It is worth pointing out, however, that the mean redshift of the sample in B18 is lower at smaller primary masses (see their Fig. 3). Thus, the satellite luminosity function may be measured to greater depth for these systems, explaining the flatter slope. B18 also point out that they may be biased towards preferentially identifying multiple systems at low stellar mass, due to their bluer colours. We conclude that our results are consistent with B18, once again despite very different methods and samples.

3.5.4 The Satellite Luminosity Function

Finally, we can use method B to estimate satellite luminosity functions directly (for comparison, results using method A are shown in Appendix C.3). In Fig. 3.13 we present the luminosity function of satellites for three subsets of primaries. The subsets were chosen such that the mean halo mass of the middle bin, $\langle M_h \rangle = 12.1 M_\odot$, is close to the estimated mass of the Milky Way (MW) or M31 (Bland-Hawthorn & Gerhard, 2016; Posti & Helmi, 2019), such that we can compare to the observed luminosity functions for these systems. For method B, we are able to measure the satellite luminosity functions reliably down to absolute magnitudes of -14. Within this magnitude range, the observed abundance of satellites around the MW and M31 are close to the average value. One exception is at the bright end of the MW satellite luminosity function, where the presence of the LMC and SMC represent a slight ($1-2 \sigma$) departure from the average. This unusual feature of the MW's satellite population has been noted and discussed extensively elsewhere (e.g. Robotham et al. 2012 – see Speller & Taylor 2014 for further references). We will consider satellite luminosity and stellar mass functions in more detail in the next chapter.

3.6 Testing for Systematic Uncertainties in the Method

While we have shown that our method produces estimates of satellite abundance consistent with previous studies using larger samples, there remain a number of choices, assumptions or free parameters in the method that could take on different values. In this section, we will perform a set of tests to understand the effects of the various assumptions and free parameters in the clustering method.




Figure 3.13: Satellite luminosity functions, for three ranges of primary halo mass, estimated using method B (see Appendix C.3 for the same results derived using method A). The mean halo mass for each range is given in parentheses. The observed luminosity functions of the MW and M31 are shown for comparison (black and purple points and lines). These magnitudes have been converted from the V -band, assuming a fixed average colour of $V - i^+ = 0.3$.




Figure 3.14: Results of a null test where the halo component of the clustering signal is measured using galaxies outside the velocity limits of the ROI. A black dashed line corresponding to no clustering is included for comparison.

3.6.1 Null Test

First, as a null test, we calculated the clustering signals around our initial sample of primaries, using only secondaries that lay within the radial cuts, but outside the redshift cuts we defined in Section 3.3.2. Following the procedure in Section 3.4, we measured the surface number density of the secondaries for the 20 bins in primary mass and redshift, and re-fit our surface density model. The satellite component from the fit is shown in Fig. 3.14.

Most of secondaries in this test should not be physically associated with the primaries, expect for a small number of real satellites that are scattered into the foreground or background by the redshift errors. Thus, we expect the clustering signal to be close to zero, and to show little dependence on the primary halo mass. This is confirmed in Fig. 3.14.

3.6.2 Parameter and Systematic Tests

We have tested for systematic effects and uncertainties in our method by varying the parameters that define the initial primary selection and the calculation of the satellite probability. The tests

include:

- Using radial cuts of $2 R_{\text{vir}}$ or $3.5 R_{\text{vir}}$ to define the radial extent of the ROI.
- Using velocity cuts of $1 \sigma_{\Delta v}$ or $3 \sigma_{\Delta v}$ to define the velocity extent of the ROI.
- Increasing/decreasing the stellar mass by 0.16 dex, which is comparable to or larger than the typical stellar mass uncertainties in our data.
- Increasing/decreasing the halo mass derived from our fiducial SHMR by 50%.
- Varying the slope of bias-corrected SHMR (see Appendix C.1) at the high mass end to 1.5 and to 2.5, with respect the original value ~ 2.1 from B13.
- Varying the definition of the virial radius, increasing or decreasing it by 20%.
- Adding 1σ scatter to the initial stellar masses before ranking them. (We perform this test three times to check the consistency of the potential effects.)
- Keeping the primary-secondary selection fixed, but adding $0.2 R_{\text{vir}}$ scatter to the coordinates of the primaries, to test the effects on potential mis-centering. (We repeat this test three times to check for the consistency of the effects.)

In each case, all other parameters and steps in the method are kept fixed. The results of these tests are summarized in Tab. 3.2.

3.6.3 Discussion

Considering the results of the individual tests in detail, changing the halo mass mainly just shifts the points horizontally on the $S_{\text{halo}}-M_h$ plot, so the fitted parameters a and b remain relatively constant. Changing the stellar mass has quite a different effect, however. As the HSMR is shallow at the low-mass end but steep at the high-mass end, increasing or decreasing the stellar mass does not change the inferred halo mass much at the low-mass end, but can produce significant change at the high-mass end. As a result, the slope (a) of the $S_{\text{halo}}-M_h$ relation is more strongly affected.

Similarly, changing the slope of the HSMR at the high-mass end will mainly affect the halo mass estimates in this range. As a result, the slope a is shifted systematically to higher or lower values. As for the mass ranking test, in addition to the random scatter in individual mass estimates and resulting variations in detailed primary selection, there is a net change in the mass function. Since there are more low-mass galaxies than high-mass ones, adding the random scatter tends to

Table 3.2: The fitting results of all tests

	a	b	c	d
F0	$0.33^{+0.09}_{-0.08}$	$0.29^{+0.08}_{-0.10}$	$79.5^{+0.6}_{-0.6}$	$21.6^{+0.5}_{-0.5}$
LMh	$0.29^{+0.08}_{-0.07}$	$0.24^{+0.09}_{-0.10}$	$79.0^{+0.5}_{-0.5}$	$21.7^{+0.5}_{-0.5}$
SMh	$0.28^{+0.08}_{-0.07}$	$0.37^{+0.06}_{-0.08}$	$79.0^{+0.5}_{-0.6}$	$21.6^{+0.5}_{-0.5}$
LMs	$0.30^{+0.09}_{-0.08}$	$0.12^{+0.09}_{-0.12}$	$78.8^{+0.5}_{-0.5}$	$21.7^{+0.5}_{-0.5}$
SMs	$0.45^{+0.12}_{-0.10}$	$0.33^{+0.08}_{-0.09}$	$80.2^{+0.6}_{-0.6}$	$22.0^{+0.5}_{-0.5}$
$\sigma_{\Delta v}$ 1	$0.32^{+0.08}_{-0.07}$	$0.22^{+0.07}_{-0.09}$	$78.2^{+0.5}_{-0.5}$	$21.2^{+0.4}_{-0.4}$
$\sigma_{\Delta v}$ 3	$0.21^{+0.10}_{-0.09}$	$0.29^{+0.08}_{-0.10}$	$79.9^{+0.7}_{-0.7}$	$21.1^{+0.6}_{-0.5}$
R2	$0.30^{+0.06}_{-0.06}$	$0.32^{+0.06}_{-0.06}$	$78.7^{+0.5}_{-0.5}$	$21.4^{+0.4}_{-0.4}$
R3.5	$0.23^{+0.09}_{-0.08}$	$0.28^{+0.06}_{-0.08}$	$82.2^{+0.5}_{-0.5}$	$21.0^{+0.4}_{-0.4}$
LR	$0.28^{+0.09}_{-0.08}$	$0.14^{+0.08}_{-0.10}$	$78.1^{+0.5}_{-0.5}$	$21.2^{+0.4}_{-0.4}$
SR	$0.26^{+0.10}_{-0.09}$	$0.28^{+0.09}_{-0.12}$	$80.7^{+0.7}_{-0.7}$	$22.1^{+0.5}_{-0.5}$
H SHMR	$0.20^{+0.06}_{-0.06}$	$0.31^{+0.05}_{-0.07}$	$80.2^{+0.5}_{-0.5}$	$21.3^{+0.3}_{-0.4}$
L SHMR	$0.34^{+0.10}_{-0.09}$	$0.28^{+0.07}_{-0.09}$	$79.4^{+0.6}_{-0.6}$	$21.6^{+0.5}_{-0.5}$
RS 1	$0.26^{+0.08}_{-0.07}$	$0.28^{+0.07}_{-0.09}$	$78.6^{+0.6}_{-0.5}$	$21.6^{+0.5}_{-0.5}$
RS 2	$0.25^{+0.08}_{-0.07}$	$0.29^{+0.07}_{-0.08}$	$78.7^{+0.6}_{-0.5}$	$21.6^{+0.5}_{-0.5}$
RS 3	$0.25^{+0.08}_{-0.07}$	$0.28^{+0.07}_{-0.09}$	$78.6^{+0.5}_{-0.5}$	$21.6^{+0.5}_{-0.6}$
CS 1	$0.34^{+0.12}_{-0.10}$	$0.05^{+0.11}_{-0.15}$	$79.4^{+0.5}_{-0.5}$	$21.8^{+0.5}_{-0.5}$
CS 2	$0.25^{+0.10}_{-0.08}$	$0.11^{+0.9}_{-0.12}$	$79.3^{+0.5}_{-0.5}$	$21.8^{+0.5}_{-0.5}$
CS 3	$0.26^{+0.08}_{-0.08}$	$0.21^{+0.08}_{-0.10}$	$79.6^{+0.6}_{-0.6}$	$21.6^{+0.5}_{-0.5}$

(Key to the tests:

F0 = original fit result;

LMh = larger halo mass; SMh = smaller halo mass;

LMs = larger stellar mass; SMs = smaller stellar mass;

$\sigma_{\Delta v}$ 1 = using $1\sigma_{\Delta v}$ velocity cut to define ROI; $\sigma_{\Delta v}$ 2 = using $3\sigma_{\Delta v}$ velocity cut for ROI;

R2 = using $2 R_{\text{vir}}$ projected separation cut for ROI; R3.5 = using $3.5 R_{\text{vir}}$ projected separation cut for ROI;

LR = larger R_{vir} (increasing the primary virial radius by 20%);

SR = smaller R_{vir} (decreasing the primary virial radius by 20%);

H SHMR = using a higher slope of HSMR at high mass end;

L SHMR = using a lower slope of HSMR at high mass end;

RS = ranking shuffle (adding $1\text{-}\sigma$ scatter to the primary stellar masses before ranking them);

CS = centering shift (adding $0.2 R_{\text{vir}}$ scatter to the coordinates of the primaries before measuring clustering).

)

increase the number of massive galaxies relative to the fiducial catalogue. This leads to slight shifts in the fitted satellite abundance, although they are less important than in the case of the mass tests.

The parameter values obtained in each test are given in Tab. 3.2. Tests that produce a variation of more than $2\text{-}\sigma$ in the fitted parameters are highlighted in bold. We note that only one test (our first re-centering test) produces a significant change in the parameters of the satellite component. Three of the 38 tests produce significant deviations in the background fit, but this is only slightly higher than the expected rate of $2\text{-}\sigma$ deviations given the random errors (8% versus 5%). Thus overall, the systematic uncertainties associated with our tests do not appear to significantly increase the random errors quoted in the fiducial model fit.

3.7 Conclusions

In this work, we have developed and tested a method for quantifying satellite abundance, using galaxy clustering. The method establishes a basic template for the radial dependence and amplitude of the satellite component of the clustering signal by using a subsample of isolated (or at least locally dominant) systems, but then applies this template iteratively to estimate the probability that any given galaxy in the field is a satellite of a nearby system. (Note that the form of the template assumes that the surface number density of background galaxies is inversely proportional to the square of angular-diameter distance; this assumption works well at low redshift, but may need modification if applying the method at higher redshift.) In that sense it is similar to crowded-field photometry, where an initial sample of isolated stars is used to determine the point spread function of the image, and that point spread function is then applied iteratively to the entire field. Using our method, we have estimated satellite abundance as a function of primary stellar and halo mass, and also measured the satellite luminosity function, over a very broad range of primary halo mass ($10^{10}\text{--}10^{13.5}M_{\odot}$). We have also tested the method for systematic uncertainties by varying the model parameters, and found variations in the final results that are generally smaller than our random error estimates.

We have compared the results of this new technique to several previous estimates of satellite abundance from the literature, that were derived using larger catalogues. Our results are fully consistent with those of [Conroy, Wechsler & Kravtsov \(2006\)](#) at the high mass end, and of [Besla et al. \(2018\)](#) at the low mass end, while covering a much larger range in primary mass overall. We have also compared our measured luminosity functions to those of the dominant LG galaxies, assuming an average halo mass of $12.1 M_{\odot}$ for these systems. The LG satellite populations seem fairly typical, with the exception of the bright satellites of the MW (the LMC and the SMC), as noted previously in the literature. The main purpose of this chapter was to describe and

validate our method; in subsequent work we will consider in more detail the properties of the detected satellite populations, including their spatial distribution, colours, star formation rates, and dependence on primary properties.

The COSMOS catalogue was chosen for this work for its deep photometry and extremely accurate photo-zs. Other deep surveys with accurate distance information will also be good candidates to apply our method. One potentially important survey is planned with SPHEREx³, an all-sky survey satellite with a wide-field spectral imager. SPHEREx is currently scheduled to launch in 2024, and will produce, during its two-year mission, four all-sky maps, with hundreds of millions of near-infrared stellar and galactic spectra (0.75–5.0 micron) (Bock & SPHEREx Science Team, 2018; Spangelo et al., 2015). The redshifts in the SPHEREx surveys are estimated by fitting template SEDs to observations, similarly to COSMOS. While COSMOS used photometry in 30 bands to derive its photo-zs, SPHEREx will produce low-resolution ($R \sim 20\text{--}100$) spectra, with a similar final redshift accuracy, as discussed in Stickley et al. (2016). While the main survey will be shallower than COSMOS, two regions at the polar caps will be visited multiple times, providing ~ 100 square degrees of coverage to a depth similar to COSMOS. Thus, SPHEREx should provide a redshift catalogue of similar accuracy to the COSMOS catalogue used here, but covering an area 50 times larger. The resulting increase in the SNR of the clustering signal would allow much finer binning in primary or secondary properties, giving a much more detailed view of the relationship between satellites and their central galaxies.

Acknowledgements

We thank N. Afshordi, A. Broderick, M. Balogh, M. Hudson, and our friends and collaborators from the COSMOS survey for their comments and advice. JET acknowledges support from the Natural Science and Engineering Research Council of Canada, through a Discovery Grant. The COSMOS 2015 catalog is based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under ESO programme ID 179.A-2005 and on data products produced by TERAPIX and the Cambridge Astronomy Survey Unit on behalf of the UltraVISTA consortium.

³see this <https://spherex.caltech.edu/> for more details

3.8 Data Availability

Most of the data underlying this article are publicly available. The COSMOS 2015 catalogue (Laigle et al., 2016) can be accessed from the COSMOS website, at <http://cosmos.astro.caltech.edu/page/photom>. A few spectroscopic redshifts that are unpublished from the COSMOS collaboration (M. Salvato, private communication) will be shared on reasonable request to the corresponding author with permission of the COSMOS collaboration. The derived data generated in this research will also be shared on reasonable request to the corresponding author.

Chapter 4

The hierarchical clustering method: abundance and properties of local satellite populations

Abstract

The faint satellites of the nearby Universe provide an important benchmark for our understanding of structure formation and galaxy formation, but satellite populations are hard to identify beyond the Local Group. We recently developed an iterative method to quantify satellite abundance using galaxy clustering, and tested it on a local sample in the COSMOS field, where accurate photometric redshifts are available for a large number of faint objects. In this chapter, we consider the properties of these satellite populations in more detail, studying the satellite stellar mass function (SSMF), the satellite-central connection, and quenching as a function of satellite and central mass and colour. Despite the limited sample size, our results show good consistency with those from much larger surveys, and extend measurements of the SSMF down to the lowest primary masses ($\log(M_h/M_\odot) \sim 10.2$, or $\log(M_*/M_\odot) \sim 8$) considered to date. We reproduce several known trends in satellite abundance and quenching, and find evidence for one new one, a dependence of the quiescent fraction on the primary-to-secondary halo mass ratio. We discuss the prospects for the clustering method in current and forthcoming surveys.

Key words: dark matter – galaxies: dwarf – galaxies: formation – galaxies: groups: general – galaxies: luminosity function, mass function – Local Group

4.1 Introduction

On large scales, the visible structure of the universe is clearly dominated by cold dark matter (CDM), or something that behaves very much like it. Given this framework, galaxies are predicted to form and grow in the centers of dense, virialized regions known as CDM halos. Halos themselves grow both by the gradual accretion of diffuse matter, and by stochastic mergers with other halos. During these mergers, smaller halos falling into larger ones will often survive as distinct, long-lived sub-structures within the final system. The central galaxies of these smaller halos will correspondingly survive as distinct “satellite” galaxies, orbiting the central galaxy of the larger halo. Dark matter substructure is intrinsically hard to detect directly, so satellite galaxies provide an important tracer of small-scale structure and halo assembly.

Galaxy formation is clearly strongly modulated on these small scales; a number of feedback mechanisms related to internal star formation (e.g. [Dekel & Silk, 1986](#); [Mashchenko, Wadsley & Couchman, 2008](#); [Governato et al., 2010](#); [Wetzel et al., 2016](#)), the photo-ionizing background (e.g. [Bullock, Kravtsov & Weinberg, 2000](#); [Gnedin & Kravtsov, 2006](#); [Lunnan et al., 2012](#); [Katz et al., 2020](#)), and/or environmental effects (e.g. [Taylor & Babul, 2001](#); [Mayer et al., 2006](#); [Lokas, Kazantzidis & Mayer, 2012](#); [Richings et al., 2020](#)) probably combine to greatly reduce the abundance of dwarf satellites, but it is possible dark matter structure itself is modified or suppressed on these scales (e.g. [Spergel & Steinhardt, 2000](#); [Macciò & Fontanot, 2010](#); [Anderhalden et al., 2013](#); [Kennedy et al., 2014](#); [Lovell et al., 2014](#); [Fry et al., 2015](#); [Elbert et al., 2015](#); [Nadler et al., 2019](#)). Satellite populations provide an important test of this rich array of physical processes.

The Local Group (LG), a composite system dominated by the Milky Way (MW) and the Andromeda Galaxy (M31), has the best-studied satellite populations. We can detect LG satellites down to much fainter magnitude and surface brightness limits than in any external group, and can resolve systems into individual stars, allowing detailed studies of their evolutionary history ([McConnachie, 2012](#)). It is not clear, however, that the LG is completely representative of the group environment or of satellite populations. Observational studies have shown that it is relatively rare for galaxies like the MW to have two bright star forming satellites like the Magellanic Clouds, for instance ([Guo et al., 2011](#); [Liu et al., 2011](#); [Strigari & Wechsler, 2012](#); [Robotham et al., 2012](#); [Speller & Taylor, 2014](#)), and numerical simulations have reached similar conclusions ([Boylan-Kolchin et al., 2010](#); [Busha et al., 2011](#); [Kang, Wang & Luo, 2016](#); [Zhang, Luo & Kang, 2019](#); [Evans et al., 2020](#)). Complete spectroscopic surveys (e.g. [Geha et al., 2017](#); [Mao et al., 2021](#)) are the ultimate tool for extending our understanding of satellite populations to other groups, but they can be very expensive in observing time, particularly for low-mass systems where genuine satellites are rare.

In previous work ([Speller & Taylor, 2014](#); [Xi et al., 2018](#), – chapter 2) we have explored the

use of galaxy clustering to quantify satellite populations around isolated primaries. In chapter 3 (published as [Xi & Taylor \(2021\)](#)) we developed and tested a new, hierarchical method for quantifying satellite abundance from galaxy clustering measurements (the ‘clustering method’ hereafter). The method uses the most obvious central galaxies, those that clearly dominate a volume in projected area and redshift around them, to establish a basic clustering template for the radial dependence and amplitude of the satellite population. This template is then applied iteratively to the entire survey, assigning probabilities that any given galaxy is a satellite of a given nearby primary galaxy. The clustering method is somewhat similar to the group-finding methods of [Yang et al. \(2007\)](#) and [Kourkchi & Tully \(2017\)](#), which also use adaptive search radii for group members, and iteratively update satellite membership and group properties. Our method does not assign definite membership for each satellite, however, but only estimates the probability of each galaxy pair being associated as primary and secondary. The method is particularly useful in crowded fields, or where distance information is limited.

We tested the method using data from the low-redshift part of the COSMOS field. COSMOS is a deep ($AB \sim 25\text{--}26$), multi-wavelength ($0.25\ \mu\text{m}\text{--}24\ \mu\text{m}$) survey covering a $2\ \text{deg}^2$ equatorial field ([Scoville et al., 2007b](#)). The data available for the COSMOS field includes high-quality photometric redshifts (photo-zs) generated from the 30+ deep bands ([Scoville et al., 2007b](#); [Ilbert et al., 2013](#); [Laigle et al., 2016](#); [Weaver et al., 2021](#)), providing distance information for a large sample of faint galaxies below the spectroscopic limit. The combination of depth and photo-z accuracy make this data set an ideal test case for the clustering method. In COSMOS, we made significant detections of the satellite population over a wide range of primary mass ($\sim 10^{10}\text{--}10^{13.5}M_{\odot}$ in halo mass). We measured the overall satellite abundance and the satellite luminosity function, as well as their dependence on primary halo mass. We also tested for a large number of possible systematic uncertainties, and showed that these are generally within the derived statistical uncertainties.

In this chapter, we explore the properties of the detected satellite populations further. The outline of this chapter is as follows. First, in section 4.2 we review the clustering method and define the base catalogue, the first-run primary sample, and the primary and satellite probabilities used in subsequent calculations. In section 4.3 we measure the satellite stellar mass function (SSMF), both in absolute terms, and relative to the stellar mass of the primary. In section 4.4 we study satellite abundance as a function of primary colour or specific star formation rate (sSFR, the star formation rate per unit stellar mass) and discuss the evidence for ‘conformity’ in our spectroscopic subsample. In section 4.5 we measure the quiescent fraction as a function of secondary stellar mass, primary stellar mass, primary colour, and primary-to-secondary halo mass ratio. Finally, we summarize our results and discuss future prospects for the clustering method in section 4.6.

4.2 Review of the method

4.2.1 The Hierarchical Clustering Method

To characterize satellite populations, we need to measure the abundance of satellite (or ‘secondary’) galaxies around central (or ‘primary’) galaxies. This is challenging, as satellites are faint and relatively rare with respect to foreground and background galaxies seen in projection. Our method uses projected clustering, relying on the fact that satellites are located close to centrals, and thus cause a slight density excess in the projected galaxy distribution. This overdensity can be used to estimate the satellite abundance statistically, giving each galaxy a probability of being associated with a given central, based on the measured clustering strength.

In chapter 3, we tested our method using data from the COSMOS field. Starting with the COSMOS2015 photometric redshift catalogue (Laigle et al., 2016), we applied the following cuts to select a low-redshift sample of galaxies with reasonably accurate redshift information:

- $i^+ < 25.5$;
- $0 < z_{\text{pdf}} < 6.9$;
- $z - 2\sigma_z < 0.3$;
- $\sigma_z < 0.5$.

These cuts give us a ‘base’ catalogue of 37,578 galaxies to work with. In most of our analysis, we also apply a stellar mass cut at $M_* > 10^{7.3} M_\odot$, as our sample is reasonably complete above this mass (see section 2.4 of chapter 3).

To define an initial sample of primaries, we make a conservative selection of systems most likely to dominate their surrounding region. First, a halo mass is assigned to each galaxy in the base catalogue as if it were the central galaxy of its own host halo, using a mean halo-to-stellar mass ratio based on Behroozi, Wechsler & Conroy (2013) (see chapter 3, section 3.1), and the virial radius R_{vir} and circular velocity at the virial radius v_{vir} corresponding to this halo mass are calculated. Each galaxy, in order of the assigned halo mass from high to low, is then checked to see if it is the most massive galaxy within its local Region of Interest (ROI), and thus a potential primary.

As discussed in chapter 3, the ROI is a cylindrical region around each primary, extending out to a maximum radius AR_{vir} in projected separation R_P (calculated at the distance of the primary) and a length $\pm Bv_{\text{vir}}$ with along the velocity-difference axis Δv . In addition, since

velocity uncertainties often exceed v_{vir} , we put an additional cut of $\Delta v > C\sigma_{\Delta v}$ on the sample of galaxies within the ROI, where $\sigma_{\Delta v}$ is the uncertainty in velocity difference Δv . Based on the clustering signal measured in stacked samples, we choose the values $(A, B, C) = (3.0, 2.0, 1.5)$ to define the limits of the ROI, as discussed in section 3.2 of chapter 3.

Primary candidates close to the boundaries of the survey volume in position or in redshift, or near heavily masked regions, could have more massive companions that were not included in the catalogue. Thus, we apply two additional cuts to the list of potential primaries, to ensure reasonably complete coverage of the ROI around each one:

- $z < 0.25$;
- area completeness > 0.65 ;

This selection produces an initial sample of 815 galaxies, which we will refer to as the ‘first-run’ primary sample. All other galaxies in the catalogue at this point are considered potential secondaries.

We measure the clustering signal around the first-run sample of primaries and model it as the sum of two components, a satellite component and a background component of constant surface density. We assume that the background component depends only on redshift, while the satellite component depends only on halo mass (given the fairly narrow redshift range of our sample). Fitting this two-component model to the data, we can estimate what fraction of all galaxies at a given distance from the primary are genuine satellites. Thus, in a second step, for every secondary in the ROI of a primary i , we calculate this probability P_{sat}^i of being a satellite. (We will also refer to the converse ‘field’ probability, $P_{\text{field}}^i = 1 - P_{\text{sat}}^i$, that a galaxy is *not* a satellite of primary i .) Given that the probability P_{sat}^i depends both on the secondary and on the primary, we will also refer to primary-secondary pairs, and associate with each a probability $P = P_{\text{sat}}^i$. We note that since ROIs can overlap, a given secondary may be a member of more than one pair, having a non-zero probability of being a satellite of more than one nearby primary. Thus, the total number of pairs exceeds the number of objects in the base catalogue.

Fig. 4.1 shows the number of pairs versus satellite probability (lower red histogram), as well as the cumulative number over a given probability, with (upper black histogram) and without (middle blue histogram) the spectroscopically confirmed pairs discussed in Section 4.2.3 below, which we assume to be 100% probable (point on the right-hand axis). We note that the pair distribution is very close to a power law in probability p , $n(p) \propto p^{-2}$.

For a galaxy to be a primary, it must not be a satellite of any other system. Thus, if a secondary is a member of more than one pair, then treating the probabilities P_{field}^i as independent, the net




Figure 4.1: Number of pairs versus satellite probability (red histogram), as well as the cumulative number over that probability, with and without the spectroscopically confirmed pairs (blue and black histograms respectively). The spectroscopic pairs are assumed to have a probability of 1.0 (black square on the right-hand axis).

primary or field probability P_{field} is the product of the individual probabilities:

$$P_{\text{field}} = \prod (P_{\text{field}}^i) = \prod (1 - P_{\text{sat}}^i). \quad (4.1)$$

Similarly, the net probability of a galaxy being a satellite of *some* system, P_{sat} , is:

$$P_{\text{sat}} = 1 - P_{\text{field}} = 1 - \prod_i P_{\text{field}}^i = 1 - \prod_i (1 - P_{\text{sat}}^i). \quad (4.2)$$

The initial primary selection is deliberately very conservative, including only those galaxies most likely to dominate their ROI. Given the non-zero background at all radii, however, no secondary is identified as a satellite with 100% probability, and thus any galaxy in the catalogue could be a potential primary. To account for this, we iterate, measuring the clustering signal stacked around *all* objects in the base catalogue, but weighting the contribution of each galaxy by its primary (i.e. field) probability. This iteration is particularly important at low masses, as these galaxies are mainly excluded from the first-run primary sample (see Fig. 3.3 of chapter 3; the first-run sample includes almost no objects below $\log(M_h/M_\odot) = 10.5$, whereas the base catalogue extends down to $\log(M_h/M_\odot) \sim 10$).

4.2.2 Calculating Distributions of Satellite Properties

The satellite probabilities described above allow us to study statistically the distribution of ancillary properties such as colour, luminosity, stellar mass, or star-formation rate (SFR). The simplest approach, which we called ‘Method A’ in our previous chapter, is to use the satellite probabilities as weights when calculating the distribution of a given ancillary property. Thus, for instance, we could estimate the satellite colour distribution by calculating the colour distribution for the entire base catalogue, weighted by the satellite probability. This assumes, however, that the ancillary property in question (colour, in this example) is uncorrelated with the clustering, which may not be true in many cases. Results calculated using Method A can be biased if the unclustered background population has a different distribution of ancillary properties than the satellite population (cf. Appendix C of chapter 3); in this case, selecting on the ancillary property first will change the clustering strength.

Ideally, we would resolve this problem by splitting the catalogue into narrow bins in the ancillary property, and recalculating the clustering signal and satellite probabilities for each of these bins individually. (We called this ‘Method C’ in the previous chapter.) Given the limited signal-to-noise ratio (SNR) of the clustering signal, however, binning our sample to this degree would increase the Poisson noise in the background estimate to unacceptable levels. In chapter 3 we described a practical alternative (‘Method B’) that works as follows: the total abundance of the

satellite and background components within the ROI are estimated from the satellite probabilities described above. We also use these probabilities to define ‘field’ galaxies, with net probabilities of being satellites of less than 1%, throughout the entire survey. This field population is used to define a (normalized) background distribution in the ancillary property, e.g. colour, stellar mass, etc., and it is assumed the background component in the ROI follows this same distribution. Thus, we can calculate the background contribution to a given bin in the ancillary property, as the fraction of field galaxies in that bin, times the total number of background galaxies in the ROI. (See sections 5.1 & 5.2 of chapter 3, as well as [Speller & Taylor \(2014\)](#), for further discussion).

4.2.3 Spectroscopically Confirmed Subsample

A total of 2,506 spectroscopic redshifts are also available for the base catalogue. Searching through this sample, we find 278 spectroscopically confirmed primary-secondary pairs with projected separations of less than 2.0 times the virial radius of the primary and redshift separations of less than 10^{-3} (roughly 300 km/s). This subsample includes 136 unique primaries and 271 unique satellites (by the definitions above, there are a few objects that could be satellites of more than one primary). On average, each primary in the spectroscopic subsample has slightly more than two confirmed satellites, but the most massive primaries contain significantly more satellites than the rest. There are 15 primaries with $\log(M_*/M_\odot) > 10.9$, with a total of 65 confirmed satellites (or 4.3 each on average). We will treat the 278 primary-secondary pairs as completely certain, and use this subsample to test some of our results. We note however that it is assembled from many different redshift sources, and is not complete nor homogeneous.

The spectroscopic sample provides an interesting test of our estimated satellite probabilities. The clustering method identifies 442 (130) secondaries with P_{sat} larger than 0.3 (0.5), of which 26 (14) have spectroscopic redshifts, and 7 (2) are spectroscopically-confirmed satellites. This is completely consistent with expectations for the $P_{\text{sat}} > 0.3$ sample (46% probability of getting $7/26 = 27\%$ or fewer confirmed satellites), but inconsistent for the $P_{\text{sat}} > 0.5$ sample (only a 0.64% probability of getting $2/14 = 14\%$ or fewer confirmed satellites, corresponding to a 2.7σ deviation.) Given the small numbers of objects involved, however, it is unclear whether the latter discrepancy is important, or simply an artifact of uneven spectroscopic follow-up (e.g. with coverage that undersamples close pairs).

We note that the high-probability samples provide an interesting set of possible targets for further spectroscopic follow-up; it would be fairly easy, for instance, to obtain spectra for magnitude-limited samples of the $(442-26) = 416$ objects with $P_{\text{sat}} > 0.3$, or the 2008 objects with $P_{\text{sat}} > 0.1$, with the expectation that significant fractions of these objects were genuine satellites.

4.3 Satellite Stellar Mass functions

Our base COSMOS catalogue includes large numbers of intrinsically faint, low-redshift galaxies, and thus it provides a good opportunity to study the low-mass end of the SSMF. In this section, we will explore the SSMF and various related quantities, and their dependence on environment (i.e. primary halo mass). We will then compare our results to previous measurements from the literature.

Laigle et al. (2016) derived stellar mass estimates for the COSMOS2015 catalogue using a library of synthetic spectra generated by the stellar population synthesis model of Bruzual & Charlot (2003), the same method also described in Ilbert et al. (2015). They also estimated the stellar mass completeness limit empirically, based on the masses of detected galaxies and their K_s magnitudes relative to the limiting magnitude in this band. For the redshift range of $z < 0 < 0.35$, they estimate a 90% completeness limit of $M_* = 10^{8.4} M_\odot$ for quiescent galaxies and $M_* = 10^{8.1} M_\odot$ for the star-forming galaxies. In chapter 3, examining number counts, we concluded that our local sample is complete down to at least $M_* = 10^{8.2} M_\odot$ for the redshift range ($z < 0.25$). For very low-redshift systems ($z < 0.07$), our sample appears to be complete down to as low as $M_* = 10^{7.2} - 10^{7.5} M_\odot$. For primary masses, these limits are amongst the deepest published, as discussed in section 4.3.3 below.

4.3.1 Absolute Stellar Mass Functions

In Figs. 4.2 and 4.3 we show the (cumulative) satellite stellar mass function (SSMF), binned in six ranges of primary stellar mass and six ranges of primary halo mass, respectively. The mass functions are calculated using method B described above. Only satellites within $1.5 R_{\text{vir}}$ in projection are included in the total.

We note that since these are cumulative plots, the error bars in different bins of secondary stellar mass are not independent; as an indication of the overall significance of the clustering signal, we can consider the measured abundances at any single value of M_* . Also, as explained in chapter 3, the halo-background decomposition used to fit the clustering signal and set the initial satellite probabilities assumes and fits a power-law dependence on halo mass; although the connection to the final measured satellite abundance is indirect, this may smooth slightly the variation of the SSMF with primary mass. From the results of chapter 3, fitting the clustering signal in each bin of primary mass independently might add an additional scatter of 20-30% to the normalization of the SSMF, particularly at lower primary masses. This is generally less than the errors shown here, however, so the net effect would be minor.




Figure 4.2: The (cumulative) satellite stellar mass function, for 6 bins of primary stellar mass. Only satellites within $1.5 R_{\text{vir}}$ in projection are included. The mass range and average mass of each primary mass bin are given in the legend.

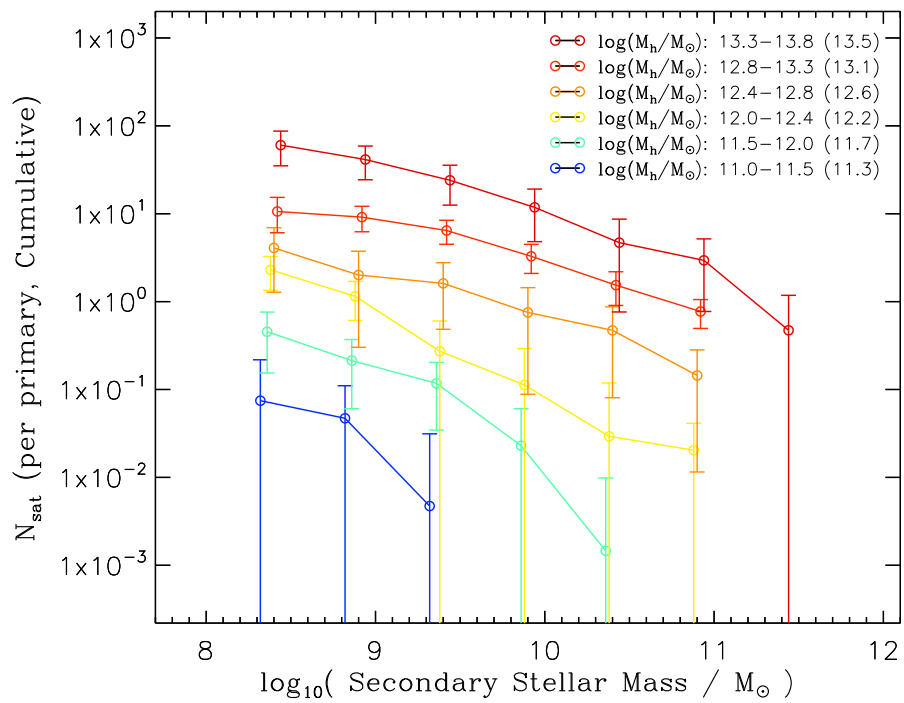


Figure 4.3: The (cumulative) satellite stellar mass function, for 6 bins of primary halo mass. Only satellites within $1.5 R_{\text{vir}}$ in projection are included.

Overall, the mass function appears Schechter-like, with a power-law dependence on secondary stellar mass and an exponential cutoff at the high-mass end. The amplitude of the SSMF also scales fairly smoothly with primary stellar or halo mass (given the caveat about smoothing discussed above). A similar conclusion was reached in chapter 3 (see e.g., Fig. 10). Previous work by Wang & White (2012) claimed a change in the slope of the (differential) SSMF with primary mass, but subsequent work (e.g. Lan, Ménard & Mo (2016); Zu & Mandelbaum (2015)) suggests that this comes from fitting a single power-law to a multi-component SSMF that was sampled over different ranges of secondary mass for different ranges of primary mass. The true SSMF probably includes Schechter-function components for red and blue satellites at the faint end, plus a separate log-normal component of red galaxies above $M_* \sim 10^{9.5}-10^{10} M_\odot$ (Lan, Ménard & Mo, 2016). While there is a faint suggestion of this log-normal component in our 1–2 most massive primary ranges, we lack the SNR to distinguish it clearly. For less massive primaries, we clearly see the exponential cutoff in the SSMF at large masses.

Satellites in the lowest stellar mass bin on Figs. 4.2 and 4.3 are approaching the stellar mass completeness limit of $10^{8.2} M_\odot$ mentioned above. This is particularly true for massive primaries (where the slope of the SSMF appears to flatten slightly below $10^9 M_\odot$), as these are rare, and tend to lie at higher redshift within our target volume (see Fig. D.2 in Appendix D.2); the highest mass bin, for instance, has no primaries below $z = 0.22$. We suspect completeness effects toward the high redshift limit of the sample may reduce the counts slightly at faint magnitudes.

4.3.2 Relative Stellar Mass Functions

The subhalo mass function predicted by theory is approximately scale invariant, that is, the number of subhalos with a given fraction of the main halo mass or circular velocity is roughly independent of primary halo mass (e.g. Moore et al., 1999a). Alternately, since the cumulative subhalo mass function goes roughly as $(M_{\text{sub}})^{-1}$, this scale-invariance also implies that the number of subhalos of a given (dark matter) mass M_{sub} per unit primary halo mass is roughly constant (e.g. Gao et al., 2004).

We can compare this prediction to the equivalent quantity in terms of stellar mass, that is the abundance of satellites of a given stellar mass, per unit primary halo mass. We refer to this as the net ‘efficiency’ of subhalo occupation on these scales, assuming a monotonic relation between subhalos and satellites. Fig. 4.4 shows this efficiency or relative abundance, for the same bins of primary halo mass as in Fig. 4.3. At the low-mass end, we see that the SSMF per unit halo mass is constant to within the uncertainties, at least for the top four primary halo mass bins. At the high-mass end, the relative mass function is truncated when the satellite mass is roughly equal to the primary mass; this is a natural consequence of our requirement that $M_*^{\text{sat}} < M_*^P$.




Figure 4.4: The SSMF per unit halo mass, or ‘efficiency’ of subhalo occupation, using the same bins of primary halo mass as in Fig. 4.3.

Overall, these results are consistent with subhalo occupation being fairly universal at the masses probed here; that is, it seems that a subhalo of a given dark matter mass will tend to host a satellite galaxy of a set average stellar mass, independent of environment (where by ‘environment’, we mean primary halo mass). We can compare these results to those of [Zu & Mandelbaum \(2015\)](#). Down to primary halo masses of $\log(M_*/M_\odot) = 12.2$, they measure a constant efficiency in the differential counts for satellite masses $\log(M_*/M_\odot) = 8.5\text{--}10.2$. The value they measure does vary with primary mass at satellite masses $\log(M_*/M_\odot) \gtrsim 10.2$. This will shift the cumulative mass functions slightly, but given the steep slope of the counts we might expect cumulative counts to agree to within $\sim 20\%$ for the top 4 bins shown in [Fig. 4.4](#), which is certainly consistent with our results. Thus, our measured efficiencies are consistent with those of [Zu & Mandelbaum \(2015\)](#), over the primary mass range where they overlap.

At the lowest primary halo masses, we do see a possible indication of a drop in efficiency, but only at $1\text{--}2\sigma$ level. To test or demonstrate universal subhalo occupation conclusively, we would need higher SNR data, as well as more information on the radial distribution of satellites in different environments (particularly around low-mass primaries), to show that subhalos of a given mass are occupied in exactly the same way.

Some evolutionary effects, such as infall due to dynamical friction, *should* depend on environment, however; specifically, they will depend on the subhalo-to-main halo mass ratio. Relatively massive satellites (i.e. small primary-secondary mass ratios) will experience stronger dynamical friction, and will merge faster, depleting the observed satellite population. We can look for evidence of such effects by measuring satellite abundance in terms of the corresponding stellar mass ratio $\mu_* \equiv M_{*,\text{main}}/M_{*,\text{sat}}$. We will call the abundance as a function of this ratio the ‘relative stellar mass function’ (RSMF).

[Fig. 4.5](#) shows the RSMF, estimated using method B, for the same primary halo mass bins used previously. Note that since we have applied a fixed stellar mass completeness cut to the base catalogue at $10^{8.2}M_\odot$, the relative depth varies with the primary stellar mass. Points where satellites around more than half the primaries would lie below this completeness limit for a given stellar mass ratio are excluded. We can see that the top mass bin has a very good SNR, and probes the RSMF down to 3.4 dex in stellar mass ratio. For smaller halos we only reach depths of 1.5–2 dex.

For the higher primary mass bins, the amplitude of the RSMF varies as expected from [Fig. 4.4](#). For the lower bins, the overall scaling is harder to determine given the low SNR, but it seems to flatten, and may even be inverted. This could be a result of the flatter stellar-to-halo mass relation (SHMR – e.g. [Leauthaud et al., 2012](#); [Behroozi, Wechsler & Conroy, 2013](#); [Grossauer et al., 2015](#), [Shuntov et al. in prep.](#)) at these halo masses (which would produce more overlap between the primary halo mass bins), but it may also indicate residual background contamination




Figure 4.5: Satellite abundance as a function of the primary-to-secondary stellar mass ratio μ_* (the RSMF). Points where 50% or more of the primaries would lie below the stellar mass completeness limit of the sample are excluded from the plot.

problems in the lowest bins.

The shape of the RSMF is similar in all bins, with a fairly shallow slope at large mass ratios (low-mass satellites), and a steeper slope at small mass ratios (massive satellites). There is a slight indication that the rise at small mass ratios is steeper for less massive primaries than for more massive ones. Re-binning the results into two bins of primary mass to increase the SNR, we find that for higher primary masses ($M_h > 12.8$), $d \ln N / d \ln \mu_* \sim 1.5 \pm 1$, whereas for ($M_h < 12.8$), $d \ln N / d \ln \mu_* \sim 3 \pm 2$. Thus, there is a difference, but it is less than $1\text{-}\sigma$ significance. Here again, this difference would be expected from the flatter slope of the SHMR at low masses. For a flatter SHMR, the range of μ_* covered by the first bin in the figure would map onto a smaller range of dark matter mass ratios. Thus, the first data point would represent a purer sample of major mergers, and dynamical friction effects would be stronger in this bin. Clearly, it would take a larger sample to confirm this. We revisit this point and show further evidence for dynamical friction effects below, however, in Section 4.5.

4.3.3 Comparison with Literature Results

The satellite luminosity function and stellar mass function have been measured many times, using different techniques to identify satellites individually or statistically. Fig. 4.6 shows one comparison of some previous measurements, in terms of their limiting secondary stellar mass sensitivity and primary halo mass coverage. Symbol types indicate measurements of the luminosity function (triangles), the stellar mass function (squares), both (pentagons), or the relative luminosity function (the number as a function of magnitude offset – circles). Note that the limits are approximate, as they have been converted from various bands and/or assume different model mass-to-light ratios. Different surveys also vary enormously in size and SNR. With these caveats, we see that our current work extends the sensitivity to the lowest primary halo masses to date. The sensitivity in secondary mass is less exceptional, although at very low redshift we do in principle have some sensitivity down to masses of $\log(M_{*,\text{sat}}/M_\odot) \sim 7.5$ (not shown here).

Given these previous results, we will compare the measured amplitude of the SSMF to the values in Wang & White (2012) and Lan, Ménard & Mo (2016), since these are the deepest published SSMFs; for the shape of the SSMF, we will convert the relative luminosity function published by Mao et al. (2021) to stellar mass.

Fig. 4.7 shows our measured amplitude of the SSMF, as a function of primary halo mass, compared to the results of (Wang & White, 2012, – blue) and (Lan, Ménard & Mo, 2016, – cyan). We have used the value of the cumulative mass function at $M_{*,\text{sat}} = 10^9 M_\odot$, as a measure of the amplitude, as this point is reasonably well sampled over the whole range of primary halo mass considered. Symbols show the halo mass range in each bin as horizontal width, and the




Figure 4.6: Limiting sensitivity, in secondary stellar mass versus primary halo mass, of a number of previous studies of satellite abundance (from top right to bottom left, Conroy, Wechsler & Kravtsov, 2006; Zu & Mandelbaum, 2015; Wang & White, 2012; Wu et al., 2021; Nierenberg et al., 2016; Lan, Ménard & Mo, 2016; Besla et al., 2018; Mao et al., 2021; Speller & Taylor, 2014; Wang et al., 2020). Triangles indicate published luminosity functions; squares indicate published mass functions; pentagons indicate both, while circles indicate relative luminosity functions (i.e. number versus magnitude offset). Values are approximate, as discussed in the text.

uncertainty as height (the uncertainties for Wang & White (2012) and Lan, Ménard & Mo (2016) were not directly reported, but are assumed to be small – thus the nominal height shown here). Note that the different studies also used slightly different limiting radii to define the satellite abundance – Wang & White (2012) used a fixed radius of 300 kpc (or 170 kpc in their lowest mass bins), Lan, Ménard & Mo (2016) used the virial radius R_{vir} , which should have a value of ~ 250 kpc for primary halo mass $\log(M_{\text{h,p}}/M_{\odot}) = 12$, while we have used $R_{\text{vir}} \sim 375$ kpc at $\log(M_{\text{h,p}}/M_{\odot}) = 12$. Testing our model with different radial cuts, we expect these differences may shift the measured amplitude by $\sim 20\%$ between Lan, Ménard & Mo (2016) and Wang & White (2012), and a further $\sim 20\%$ between Wang & White (2012) and ourselves. These shifts would slightly improve the agreement between our results and Wang & White (2012), particularly at the high-mass end, but they are relatively small compared to the uncertainties, so we have not included them here.

Overall, we see excellent consistency between our measured amplitude of the SSMF and the values found in previous work; while there appears to be a slight offset between results around $\log(M_{*,\text{p}}/M_{\odot}) = 10.8$, it is well within our uncertainties. The amplitude appears to vary as a simple power law in $\log M_{*,\text{p}}$, over four decades in primary stellar mass.

We can also test the shape of the SSMF, or equivalently the RSMF, at a given primary mass. Here we compare to Mao et al. (2021), who recently presented the data release from the SAGA II spectroscopic survey, as well as some initial science results, including a relative luminosity function. Estimating stellar masses for the SAGA sample as described in appendix D.3, we can use their satellite and host catalogues to construct a RSMF for MW-like hosts.

The comparison is shown in Fig.4.8. Once again, the agreement is excellent; over the range of mass ratios $\log(M_{*,\text{p}}/M_{*,\text{s}}) = 0.5\text{--}2$, the SAGA results lie well within the uncertainties of our RSMF measurements. Our results appear to roll over, becoming incomplete around mass ratios of $\mu_* = 2\text{--}2.5$, depending on the primary mass bin; the SAGA results extend deeper, with rising counts down to a mass ratio of around $\mu_* = 3$. (Given the SAGA II magnitude limit is about $M_r = -12.3$, and their hosts are selected in the range $-23 > M_K > -24.6$, which corresponds approximately to $-20.5 > M_r > -22.1$, the SAGA results should be relatively complete down to the last few points plotted, or $\mu_* \sim 3.2$.)

We note that Mao et al. (2021) include only satellites at separations $R_p < 300$ kpc and velocity offsets $\Delta v < 250 \text{ km s}^{-1}$, which are slightly more restrictive than our limits. Judging from Figure 10 of their first paper (Geha et al., 2017), this might reduce their measured satellite abundance by $\sim 30\%$ relative to ours, but this is still consistent with our results (and in fact a slightly better match at low mass ratios/large satellite masses).

Overall, we conclude that for the ranges of primary and secondary mass where we can directly compare them, both the normalization and shape of our SSMF (and/or RSMF) are in excellent




Figure 4.7: Amplitude of the cumulative stellar mass function at $M_* = 10^9 M_\odot$, for the current study (red points and squares), compared to previous results from Wang & White (2012) and Lan, Ménard & Mo (2016). The horizontal extent of the symbols indicates the bin width, while the vertical extent indicates the approximate uncertainty in the central value (red points).




Figure 4.8: The RSMF for MW-mass hosts, compared to the SAGA II results of [Mao et al. \(2021\)](#). We have included two of our primary halo mass bins that bracket the range covered by the SAGA sample. The legend specifies the primary halo mass range, together with the median halo mass and the median COSMOS K_s band magnitude for each bin. (Note COSMOS uses AB magnitudes, while SAGA II uses Vega magnitudes.)

agreement with previous results. The clustering method has allowed us to extend these previous measurements to very low primary masses, however, despite the relatively small COSMOS field, and thus shows great potential for future surveys.

4.4 Satellite-Central connection

Given the overall consistency of our SSMF measurements with previous results, we will proceed to consider satellite populations as a function of primary properties. In particular, there is already considerable evidence that primary morphology influences the abundance and properties of satellites:

- Several previous studies find that red/blue primaries have a higher/lower abundance of satellites at fixed stellar mass (e.g. [Wang & White, 2012](#); [Mandelbaum et al., 2016](#)).
- Central galaxies of different morphological type may follow different SHMRs ([Wojtak & Mamon, 2013](#); [Hudson et al., 2015](#); [Mandelbaum et al., 2016](#); [Correa & Schaye, 2020](#), Spitzer et al. in prep.), which might or might not explain the first point.
- Satellite populations also show “conformity” in colour, that is red primaries tend to have more red satellites, while blue primaries have more blue satellites (e.g. [Weinmann et al., 2006](#); [Wang & White, 2012](#); [Hartley et al., 2015](#); [Knobel et al., 2015](#)).

These patterns could reflect assembly bias ([Gao, Springel & White, 2005](#)), red primary galaxies lying preferentially in denser regions where structure has formed earlier and dwarf galaxies are more abundant, but satellites tend to be quenched. Alternately, they could be due to effects at the single-halo scale, including a different SHMR for central galaxies with different colours or morphologies. If central morphology depends on the detailed merger history of the system, this may also influence the satellite population; a simple example is fossil groups ([Ponman et al., 1994](#)), where recent mergers appear to have depleted the bright end of the satellite luminosity function.

To test for these effects, we consider satellite populations around subsamples of primaries split by colour or by sSFR, compare our results to previous measurements from the literature, and also consider the question of conformity.

4.4.1 Satellite Populations Split by Central Colour or sSFR

First, we consider the division of primaries into subsamples. For colour, we choose to split on the index $c = (B - i^+)$, as explained in Appendix D.4. We use a colour cut at $c = 1.6$ for massive primaries ($\log(M_*/M_\odot) > 10$), while for less massive primaries, we move this to $c = 1.5$, to reflect the evolution of the red sequence at lower stellar mass. For sSFR, we split at the value -11 , which produces roughly comparable subsamples of star-forming and passive galaxies, as explained in Appendix D.4.

Fig. 4.9 shows the SSMF in subsamples split by colour. Note that the signal of the highest primary mass bin is entirely from red primaries, as the primary sample has no blue galaxies in this mass range. While our results are consistent with those mentioned above (the SSMF for red primaries lies above the SSMF for blue primaries), we lack the SNR to reach significant conclusions. In the intermediate primary mass bin, there is a significant detection of massive satellites ($\log(M_{*,\text{sat}}/M_\odot) > 10$) around red primaries, versus no detection around blue primaries; thus the two populations differ at the $\sim 1.5\sigma$ level. (This measured difference is also consistent with earlier results, e.g. Wang & White 2012.)

A number of previous studies have indicated that the host halos of red primaries are, on average, more massive than those of blue ones (e.g. Wojtak & Mamon, 2013; Mandelbaum et al., 2016; Correa & Schaye, 2020). Correa & Schaye (2020), for instance, find that disk galaxies have stellar masses up to 1.5 times larger at fixed halo mass, at the high-mass end of the galaxy population. Given the slope of the SHMR, this is equivalent to a factor of ~ 2 difference in halo mass at fixed stellar mass. Earlier work by Mandelbaum et al. (2016) found a similar result: for primary stellar masses $\log(M_*/M_\odot) = 10.3\text{--}11.6$, the halos of passive centrals are at least twice as massive as those of star-forming centrals of the same stellar mass. If we assume that total satellite number is directly proportional to host halo mass (as suggested by Fig. 4.4), we would expect a higher number of satellites around red primaries than around blue primaries of equal stellar mass. The amplitude of this effect is only a factor of ~ 2 , however, which is within the uncertainties in Fig. 4.9. There could also be other factors that affect the normalization of the SSMF; if the satellites of red or passive primaries are more often quenched themselves (‘conformity’), they may have a different *subhalo* SHMR, which could cancel out some or all of the effect of the primary halo mass difference. We consider conformity below, and quenching rates in Section 4.5.

We have also measured the SSMF in subsamples split by sSFR, as shown in Fig. 4.10, to test whether they showed more evidence of systematic differences in satellite populations. At the highest mass bins, we find no significant difference between the satellite abundance around high-sSFR primaries and low-sSFR ones. At intermediate and low primary mass, low-sSFR primaries seem to have more satellites, consistent with previous results, but once again the SNR




Figure 4.9: The SSMF for red/blue primaries in three primary stellar mass ranges. The bins in primary mass are adjusted to give the red and blue subsamples comparable mean stellar mass. The top dashed line is the SSMF for primaries with stellar mass larger than $10^{10.9} M_{\odot}$, the range in which there are no blue primaries.




Figure 4.10: As Fig. 4.9, but for primary samples split by sSFR.

of the detection is only about 1σ .

4.4.2 Conformity in the Spectroscopic Sample

While we lack the SNR to properly test conformity with the clustering method (our attempts to measure it were inconclusive), the spectroscopic sample does show some evidence for this effect. First, Fig. 4.11 shows the distribution of colour versus stellar mass for the spectroscopically confirmed secondaries (a similar plot is shown for the primaries in Appendix D.4). We see a strong correlation between the stellar mass and colour, with a well-defined red sequence visible at the high-mass end.

Fig. 4.12 then compares the colour indices of primary and secondary galaxies. The horizontal line indicates the division between red and blue primaries (at high primary mass), while the horizontal line indicates the median colour of the secondary sample. Examining each quadrant,




Figure 4.11: Colour index versus stellar mass for the spectroscopically confirmed satellites; points are coloured by the redshift.




Figure 4.12: The colour index of the spectroscopically confirmed satellites versus the colour index of their primaries. The dashed vertical line marks the division between red and blue primaries (at low stellar mass); the horizontal line marks the median colour of the secondary sample.

we see that red secondaries are rare around blue primaries. Thus, red satellites occur mainly around red primaries, while blue satellites occur around a range of primaries; also the satellites of blue primaries tend to be blue, while the satellites of red primaries have a range of colours. While this pattern is broadly consistent with previous measurements of conformity (e.g. [Weinmann et al., 2006](#); [Wang & White, 2012](#); [Hartley et al., 2015](#); [Knobel et al., 2015](#)), we caution that the spectroscopic sample is inhomogenous and incomplete, particularly for the more massive primaries, which tend to lie at higher redshift (see Appendix D.2). Thus it is possible that we are missing some faint blue satellites around distant, massive red primaries, (although it seems less likely that we are missing red satellites around nearby, low-mass blue primaries).

4.5 Environmental Quenching

The quenching of star formation in a galaxy can be driven by internal processes (e.g. rapid gas loss via violent starbursts), or environmental ones (e.g. ram pressure stripping by the intra-cluster medium). There are clear indications of environmental quenching even on small scales; in the Local Group, for instance, most nearby dwarfs are quiescent, and star-forming dwarfs are generally distant, with the notable exception of the Magellanic clouds (McConnachie, 2012). Studies of environmental quenching consider either the total quiescent fraction in the satellite population, or the “environmentally quenched fraction”. This is the fraction of those galaxies that would be expected to be star-forming in the field, but that are observed to be quiescent in satellite populations. It can be calculated as

$$f_{\text{qe}} = (f_s - f_f)/(1 - f_f) \quad (4.3)$$

where f_s refers to the quiescent fraction of satellites and f_f refers to the quiescent fraction of field galaxies. Thus for instance, if $f_f = 0.2$ and $f_s = 0.8$, we conclude that the fraction of (star-forming) systems quenched by their environment is $0.6/0.8 = 75\%$.

The environmentally quenched fraction has previously been found to depend on secondary stellar mass, with hints of an abrupt change around a secondary stellar mass $\log(M_*/M_\odot) \sim 8$ (Wheeler et al., 2014), but measurements in the LG (Wetzel, Tollerud & Weisz, 2015) and in the SAGA II survey (Mao et al., 2021) do not confirm this feature, and in general there is limited information in the secondary mass range $10^7\text{--}10^9 M_\odot$. Quenching may also vary as a function of primary properties, as mentioned above, so we will consider it in the various primary subsamples defined in Section 4.4.

To distinguish quiescent galaxies from star-forming ones, we use the ‘CLASS’ flag in the COSMOS2015 catalogue, which classifies galaxies based on their location in the (NUV- r)-($r - J$) colour-colour plane (Laigle et al., 2016). We noticed, however, that about 30% of the galaxies in the base catalogue do not have valid NUV-band measurements. In principle, these galaxies should not have been classified, due to the lack of NUV- r colour. In practice, however, they were classified as ‘CLASS=1’ (or star-forming) in the COSMOS2105 catalogue. To correct this, for these galaxies we apply our own classification, similar to that of Laigle et al. (2016), but replacing the missing NUV- r colour with ‘MNUV_MR’, an estimated colour given in the catalogue that appears to be available for most objects, and is also corrected for dust-extinction. Using this colour ensures full completeness for our base catalogue, and for galaxies where valid NUV measurements are available, the MNUV_MR-($r - J$) colour-colour distribution appears to be roughly equivalent to the (NUV- r)-($r - J$) distribution.

The new classification ‘CLASS_XI’ is defined as:

1. 'CLASS_XI=0' (Quiescent) if ['CLASS=0' OR ('MNUV_MR> 3.1' AND 'MNUV_MR> 3 × (r - J) + 1')];
2. 'CLASS_XI'=1 (Star-forming) for the rest.

4.5.1 Quenching versus Secondary Stellar Mass

To estimate the environmentally quenched satellite fraction, we can proceed in two ways. First, we can consider the whole catalogue (binned by stellar mass) and use the previously calculated overall satellite/field probabilities, P_{sat} and P_{field} , without reference to particular primaries.

Suppose the catalogue contains $N_{\text{q,tot}}$ quiescent galaxies in a given stellar mass bin. We can divide this into satellite and field or primary populations, such that

$$N_{\text{q,sat}} = N_{\text{q,tot}} - N_{\text{q,field}}. \quad (4.4)$$

On the other hand, the total number of satellites can be written:

$$N_{\text{sat}} = \sum_i P_{\text{sat}}^i, \quad (4.5)$$

and we can estimate the number of quiescent field galaxies by assuming an universal quiescent fraction for field galaxies within the given mass bin, $f_{\text{q,field}}$, such that

$$N_{\text{q,field}} = f_{\text{q,field}} N_{\text{field}} = f_{\text{q,field}} \sum P_{\text{field}}^i. \quad (4.6)$$

Combining these expressions, the quiescent satellite fraction $f_{\text{q,sat}}$ is:

$$f_{\text{q,sat}} = \frac{N_{\text{q,sat}}}{N_{\text{sat}}} = \frac{N_{\text{q,tot}} - f_{\text{q,field}} \sum P_{\text{field}}^i}{\sum_i P_{\text{sat}}^i} \quad (4.7)$$

To estimate $f_{\text{q,field}}$, we select objects with $P_{\text{field}} > 0.99$ as the field galaxy subset, and measure the field quiescent fraction as a function of stellar mass, based on this subset. For the satellite sample, including galaxies with very low P_{sat} will lower the SNR of satellite quiescent fraction, but restricting the sample to objects with the highest values of P_{sat} will limit the sample size. As a compromise, we choose to limit the sample to objects with $P_{\text{sat}} > 0.8$ for the calculation in Eqn. 4.7.

Fig. 4.13 compares the quiescent fractions as a function of stellar mass for these different galaxy populations. In the legend, “Total” (filled black circles) indicates the whole base catalogue.




Figure 4.13: Quiescent fraction versus stellar mass. Results are shown for the three different samples defined in the text, the whole base catalogue (black circles), field galaxies ($P_{\text{field}} > 0.99$; blue squares), and satellites ($P_{\text{sat}} > 0.8$; red triangles). For comparison, the SAGA II (Mao et al., 2021) and LG (Wetzel, Tollerud & Weisz, 2015) results are shown as green and orange shaded regions respectively.

"Field/Central" (blue squares) indicates objects likely to be field galaxies, defined as above, while "Satellites" (red triangles) indicates objects likely to be satellites. The quiescent fraction for the satellites is estimated using Eqn. 4.7.

Overall, the quiescent fraction of all galaxies in the base catalogue is low, but increases both towards the high-mass end and towards the low-mass end. The trend at the high-mass end is expected as a result of the increasing mass quenching effect for high-mass field galaxies. We suspect the increase at low masses is largely contributed by satellites. Note that we do also see a small increasing trend at the low-mass end for field galaxies as well, which may come with an increasing fraction of mis-classification between field galaxies and satellites at very low masses. The field galaxies have significantly lower quiescent fractions than the average of all galaxies, implying a high quiescent fraction for satellites. The significant difference between the satellites and field galaxies provides strong evidence for environmental quenching effects.

Note that this discussion assumes the COSMOS2015 definition of quiescence; if we consider instead the red fraction with $c > 1.5$ or 1.6 , as defined in Section 4.4.1, we find red fractions of 33% 54%, and 93% (or 22.5/54/93% for a uniform cut at $c = 1.6$) for the top three stellar mass bins shown in Fig. 4.13. These are $\sim 20\%$ higher than the red fractions measured for isolated samples (e.g. Wang & White, 2012; Geha et al., 2012), but closer to the fractions measured in Wang & White (2012) for the SDSS main sample. The difference may be partly due to the colour index we use, but it probably reflects the fact that the isolation cuts in these studies were quite strict, whereas our field probability calculation includes some primaries in denser regions.

For satellites, the quiescent fraction is fairly high, around 0.4–0.9. This is comparable to the result of Wetzel, Tollerud & Weisz (2015) but systematically higher than the fraction reported in Mao et al. (2021) except for the lowest stellar mass bin around $10^{7.5} M_{\odot}$. We will consider a few possible explanations for this discrepancy below. Beyond $\log(M_*/M_{\odot}) = 10$, the measured fractions are very uncertain, as we have very few pairs in this mass range with satellite probabilities of 0.8 or more. Below $\log(M_*/M_{\odot}) = 8$, the quiescent fraction drops significantly, it is likely due to a higher fraction of misclassified field/satellites at the low-mass end.

4.5.2 Quiescent Fraction versus Primary Morphology and Stellar Mass

As noted earlier, quenching may be correlated with central properties, so we will also consider the satellite quiescent fraction for the red and blue primary subsamples defined in Section 4.4, as well as for several bins of primary stellar mass. To calculate these fractions, we follow the procedure described above (Eqn. 4.7), but first split the satellite samples by primary properties, while still using the whole sample to estimate $f_{q,\text{field}}$. For instance, we can calculate the satellite

quiescent fraction around red primaries as

$$f_{\text{sat}}^{\text{q,RP}} = \frac{N_{\text{q,sat}}^{\text{RP}}}{N_{\text{sat}}^{\text{RP}}} = \frac{N_{\text{q,tot}}^{\text{RP}} - f_{\text{q,field}} \sum P_{\text{field}}^{i,\text{RP}}}{\sum_i P_{\text{sat}}^{i,\text{RP}}}, \quad (4.8)$$

where the notation ‘‘RP’’ indicates quantities measured around red primaries. A single galaxy will often be a potential satellite of several primaries. In these cases, we associate the satellite with central for which it has the largest satellite probability, although this may introduce some noise into our split by primary properties, by mixing satellites from the two subsamples.

Fig. 4.14 shows the quiescent fraction for red and blue primary subsamples, as well as for the entire base catalogue. We note that the SNR of the blue sample is lower, due to the smaller sample size. Despite the large uncertainties, we can see that the satellites around blue primaries show a significantly lower quiescent fraction than the ones around red primaries, over the stellar mass range of $\log(M_*/M_\odot) = 8.5\text{--}10$. We note that a difference in the quiescent fraction of satellites around red and blue primaries might affect the normalization of the SSMF, as discussed above. The contribution of the high satellite quiescent fraction of red primaries may also explain why our measurements are higher than the MW results from [Wetzel, Tollerud & Weisz \(2015\)](#) at the stellar mass of $\log(M_*/M_\odot) \sim 9$.

Finally, in Fig. 4.15 we test the dependence of the quiescent fraction on primary stellar mass. For the mass bins plotted, low-mass primaries generally have lower quiescent fractions, but the SNR is too poor to reach a definitive conclusion. If we bin together the results for the secondary mass range $\log(M_*/M_\odot) = 7.5\text{--}10$, we find fractions 0.64 ± 0.04 for the higher primary mass range, versus 0.52 ± 0.15 for the intermediate mass range, a difference of $\sim 0.8\sigma$. For the lowest mass range, there is too little SNR to reach any conclusion. Nonetheless, the suggestion that the quiescent fraction may be lower around low-mass primaries motivates us to consider its dependence on the primary-to-secondary mass *ratio*; we will discuss this next.

4.5.3 Quiescent fraction vs relative mass

Dynamical friction will drag any satellites that are comparable in mass to the primary into the centre of the main halo on a very short timescale, equivalent to a few orbits (e.g. [Colpi, Mayer & Governato, 1999](#); [Taylor & Babul, 2004](#)). By implication, the most massive surviving satellites are necessarily recent mergers. We anticipate that this could produce a selection effect, whereby satellites with large relative masses are still star-forming, and the quiescent fraction depends on the primary-to-secondary mass ratio.




Figure 4.14: Quiescent fraction as function of secondary stellar mass, for the red and blue primary subsamples (red and blue points), as well as the whole base catalogue (black points).




Figure 4.15: Quiescent fraction as function of secondary stellar mass for three primary stellar mass bins (red, green and blue points), and for the whole base catalogue (black points).

Colpi, Mayer & Governato (1999) have estimated the timescale for infall due to dynamical friction, on an orbit of a given initial energy and angular momentum, as

$$\begin{aligned}
T_{\text{DF}} &= 1.2 \frac{J_{\text{circ}} r_{\text{circ}}}{(GM_{\text{h,s}}/e) \log(M_{\text{h,m}}/M_{\text{h,s}})} \epsilon^{0.4} \\
&= 0.519 P_{\text{vir}} \frac{\mu}{\log(\mu)} \epsilon^{0.4} \\
&= 0.5 P_{\text{rad}} \frac{\mu}{\log(\mu)} \left(\frac{\epsilon}{0.5} \right)^{0.4}
\end{aligned} \tag{4.9}$$

where J_{circ} and r_{circ} are the angular momentum and radius of a circular orbit of the same energy, $M_{\text{h,m}}$ and $M_{\text{h,s}}$ are the mass of the primary halo and the mass of the satellite halo at the initial time, $\mu = M_{\text{h,m}}/M_{\text{h,s}}$ is the ratio of the two, and $\epsilon = J/J_{\text{circ}}$ is the initial circularity of the orbit (e in the equation is simply Euler's number). The second equality assumes the energy of the orbit is equal to that of a circular orbit at the virial radius (with period P_{vir}), while in the final equality, we have substituted the radial orbital period P_{rad} (Binney & Tremaine, 2008). By comparison, the first and second pericentric passages, where tidal effects are strongest and quenching through triggered starbursts is likely to occur, take place after approximately $0.2 P_{\text{rad}}$ and $1.2 P_{\text{rad}}$, respectively.

Suppose satellites are all star-forming on initial infall into the main system (as indicated by Fig 4.13). If systems are completely quenched at the first pericentric passage, and then fall into the centre of the main halo and merge by τ_{DF} , then assuming a uniform distribution of infall times, the quiescent fraction of the surviving satellites will be $(\tau - 0.2)/\tau = 1 - 0.2/\tau$, where $\tau \equiv T_{\text{DF}}/P_{\text{rad}}$. Similarly, if systems are only quenched after the second pericentric passage, the quiescent fraction should be $1 - 1.2/\tau$.

Figure 4.16 shows the quiescent fraction for our satellite sample, binned by mass ratio (points with errorbars). We see a clear difference between the first data point and the subsequent ones; it is $\sim 2\sigma$ below the next bin, and deviates from the average of the others by almost 3σ . The smooth curves indicate the expected fraction as a function of mass ratio, if quenching is 100% efficient at the first pericentric passage (short dashed lines), 100% efficient at the second pericentric passage (solid lines), or 50% quenching efficiency at the first and 30% efficiency at the second (long-dashed line). In the first two cases, the three curves are for orbital circularities $\epsilon = 0.9, 0.5, 0.1$ from top to bottom; in the final case, for clarity we show only the results for $\epsilon = 0.5$. While we have not fit the data explicitly to any of these models, a model with partial quenching at each of the first two pericentric passages clearly matches the general trend in our results.

In the limit of large mass ratios, the quiescent fraction goes to ~ 0.65 , rather than 100%. This could be because a significant fraction of all satellites have fallen in recently (e.g. Taylor & Babul, 2005, suggest the median number of orbits spent in the main halo is around 2 for low-mass

satellites); it could be because some satellites are on circular orbits with large pericentres, and don't experience strong tidal triggering, or it could be because quenching/triggering is not 100% effective at quenching galaxies, even on radial orbits. We also note that the lowest bin may be affected by misclassification of near-equal-mass secondaries. Since the scatter in the SHMR is ~ 0.16 dex (cf. chapter 3), the uncertainty in the primary-secondary mass ratio is approximately 0.25 dex, so a $+1\sigma$ deviation could result in an primary and secondary with the median mass ratio of the first bin being switched. Assuming this positive deviation happens 16% of the time, and that the field quenched fraction is low (cf. Fig. 4.13), this will reduce the quiescent fraction measured in this bin by up to 16%. The red diamond shows the value after correcting for this effect.

4.5.4 Quenching: Summary

We can summarize our results on environmental quenching as follows:

1. The quiescent fraction in our satellite populations does not appear to depend strongly on secondary mass (Fig. 4.13), at least not below secondary stellar masses $\log(M_*/M_\odot) \sim 10$ where we can measure it reliably.
2. The quiescent fraction may depend on primary colour (cf. Fig. 4.14); blue primaries appear to have lower quiescent fractions than red primaries, though the difference is only significant at 1σ . This is consistent with the idea of "conformity", for which there is clear evidence in the spectroscopic sample (cf. Section 4.4.2).
3. The quiescent fraction may depend on primary mass, with a higher fraction in more massive halos (cf. Fig. 4.15), though we lack the SNR to establish this conclusively.
4. There is stronger evidence for the quiescent fraction depending on the primary-to-secondary mass ratio (cf. Fig. 4.16). This is consistent with a model where the quiescent fraction reflects the mean accretion time of satellites, and massive satellites that have only recently fallen in to the main halo are not yet quenched (see Section 4.5.3).

Given the complex, multi-variate dependence of quiescence on primary and secondary properties, we clearly need more data to map out these trends in detail. For low-mass groups, dedicated spectroscopic campaigns targeting local examples, such as Mao et al. (2021), are probably the most promising approach, as they provide additional information about satellite orbits, quenching times and infall times.




Figure 4.16: Satellite quiescent fraction versus the mass ratio of the primary halo mass to the satellite infall halo mass (black points with errorbars). Curves indicate the expected fraction, assuming 100% quenching efficiency at the first pericentric passage (short dashed lines), 100% quenching efficiency at the second pericentric passage (solid lines), or 50% quenching efficiency at the first and 30% efficiency at the second (long-dashed line). In the first two cases, the three curves are for orbital circularities $\epsilon = 0.9, 0.5, 0.1$ from top to bottom; in the final case, for clarity we show only the results for $\epsilon = 0.5$. Note the lowest bin may be affected by misclassification of near-equal mass secondaries; the red diamond shows the corrected value if we assume these artificially reduce the quiescent fraction by 16%.

4.6 Summary and Outlook

In chapter 3, we established a new, iterative method for quantifying satellite abundance using clustering in catalogues with accurate photo-zs. The method allows us to measure the clustering signal even in crowded fields, avoiding biases that may arise from selecting only the most isolated systems. We tested the method using the COSMOS2015 photo-z catalogue of [Laigle et al. \(2016\)](#), which has excellent photo-z accuracy, but covers only a very small field.

In this chapter, we have explored the properties of the satellite populations detected in the COSMOS field. We measure the amplitude and shape of the SSMF, and find results for both that are consistent with previous measurements by [Wang & White \(2012\)](#) and [Lan, Ménard & Mo \(2016\)](#), but extend these down to primary masses $\log(M_*/M_\odot) \sim 10.2$. We also measure the SSMF per unit halo mass, an indicator of the net efficiency of galaxy formation in different environments, and find fairly constant efficiency at primary halo masses $\log(M_h/M_\odot) \gtrsim 12$, consistent with previous studies (e.g. [Zu & Mandelbaum, 2015](#)), while at lower primary halo masses there is some marginal evidence for reduced efficiency. Expressing the SSMF as a function of the primary-to-secondary mass ratio (the ‘RSMF’), we find marginal evidence for dynamical friction effects depleting the low-mass ratio end of the satellite population (i.e. relatively massive satellites). Splitting the primary sample by colour, we find some evidence for greater satellite abundance around red primaries, particularly for large secondary masses. Here again, this is consistent with previous results where we overlap in primary mass (e.g. [Wang & White, 2012](#)). Examining the spectroscopic subsample within our data, we also see evidence for conformity in colour; in particular, red satellites appear rarer around blue galaxies. Finally, we study quenching as a function of secondary mass and primary properties, seeing evidence for a number of trends. The most significant one (and previously unreported, to our knowledge), is a selection effect whereby the satellites with the largest relative masses have lower quiescent fractions (cf. [Fig. 4.16](#)), because they have only merged into the main halo recently, and have not yet been quenched by pericentric passages.

Overall, our results show good consistency with previous studies, and extend these to lower primary mass; they are limited, however, by low SNRs and the small size of the COSMOS field. Many large surveys currently underway, or planned for the near future, will greatly increase the deep imaging data available; on the other hand, most will also have much larger photo-z errors, relative to COSMOS. One of the most promising surveys for applying our method is the deep polar-cap component of SPHEREx¹, an all-sky survey satellite with a wide-field spectral imager that will produce low-resolution ($R \sim 20\text{--}100$) spectra, with a final redshift accuracy similar to COSMOS, but over an area 50 times larger (see chapter 3). The resulting increase in the SNR of

¹see this <https://spherex.caltech.edu/> for more details

the clustering signal should confirm or rule out many of the marginal trends seen in the current work, and further clarify the complex relationship between satellite and central galaxies.

Acknowledgements

We thank N. Afshordi, A. Broderick, M. Balogh, M. Hudson, and our friends and collaborators from the COSMOS survey for their comments and advice. JET acknowledges support from the Natural Sciences and Engineering Research Council of Canada, through a Discovery Grant. The COSMOS2015 catalogue is based on data products from observations made with European Southern Observatory (ESO) Telescopes at the La Silla Paranal Observatory under ESO programme ID 179. A-2005, and on data products produced by TERAPIX and the Cambridge Astronomy Survey Unit on behalf of the UltraVISTA consortium.

Data Availability

Most of the basic data presented in this article are publicly available. The COSMOS2015 catalogue (Laigle et al., 2016) can be accessed from the COSMOS website, at <http://cosmos.astro.caltech.edu/page/photom>. A few spectroscopic redshifts that are unpublished from the COSMOS collaboration (M. Salvato, private communication) will be shared on reasonable request to the corresponding author with permission of the COSMOS collaboration. The derived data generated for this chapter will also be shared on reasonable request to the corresponding author.

Chapter 5

Conclusion

We will conclude the thesis by summarizing the key findings of each chapter and discussing the significance of each in relation to our initial research aims. We will also review the limitations of our work, and propose opportunities for future research.

The properties of dwarf galaxies provide important tests for the theories of galaxy formation, cosmic structure formation, and the nature of dark matter. While the properties of bright galaxies are relatively well determined, those of faint dwarf galaxies are much less well known. To understand the Local Group relative to other similar environments, however, it is precisely these objects we have to identify and study. To do so efficiently requires new methods for estimating distances and group membership, either individually or statistically. This thesis presented three papers that aim to advance techniques, in that area, and open up opportunities for future research with new samples.

In chapter 2, we tested the effectiveness of using structural properties to identify nearby dwarf populations statistically, selecting samples using cuts in magnitude, surface brightness and size, and confirming the fraction of galaxies that are indeed nearby by measuring their clustering with respect to bright galaxies with known redshifts. Although we considered only a relatively small field in the sky, (the 1.5 square degree COSMOS field), we demonstrated that structural selection works surprisingly well out to a redshift of 0.15. The sample selected by this method is not complete nor unbiased, but can nevertheless be very useful in placing a lower limit on mean satellite abundance. In the best case, we are able to recover about two-thirds of the clustering signal measured from samples selected using the extremely high-quality COSMOS photo-zs, with 80% of the SNR, and a net purity of 33% on the final sample.

We deliberately tested our method using SDSS photometry data within the COSMOS field. The SDSS photometry has fairly bright surface brightness limits ($\mu_{\text{eff}} \leq 25\text{--}27$) and limited spatial

resolution. New generations of surveys with deeper photometry and better image resolution may push the detection limits much deeper, and therefore explore a much larger parameter space in structural properties. We are able to find hundreds of MW-like systems in our analysis of a very small sky field. Large-area surveys such as Euclid and WFIRST will identify hundreds of thousands of such systems, and thus greatly improve the statistics of our detections.

In chapter 3, we developed and tested a more complex method of using clustering to quantify satellite abundance. This method allows us to measure satellite abundance around centrals even in crowded fields, and thus it can help to reduce biases introduced by selecting only the most isolated systems. The method relies on high-precision photo-zs, rather than spectroscopic redshifts, and thus it can push satellite detection to greater depths with a relatively modest investment of observing time.

With this method, we are able to estimate luminosity functions and various other properties of satellite populations over a very broad range of primary halo mass (10^{10} – $10^{13.5} M_{\odot}$). We compared our estimate of satellite abundance to several previous results from the literature that were derived from larger data sets, including [Conroy, Wechsler & Kravtsov \(2006\)](#) for high-mass primaries and [Besla et al. \(2018\)](#) for low-mass primaries. In both cases, we found excellent consistency, while our results cover a much larger range of primary mass overall. We also compared our results for primaries with a halo mass of around $12.1 M_{\odot}$ to the luminosity function of the two main LG primaries, the MW and M31. The LG luminosity functions lie close to the measured average, except for the two brightest satellites of MW (the LMC and SMC). This confirms that the LG is statistically rare in having two bright satellites like the LMC and SMC. Finally, testing the method for systematic uncertainties by varying the model parameters used in the analysis, we confirmed that systematic uncertainties in the method are generally smaller than our random error estimates.

In chapter 4, we used the method established in the previous chapter to explore the properties of the local satellite populations detected within the COSMOS field. We measured the shape and amplitude of the satellite stellar mass function (SSMF) and find results that are consistent with previous literature (e.g. [Wang & White, 2012](#); [Lan, Ménard & Mo, 2016](#)). At the same time, we were able to extend estimates of the SSMF down to a primary halo mass of $10^{10.2} M_{\odot}$, which is close to the lowest limit in the literature. As an indicator of the net efficiency of galaxy formation in different environments, we calculated the SSMF per unit halo mass. We found fairly constant efficiency for larger primary halo masses ($\log(M_h/M_{\odot}) \leq 12$), which is also consistent with previous studies ([Zu & Mandelbaum, 2015](#), e.g.). At lower primary halo mass, our results show some evidence for reduced efficiency, but only with marginal significance. We also calculated the RSMF, that is the SSMF as a function of primary-secondary halo mass ratio, and found some slight evidence that dynamical friction effects deplete the low-mass ratio end of the satellite population (i.e. relatively massive satellites).

Splitting the SSMR by primary colour, we found that red primaries have marginally higher satellite abundance, especially at the high secondary-mass end (relatively massive satellites). This is also consistent with previous studies (e.g. [Wang & White, 2012](#)), at least in the range where we overlap in primary mass. While examining the spectroscopic subsample within our data, we found some evidence for colour conformity, in that red satellites seem particularly rare around blue primaries.

Finally, we studied quenching as a function of primary and secondary mass, noticing several trends. The most significant one is that relatively massive satellites (i.e. those with small primary-to-secondary mass ratios) have a lower quiescent fraction. We speculate that this is because satellites in this relative mass range must have merged recently, since they will not survive long in the main system due to dynamical friction effects; thus a large fraction has not yet been quenched by pericentric passages, and continue to form stars.

5.1 Discussion

The work presented in the thesis has a number of shortcomings and limitations, of which we summarize the main ones below.

- Our work has focused on developing and testing new methods. In practice, the significance of many of our measurements is limited by the small volume probed by the COSMOS field at low redshift. With data from large-area surveys, we could not only improve the statistics of our current measurements, but focus on the lowest redshift range, where our methods are most effective. For the moment, however, most current and forthcoming surveys have limited quality photo-zs, so the techniques of chapters 3 & 4 will be less effective.
- In chapter 2, we only considered the simplest linear cuts in size and magnitude or surface brightness and magnitude. This was partly due to the limited SNR of the clustering signal, which would make the effectiveness of more complicated cuts harder to assess. With larger data sets and higher SNR, however, we could explore the problem in a larger number of dimensions and test more sophisticated selection criteria.
- The method developed in chapters 3 & 4 relies on high-quality photo-zs, as mentioned previously. We have not tested it on samples with the lower precision photo-zs expected from 5-band photometry. Redshift uncertainties will increase the depth of the ROI around each primary, as described in section [3.3.2](#). This will not only introduce more noise from foreground/background populations, decreasing SNR of any clustering detection, but it may

also raise issues of sample completeness. Completeness is strongly related to distance; if the ROI gets too large as a result of poor photo-zs, the completeness limit at the backs of the ROI may be significantly different from that at the front. Additional precautions might then be needed, such as applying more aggressive cuts in magnitude or modelling the variation in completeness along the line of sight. We would recommend testing these effects extensively with mock catalogues, to develop a better understanding of this issue.

- When selecting the primaries (section 2.4.1), we excluded candidates close to the field boundaries, as they may be satellites of more massive galaxies outside the field. This means that we lose a certain percentage of the effective survey area, especially at low redshift, where the radius of the ROI corresponds to a large angle. This would be less of an issue for contiguous, large-area surveys, where the affected area makes up a smaller fraction of the total sky coverage. One might also address this issue by identifying large galaxies near the main field in shallower ancillary data.
- In section 3.4.2, given the relatively low SNR of the clustering signal, we used a simple, approximate model (Eqn. 3.5) to fit its radial dependence profile of the clustering signal. With larger sets of data and better SNR, one should use a projected NFW or Einasto profile, and consider varying the concentration parameter with primary mass.
- When fitting the clustering signal, we ignored satellites at the smallest projected separations (see section 3.4 and section 3.4.2), as artifacts such as blending between galaxies or fragmentation of single galaxies occur more frequently on small scales. This means omitting a part of the satellite population that lies very close to the primary. The effect is small overall but may lead to bias in certain cases. For example, galaxies with distant H_2 regions can block or confuse satellite detections at small separations, relative to early-type galaxies. One could try to model and/or correct for these effects, so as to capture the innermost satellite population.
- We modelled the bias of SHMR with a simple linear fit at the high mass end (see section C.1). This helped us to avoid overestimating the halo mass associated with large galaxies. Within the mass that is considered in the work of this thesis, our simple model actually works very well and shows good consistency with the literature (e.g. Zu & Mandelbaum, 2015, Figure 11). However, a more sophisticated model may work more consistently within a larger range of mass. More importantly, the potential uncertainty of the fitting parameters was not considered/propagated into subsequent calculations. We tested and demonstrated that the potential effects of SHMR uncertainties are not major, but properly quantifying and utilizing these uncertainties, e.g. by using Markov Chain Monte Carlo methods, would make our results more robust.

- In section 3.5, we introduced two background-correction methods, “method A” and “method B” to estimate satellite abundance given secondary cuts on satellite properties. While “Method B” is an improvement over “method A”, the optimal approach (which we referred to as “method C”) would be to divide the sample by the secondary property first (e.g. by colour into red/blue galaxies) and then measure the clustering amplitude separately on each subsample. This would mitigate any potential bias where the secondary property is correlated with clustering strength. We did not attempt “method C” in COSMOS, as we did not have the SNR to do so. This should be an obvious improvement when applying our method to a significantly larger data set.
- We relied on the “CLASS” flag provided by the COSMOS2015 catalogue to distinguish quiescent galaxies from star-forming ones (see section 4.5). As we pointed out in chapter 4, the cuts the catalogues used to determine quiescent galaxies are very conservative. This makes it harder to compare our results to quenching to the literature. Furthermore, the quiescent cuts in COSMOS2015 were optimized for the whole redshift range of the catalogue up to $z \sim 6$. Since our focus is only on nearby systems at $z < 0.25$, it might be worth trying a dedicated analysis optimized to this low redshift range.

Despite these limitations, our work provides useful tools for exploring faint satellite populations outside the LG. Both the morphological selection developed in chapter 2 and the satellite probability estimate developed in chapters 3 and 4 can also be useful in providing nearby dwarf candidates for spectroscopic follow-up, greatly improving the efficiency of such surveys.

Our work in chapters 3 and 4 explored the potential of using high-quality photo-zs to search for nearby dwarf galaxies, rather than depending on spectroscopic redshifts. Although spectroscopic measurements will always be required to confirm individual satellites, photo-zs provide a cheap way to obtain statistical information down to much fainter magnitudes. Currently, COSMOS is one of the few datasets with sub-percent precision photo-zs. An important forthcoming survey is SPHEREx (Stickley et al., 2016), which features the innovative approach of imaging with rectangular linear variable filters (LVFs). This technique allows the instrument to scan and perform photometry in multiple bands simultaneously, and then stack them into a full sky, low-resolution spectral (0.75–4.2 μm at $R = 41$; and 4.2–5 μm at $R = 135$) survey. The net quality of the redshifts produced by this technique is expected to be comparable to COSMOS photo-zs or better. While most of the survey area will be shallower than COSMOS, the two polar regions will be visited multiple times, producing a final depth comparable to COSMOS over an area 50 times larger (about 100 square degrees). With the future data from SPHEREx, we can overcome many of our current limitations by achieving ~ 7 times greater SNR in the clustering signal. This should allow us to split primary and secondary samples further, and study their properties in exquisite detail. More importantly, this mission will highlight the value of high-quality photo-zs

in observational studies of galaxy evolution and cosmology, which we hope may influence the observation strategy of future astronomical surveys.

5.2 Future Research Directions

As discussed above, the work presented in this thesis suffers from a number of limitations and compromises, many of which came from the limited survey field and sample size. With the new generation of surveys, there are many opportunities for improvements for future studies.

In chapter 2, we developed a method to identify nearby dwarfs, based on structural properties including magnitude, size and surface density. One key idea was to use the known distribution of nearby dwarfs in this parameter space to identify other nearby systems. For future work, it might be interesting to apply more sophisticated techniques, such as machine learning, to explore this or other spaces of structural parameters. In going to higher-dimensional spaces, there is of course a risk of over-fitting the sample data, so some caution will be required.

In general, there is a trade-off in the clustering method between SNR and how deep we can go in brightness and stellar mass. If we want to put more emphasis on the low-mass end of the satellite population, one obvious solution is to focus on the very nearest systems, but this will lower the overall sample size and lead to a lower SNR of clustering detection. In practice, we found we were not able to measure a significant signal for primaries at $z < 0.07$ in COSMOS, given the size of the field. With surveys with similar depth but a much larger field, it would be quite interesting to study this very low redshift range.

As mentioned above, we used several simplified models, such as a two-parameter fit for the SHMR at the high mass end, and a simple model for the projected halo density profile. These worked sufficiently well at the SNR that we were able to achieve, but could be improved in future work with larger samples. Very high SNR data should even allow us to study the true radial distribution of satellites, as a function of primary mass, morphology and other properties.

Thus, there remains much interesting science to explore by combining the methods we have developed with new, larger datasets.

References

- Abraham R. G., van Dokkum P. G., 2014, *PASP*, 126, 55
- Aihara H. et al., 2018, *PASJ*, 70, S4
- Albareti F. D. et al., 2017, *ApJS*, 233, 25
- Anderhalden D., Schneider A., Macciò A. V., Diemand J., Bertone G., 2013, *JCAP*, 2013, 014
- Anderson M. E., Gaspari M., White S. D. M., Wang W., Dai X., 2015, *MNRAS*, 449, 3806
- Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, *MNRAS*, 426, 2046
- Babcock H. W., 1939, *Lick Observatory Bulletin*, 19, 41
- Bailin J., Steinmetz M., 2005, *ApJ*, 627, 647
- Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Ž., Lupton R. H., Nichol R. C., Szalay A. S., 2004, *ApJ*, 600, 681
- Balogh M. L., Baldry I. K., Nichol R., Miller C., Bower R., Glazebrook K., 2004, *ApJ*, 615, L101
- Barro G. et al., 2017, *ApJ*, 840, 47
- Bechtol K. et al., 2015, *ApJ*, 807, 50
- Behroozi P. S., Conroy C., Wechsler R. H., 2010, *ApJ*, 717, 379
- Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
- Bell E. F. et al., 2004, *ApJ*, 608, 752
- Benson A. J., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2002, *MNRAS*, 333, 177

Berlind A. A., Weinberg D. H., 2002, *ApJ*, 575, 587

Besla G. et al., 2018, *MNRAS*, 480, 3376

Bett P., Eke V., Frenk C. S., Jenkins A., Helly J., Navarro J., 2007, *MNRAS*, 376, 215

Binney J., Tremaine S., 2008, *Galactic Dynamics: Second Edition*

Bland-Hawthorn J., Gerhard O., 2016, *ARA&A*, 54, 529

Blanton M. R. et al., 2003, *ApJ*, 594, 186

Blanton M. R., Lupton R. H., Schlegel D. J., Strauss M. A., Brinkmann J., Fukugita M., Loveday J., 2005, *ApJ*, 631, 208

Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, *Nature*, 311, 517

Bock J., SPHEREx Science Team, 2018, in *American Astronomical Society Meeting Abstracts*, Vol. 231, *American Astronomical Society Meeting Abstracts #231*, p. 354.21

Bode P., Ostriker J. P., Turok N., 2001, *ApJ*, 556, 93

Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, *MNRAS*, 415, L40

Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., 2010, *Monthly Notices of the Royal Astronomical Society*, 406, 896

Brainerd T. G., Blandford R. D., Smail I., 1996, *ApJ*, 466, 623

Bremer M. N. et al., 2018, *MNRAS*, 476, 12

Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151

Brooks A. M., Zolotov A., 2014, *ApJ*, 786, 87

Brown T. et al., 2017, *MNRAS*, 466, 1275

Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000

Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80

Bullock J. S., Boylan-Kolchin M., 2017, *ARA&A*, 55, 343

Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, *ApJ*, 539, 517

Busha M. T., Wechsler R. H., Behroozi P. S., Gerke B. F., Klypin A. A., Primack J. R., 2011, *The Astrophysical Journal*, 743, 117

Buta R. J., 2013, in *Planets, Stars and Stellar Systems. Volume 6: Extragalactic Astronomy and Cosmology*, Oswalt T. D., Keel W. C., eds., Vol. 6, p. 1

Callingham T. et al., 2018, arXiv e-prints, arXiv:1808.10456

Capak P. et al., 2007, *ApJS*, 172, 99

Carlin J. L. et al., 2016, *ApJ*, 828, L5

Carlsten S. G., Beaton R. L., Greco J. P., Greene J. E., 2019, *ApJ*, 879, 13

Chiboucas K., Jacobs B. A., Tully R. B., Karachentsev I. D., 2013, *AJ*, 146, 126

Chilingarian I. V., Melchior A.-L., Zolotukhin I. Y., 2010, *MNRAS*, 405, 1409

Chua K. T. E., Pillepich A., Rodriguez-Gomez V., Vogelsberger M., Bird S., Hernquist L., 2017, *MNRAS*, 472, 4343

Cohen Y. et al., 2018, *ApJ*, 868, 96

Cole S., Lacey C., 1996, *MNRAS*, 281, 716

Colpi M., Mayer L., Governato F., 1999, *ApJ*, 525, 720

Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201

Conselice C. J., 2014, *ARA&A*, 52, 291

Contini E., Gu Q., Ge X., Rhee J., Yi S. K., Kang X., 2020, arXiv e-prints, arXiv:2001.01369

Cooray A., Sheth R., 2002, *Phys. Rep.*, 372, 1

Correa C. A., Schaye J., 2020, *MNRAS*, 499, 3578

Coupon J. et al., 2012, *A&A*, 542, A5

Crain R. A. et al., 2015, *MNRAS*, 450, 1937

Crnojević D. et al., 2016, *ApJ*, 823, 19

Curtis H. D., 1988, PASP, 100, 6

Dai X., Bregman J. N., Kochanek C. S., Rasia E., 2010, ApJ, 719, 119

Danieli S., van Dokkum P., Abraham R., Conroy C., Dolphin A. E., Romanowsky A. J., 2019, arXiv e-prints, arXiv:1910.07529

Danieli S., van Dokkum P., Conroy C., 2018, ApJ, 856, 69

Danieli S., van Dokkum P., Merritt A., Abraham R., Zhang J., Karachentsev I. D., Makarova L. N., 2017, ApJ, 837, 136

Davies L. J. M. et al., 2016, MNRAS, 455, 4013

Davies L. J. M. et al., 2019, arXiv e-prints

Dekel A., Silk J., 1986, ApJ, 303, 39

Diemand J., Kuhlen M., Madau P., Zemp M., Moore B., Potter D., Stadel J., 2008, Nature, 454, 735

Dolag K., Komatsu E., Sunyaev R., 2016, MNRAS, 463, 1797

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

Drlica-Wagner A. et al., 2020, ApJ, 893, 47

Drlica-Wagner A. et al., 2015, ApJ, 813, 109

Eadie G. M., Harris W. E., 2016, ApJ, 829, 108

Einasto J., Saar E., Kaasik A., Chernin A. D., 1974, Nature, 252, 111

Elbert O. D., Bullock J. S., Garrison-Kimmel S., Rocha M., Oñorbe J., Peter A. H. G., 2015, MNRAS, 453, 29

Evans J., 1998, The history & practice of ancient astronomy

Evans T. A., Fattahi A., Deason A. J., Frenk C. S., 2020, MNRAS, 497, 4311

Faber S. M. et al., 2007, ApJ, 665, 265

Fattahi A., Navarro J. F., Frenk C. S., Oman K. A., Sawala T., Schaller M., 2018, MNRAS, 476, 3816

Ferrarese L. et al., 2012, ApJS, 200, 4

Fitts A. et al., 2017, MNRAS, 471, 3547

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306

Frenk C. S., White S. D. M., 2012, Annalen der Physik, 524, 507

Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, ApJ, 327, 507

Fritz T. K., Battaglia G., Pawlowski M. S., Kallivayalil N., van der Marel R., Sohn S. T., Brook C., Besla G., 2018, A&A, 619, A103

Fry A. B. et al., 2015, MNRAS, 452, 1468

Gao L., Springel V., White S. D. M., 2005, MNRAS, 363, L66

Gao L., White S. D. M., Jenkins A., Stoehr F., Springel V., 2004, MNRAS, 355, 819

Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., Lee K., 2014, MNRAS, 438, 2578

Geha M., Blanton M. R., Yan R., Tinker J. L., 2012, ApJ, 757, 85

Geha M. et al., 2017, ApJ, 847, 4

Genel S. et al., 2014, MNRAS, 445, 175

Ghigna S., Moore B., Governato F., Lake G., Quinn T., Stadel J., 1998, MNRAS, 300, 146

Gillessen S. et al., 2017, ApJ, 837, 30

Giodini S. et al., 2009, ApJ, 703, 982

Gnedin N. Y., Kravtsov A. V., 2006, ApJ, 645, 1054

Governato F. et al., 2010, Nature, 463, 203

Greco J. P. et al., 2018, ApJ, 857, 104

Grossauer J. et al., 2015, ApJ, 807, 88

Gunn J. E., Gott, J. Richard I., 1972, ApJ, 176, 1

Guo H., Zehavi I., Zheng Z., 2012, ApJ, 756, 127

Guo Q., Cole S., Eke V., Frenk C., 2011, MNRAS, 417, 370

Guo Q., Cole S., Eke V., Frenk C., 2012, MNRAS, 427, 428

Hargis J. R., Willman B., Peter A. H. G., 2014, ApJ, 795, L13

Hartley W. G., Conselice C. J., Mortlock A., Foucaud S., Simpson C., 2015, MNRAS, 451, 1613

Hayashi E., Navarro J. F., Springel V., 2007, MNRAS, 377, 50

Hearin A. P., Zentner A. R., Berlind A. A., Newman J. A., 2013, MNRAS, 433, 659

Heitmann K. et al., 2015, ApJS, 219, 34

Holmberg E., 1941, ApJ, 94, 385

Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, MNRAS, 445, 581

Hubble E. P., 1927, The Observatory, 50, 276

Hubble E. P., 1929, ApJ, 69

Hudson M. J. et al., 2015, MNRAS, 447, 298

Ilbert O. et al., 2015, A&A, 579, A2

Ilbert O. et al., 2009, ApJ, 690, 1236

Ilbert O. et al., 2013, A&A, 556, A55

Impey C. D., Sprayberry D., Irwin M. J., Bothun G. D., 1996, ApJS, 105, 209

Javanmardi B. et al., 2016, A&A, 588, A89

Jiang F., van den Bosch F. C., 2017, MNRAS, 472, 657

Jing Y. P., Mo H. J., Börner G., 1998, ApJ, 494, 1

Jing Y. P., Suto Y., 2002, ApJ, 574, 538

Kang X., Wang L., Luo Y., 2016, Monthly Notices of the Royal Astronomical Society, 460, 2152

Karachentsev I. D., Karachentseva V. E., Huchtmeier W. K., Makarov D. I., 2004, AJ, 127, 2031

Karachentsev I. D., Kashibadze O. G., 2006, *Astrophysics*, 49, 3

Karachentsev I. D., Makarov D. I., Kaisina E. I., 2013, *AJ*, 145, 101

Katz H. et al., 2020, *MNRAS*, 494, 2200

Kauffmann G. et al., 2003a, *MNRAS*, 346, 1055

Kauffmann G. et al., 2003b, *MNRAS*, 341, 33

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

Kauffmann G., White S. D. M., Heckman T. M., Ménard B., Brinchmann J., Charlot S., Tremonti C., Brinkmann J., 2004, *MNRAS*, 353, 713

Kennedy R., Frenk C., Cole S., Benson A., 2014, *Monthly Notices of the Royal Astronomical Society*, 442, 2487

Klypin A., Gottlöber S., Kravtsov A. V., Khokhlov A. M., 1999a, *ApJ*, 516, 530

Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999b, *ApJ*, 522, 82

Klypin A., Zhao H., Somerville R. S., 2002, *ApJ*, 573, 597

Kniazev A. Y., Grebel E. K., Pustilnik S. A., Pramskij A. G., Kniazeva T. F., Prada F., Harbeck D., 2004, *AJ*, 127, 704

Knobel C., Lilly S. J., Woo J., Kovač K., 2015, *ApJ*, 800, 24

Koekemoer A. M. et al., 2007, *ApJS*, 172, 196

Koekemoer A. M. et al., 2011, *ApJS*, 197, 36

Koposov S. E., Belokurov V., Torrealba G., Evans N. W., 2015, *ApJ*, 805, 130

Kourkchi E., Tully R. B., 2017, *ApJ*, 843, 16

Kravtsov A. V., Berlind A. A., Wechsler R. H., Klypin A. A., Gottlöber S., Allgood B., Primack J. R., 2004, *ApJ*, 609, 35

Kravtsov A. V., Gnedin O. Y., Klypin A. A., 2004, *ApJ*, 609, 482

Kuhlen M., Diemand J., Madau P., 2007, *ApJ*, 671, 1135

Laevens B. P. M. et al., 2015a, ApJ, 813, 44

Laevens B. P. M. et al., 2015b, ApJ, 802, L18

Laigle C. et al., 2016, ApJS, 224, 24

Lan T.-W., Ménard B., Mo H., 2016, MNRAS, 459, 3998

Lang P. et al., 2014, ApJ, 788, 11

Larson R. B., Tinsley B. M., Caldwell C. N., 1980, ApJ, 237, 692

Leauthaud A. et al., 2012, ApJ, 744, 159

Lilly S. J. et al., 2007, ApJS, 172, 70

Liu L., Gerke B. F., Wechsler R. H., Behroozi P. S., Busha M. T., 2011, ApJ, 733, 62

Lokas E. L., Kazantzidis S., Mayer L., 2012, ApJ, 751, L15

Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, MNRAS, 439, 300

Lunnan R., Vogelsberger M., Frebel A., Hernquist L., Lidz A., Boylan-Kolchin M., 2012, ApJ, 746, 109

Macciò A. V., Fontanot F., 2010, MNRAS, 404, L16

Madau P., Shen S., Governato F., 2014, ApJ, 789, L17

Mandelbaum R., Seljak U., Kauffmann G., Hirata C. M., Brinkmann J., 2006, MNRAS, 368, 715

Mandelbaum R., Wang W., Zu Y., White S., Henriques B., More S., 2016, MNRAS, 457, 3200

Mao Y.-Y., Geha M., Wechsler R. H., Weiner B., Tollerud E. J., Nadler E. O., Kallivayalil N., 2021, ApJ, 907, 85

Marinacci F. et al., 2018, MNRAS, 480, 5113

Marinoni C., Hudson M. J., 2002, ApJ, 569, 101

Martin D. C. et al., 2007, ApJS, 173, 342

Mashchenko S., Wadsley J., Couchman H. M. P., 2008, Science, 319, 174

Masters D. C., Stern D. K., Cohen J. G., Capak P. L., Rhodes J. D., Castander F. J., Paltani S., 2017, *ApJ*, 841, 111

Mayer L., Mastrogiuseppe C., Wadsley J., Stadel J., Moore B., 2006, *MNRAS*, 369, 1021

McConnachie A. W., 2012, *AJ*, 144, 4

McCracken H. J. et al., 2012, *A&A*, 544, A156

Ménard B., Scranton R., Schmidt S., Morrison C., Jeong D., Budavari T., Rahman M., 2013, *ArXiv e-prints*

Merchán M. E., Zandivarez A., 2005, *ApJ*, 630, 759

Merritt A., van Dokkum P., Abraham R., 2014, *ApJ*, 787, L37

Miyazaki S. et al., 2012, in *Proc. SPIE*, Vol. 8446, *Ground-based and Airborne Instrumentation for Astronomy IV*, p. 84460Z

Mo H., van den Bosch F. C., White S., 2010, *Galaxy Formation and Evolution*

Mobasher B. et al., 2007, *The Astrophysical Journal Supplement Series*, 172, 117

Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999a, *ApJ*, 524, L19

Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, *Nature*, 379, 613

Moore B., Lake G., Quinn T., Stadel J., 1999b, *MNRAS*, 304, 465

More S., van den Bosch F. C., Cacciato M., Mo H. J., Yang X., Li R., 2009, *MNRAS*, 392, 801

More S., van den Bosch F. C., Cacciato M., Skibba R., Mo H. J., Yang X., 2011, *MNRAS*, 410, 210

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

Moustakas J. et al., 2013, *ApJ*, 767, 50

Müller O., Scalera R., Binggeli B., Jerjen H., 2017, *A&A*, 602, A119

Nadler E. O., Gluscevic V., Boddy K. K., Wechsler R. H., 2019, *ApJ*, 878, L32

Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563

Navarro J. F. et al., 2004, MNRAS, 349, 1039

Nelson D. et al., 2015, Astronomy and Computing, 13, 12

Newman J. A. et al., 2013, ApJS, 208, 5

Newton O., Cautun M., Jenkins A., Frenk C. S., Helly J. C., 2018, MNRAS, 479, 2853

Nichols M., Bland-Hawthorn J., 2011, ApJ, 732, 17

Nierenberg A. M., Treu T., Menci N., Lu Y., Torrey P., Vogelsberger M., 2016, MNRAS, 462, 4473

Oñorbe J., Boylan-Kolchin M., Bullock J. S., Hopkins P. F., Kereš D., Faucher-Giguère C.-A., Quataert E., Murray N., 2015, MNRAS, 454, 2092

Ostriker J. P., Peebles P. J. E., 1973, ApJ, 186, 467

Ostriker J. P., Peebles P. J. E., Yahil A., 1974, ApJ, 193, L1

Papastergis E., Ponomareva A. A., 2017, A&A, 601, A1

Patel E., Besla G., Mandel K., Sohn S. T., 2018, ApJ, 857, 78

Patton D. R., Ellison S. L., Simard L., McConnachie A. W., Mendel J. T., 2011, MNRAS, 412, 591

Pawlowski M. S., Famaey B., Merritt D., Kroupa P., 2015, ApJ, 815, 19

Peacock J. A., Smith R. E., 2000, MNRAS, 318, 1144

Peng Y., Maiolino R., Cochrane R., 2015, Nature, 521, 192

Peng Y.-j. et al., 2010, ApJ, 721, 193

Peng Y.-j., Lilly S. J., Renzini A., Carollo M., 2012, ApJ, 757, 4

Planck Collaboration et al., 2014, A&A, 571, A16

Planck Collaboration et al., 2016, A&A, 594, A13

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

Poggianti B. M. et al., 2017, ApJ, 844, 48

Polletta M. et al., 2007, ApJ, 663, 81

Ponman T. J., Allan D. J., Jones L. R., Merrifield M., McHardy I. M., Lehto H. J., Luppino G. A., 1994, Nature, 369, 462

Posti L., Helmi A., 2019, A&A, 621, A56

Press W. H., Schechter P., 1974, ApJ, 187, 425

Rahman M., Ménard B., Scranton R., Schmidt S. J., Morrison C. B., 2015, MNRAS, 447, 3500

Rahman M., Mendez A. J., Ménard B., Scranton R., Schmidt S. J., Morrison C. B., Budavári T., 2016, MNRAS, 460, 163

Read J. I., Agertz O., Collins M. L. M., 2016, MNRAS, 459, 2573

Reddick R. M., Wechsler R. H., Tinker J. L., Behroozi P. S., 2013, ApJ, 771, 30

Ribas I., Jordi C., Vilardell F., Fitzpatrick E. L., Hilditch R. W., Guinan E. F., 2005, ApJ, 635, L37

Richings J. et al., 2020, MNRAS, 492, 5780

Robotham A. S. G. et al., 2012, MNRAS, 424, 1448

Robotham A. S. G. et al., 2014, MNRAS, 444, 3986

Rubin V. C., Ford, W. K. J., Thonnard N., 1980, ApJ, 238, 471

Rubin V. C., Ford, W. Kent J., 1970, ApJ, 159, 379

Sales L. V., Wang W., White S. D. M., Navarro J. F., 2013, MNRAS, 428, 573

Salim S., 2014, Serbian Astronomical Journal, 189, 1

Sand D. J. et al., 2014, ApJ, 793, L7

Sawala T. et al., 2015, MNRAS, 448, 2941

Schaefer A. L. et al., 2017, MNRAS, 464, 121

Schawinski K. et al., 2014, MNRAS, 440, 889

Schaye J. et al., 2015, MNRAS, 446, 521

Scoville N. et al., 2013, *ApJS*, 206, 3

Scoville N. et al., 2007a, *ApJS*, 172, 150

Scoville N. et al., 2007b, *ApJS*, 172, 1

Seljak U., 2000, *MNRAS*, 318, 203

Seljak U. et al., 2005, *Phys. Rev. D*, 71, 103515

Silk J., 2013, *ApJ*, 772, 112

Slipher V. M., 1913, *Lowell Observatory Bulletin*, 1, 56

Slipher V. M., 1915, *Popular Astronomy*, 23, 21

Smith A. et al., 2019, *MNRAS*, 484, 1285

Sohn S. T., Watkins L. L., Fardal M. A., van der Marel R. P., Deason A. J., Besla G., Bellini A., 2018, *ApJ*, 862, 52

Somerville R. S., 2002, *ApJ*, 572, L23

Somerville R. S., Lee K., Ferguson H. C., Gardner J. P., Moustakas L. A., Giavalisco M., 2004, *ApJ*, 600, L171

Spangelo S. C., Katti R. M., Unwin S. C., Bock J. J., 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 037001

Speller R., Taylor J. E., 2014, *ApJ*, 788, 188

Spergel D. N., Steinhardt P. J., 2000, *Phys. Rev. Lett.*, 84, 3760

Springel V. et al., 2018, *MNRAS*, 475, 676

Springel V. et al., 2008, *MNRAS*, 391, 1685

Springel V. et al., 2005, *Nature*, 435, 629

Stickley N. R., Capak P., Masters D., de Putter R., Doré O., Bock J., 2016, *arXiv e-prints*, arXiv:1606.06374

Strigari L. E., Wechsler R. H., 2012, *ApJ*, 749, 75

Taylor C. L., Thomas D. L., Brinks E., Skillman E. D., 1996, *ApJS*, 107, 143

Taylor E. N. et al., 2015, *MNRAS*, 446, 2144

Taylor J. E., Babul A., 2001, *ApJ*, 559, 716

Taylor J. E., Babul A., 2004, *MNRAS*, 348, 811

Taylor J. E., Babul A., 2005, *MNRAS*, 364, 515

Thomas P. A. et al., 1998, *MNRAS*, 296, 1061

Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, *ApJ*, 688, 709

Tinker J. L., Norberg P., Weinberg D. H., Warren M. S., 2007, *ApJ*, 659, 877

Tolstoy E., Hill V., Tosi M., 2009, *Annual Review of Astronomy and Astrophysics*, 47, 371

Trussler J., Maiolino R., Maraston C., Peng Y., Thomas D., Goddard D., Lian J., 2019, *MNRAS*, 2878

Tully R. B., Courtois H. M., Sorce J. G., 2016, *AJ*, 152, 50

Tully R. B., de Marseille O., Fisher J. R., 1975, in *BAAS*, Vol. 7, *Bulletin of the American Astronomical Society*, p. 426

Tully R. B., Rizzi L., Shaya E. J., Courtois H. M., Makarov D. I., Jacobs B. A., 2009, *AJ*, 138, 323

Vale A., Ostriker J. P., 2004, *MNRAS*, 353, 189

van den Bergh S., 2000, *The Galaxies of the Local Group*

van den Bosch F. C., Aquino D., Yang X., Mo H. J., Pasquali A., McIntosh D. H., Weinmann S. M., Kang X., 2008, *MNRAS*, 387, 79

van den Bosch F. C., Norberg P., Mo H. J., Yang X., 2004, *MNRAS*, 352, 1302

van Dokkum P. et al., 2017, *ApJ*, 844, L11

van Dokkum P., Danieli S., Cohen Y., Romanowsky A. J., Conroy C., 2018, *ApJ*, 864, L18

van Dokkum P. G., Abraham R., Merritt A., Zhang J., Geha M., Conroy C., 2015, *ApJ*, 798, L45

Velazquez H., White S. D. M., 1999, MNRAS, 304, 254

Vera-Ciro C. A., Sales L. V., Helmi A., Frenk C. S., Navarro J. F., Springel V., Vogelsberger M., White S. D. M., 2011, MNRAS, 416, 1377

Vogelsberger M. et al., 2014, MNRAS, 444, 1518

Wake D. A., van Dokkum P. G., Franx M., 2012, ApJ, 751, L44

Wang J., Frenk C. S., Navarro J. F., Gao L., Sawala T., 2012, MNRAS, 424, 2715

Wang W., Han J., Cooper A. P., Cole S., Frenk C., Lowing B., 2015, Monthly Notices of the Royal Astronomical Society, 453, 377

Wang W., Sales L. V., Henriques B. M. B., White S. D. M., 2014, MNRAS, 442, 1363

Wang W. et al., 2020, arXiv e-prints, arXiv:2009.06882

Wang W., White S. D. M., 2012, MNRAS, 424, 2574

Watkins L. L., van der Marel R. P., Sohn S. T., Evans N. W., 2018, arXiv e-prints, arXiv:1804.11348

Weaver J. R. et al., 2021, arXiv e-prints, arXiv:2110.13923

Wechsler R. H., Tinker J. L., 2018, ARA&A, 56, 435

Weinmann S. M., van den Bosch F. C., Yang X., Mo H. J., 2006, MNRAS, 366, 2

Weisz D. R. et al., 2011, ApJ, 743, 8

Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, ApJ, 827, L23

Wetzel A. R., Tinker J. L., Conroy C., 2012, MNRAS, 424, 232

Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013, MNRAS, 432, 336

Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2014, MNRAS, 439, 2687

Wetzel A. R., Tollerud E. J., Weisz D. R., 2015, ApJ, 808, L27

Wheeler C., Phillips J. I., Cooper M. C., Boylan-Kolchin M., Bullock J. S., 2014, MNRAS, 442, 1396

White S. D. M., Rees M. J., 1978, MNRAS, 183, 341

Willman B., 2010, Advances in Astronomy, 2010, 285454

Wojtak R., Mamon G. A., 2013, MNRAS, 428, 2407

Wolf C. et al., 2009, MNRAS, 393, 1302

Woo J., Dekel A., Faber S. M., Koo D. C., 2015, MNRAS, 448, 237

Wu J. F. et al., 2021, arXiv e-prints, arXiv:2112.01542

Wyder T. K. et al., 2007, ApJS, 173, 293

Xi C., Taylor J. E., 2021, MNRAS

Xi C., Taylor J. E., Massey R. J., Rhodes J., Koekemoer A., Salvato M., 2018, MNRAS, 478, 5336

Yang X., Mo H. J., van den Bosch F. C., 2003, MNRAS, 339, 1057

Yang X., Mo H. J., van den Bosch F. C., 2008, ApJ, 676, 248

Yang X., Mo H. J., van den Bosch F. C., 2009, ApJ, 695, 900

Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, ApJ, 671, 153

York D. G. et al., 2000, AJ, 120, 1579

Zaritsky D., Conroy C., Zhang H., Bonaca A., Caldwell N., Cargile P. A., Johnson B. D., Naidu R. P., 2019, arXiv e-prints, arXiv:1909.02025

Zehavi I. et al., 2002, ApJ, 571, 172

Zhang D., Luo Y., Kang X., 2019, MNRAS, 486, 2440

Zu Y., Mandelbaum R., 2015, MNRAS, 454, 1161

Zwicky F., 1933, Helvetica Physica Acta, 6, 110

APPENDICES

Appendix A

Additional Materials for Introduction

A.1 Additional references for the total mass of MW

The following are estimates of the total mass of the MW from recent studies. Note that there are mainly two types of mass definition – M_{vir} and M_{200c} – which refer to the virial mass and the mass within a volume that has an average density of 200 times of the critical density of the Universe, respectively.

- [Zaritsky et al. \(2019\)](#) suggests a lower limit of $M_{200} > 0.91 \times 10^{12} M_{\odot}$, and a best estimate at $1.4 \times 10^{12} M_{\odot}$. It uses 32 far-side stars from the H3 survey, which have galactic radii $R > 60$ kpc.
- [Posti & Helmi \(2019\)](#) derives $M_{\text{vir}} = 1.3 \pm 0.3 \times 10^{12} M_{\odot}$ by analyzing the phase distribution of 75 globular clusters from Gaia and Hubble data, with an assumption of a concentration-mass relation. They also estimate a stellar mass of $M(< 20 \text{ kpc}) = 1.91^{+0.18}_{-0.17} \times 10^{11} M_{\odot}$ and a central dark matter mass $M_{DM}(< 20 \text{ kpc}) = 1.37^{+0.18}_{-0.17} \times 10^{11} M_{\odot}$.
- [Fritz et al. \(2018\)](#) estimates $M_{MW} = 0.8$ to $1.6 \times 10^{12} M_{\odot}$ from the proper motions of 39 dwarf galaxies within 420 kpc, based on Gaia DR2 data.
- [Sohn et al. \(2018\)](#) estimates $M_{\text{vir}} = 2.05^{+0.97}_{-0.79} \times 10^{12} M_{\odot}$ using HST to measure the proper motion of 20 globular clusters at 10–100 kpc. (looking for the relative displacement between images at two epochs, with a time separations that vary from 6 to 10 years). They also estimate an enclosed mass $M(< 39.5 \text{ kpc}) = 0.61^{+0.18}_{-0.12} \times 10^{12} M_{\odot}$.

- [Callingham et al. \(2018\)](#) derive $M_{200} = 1.17^{+0.21}_{-0.15} \times 10^{12} M_{\odot}$ by comparing the dynamics of MW satellites to the EAGLE hydrodynamic simulations. Fig. 7 of this paper is a good review of literature of MW mass estimates. (Although it appears to include a few some mistakes in the cited values. e.g. Patel18 and Eadie16.)
- [Watkins et al. \(2018\)](#) estimates $M_{\text{vir}} = 1.41^{+1.99}_{-0.52} \times 10^{12} M_{\odot}$, using the kinematics of 34 globular clusters within 2.0–21.1 kpc. Combining this with another similar result using 46 HST globular cluster proper motion measurements, they conclude that $M_{\text{vir}} = 1.67^{+0.79}_{-0.50} \times 10^{12} M_{\odot}$.
- [Patel et al. \(2018\)](#) estimates $M_{\text{vir}} = 1.19^{+0.19}_{-0.21} \times 10^{12} M_{\odot}$ by analyzing the proper motions of the classical satellites. Note that the result cited here excludes the Sagittarius dSph, which reportedly biases the result towards lower masses.
- [Eadie & Harris \(2016\)](#) estimates $0.91 \times 10^{12} M_{\odot}$ ($0.57\text{--}1.09$, $r_{\text{vir}} = 198^{+19}_{-24}$ kpc, only using the GCs beyond 10 kpc); Another “conservative” estimate is $0.68 \times 10^{12} M_{\odot}$ ($0.61\text{--}0.75$, $r_{\text{vir}} = 185^{+7}_{-7}$ kpc, all GCs). This paper uses a Bayesian estimates based on phase-space distribution function of globular clusters (89 GCs after several selection cuts). It also discusses how using subsamples with different radial distributions affects the estimates. This paper also contains a good review of the literature.
- [Bland-Hawthorn & Gerhard \(2016\)](#) is a literature review of MW properties. Its section 6.3 is a good summary of the recent studies on MW total mass estimates. Overall, this review suggests $M_{200} = 1.1 \pm 0.3 \times 10^{12} M_{\odot}$ or $M_{\text{vir}} = 1.3 \pm 0.3 \times 10^{12} M_{\odot}$.
- [Wang et al. \(2015\)](#): Fig. 1 of this paper also gives a review of the MW mass estimates, using various dynamical tracers and several different analysis methods. Those results include several estimates range from 0.5 to $2.5 \times 10^{12} M_{\odot}$). This paper also tests various methods using mock samples and estimating bias in the results.

Appendix B

Appendix for Chapter 2

B.1 The Serendipitous Catalogue

In Table [B.1](#) we list the serendipitous catalogue of nearby objects. Columns are visual class (as explained in section [2.7](#)), ID from the COSMOS 2015 catalogue (where available), coordinates, redshift, redshift error (for objects with photometric redshifts only), apparent i^+ -band magnitude, approximate absolute magnitude in the same band (assuming a distance $D = cz/H_0$ with $H_0 = 0.678$), and any comments. As noted previously, repeated visual searches suggest classes 1 & 2 are reasonably complete, while classes 3–7 contain only a few representative examples of the many objects of this kind.

Table B.1: The Serendipitous Catalog

Class	COSMOS ID (Laigle et al. 2016)	R.A. (J2000)	Decl. (J2000)	z	σ_z^*	i^+ (mag)	M_{i^+} (mag)	Comments
1	213165	150.6950	1.6139	0.030		18.02	-17.5	conflicting redshift 0.1529
	260583	149.6202	1.6936	0.006		17.90	-14.2	part of group at 26 Mpc
	331749	150.3456	1.7936	0.019		18.96	-15.6	
	401988	150.0245	1.9110	0.006		17.15	-15.0	part of group at 26 Mpc
	458976	149.8663	2.0071	0.013		18.59	-15.2	
	549719	150.1254	2.1498	0.005	0.008	19.00	-12.7	
	551648	150.0433	2.1560	0.006		14.24	-17.8	appears more distant?
	561851	150.6131	2.1668	0.006		18.02	-14.1	part of group at 26 Mpc
	653748	150.3134	2.3064	0.027		17.42	-17.9	
	677414	149.6951	2.3477	0.006		17.37	-14.6	part of group at 26 Mpc
	686606	150.3666	2.3404	0.007		20.28	-12.2	part of group at 26 Mpc
	709026	150.0284	2.3793	0.012		19.45	-14.1	size incorrect in catalogue?
	733922	150.4743	2.4138	0.007		17.40	-14.9	part of group at 26 Mpc
	2	219550	149.8758	1.6103	0.040	0.034	19.87	-16.3
221686		149.5820	1.6156	0.043	0.035	21.77	-14.6	
259971		149.4614	1.6750	0.010		20.20	-13.0	conflicting redshift of 0.8058
279307		149.9644	1.7067	0.025		19.60	-15.6	
300323		150.4282	1.7425	0.045	0.045	22.28	-14.1	
316142		149.4853	1.7645	0.018	0.029	20.19	-14.2	
424575		149.5127	1.9533	0.005	0.008	17.97	-13.7	
556961		149.6577	2.1597	0.005	0.008	20.84	-10.8	
589205		149.8118	2.1923	0.025		19.92	-15.2	
627637		149.7679	2.2548	0.025		19.07	-16.1	
642238		149.4566	2.2722	0.005	0.008	19.93	-11.8	
689831		150.6784	2.3433	0.005	0.008	21.32	-10.3	
880363		149.9964	2.6334	0.060	0.040	21.10	-16.0	
918161		150.3921	2.6917	0.012		20.25	-13.3	
989145		150.4089	2.8052	0.044		20.26	-16.1	
997756		149.6831	2.8163	0.023	0.023	20.19	-14.8	
3	183741	149.5938	1.5848	0.028		15.89	-19.5	
	246757	149.4982	1.6542	0.022		19.83	-15.0	
	460674	150.5469	2.0216	0.021		16.21	-18.6	
	532836	150.5065	2.1134	0.046		18.78	-17.7	
	534651	150.1830	2.1148	0.100	0.060	21.62	-16.5	
	538389	150.0464	2.1188	0.029		21.86	-13.6	
	622498	150.1930	2.2445	0.677		21.30	-21.0	
	660791	149.9128	2.3040	0.705		21.91	-20.5	
	706494	150.2301	2.3955	0.045		15.44	-21.0	
	718332	149.8389	2.3875	0.028		22.17	-13.2	

* redshift error, listed only for objects with photometric redshifts

Table B.2: The Serendipitous Catalog

Class	COSMOS ID (Laigle et al. 2016)	R.A. (J2000)	Decl. (J2000)	z	σ_z^*	i^+ (mag)	M_{i^+} (mag)	Comments
	824852	149.7570	2.5499	0.029		18.66	-16.8	
	905622	150.4302	2.6859	0.047		17.31	-19.2	
	915194	149.8467	2.6938	0.048		16.78	-19.8	
	923647	150.0386	2.7132	0.033		16.77	-19.0	
	955856	150.0338	2.7651	0.029		15.41	-20.1	
4	261496	149.5315	1.6786	0.021	0.026	20.63	-14.1	
	282078	149.8230	1.7285	0.055	0.183	20.18	-16.7	
	643833	149.9028	2.2784	0.005	0.008	19.86	-11.8	
	733610	150.1712	2.4130	0.043		19.50	-16.9	
	771819	150.3126	2.4689	0.005	0.008	20.92	-10.7	
	1038253	149.8371	2.8744	0.050	0.035	20.08	-16.6	
5	377112	150.1917	1.8634	0.027		20.84	-14.5	
	484608	150.4819	2.0372	0.005	0.008	20.66	-11.0	
	494700	150.4874	2.0533	0.093		20.86	-17.1	
	648571	150.3759	2.2856	0.051	0.036	20.76	-15.9	
	864285	150.6092	2.6075	0.104		20.68	-17.6	
6	380820	150.0600	1.8665	0.024		20.04	-15.0	
	423926	150.3431	1.9400	0.046		20.08	-16.4	
	532809	150.7758	2.1105	0.005	0.008	18.45	-13.3	
	840592	150.7351	2.5780	0.005	0.008	17.78	-14.0	
	880547	150.0023	2.6332	0.024		20.20	-14.9	
7	216843	149.6873	1.6104	0.050	0.035	19.93	-16.7	
	349181	149.8123	1.8196	0.081	0.051	19.39	-18.3	
	400833	150.7306	1.9004	0.005	0.008	21.05	-10.7	
	516283	150.6366	2.0837	0.093		20.92	-17.1	
	518816	150.7234	2.0883	0.069		21.37	-16.0	
	523477	150.4045	2.1067	0.054	0.040	18.10	-18.7	
	731241	150.1731	2.4042	0.036	0.032	19.32	-16.6	
	757311	150.0542	2.4513	0.082		20.38	-17.4	
	837992	150.6170	2.5750	0.007	0.024	18.84	-13.5	
	840823	150.4003	2.5727	0.092		21.09	-16.9	
	862172	149.7740	2.6061	0.188		21.47	-18.0	
	908277	150.7575	2.6765	0.005	0.008	21.56	-10.1	
	943231	150.0981	2.7438	0.059	0.040	19.21	-17.8	
	(masked)	150.5946	2.4223	0.090	0.090	21.46 [†]	-16.5	SDSS J100222.70+022520.3 [†] SDSS i -band <i>model</i> magnitude

* redshift error, listed only for objects with photometric redshifts

Appendix C

Appendix for Chapter 3

C.1 Bias in Halo Masses Derived from the SHMR

Throughout this work we assume the SHMR proposed by [Behroozi, Wechsler & Conroy \(2013\)](#):

$$\log_{10}(M_*(M_h)) = \log_{10}(\epsilon M_1) + f\left(\log_{10}\left(\frac{M_h}{M_1}\right)\right) - f(0), \quad (\text{C.1})$$

where the function $f(x)$ is defined as:

$$f(x) = -\log_{10}(10^{\alpha x} + 1) + \delta \frac{(\log_{10}(1 + \exp(x)))^\gamma}{1 + \exp(10^{-x})}. \quad (\text{C.2})$$

The free parameters vary with redshift as follows:

$$\begin{aligned} \log_{10}(M_1) &= M_{1,0} + (M_{1,a}(a-1) + M_{1,z}z) \exp(-4a^2) \\ \log_{10}(\epsilon) &= \epsilon_0 + (\epsilon_a(a-1) + \epsilon_z z) \exp(-4a^2) + \epsilon_{a,2}(a-1) \\ \alpha &= \alpha_0 + (\alpha_a(a-1)) \exp(-4a^2) \\ \delta &= \delta_0 + (\delta_a(a-1) + \delta_z z) \exp(-4a^2) \\ \gamma &= \gamma_0 + (\gamma_a(a-1) + \gamma_z z) \exp(-4a^2). \end{aligned} \quad (\text{C.3})$$

where $a = 1/(1+z)$ is the scale factor. (The $1-\sigma$ uncertainty range in these parameter values is listed on p.9 of [Behroozi, Wechsler & Conroy \(2013\)](#).) In our case, since our primary sample covers the fairly narrow redshift range of $z = 0-0.25$, little variation is predicted in the SHMR. Thus, we simply use an intermediate redshift of $z = 0.15$ for the analysis below.




Figure C.1: Stellar mass errors versus stellar mass in the mock sample (for clarity, only a third of the data points are shown), together with the mean relation in bins of stellar mass (red points and errorbars).

The (Behroozi, Wechsler & Conroy, 2013) relations are theoretical, unbiased mean values of M_* , given a specific halo mass M_h . In any real survey, this relationship will be biased by intrinsic scatter and observational errors (B13, Leauthaud et al. (2012)). To quantify this bias for the COSMOS catalogue, we generated a random Monte-Carlo sample of 12,000 halos selected from the halo mass function given in Tinker et al. (2008). “True” stellar masses were calculated for these objects using the SHMR given above. We then added intrinsic scatter and random errors to each stellar mass 50 times independently, to simulate an “observed” stellar mass sample. The intrinsic scatter in the SHMR is about 0.14–0.2 dex at redshift of 0 (More et al., 2009; Yang, Mo & van den Bosch, 2009; Reddick et al., 2013), and there is no evidence for any trend with mass, at least down to halo masses of $10^{12} M_\odot$ (Reddick et al., 2013; Behroozi, Wechsler & Conroy, 2013). Thus we added an intrinsic scatter of 0.15 to all stellar masses derived for our “observed” sample. Average observational errors, as a function of stellar mass, were estimated directly from the COSMOS 2015 catalogue, as shown in Fig. C.1. We added these in quadrature to determine the final stellar masses of the mock sample.

The light blue points in Fig. C.2 show the “observed” stellar masses of the mock sample, after




Figure C.2: Halo-to-stellar mass relation of the Monte-Carlo samples. The dark blue dots are our Monte-Carlo halo mass sample with stellar masses assigned using B13. The light blue points show the “observed” stellar masses of the sample after adding intrinsic scatter and observational errors. (Note that only 15% of the mock sample is shown for clarity.) The red crosses with error bars are the average halo mass in each bin of “observed” stellar mass. The solid black line is a linear fit to the red crosses (in $\log(M_*/M_\odot)$), over the range 10.7–11.8.

adding the intrinsic scatter and observational errors. The dark blue points show the underlying “true” stellar masses, while the black dashed line shows the theoretical SHMR from B13. The red crosses show the mean halo mass in each “observed” stellar mass bin. We can see that the “observed” SHMR follows the theoretical SHMR reasonably well at low masses, but departs from it at the high masses. Given this pattern, we will use a two-part SHMR to assign halo masses in our work. For stellar masses of $10^{10.7} M_\odot$ or less, we use an unmodified B13 SHMR, while for masses greater than $10^{10.7}$, we use the linear fit to the average “observed” values listed and shown as a black solid line on the plot.

C.2 Measuring the Masking and field boundaries

Regions of the COSMOS field have poor photometry in one or more bands, due to contamination from bright stars, internal reflections, or other artifacts. Data from these regions are tagged with a “masking” flag (“FLAG_PETER” in the COSMOS 2015 catalogue), which can be used to exclude those data from further analysis. The shape of these masked regions, together with the field boundaries, needs to be measured to determine the area completeness η around any given primary. Although detailed mask files are available for the COSMOS field, we found it less computationally demanding for our work to use a single, approximate mask image with coarser spatial sampling. We use the method described in Chapter 2 to generate this global mask. First, a coarse map consisting of 390×390 cells is constructed, covering the whole COSMOS field. We search for objects in each cell of this map, to determine whether it should be included or masked out. In the first round, any cell with one or fewer objects counts is selected as a potential masked region. In a second round, these candidates are confirmed as masked if they have one or more neighbouring cells with no counts. This two-step selection process reduces to 0.0026% the probability of artificially eliminating cells due to Poisson fluctuations in their object counts. The map resolution and count threshold were determined empirically after testing various resolutions from 200×200 to 600×600 , with different thresholds in each case. We found that the effect on the clustering signal of variations in the masking parameters is small, producing variations in S_{halo} of roughly 5% or less. The final resolution was selected to provide the most accurate overall mask, relative to the full images.

Given a single global mask for the COSMOS field, we then generated a large, random sample of points, and used the distribution of the points around each primary to estimate its area completeness as a function of projected separation. For each galaxy, we counted the number of random samples in projected radial bins with and without applying the masks and boundaries. Each bin had a size of $0.2 R_{\text{vir}}$ of the galaxy, up to $3.6 R_{\text{vir}}$. The area completeness is then

$$\eta(R_P) \equiv \frac{A_M(R_P)}{A_T(R_P)} = \frac{N_M(R_P)}{N_T(R_P)} \quad (\text{C.4})$$

where A_M and A_T are the masked and total areas, and N_M and N_T are the random counts with and without masking.

Besides the area completeness in individual radial bins, we also measured the total area completeness of each primary within $3.0 R_{\text{vir}}$. Galaxies with poor completeness were excluded from the primary sample, as described in the main text.

C.3 Comparing Background Estimation Methods A and B

While Method A is simple to implement, it may introduce systematic biases in the inferred satellite properties, as described in section 3.5.1. Method B removes the contribution from the background statistically, and should produce less biased, albeit noisier, results.

Fig. C.3 shows the (cumulative) satellite luminosity function for three sets of primaries with different halo mass ranges, using methods A (dashed lines) and B (solid lines). Overall, method A produces a luminosity function with a steeper slope, that continues to rise at faint magnitudes, whereas for method B, the cumulative luminosity function flattens. As shown in Fig. 3.13, the method B results are in better agreement with Local Group data.




Figure C.3: Cumulative satellite luminosity functions estimated using methods A (dashed) and B (solid), for primaries in three halo mass bins (as in Fig. 3.13).

Appendix D

Appendix for Chapter 4

D.1 Catalogue completeness

In this chapter, we used several quantities from the COSMOS 2015 catalogue that are not available for all objects. This incompleteness could introduce systematic errors in some of our analysis. The most important examples of incomplete quantities are stellar mass, sSFR and colour. To test for effects related to incompleteness, we constructed a test sample from the COSMOS 2015 catalogue that includes galaxies within the same redshift range as our base catalogue, but with no magnitude cut applied (though masking was still applied). From this sample, we identified objects with missing or invalid values; for instance, a galaxy without a reasonable B-band or i^+ -band magnitude was identified as lacking a valid B- i^+ colour. Fig. D.1 shows the number of galaxies lacking particular measurements, versus i^+ -band magnitude. By comparing these distributions to the magnitude distribution of the whole test sample, we can see that the samples without valid values comprise only a small tiny fraction of the total. In the worst case, galaxies that lack a proper colour only make up less than 0.5% of the whole sample near our magnitude cut. Thus, we conclude that the potential effects of incompleteness in these quantities on our analysis are negligible.

D.2 Redshift distribution for the six primary mass bins

We split our first-run primaries into six bins over their halo mass in some of our analysis. The redshift distributions the six bins are shown in Fig. D.2. The dashed vertical lines indicate the




Figure D.1: Histogram of the i^+ -band magnitude for the whole test sample and the galaxies lack valid values for certain quantities. The black line shows the whole test sample; the blue line shows the sample without a valid B- i^+ colour; red line shows the sample without a valid stellar mass measurement; the orange line shows the sample without a valid sSFR measurement. The vertical dash line shows the magnitude cut we used for our base catalogue.




Figure D.2: Cumulative redshift distributions for the six primary halo mass bins used in the SSMF calculation. The dashed lines indicate the median redshift for each bin. The solid black line shows the upper redshift cut for the first-run primary selection.

median redshift for each bin. Note that the median redshift clearly increases with primary halo mass, as discussed in the main text.

D.3 Comparison to SAGA II (Mao et al. 2021)

To construct the RSMF for the SAGA II survey, we need stellar mass estimates for both primaries and secondaries. A full set of primary stellar masses were not published in Mao et al. (2021), though masses are given for a subset of the same sample in Geha et al. (2017). We estimate the remaining masses from the published K-band absolute magnitudes, using a mean mass-to-light

ratio derived from the published values:

$$\log[M_*/M_\odot] = -0.4M_K + C_{S1}, \quad (\text{D.1})$$

where $C_{S1} = 1.042$ is the median of $\log[M_*/M_\odot] + 0.4M_K$ for the subsample in [Geha et al. \(2017\)](#).

The satellite stellar masses published in [Mao et al. \(2021\)](#) were instead estimated by using the r -band magnitude and $g - r$ colour:

$$\log[M_*/M_\odot] = 1.254 + 1.098(g - r)_o - 0.4M_{r,o}, \quad (\text{D.2})$$

where $M_{r,o}$ refers to K-corrected r -band absolute magnitude and $(g - r)_o$ is the K-corrected $g - r$ colour.

D.4 Primary Subsamples in Colour and sSFR

To separate red and blue primaries, we use $c = B - i^+$ as our colour index. Among the various photometric bands for the COSMOS2015 catalogue, B -band is the deepest Subaru broad band that covers the whole field. For the other band, we select i^+ for the best combination of wavelength baseline, depth, and noise properties (cf. [Laigle et al. 2016](#), Table 3). For consistency, we use the magnitudes measured using a fixed $3''$ aperture in each case. We also applied K-corrections to our colour indices, using the public tool available at this URL¹, which based on the method in [Chilingarian, Melchior & Zolotukhin \(2010\)](#). As our filters are not included in this tool, we assume the nearest equivalents in colour; since our objects are fairly low-redshift, the error introduced by the conversion should be minor.

[Fig. D.3](#) shows primary colour versus stellar mass. There is a clear correlation, as well as clustering of points at $c = B - i^+ > 1.6$ that we identify as the red sequence. We will use this colour cut to define the red primary sample for $\log(M_*/M_\odot) > 10$. For lower stellar masses, the red sequence moves bluewards due to metallicity effects, so we move our cut to $c = 1.5$. Note that by these definitions, all primaries with a stellar mass over $\log(M_*/M_\odot) = 10.9$ are red.

The star formation rate (SFR) and sSFR in the COSMOS2015 catalogue were derived along with the redshift and stellar mass, by using the best-fit templates from [Polletta et al. \(2007\)](#) and [Bruzual & Charlot \(2003\)](#). The catalogue includes median best-fit estimates, as well as upper and lower limits. [Fig. D.4](#) shows the sSFR versus colour for the first-run primary sample. To define a cut in sSFR that roughly corresponds to our colour selection, we choose the median value for the red sample, which is approximately $sSFR = 10^{-11} \text{yr}^{-1}$.

¹<http://kcor.sai.msu.ru/getthecode/>




Figure D.3: The $B - i^+$ colour versus stellar mass distribution for the primaries. The two dashed lines show the two colour cuts used in our analysis.




Figure D.4: sSFR versus $B - i^+$ colour for the primaries. The vertical dashed line indicates the division between red and blue primaries; the horizontal dashed line indicates the division between star-forming and passive primaries.

Glossary

***N*-body simulation** In astronomy, an *N*-body simulation is a simulation of a dynamical system of particles, usually under the influence of physical forces, such as gravity. [7](#)

baryon A type of composite subatomic particle which contains an odd number of valence quarks (at least 3), such as protons and neutrons. [5](#)

Cepheid variable A Cepheid variable star is a type of star that pulsates radially, varying in both diameter and temperature and producing changes in brightness with a well-defined stable period and amplitude. The strong direct relationship between a Cepheid variable's luminosity and pulsation period allow it to be used as a distance estimator. [2](#)

Cosmic Microwave Background Electromagnetic radiation which is a remnant from an early stage of the universe. [5](#)

Dark Matter Matter that does not interact with light. [5](#)

electromagnetic spectrum The full range of all frequencies of electromagnetic radiation and also to the characteristic distribution of electromagnetic radiation emitted or absorbed by that particular object. [5](#)

galaxies A galaxy is a gravitationally bound system of stars, stellar remnants, interstellar gas, dust, and dark matter. [5](#)

galaxy clusters a galaxy cluster is a structure that consists of anywhere from hundreds to thousands of galaxies that are bound together by gravity, with typical total masses ranging from 10^{14} – 10^{15} solar masses. [5](#)

gravitational instabilities The distribution of material is unstable when only under the effect of gravity. Material is set to be brought together to form structures due to gravitational instabilities. [5](#)

parsec The parsec (symbol: pc) is a unit of length used to measure the large distances to astronomical objects outside the Solar System, approximately equal to 3.26 light-years or 206,000 astronomical units (au). [14](#)

primordial fluctuations Primordial fluctuations are density variations in the early universe which are considered the seeds of all structure in the universe. Currently, the most widely accepted explanation for their origin is in the context of cosmic inflation. [6](#)

redshift A redshift is an increase in wavelength of electromagnetic radiation, or a decrease in frequency. [5](#)

wavenumber In physics, the wavenumber is defined as the number of wavelengths per unit distance of the spacial wave frequency and is known as spatial frequency. [6](#)

Abbreviations

Λ CDM Dark energy plus Cold Dark Matter [5](#)

AGN Active Galactic Nucleus [10](#)

AM Abundance matching [11](#)

CMB Cosmic Microwave Background [5, 6](#)

DE Dark Energy [5](#)

DM Dark Matter [5](#)

HDM Hot Dark Matter [17](#)

HOD Halo Occupation Distribution [11](#)

kpc kilo-parsec [3](#)

LG Local Group [4, 13](#)

LMC Large Magellanic Cloud [2](#)

LVC Local Volume Catalogue [27](#)

ly light year [3](#)

M31 Messier 31 [2, 13](#)

Mpc Mega-parsec [14](#)

MSP Missing satellite Problem [16](#)

MW Milky Way [1](#)

ROI Region of Interest [82](#)

SAGA Satellites Around Galactic Analogs [20](#)

SDSS Sloan Digital Sky Survey [5](#)

SFR Star Formation Rate [3](#)

SHAM Sub-halo Abundance Matching [11](#)

SHMR Stellar-to-Halo mass relation [11](#)

SIMBAD the Set of Identifications, Measurements and Bibliography for Astronomical Data [14](#)

SMBH Suppermassive black hole [10](#)

SMC Small Magellanic Cloud [2](#)

SSCs Small-Scale Challenges [16](#)

TBTF Too Big To Fail [16](#)

TRGB Tip of the Red Giant Branch [20](#)

WDM Warm Dark Matter [17](#)

List of Symbols

- δ Density fluctuation. Defined as $\delta \equiv \Delta\rho/\bar{\rho}$, where $\bar{\rho}$ is the mean mass density and $\Delta\rho$ is the excess density. 6
- M_h Halo mass. The total mass of dark matter within a halo. 3
- M_* Stellar mass. The total mass that is in the form of stars. 3
- M_\odot Solar mass. A commonly used mass unit, which equals to the mass of the Sun. 4
- M_{vir} Virial mass. It is one of the common descriptions for halo mass. 8
- r_{vir} Virial radius. It is one of the common descriptions for halo size. 8
- a Scale factor of the Universe. It describes the relative expansion of the Universe. 6