

Globular Cluster System Properties by Host
Galaxy Environment

GLOBAL CLUSTER SYSTEM PROPERTIES BY HOST GALAXY
ENVIRONMENT

By Kate HARTMAN,
B.A., Pomona College (Claremont, CA)

*A Thesis Submitted to the School of Graduate Studies in the Partial Fulfillment
of the Requirements for the Degree Master of Science*

McMaster University © Copyright by Kate HARTMAN August 10, 2022

McMaster University

Master of Science (2022)

Hamilton, Ontario (Department of Physics & Astronomy)

TITLE: Globular Cluster System Properties by Host Galaxy Environment

AUTHOR: Kate HARTMAN (McMaster University)

SUPERVISOR: Dr. Bill HARRIS

NUMBER OF PAGES: x, 76

Abstract

I present *Hubble Space Telescope* photometry in optical (F475X, 475 nm) and near-infrared (F110W, 1.1 μm) bands of the globular cluster (GC) system of the inner haloes of a sample of 15 brightest cluster galaxies (BCGs). I also present a quantitative model of the relation between (F475X - F110W) colour and cluster metallicity, using simulated GCs. The sample comprises massive elliptical galaxies in a range of environments, from sparsely populated groups to dense clusters. Because the material available for large galaxies to accrete varies with environment and GC systems of such galaxies are built up through accretion, I expect the metallicity distribution of the GC systems in my sample to vary with galaxy environment. GC systems in massive elliptical galaxies tend to follow a bimodal colour distribution, with two subpopulations of blue (metal-poor) and red (metal-rich) clusters. The photometry is used to create a completeness-corrected metallicity histogram for each galaxy in my sample, and to fit a double Gaussian curve to each histogram in order to model the two subpopulations. Finally, the properties of the GC metallicity distribution are correlated against each BCG environment. I found that almost no GCS properties and host galaxy environmental properties are correlated, with the exception of weak but consistent correlations between number of GCs and n th-nearest neighbour surface density and between blue fraction and n th-nearest neighbour surface density.

Acknowledgements

Thank you first and foremost to Bill—you've been everything I could ever have asked for in an advisor, and I feel so lucky to have landed at McMaster. Thank you also to Alison and Laura, my other committee members, for all of your help and advice, and to Ralph, Chris, James, Marta, and everyone else I've learned from at Mac.

Thank you to my parents for everything—the music, the sports, the love of education, the cool things we've done, and all the other stuff I couldn't possibly fit on this page. You've supported me every step of the way.

Thank you to my cousins—Doug, Adam, and Mo—for all the late night Netflix sessions, and thank you to all the extended family—Grandma, Grandpa, Aunt Kim, Uncle Marc, and Aunt Kori—for being my biggest fans. Thank you also to all the family friends who have supported me, especially David and Holly.

Thank you to Claude, Veronika, Jeremy, Megan, Taavishi, and Michelle for helping me solve problems, going on snack runs, offering moral support, and being the best officemates (and thank you to Blake, Hao, and Nathan for the same thing before I moved across the hall).

Thank you to my teammates at Mac, especially Alyson, Michelle, Rachel, Oscar, Esther, Natalie, and Shaleesa, for making practices fun and for venting with me.

And thank you to Gianfranc, Ali, Joey, Pedro, and everyone else at West End Swords for keeping me sane and making it all worth it. You've created a fantastic community in WES (and as a bonus, I'm a much better fencer now). Three times in the middle!

CONTENTS

Abstract	iii
Acknowledgements	iv
Acronyms and Abbreviations	x
1 Introduction	1
1.1 Globular Clusters: the Basics	1
1.2 Giant Elliptical Galaxies: Properties and Formation	2
1.3 Globular Cluster Systems: What We Know So Far	3
1.4 Motivation	5
2 Data and Photometry	8
2.1 Data	8
2.2 Principles of Photometry	9
2.3 Photometry in This Work	13
2.4 Post-Photometry Cleanup	14
3 Completeness	18
3.1 Artificial Star Tests	18
3.2 Principles of Completeness	22
3.3 Correcting for Completeness as a Function of Surface Brightness	27
4 Metallicity	32
4.1 Colour vs. Metallicity for Optical-Infrared Colour Indices	32
4.2 Creating a Colour-Metallicity Conversion	33
5 Quantifying the Metallicity Distributions	37

5.1	Applying Completeness Corrections and Metallicity Conversions	37
5.2	Double Gaussian Fitting	41
6	Analysis (or the Chapter of a Thousand Plots)	46
6.1	Host Galaxy Environment Data	46
6.2	GCS-Environment Comparisons	53
7	Discussion	63
7.1	GCS Properties by Host Galaxy Environment?	63
7.2	Conclusions	63
7.3	Future Work	64
A	Sample <i>calcsky</i> script	67
B	Sample DOLPHOT parameter file	68
C	Table of fitted double Gaussian parameters	72
	Bibliography	73

LIST OF FIGURES

1.1	Example: dark matter merger tree	2
1.2	Example: simulated galaxy environments	4
1.3	Example: GCS-galaxy relations	6
2.1	NGC 777, F110W	9
2.2	NGC 777, F475X	10
2.3	Initial CMDs	15
2.4	Early cleanup steps	16
3.1	Artificial star CMDs	19
3.2	Artificial star delta magnitude vs. magnitude	19
3.3	Artificial star delta x or y vs. magnitude	20
3.4	Artificial star delta x or y vs. x or y	20
3.5	Artificial star sharpness	21
3.6	Artificial star chi	22
3.7	Artificial star S/N	23
3.8	Artificial star CMD vs. GC CMD	24
3.9	Example completeness curve	25
3.10	Sky brightness	25
3.11	Artificial star zones	26
3.12	Sigmoid parameters vs. sky brightness	27
3.13	Completeness curves by zone	28
3.14	m_0 vs. sky brightness	29
3.15	m_0 vs. sky brightness - all three case study galaxies	30
4.1	Simulated GCs age comparison	34
4.2	Simulated GCs fit comparison	34
4.3	Simulated GCs colour comparison	35

4.4	Simulated GCs colour-colour plot	36
4.5	[M/H] vs. [Fe/H]	36
5.1	CMDs with completeness cut	38
5.4	Colour histogram	38
5.2	F475X CMDs	39
5.3	F110W CMDs	40
5.5	Colour and metallicity histograms	41
5.6	Metallicity histogram with double Gaussian	42
5.7	Final GCS overlay	43
5.9	NGC 741	44
5.8	Post-cut metallicity histogram with double Gaussian	44
6.1	Intra-variable checks	48
6.2	Distance checks	50
6.3	Stellar mass checks	51
6.4	Environmental checks	52
6.5	Metallicity histograms - full dataset	54
6.6	GCs vs. environment	56
6.7	Mean metallicity vs. environment	56
6.8	Mode difference vs. environment	57
6.9	Peak separation D vs. environment	57
6.10	Amplitude vs. environment	58
6.11	Mode vs. environment	59
6.12	Sigma vs. environment	60
6.13	Blue fraction vs. environment	61
6.14	GCs vs. environment with fitting	61
6.15	Blue fraction vs. environment with fitting	62
6.16	n th-nearest neighbour linear fit comparison	62
7.1	GC accreted fraction vs. stellar mass	65

LIST OF TABLES

2.1	Observational information	11
2.2	Basic galaxy data	12
5.1	Example double Gaussian parameters	44
6.1	MASSIVE data	47
6.2	CHM data	48
6.3	n th nearest neighbour metrics	49
6.4	Double Gaussian fits	53
6.5	GC tallies	55
A3.1	Fitted double Gaussian parameters	72

Acronyms and Abbreviations

ACS	Advanced Camera for Surveys
BCG	Brightest cluster galaxy
CDF	Colour distribution function
CMD	Colour-magnitude diagram
FWHM	Full width half maximum
GC	Globular cluster
GCS	Globular cluster system
GMM	Gaussian mixture modelling
HDC	High-density contrast
<i>HST</i>	<i>Hubble Space Telescope</i>
IMF	Initial mass function
LDC	Low-density contrast
MDF	Metallicity distribution function
NED	NASA/IPAC Extragalactic Database
NGC	New General Catalog
NIR	Near infrared
PSF	Point spread function
S/N	Signal-to-noise
WFC3	Wide Field Camera 3

1.1 Globular Clusters: the Basics

Globular clusters (GCs) are massive, dense, gravitationally bound systems of stars. The median GC is of order 10^5 Solar masses, with a half-light radius of order 3 pc, and it is very old. In the Milky Way, most GCs are at least 10 Gyr old, with some forming as far back as ~ 13 Gyr, and the same is true for GCs in other galaxies (Beasley 2020, Harris 2010). GCs are also relatively metal-poor (Beasley 2020, for example, sets the cutoff between metal-poor and metal-rich GCs at 1.0 dex below Solar metallicity), because they formed before the available star-forming material reached its present level of enrichment.

There is, however, variation in GC metallicity. The most metal-poor GC in the Milky Way is M15, at $[\text{Fe}/\text{H}] \sim -2.5$ dex, and GCs in other galaxies bottom out around that value as well. (Note, however, Larsen et al. 2020, which measured an M31 GC at $[\text{Fe}/\text{H}] \sim -2.9$ dex.) On the other end of the scale are the most metal-rich GCs, at or near Solar metallicity (Beasley 2020).

Because of GCs' near ubiquity in galaxies, their apparent connection to host galaxy properties (more on GC scaling relations later in this chapter), and their age, which makes them as old as their host galaxies (Harris 2010), GC formation has been of great interest to theorists. Recent simulations such as E-MOSAICS (Pfeffer et al. 2018, Kruijssen et al. 2019) have tried to model galaxy growth and GC formation simultaneously, a computationally challenging task given the wide range of scales involved. These simulations generally have to rely on subgrid models—i.e. approximate analytic models of the physics—on GC spatial and mass scales because of the lack of better resolution (simulators have to compromise on resolution in order for their simulations to finish running in a reasonable amount of time), but by comparing model GC properties to observed GC properties, simulators can fine-tune their models and inch closer to accurate GC formation and evolution physics.

GCs and their properties are also important tools for probing galaxy formation and evolution for observers. Studies such as Harris et al. 2013, Beasley et al. 2018, and Harris et al. 2017 have used GC and globular cluster system (GCS) properties to trace the history of the GCs’ host galaxies, working backwards from the galaxies’ present-day states to what may have happened to get them there. This approach is particularly promising for massive elliptical galaxies, which tend to have large GCSs.

1.2 Giant Elliptical Galaxies: Properties and Formation

This work focuses on massive elliptical galaxies, as typified by the objects from the MASSIVE survey (Ma et al. 2014; more on this survey in the following chapter). The most massive galaxies tend to be elliptical (E) or occasionally lenticular (S0), and they have globular cluster systems (GCSs) with GCs numbering in the thousands (see Harris et al. 2017 for examples of large GCSs).

Simulations support a hierarchical merger model of galaxy growth in which massive elliptical galaxies grow to their present sizes by merging with smaller satellite galaxies (and sometimes other large ellipticals), and like GC-focused modelling, recent work such as Pillepich et al. 2018 (IllustrisTNG) and El-Badry et al. 2019 have improved the resolution, allowing for more accurate modelling of galaxy formation and evolution. These simulations are built on dark matter merger trees (see Figure 1.1 for an example), with baryonic matter attached to each dark matter halo.

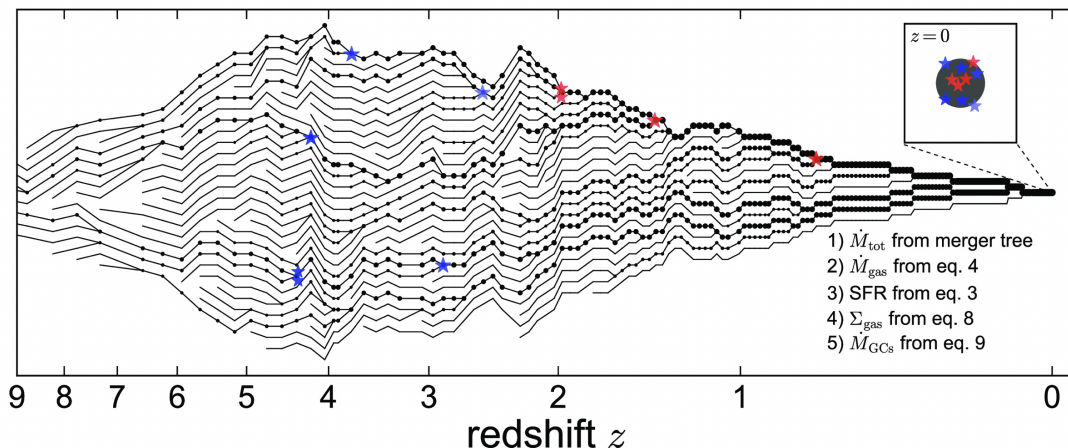


FIGURE 1.1: Figure 1 from El-Badry et al. 2019 (*Monthly Notices of the Royal Astronomical Society*), showing an example of a merger tree. As time progresses (moving right on the x axis), smaller haloes merge to become a single large halo and bring their associated baryonic matter along with them.

Some studies, such as Choksi et al. 2018, El-Badry et al. 2019, and Choksi and Gnedin 2019, focus on the GCs in individual galaxies (more on that work later), while others,

such as IllustrisTNG (Pillepich et al. 2018), encompass a wider volume and model the variety of environments in which galaxies can reside. Figure 1.2, from Pillepich et al. 2018, showcases several different environments found in IllustrisTNG, from sparse areas that only allow for infrequent mergers (such as the top left inset), to rich clusters with ample merging opportunity (such as the second-from-bottom left inset). Rich galaxy clusters like the one in the second-from-bottom left inset tend to be organized as shown in the simulation, with smaller satellite galaxies in the outskirts of the cluster, and the centre of the cluster dominated by a massive elliptical galaxy, sometimes called the brightest cluster galaxy (BCG). By virtue of being the most massive object in the area, BCGs draw in satellite galaxies, growing their own stellar and halo masses and adding the satellites' GCs to their own GCSs. On the observational side, studies such as Harris et al. 2017 pick out BCGs specifically from the larger population of elliptical galaxies and look at the GCS properties that result from a lifetime of merging with smaller satellite galaxies.

Meanwhile, galaxies in low-density areas evolve differently—they have access to far less accretable material than a BCG in a rich galaxy cluster. These galaxies presumably build up a far larger proportion of their GCs locally, rather than bringing them in from accreted satellites. Studies such as Cho et al. 2012, Ennis et al. 2019, and Salinas et al. 2015 focus on these isolated galaxies and their GCS properties, but the target galaxies in these studies are less luminous and less massive than those in this work.

1.3 Globular Cluster Systems: What We Know So Far

Because GCSs form concurrently with their host galaxies (Beasley 2020), it stands to reason that GCS properties may be related to galaxy properties—a conjecture borne out by both observations and simulations.

GCSs exhibit tight scaling relations with host galaxy halo mass and host galaxy stellar mass (which are themselves related). In a large observational study, Harris et al. 2013 quantified various relations between the number of GCs in a GCS and galaxy stellar mass, luminosity, and velocity dispersion (all three galaxy metrics are themselves related). The study also points convincingly to a simple linear relation between GCS mass and host galaxy halo mass (see also Blakeslee et al. 1997, Spitler and Forbes 2009, and Hudson et al. 2014). Following up on their previous work, Harris et al. 2017 focus on GCS metallicity, exploring GCS metallicity distribution function (MDF) bimodality, which appears in many galaxies. In galaxies with a bimodal MDF, the GCS can be divided into a blue, metal-poor subpopulation, which tends to be built up through accretion onto massive galaxies, and a red, metal-rich subpopulation, which requires the chemical enrichment associated with large haloes to account for its metal content. Harris et al. 2017 find that while many BCGs exhibit the typical MDF bimodality, there is variation from galaxy to galaxy in the shape of the MDF, with some GCSs edging toward unimodality as the red and blue subpopulations overlap.




FIGURE 1.2: Figure 1 from Pillepich et al. 2018 (*Monthly Notices of the Royal Astronomical Society*), showing the dark matter web from two IllustrisTNG simulations. In the insets are examples of the different types of galaxies and galactic environments produced in IllustrisTNG.

Simulators have been able to reproduce the observed mass scaling relations using GC and galaxy assembly physics. El-Badry et al. 2019 used a semi-analytical model to track GCs from high redshift to the present day, through formation and galaxy assembly; they find that the tight GCS mass-halo mass relation is a consequence of hierarchical merging and the central limit theorem. Using the hydrodynamic E-MOSAICS simulations, Bastian et al. 2020 produced the same result. Meanwhile, Choksi et al. 2018 tracked

several GCS relations through cosmic time, finding connections between halo mass and mean GCS metallicity, GCS metallicity dispersion, GCS mass, metallicity distribution bimodality, and mean metallicity of metal-poor and metal-rich GC subpopulations.

Many of the studies cited here also address a GC age-metallicity relation (see e.g. Choksi et al. 2018, El-Badry et al. 2019, Horta et al. 2021, and Kruijssen et al. 2019)—metal-rich GCs are, on average, 2 to 3 Gyr younger than their metal-poor counterparts.

Considering the evidence in favour of a hierarchical merger model of galaxy formation in which massive elliptical galaxies build themselves up by absorbing nearby satellites, neighbouring galaxies are conspicuously absent from the list of GCS-galaxy relations. Choksi and Gnedin 2019 addressed this problem in a coarse way by comparing full merger trees (including accreted satellites) to the main progenitor branch only (excluding accretion), finding that only the full set of accretions matches observations (see Figure 1.3) but their all-or-nothing approach leaves us unable to tease out the detail in any existing GCS-neighbour relations. On the observational side, studies such as Cho et al. 2012, Salinas et al. 2015, and Ennis et al. 2019 focus on galaxies in low-density environments rather than the rich galaxy clusters where BCGs are found, but do not compare them to their more crowded peer galaxies. Meanwhile, De Bortoli et al. 2022 addresses connections between GCS properties and environment, but it is not the environment that we are looking for—that study looks at environment within a halo rather than extending its reach out into each galaxy’s surrounding neighbourhood. So far, studies have neglected the question of what material is, and which GCs are, available to accrete during galaxy assembly—an omission that this work will remedy.

1.4 Motivation

To summarize, let’s explore! A large portion of our motivation stems from the environmental parameter space being not well-studied to date. We cannot incorporate environmental differences into our understanding of cluster and galaxy formation and evolution if we do not know what is there. More specifically, it may prove possible to tie any existing GCS-galaxy environment relations to GCS and galaxy formation scenarios.

In growing to their present-day sizes, BGCs accrete whatever smaller galaxies are available to accrete—i.e. their accretion history, and the accretion history of their GCSs, depends on the contents of their environment. A BCG in a crowded environment with many small satellites would have had ample opportunity to accrete the blue, metal-poor GCs that dwarf galaxies preferentially form, while a BGC in a sparse environment would fill out its GCS mainly with its own red, metal-rich GCs. If that scenario accurately describes BCG assembly, then we might expect BGCs in crowded environments to have a higher proportion of blue GCs than otherwise similar galaxies in low-density environments. That said, this is primarily exploratory work.

In the following chapters, I will describe the BCG data and photometry used in this work (Chapter 2), a method for correcting for incompleteness (Chapter 3), and a method for deriving GC metallicity from a photometric colour index (Chapter 4).




FIGURE 1.3: Figure 4 from Choksi and Gnedin 2019 (*Monthly Notices of the Royal Astronomical Society*), showing the effect in including or excluding accreted satellite galaxies and their GCs in the mean GCS metallicity-halo mass relation. Black points show observations, the grey region shows model predictions including accreted satellites, and the pink region shows predictions without accretion (labeled Main Progenitor Branch in the figure).

Chapter 5 will look in detail at the GCS metallicity distributions for the galaxies in this work, and Chapter 6 will bring environmental variables into play in the first systematic

*Master of Science—Kate HARTMAN;
McMaster University—Department of Physics & Astronomy*

investigation of the effect of local environment on GCS properties in giant galaxies. Finally, I will summarize the findings of this work in Chapter 7, and discuss where we might go next.

CHAPTER 2

DATA AND PHOTOMETRY

2.1 Data

The strong scaling relations between GCS properties and host galaxy stellar mass put constraints on how the sample of galaxies for this work could be constructed. In an unrestricted sample of galaxies, the relations involving stellar mass drown out any signal from second-order relations such as those involving environment (Harris et al. 2013). The sample had to be constructed to minimize differences in stellar mass, and therefore to minimize differences in GCS properties related to stellar mass.

To that end, the sample includes fifteen galaxies from the MASSIVE survey sample. The MASSIVE survey targeted massive nearby galaxies, all within about 110 Mpc and all with a stellar mass greater than $10^{11.5}M_{\odot}$ (Ma et al. 2014). Because it comes from MASSIVE, which observed a very narrow stellar mass range, the sample for this work also represents a very narrow stellar mass range, effectively controlling for that variable. The other major benefit of the MASSIVE sample is good statistics; massive galaxies have large GC systems, which produce large sample sizes in this work. Furthermore, focusing on very massive ellipticals was an attractive option from an evolutionary standpoint; as discussed in Chapter 1, these galaxies display the greatest differences in environment and merger history under a hierarchical merger model of galaxy growth. Additionally, the MASSIVE distance criterion proved to be beneficial—it is possible to acquire high-quality observations of nearby galaxies with less telescope time than would be needed for such observations for farther galaxies while maintaining an adequately sized field of view; *HST* cameras cover about 20 kpc at the distances in this sample.

This work makes use of *HST* archival data from previous work targeting MASSIVE galaxies (Goullaud et al. 2018) and supplements it with later companion observations. The original 2018 observations comprise Wide Field Camera 3 (WFC3) images in the F110W filter (1.1 μm) for thirty-four MASSIVE galaxies; see Figure 2.1 for an example




FIGURE 2.1: NGC 777 at 1.1 μm . This image is from the original 2018 dataset (Goullaud et al. 2018)

F110W image from that dataset. Nine of the fifteen galaxies in the sample come from that set of observations, and then added to it are F110W images for an additional six galaxies and F475X (475 nm) images for the entire sample; see Figure 2.2 for an example F475X image. F475X exposures ranged in duration from 2382 s to 2658 s, and F110W exposures from 2496 s to 2812 s, i.e. one orbit per filter. See Table 2.1 for a list of galaxies in our sample and observational details, and Table 2.2 for basic galaxy properties.

In GCs, colour is tied to metallicity (Brodie and Strader 2006). This relation is nonlinear at all wavelengths, so the main consideration when choosing WFC3 filters was sensitivity to metallicity. Because wider colour ranges provide more sensitivity, the (F475X – F110W) index, in the (B – I) range, was a better choice than any of the (V – I) options. Despite the nonlinearity, this relation can be modelled (see Chapter 4 for more on this), which enabled estimation of the metallicity of each GC in a system by first observing its colour index.

2.2 Principles of Photometry

The colour index is a basic photometric difference measurement—how bright is an object at one wavelength compared to another wavelength? This work uses the (F475X – F110W) colour index. Because of their old stellar populations, GCs tend to be brighter




FIGURE 2.2: NGC 777 at 475 nm. This image is from the followup program.

at red wavelengths than they are at blue wavelengths, but the size of the difference can vary; GCs with a redder colour index (higher numbers) are much brighter at red wavelengths than blue and are referred to as "red," whereas GCs with a bluer index (lower numbers) are only slightly brighter at red wavelengths and are called "blue." High-precision measurements of metallicity require spectroscopy, but the link between GC metallicity and colour allows researchers to make good estimates of GC metallicity with single-orbit images and just two photometric measurements.

The earliest photometry that I did for this work involved the photometry program DAOPHOT (Stetson 1987). DAOPHOT was designed for stellar photometry in crowded fields, and the interactive nature of the program prevents photometry from becoming a black box and gives newer researchers a good look at each of the steps in the process. Although DAOPHOT was originally intended to measure stellar magnitudes, it also works for GCs at extragalactic scales—at the distances found in the sample of galaxies for this work, GCs appear as point sources in a CCD image, just as stars do.

TABLE 2.1: Program numbers and exposure times for each image in the dataset.

Galaxy	F475X program number	F110W program number	F475X exposure (s)	F110W exposure (s)
NGC 57	15265	14219	2394	2496
NGC 410	15265	14219	2421	2496
NGC 533	15265	14219	2382	2496
NGC 741	15265	14219	2385	2496
NGC 777	15265	14219	2421	2496
NGC 1016	15265	14219	2382	2496
NGC 1129	15265	14219	2478	2496
NGC 1600	15265	14219	2385	2496
NGC 2340	15265	15265	2658	2812
NGC 3158	15265	15265	2538	2612
NGC 3842	15265	15265	2496	2612
NGC 4073	15265	15265	2484	2612
NGC 4839	15265	15265	2505	2612
NGC 4914	15265	14219	2439	2496
NGC 7242	15265	15265	2538	2612

I used DAOPHOT’s point spread function (PSF) fitting capabilities in the first round of photometry for this work. PSF fitting takes advantage of the point source nature of the stars (or GCs) being observed; the shape and brightness of such sources across the few pixels of CCD on which they appear depends only on telescope optics (and atmospheric effects, for ground-based telescopes). This means that each object in an image can be fitted with a standard point source profile, and the scaling of that profile between objects indicates their relative brightness (Harris 1998).

DAOPHOT requires continual input from the user and gives the user minute control over the routines involved in photometry. I used the IRAF implementation of DAOPHOT, and most of the initial setup work involves retrieving telescope, camera, and filter information from the image headers and passing that to a sequence of files to be used later during photometry. At this point in the process, I made choices about which PSF model to use (DAOPHOT has six PSFs and the option to try all of them), how large an area to use for PSF fitting, and how large an area to use for sky brightness measurements. Once I had set up the photometry routines, I ran the routine *daofind*, which creates a catalogue of objects to be measured and asks the user to choose the minimum amount of deviation from the sky background required for an object to be added to the list (Stetson 1987). Step three is to choose aperture photometry parameters—an area within which to measure the total amount of light, and an area of sky brightness to subtract off—and to run the routine *phot*, which performs aperture photometry on all the objects in the catalogue and provides a good initial guess at their magnitudes (Stetson 1987).

TABLE 2.2: Basic data for each galaxy in the dataset. Distances calculated from NED overview data, absolute V-band magnitudes from de Vaucouleurs et al. 1991, absolute K-band magnitudes from Ma et al. 2014, group numbers and number of group members from Crook et al. 2007, and extinction magnitudes from Schlafly and Finkbeiner 2011. Galaxies listed with an "L" group (low density contrast) have no "H" group (high density contrast).

Galaxy	Distance (Mpc)	M_V	M_K	Group number	Number of members	E_{F475X}	E_{F110W}
NGC 57	72.69	-22.8	-25.75	L11	4	0.241	0.068
NGC 410	71.46	-23.1	-25.90	H53	29	0.183	0.052
NGC 533	74.86	-22.8	-26.05	H74	3	0.095	0.027
NGC 741	75.43	-23.1	-26.06	H102	5	0.164	0.046
NGC 777	67.97	-22.8	-25.94	H109	7	0.145	0.041
NGC 1016	91.83	-22.9	-26.33	H159	8	0.097	0.027
NGC 1129	71.54	-21.9	-26.14	H185	33	0.352	0.099
NGC 1600	66.01	-22.6	-25.99	H294	16	0.134	0.038
NGC 2340	85.83	-23.1	-25.90	H426	18	0.231	0.065
NGC 3158	103.41	-23.3	-26.28	H570	6	0.041	0.012
NGC 3842	94.90	-22.9	-25.91	H672	42	0.067	0.019
NGC 4073	88.97	-23.3	-26.33	H692	10	0.084	0.024
NGC 4839	109.06	-23.2	-25.85	H745	49	0.031	0.009
NGC 4914	69.94	-22.7	-25.72	L929	3	0.042	0.012
NGC 7242	78.26	-21.5	-26.34	H1186	15	0.472	0.133

The next step in DAOPHOT was to build a PSF model for my specific images. This was the most interactive portion of the DAOPHOT process, and while it provided a lot of insight into how point sources appear on a CCD chip, it was also the greatest source of inconsistency for this batch of photometric measurements. I ran the routine *pstselect*, which selects bright, isolated objects from the larger catalogue, and then looked at the profile of each candidate object and decided whether to include it in the PSF model based on how closely it conformed to a standard point source profile (a bright core with wings) (Stetson 1987). Because each image contains different objects, this process produced slightly different PSF models for each galaxy in my sample.

The last major DAOPHOT step is to run *allstar*, which collects the original image, the object catalogue, the magnitude file, and the PSF model and fits the model to each object to produce relative magnitudes. After that, a researcher using DAOPHOT would still need to calibrate the *allstar* output (those magnitudes are relative, not instrumental), and DAOPHOT also has options for running additional tests, such as the artificial star routine *addstar* (Stetson 1987).

2.3 Photometry in This Work

Although DAOPHOT is a good tool for those who want to look inside the black box of photometry, its interactive nature can introduce inconsistencies when a user is working with multiple images. My DAOPHOT photometry involved building a separate PSF for each galaxy, which meant that each GCS was being measured against a slightly different baseline. In order to avoid such inconsistencies, I performed the photometry used in this work with DOLPHOT.

DOLPHOT is a photometry package designed specifically for use with *HST*'s Advanced Camera for Surveys (ACS) and WFC3 images (Dolphin 2000). Unlike DAOPHOT, all of the user input comes at the beginning of the process—I set up a large parameter file and run a few routines on the images that I want to measure, but once I start the photometry itself, it runs on its own to completion. Runtime varies widely depending on how many images are being analyzed; for this work, photometry took about twenty minutes to run.

One of DOLPHOT's key features is a library of PSFs for all the *HST* ACS and WFC3 filters (and DOLPHOT was updated in 2017 with a set of PSFs for the upcoming *Roman* telescope). Rather than building a PSF by hand for each image, DOLPHOT takes the PSF step out of the user's hands, instead taking filter information from the image headers and choosing the correct PSF from its library (Dolphin 2000). Doing photometry with DOLPHOT allowed me to use an identical set of measurement parameters for every galaxy in my sample and avoid the uncertainty introduced by manual PSF building in DAOPHOT. (At the same time, having learned to do photometry with DAOPHOT gave me a backup option should I have been unable to solve any problems with DOLPHOT.)

My images required four setup steps before I could run DOLPHOT.

First, I ran DOLPHOT's WFC3 pixel masking routine, *wfc3mask*, on all individual exposures plus the combined images (the exposures have been drizzled—offset from each other by less than a pixel and then combined to reduce numerical effects). This command flags known problem pixels as bad pixels in the images (Dolphin 2013). Second, I ran the command *splitgroups* on all individual exposures plus the drizzled images. This routine separates FITS files with multiple extensions into single-extension files (Dolphin 2013). Each of my F110W single exposures remained a single fits file after *splitgroups*, as did the drizzled images in both filters, but my F475X single exposures were split into two files, each containing half of the total CCD area. Third, I put together and ran a script containing *calcsky* commands for all individual exposures plus the drizzled images. This routine creates a sky image by measuring pixel values in an annulus, calculating the mean and standard deviation, rejecting measurements that fall outside of the user's specified parameters, and repeating until no pixels are rejected (Dolphin 2013). The command takes the arguments $\langle r_{in} \rangle$ and $\langle r_{out} \rangle$, the inner and outer bounds, respectively, of the annulus in pixels; $\langle step \rangle$, the number of pixels between sampling points within the annulus; and $\langle \sigma_{low} \rangle$ and $\langle \sigma_{high} \rangle$, the minimum and maximum allowed standard

deviation values, respectively; all in that order. See Appendix A for an example *calcsky* script.

The last step in the pre-DOLPHOT process is setting up the parameter file. The first section of the parameter file contains a list and count of images (the single exposures), plus one reference image listed as "image 0" and not included in the image count (I used the drizzled images as references because they are deeper than any of the individual exposures). I used the drizzled F110W image as my reference image for thirteen of the fifteen galaxies, but because of a poor coordinate solution when using that image, I had to use the F475X drizzled image as a reference for NGC 3158 and NGC 4914. The reference image is only used for locating point sources and aligning the single exposures and not for making measurements, so after my initial data cleanup work, this did not affect my results. The second section of the parameter file deals with coordinate solutions, the PSF, the aperture used for photometry, and the apertures used for calculating other values such as sky brightness. The third portion of the file contains photometry parameters such as detection threshold, PSF analytic forms, settings for sky fitting, and signal-to-noise (S/N) and full width half maximum (FWHM) requirements for point source PSF fitting. The last section contains miscellaneous options relating to coordinate solutions, automatically generated plots, camera-specific settings, and DOLPHOT's artificial star mode (more on artificial star mode later). A partial list of parameters in this file and explanations of what they do can be found in the DOLPHOT User's Guide (Dolphin 2013), and an example parameter file appears in Appendix B.

Once the pre-DOLPHOT routines had been run and the parameter file had been set up, I could run the routine *dolphot*, passing it the parameter file and specifying the names of various output files. This routine identifies point sources, performs photometry, and produces a large data file with photometry and diagnostic measurements for both individual exposures and all exposures in a given filter combined, as well as a suite of ancillary data files including a list of columns in the main output file, PSF figures, and a file summarizing important alignment and aperture correction information (Dolphin 2013).

2.4 Post-Photometry Cleanup

Because they contain columns for each individual exposure, DOLPHOT output files can be very large (mine had up to 180 columns). I only needed some of the combined photometry metrics, so after photometry was complete, I extracted the relevant columns from the output files. I pulled the columns for x coordinates, y coordinates, object type flag, and sky brightness, VEGAMAG (the magnitude system in which Vega has magnitude zero in all filters), magnitude error, S/N, chi (a statistical goodness of fit metric; lower chi values indicate a better fit), sharpness (a PSF-fitting metric; positive values are narrower than a PSF and negative values are broader), roundness, and crowding parameters for both filters (I ultimately did not use roundness or crowding).




FIGURE 2.3: Objects from NGC 777. The lighter points are all objects, and the darker points are objects remaining in the dataset after the major culling steps (object type, magnitude, sharpness, and chi). The remaining objects appear as expected in CMDs. *Left:* blue (F475X) CMD. *Right:* red (F110W) CMD.

After extracting the metrics that I would need, I cleaned the photometry from each galaxy of non-GC objects. The biggest reduction came from the object flag—flags 2 through 5 denote objects that are too faint to be measured and objects that are not star-like in shape (Dolphin 2013); by keeping only objects with type flag 1 (point sources), I weeded out things such as cosmic ray strikes and small background galaxies and kept only candidate GCs. I then made the following measurement quality and GC-like object shape cuts:

- Magnitude < 90 in both filters, to eliminate objects too faint for valid measurement
- $S/N \geq 4$ in both filters, to ensure high-quality magnitude measurements
- $-0.15 < \text{F475X sharpness} < 0.08$, to capture point sources
- $|\text{F110W sharpness}| < 2.3 * (0.3608 - 0.0363 * (\text{F110W}) + 0.000938 * (\text{F110W})^2)$, to capture point sources with room for scatter at the faint end
- $\text{Chi} \leq 1.5$ in both filters, to ensure high-quality magnitude measurements

The F110W sharpness criterion is a pair of quadratic curves motivated by the shape of the data. There was little difference in the datasets produced by a quadratic cutoff and a constant cutoff. The F475X sharpness criterion is slightly skewed toward negative values (i.e. objects broader than the PSF). I also set the F475X sharpness bounds

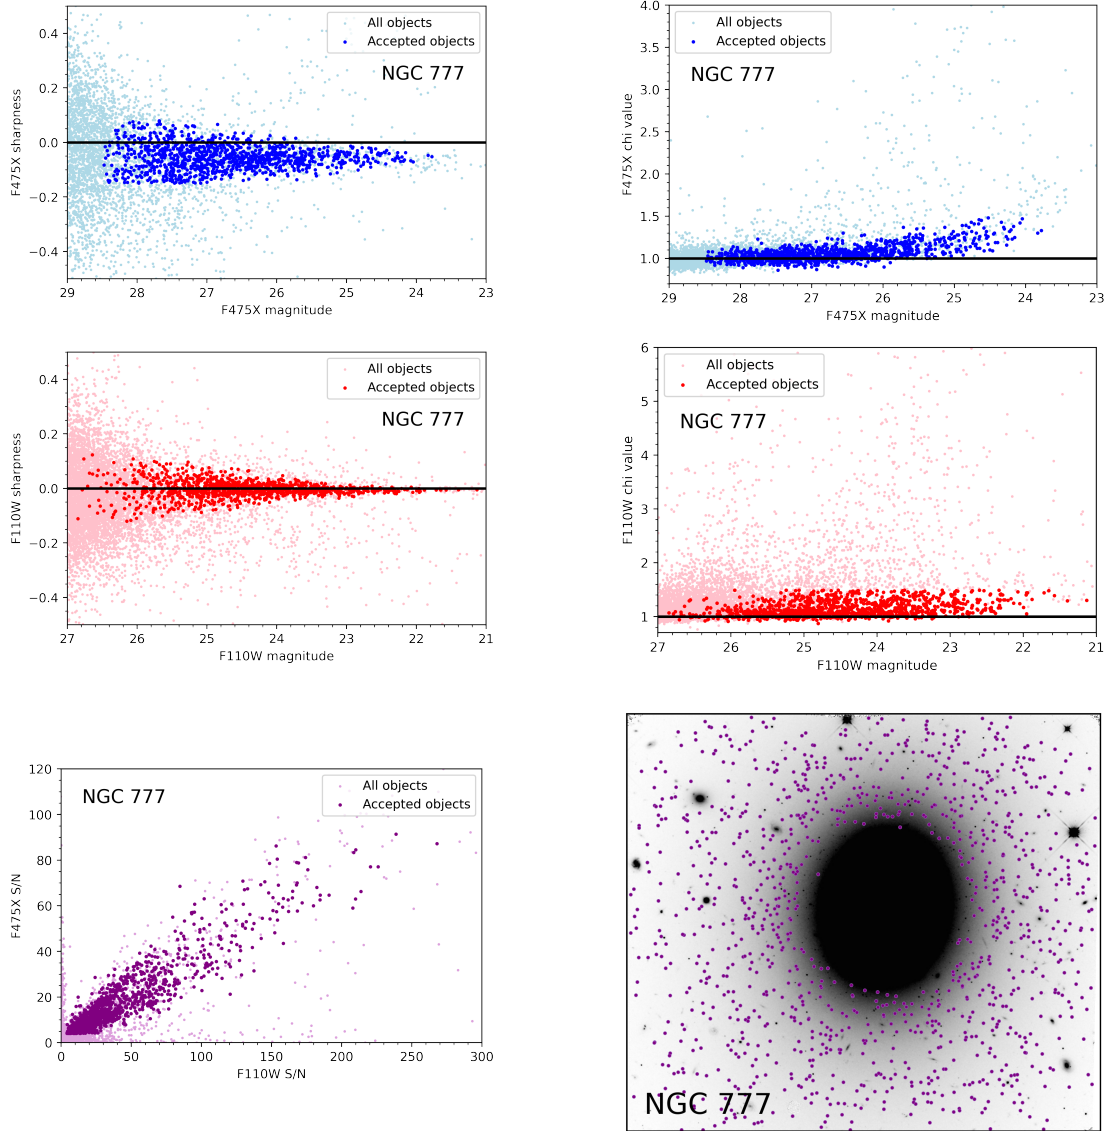


FIGURE 2.4: Objects from NGC 777. *Top left:* sharpness in the F475X filter. *Top right:* chi in the F475X filter. *Middle left:* sharpness in the F110W filter. *Middle right:* chi in the F110W filter. *Bottom left:* S/N in both filters. *Bottom right:* an overlay of objects remaining after the major culling steps on NGC 777 as seen in the F110W filter.

based on how the entire dataset at that point filled the parameter space; physically, GCs appear slightly broader than a point source in the blue filter because the CCD chips used with that filter have a higher resolution than that used with F110W, and at my galaxies' distance, GCs are extended enough to be very slightly resolved. See Figure 2.3 and Figure 2.4 for a typical example of a galaxy before and after this process, as seen in

colour-magnitude diagrams (CMDs) and the metrics used to make cuts to the dataset.

After this series of big cuts, I had for each galaxy in my sample a set of GC candidates largely devoid of contaminants. The next cleanup steps were more involved and oriented toward maximizing the statistical quality of each dataset.

As shown in the bottom right panel of Figure 2.4, the major culling steps removed from the dataset foreground stars, extended background sources, and objects lying too close to the centre of NGC 777. Unlike the first two cases, objects in the last case may indeed be GCs, but their location on the image does not allow for acceptably precise photometric measurements. The background light from NGC 777 in the centre of the image not only prevents DOLPHOT from making good measurements, but also washes out most of the GCs in that area altogether. Only a small fraction of the objects in such areas of the image can be detected in the first place, and the few objects that are detected may not be a representative sample of the GC population. This problem becomes less dire away from the centre of the image, where there is less background light from the galaxy, but it does not disappear completely. In order to properly account for GCs that DOLPHOT was not detecting, I had to understand the relation between background light and the fraction of existing objects being detected; this motivated the investigation of completeness described in Chapter 3.

3.1 Artificial Star Tests

Aside from performing photometry, the other key feature of programs like DOLPHOT is the ability to perform artificial star tests. Artificial star tests involve inserting artificial stars—i.e. standard PSF profiles scaled by any amount, to mimic a star of any magnitude—into an image and then doing photometry. All the artificial stars have known positions and magnitudes, which can then be compared to their measured positions and magnitudes. Artificial star tests allow researchers to characterize the performance of their particular photometry programs on their particular images, and to find potential sources of systematic measurement error.

To make sure that DOLPHOT was performing well, I ran a 5000-star artificial star test on every galaxy in the sample and compared the artificial star photometry to what was expected and to the GC photometry. The runtime for a 5000-star artificial star test was approximately two hours; larger tests such as the 42,000-star tests described later in this chapter take significantly more time to complete.

Figure 3.1 shows input and output CMDs for the artificial stars in this test. The input CMD provides a good visualization of artificial star input parameters—in DOLPHOT, the user chooses the number of artificial stars to be added to an image, the allowed magnitude range in the bluer filter, and the allowed colour index range, and DOLPHOT then randomly assigns magnitudes and positions to that number of stars (Dolphin 2013). This results in a rectangular input CMD, as the collection of artificial stars fills each magnitude range uniformly. Then, when the stars are measured, those that fall within areas of the image with more background light are more likely to be washed out and not detected, and less likely to be measured accurately if they are detected, resulting in the flaring and the blank space at the bottom of the right panel in Figure 3.1.

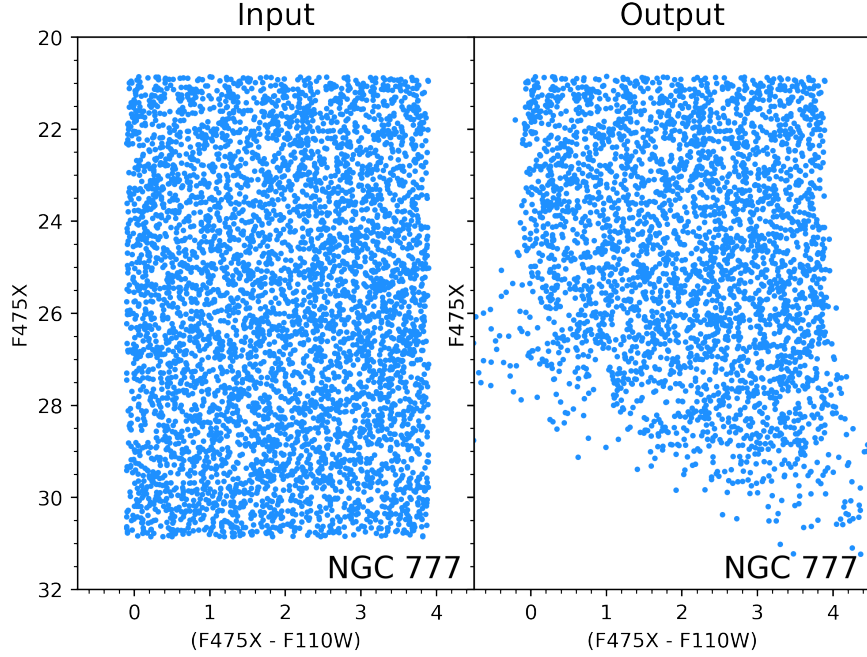


FIGURE 3.1: CMDs of artificial stars. *Left:* CMD of input stars. *Right:* CMD of recovered stars. Fainter stars are less likely to be detected (right panel, bottom left) and less likely to be measured accurately if they are detected (right panel, extreme colours).

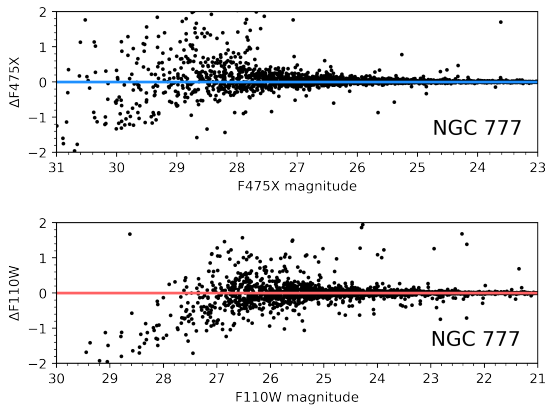


FIGURE 3.2: Difference between input and measured magnitude vs. input magnitude. Measurement error is larger on average for fainter stars. *Top:* F475X. *Bottom:* F110W.

There are a few more quality checks that can be done with only the artificial star data. First, I looked at the difference between input magnitude and measured magnitude vs. input magnitude for each artificial star. As expected and as seen in Figure 3.2, brighter stars were more likely to be measured accurately than fainter stars and show less scatter around the zero difference line, seen in blue for F475X and red for F110W. The bottom panel, for F110W, has a blank spot in the upper left-hand corner most likely because objects that would have fallen in that area were too blue to be detected in F110W.

It is also important to make sure that there is no systematic over- or underestimation of position, as the images in this work contain widely varying amounts of background light, and misplacing an

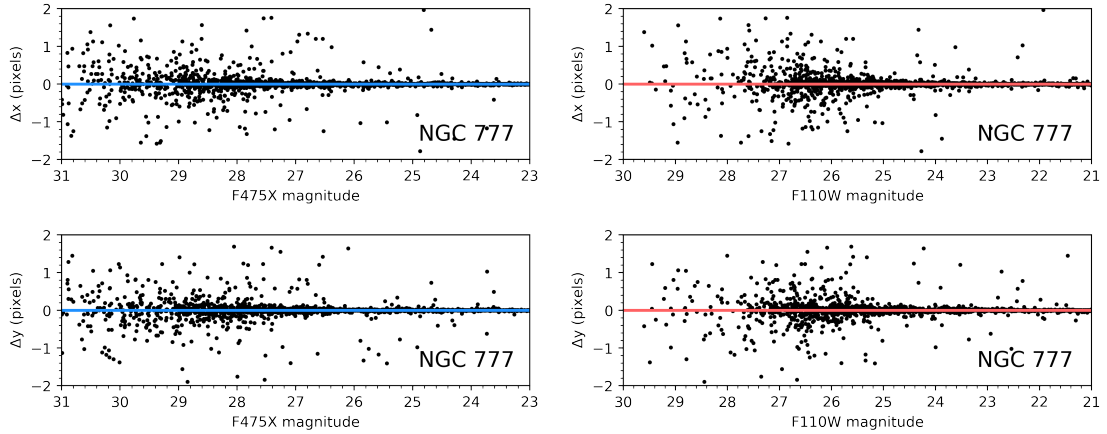


FIGURE 3.3: Difference between input and measured coordinates vs. input magnitude. Measurement error is larger on average for fainter stars. *Top left:* x coordinates in F475X. *Top right:* x coordinates in F110W. *Bottom left:* y coordinates in F475X. *Bottom right:* y coordinates in F110W.

object could lead to an inaccurate idea of how much background light is around it. To that end, I looked at the difference between input x and y position on the image and measured x and y position vs. magnitude, as seen in Figure 3.3. As expected, fainter objects have more positional scatter than brighter objects, although most scatter is within two pixels in both the x and y direction. Comparing difference between input and measured x and y coordinates with the original input coordinates also shows no adverse effects—as seen in Figure 3.4, position measurement accuracy is not dependent on the position of an object in the image.

In addition to being used on its own to look for potential problems in the photometry process, artificial star data can be compared to real photometry from the user’s images to make sure that the photometry program continues to perform well for real objects. Figure 3.5 and Figure 3.6 show good agreement where expected. As mentioned in Chapter 2, GCs are marginally resolved in F475X, which may contribute to the slightly negative (i.e. slightly wider than a point source) sharpness values in that filter.

Beyond measures of point source resemblance, the artificial stars and the real GCs begin to differ. A S/N plot

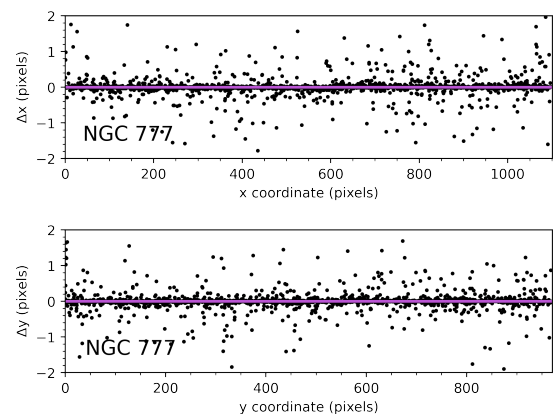


FIGURE 3.4: Difference between input coordinates and measured coordinates. *Top:* x coordinates. *Bottom:* y coordinates.

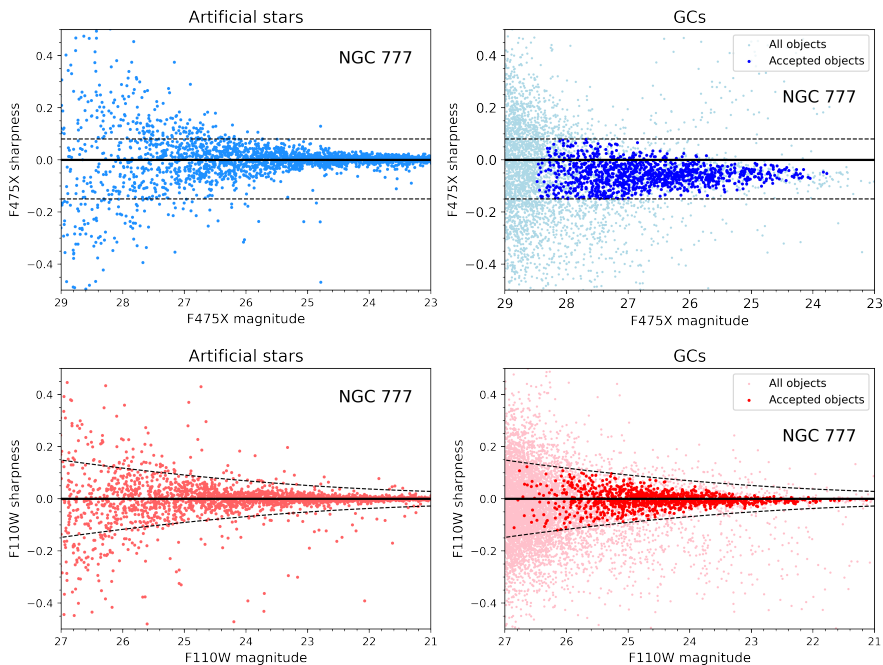


FIGURE 3.5: A comparison between artificial star and real GC sharpness measurements. Thin dashed lines denote the cutoffs described in Chapter 2. *Top left:* artificial star sharpness in F475X. *Top right:* real GC sharpness in F475X. *Bottom left:* artificial star sharpness in F110W. *Bottom right:* real GC sharpness in F110W.

of artificial stars (see the left panel of Figure 3.7) with F475X S/N vs. F110W S/N shows thorough coverage of the parameter space; because the artificial star magnitudes form a random uniform distribution, it is possible for an object's S/N in one filter to be significantly higher than its S/N in the other filter. In actuality, GC magnitudes are not random and uniform; although they have a range of colours, GCs that are brighter in one filter also tend to be brighter in the other filter, resulting in the tightened S/N distribution seen in the right panel of Figure 3.7.

CMDs are even more starkly different, as seen in Figure 3.8. The artificial star CMD is sharply rectangular at bright magnitudes and uniformly distributed through more than half of the parameter space, because of the way that artificial star magnitudes and colours are assigned. Real GCs do not have randomly assigned magnitudes and colours, so there is no rectangle in their CMD. Instead, two distinctly-coloured subpopulations appear, having a narrow range of colours and extending vertically across magnitudes—a blue subpopulation at $(F475X - F110W) \sim 2$, and a red one at $(F475X - F110W) \sim 2.5$. The two CMDs resemble each other only in the bottom left corner, where objects are too faint in one or both filters to be detected reliably. (Although "objects" are detected in that area of the real CMD in DOLPHOT's photometry, they are quickly ruled out by the initial culling steps described in Chapter 2, which are designed to limit the dataset

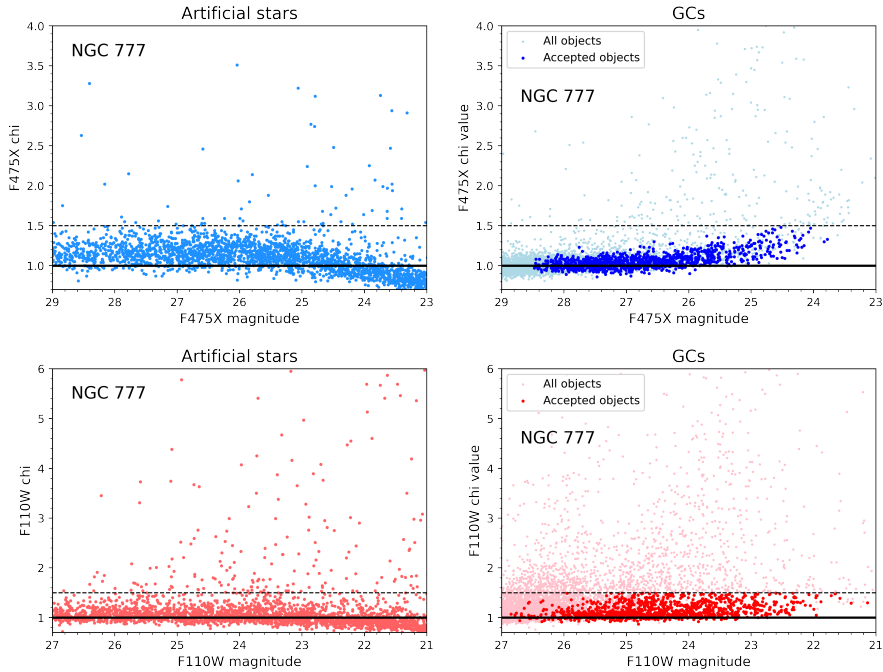


FIGURE 3.6: A comparison between artificial star chi values and real GC chi values. Thin dashed lines denote the cutoffs described in Chapter 2. The real and artificial values are comparable within those boundaries. *Top left:* artificial star chi in F475X. *To right:* real GC chi in F475X. *Bottom left:* artificial star chi in F110W. *Bottom right:* real GC chi in F110W.

to well-measured GCs.)

In addition to basic reality checks, it is possible to use artificial star tests in more complex ways to characterize the area of the CMD where objects become too faint compared to the surrounding light to be measured accurately. This is not a hard cutoff line, but a grey area where objects may or may not be detected, and it is especially apparent in the left panel of Figure 3.8, where artificial stars become less and less densely populated before fading away completely. It is possible to determine what fraction of the total number of objects is being detected, and that value is referred to as completeness—how complete is our set of detections compared to what is actually there?

3.2 Principles of Completeness

Failure to detect faint, washed-out GCs leads to a reduction in completeness—a researcher doing photometry only detects a fraction of the objects that exist in an image and obtains an incomplete dataset. When characterizing systems of objects, it is important to get accurate object counts, which incompleteness can prevent. Additionally, as




FIGURE 3.7: A comparison between artificial star S/N measurements and real GC S/N measurements. Noise from the cameras is approximately constant, while signal varies with source magnitude (i.e. a brighter object produces more signal). Because artificial star magnitudes were chosen randomly in both filters, artificial star S/N in one filter can be very different from the other filter, whereas brighter or fainter GCs tend to be brighter or fainter in both filters. *Left:* Artificial star S/N. *Right:* Real GC S/N.

a dataset captures fewer and fewer of the objects in an image, it becomes less and less advisable to use that dataset for statistics.

Completeness is quantified as a unitless number between 0 and 1 attached to each object, with 1 denoting 100% of existing objects being detected, and lower numbers indicating fractional completeness. Objects with a completeness value less than 0.5 (i.e. 50% of existing objects being detected) should not be used for analysis, because we cannot be sure that detections in that regime are representative of the full population. Completeness is also characterized with the magnitude m_0 at which the completeness fraction falls below 0.5.

The left panel of Figure 3.9 shows how we can use artificial star tests to determine m_0 for their images. To find the magnitude at which GCs are no longer reliably detected, I sorted each artificial star in DOLPHOT’s artificial star input file and each in the output file into 0.2-magnitude bins, and then compared the number of output stars (detections) to the number of input stars (objects actually there) for each magnitude bin. There is a narrow magnitude range in each filter in which the completeness value drops from ~ 1 to ~ 0 (between magnitudes 27 and 30 for F475X, and 25 and 28 for F110W). Above this dropoff, DOLPHOT detects almost all of the objects that are actually there, and below the dropoff, it detects almost none of them.

Although completeness is quantified as a fraction, detection is essentially a binary variable, i.e. a yes-or-no question—did I detect this object or not? In Eadie et al. 2022, the authors looked at the binary variable of GC presence in very low mass galaxies (does

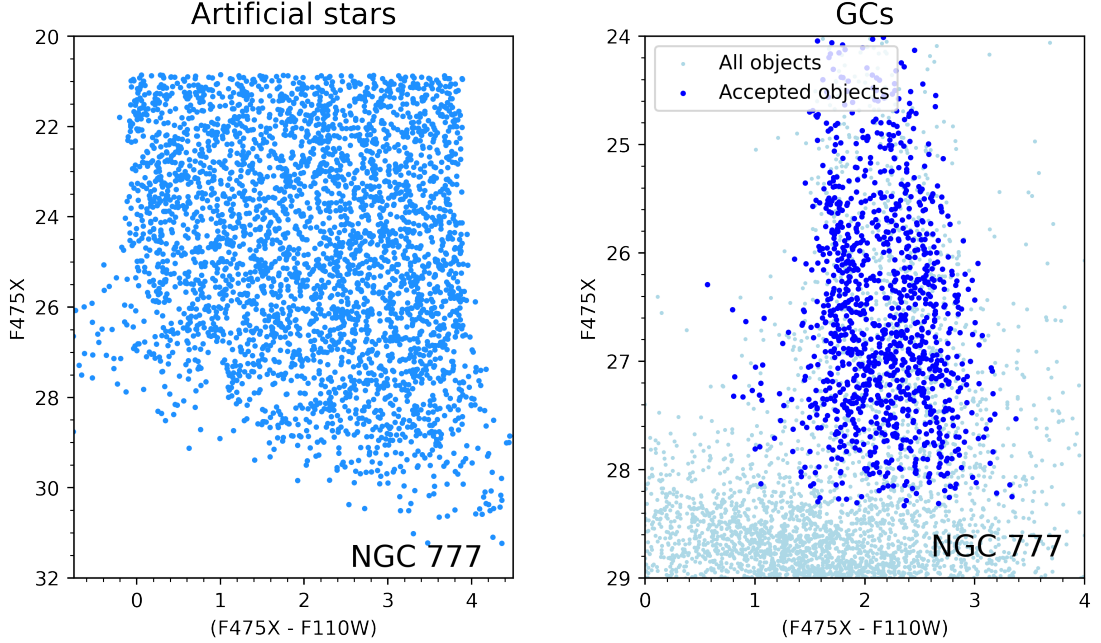


FIGURE 3.8: A comparison between the artificial star CMD and the real GC CMD, with F475X magnitudes. Artificial star magnitudes and colours are determined randomly and uniformly, which is not the case for real GC magnitudes and colours. *Left*: Artificial star output CMD. *Right*: Real GC CMD.

this galaxy have at least one GC, or does it have none?). They modelled the probability of finding at least one GC in a galaxy of a given mass using logistic regression, a statistical technique suited to the transition regions of such binary variables. Completeness takes the same approach—it is the probability of answering "yes" to the yes-or-no question of detection for objects of a given magnitude. Completeness vs. magnitude can be modelled as a sigmoid curve, described in Equation 3.1 (Harris et al. 2016):

$$completeness = \frac{1}{1 + \exp(\alpha(m - m_0))} \quad (3.1)$$

This equation captures the difference between the magnitude of a given object m and the 50% completeness cutoff magnitude m_0 , and the slope of the transition region α .

Figure 3.9 models completeness for the entirety of the NGC 777 images, which is adequate as a rough estimate for where faint GCs become unreliable for analysis—but if incompleteness is a consequence of GCs having to compete with background sky brightness, then what happens with images such as the ones from this work that have small areas of intense sky brightness? In those cases, a single estimate of m_0 is not enough to accurately characterize completeness.

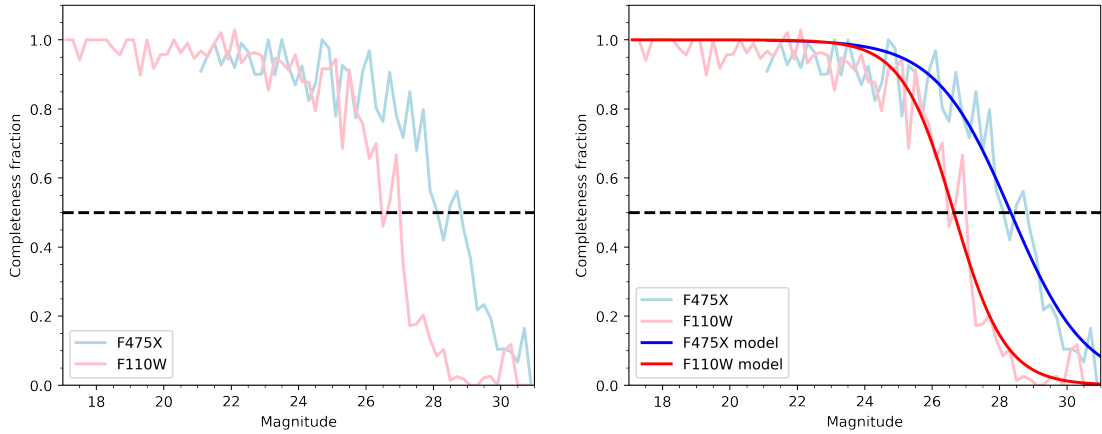


FIGURE 3.9: Completeness curves for NGC 777 in both filters. The black dotted line shows a completeness value of 0.5. *Left*: artificial star data binned by magnitude. *Right*: with sigmoid modelling overlaid.

Figure 3.10 illustrates the problem. In areas of an image with low background sky brightness, it is easy to detect faint GCs—in the left panel of the figure, GCs circled in red are significantly fainter than those circled in blue, but still visible above the background light. In areas with high sky brightness counts like that in the right panel, however, faint




FIGURE 3.10: Sky brightness as seen in portions of the F110W image for NGC 777. Bright GCs, circled in blue, would be visible in most areas of the images in this work, but fainter GCs, circled in red, would be washed out near the centre of their host galaxy. *Left*: an area with low sky brightness. *Right*: an area with high sky brightness.

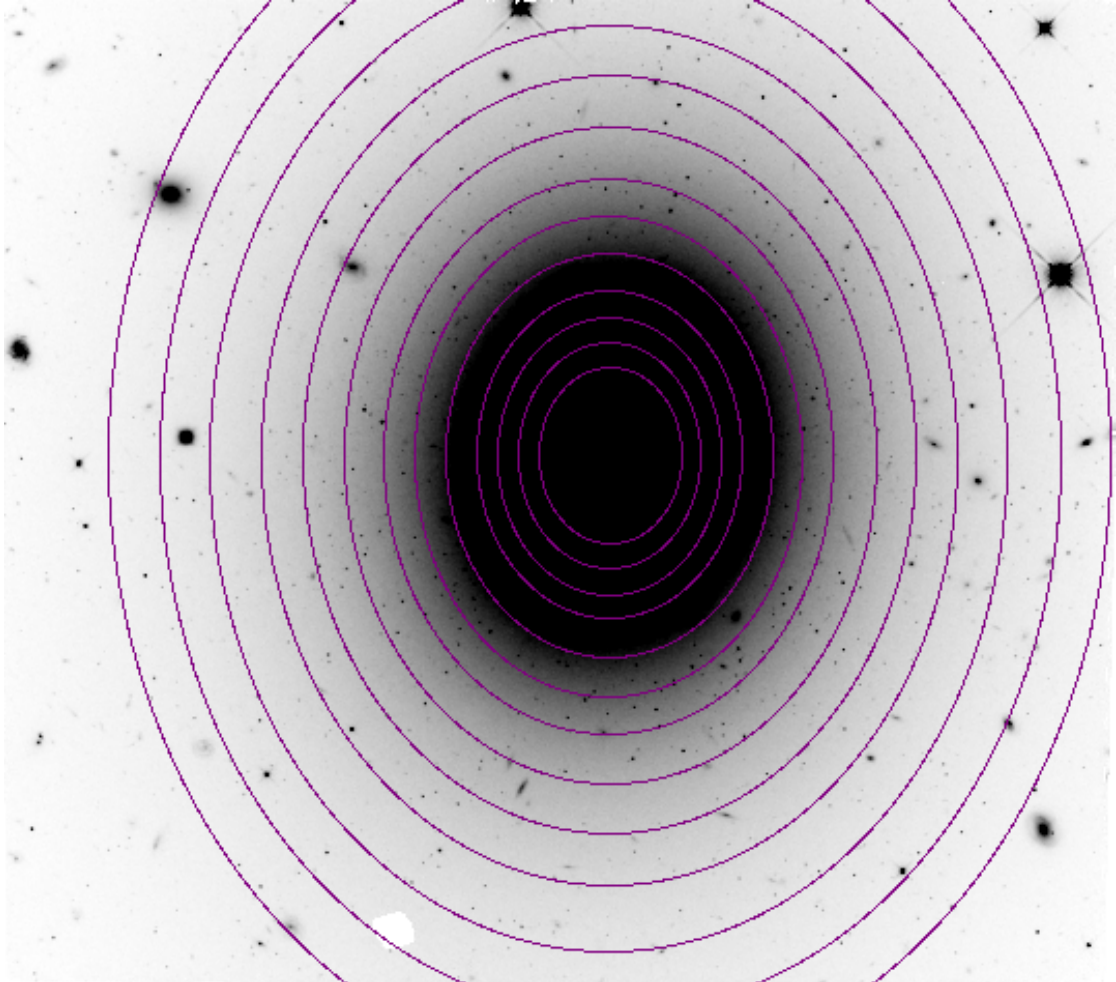


FIGURE 3.11: The twelve elliptical zones for NGC 777 overlaid on the galaxy. Zone edges have a constant eccentricity, and each zone received the same number of input artificial stars. Zone 1 is the innermost annulus, and zone 12 is the outermost annulus.

GCs fade into the background noise, and only their brighter counterparts are detectable. With that amount of sky brightness, even the brighter GCs can be difficult to see.

Because accurate object counts are critical and because those GCs in the right panel of Figure 3.10 near the centre of the galaxy are just as important as their left panel counterparts, the images in this work required a more detailed treatment of completeness—a treatment that takes into account background sky brightness.

3.3 Correcting for Completeness as a Function of Surface Brightness

I used three galaxies—NGC 57, NGC 777, and NGC 1016—as case studies, and looked at completeness value as a function of background sky brightness. If the results from each of these galaxies were similar, then it would be safe to assume that the other twelve galaxies in the sample would behave in the same way.

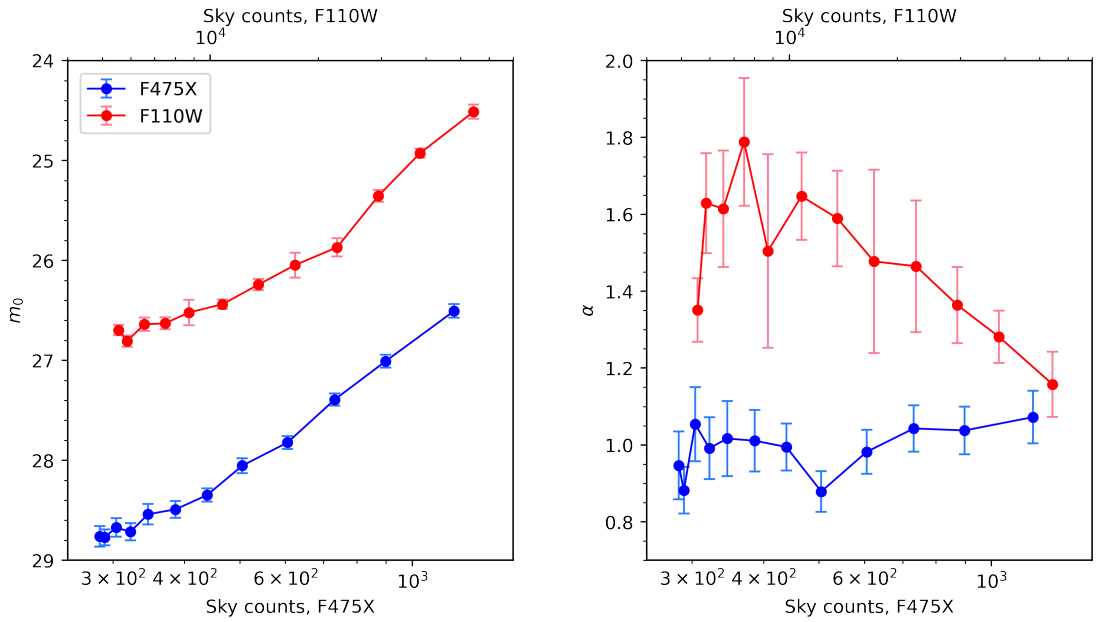


FIGURE 3.12: Fitted sigmoid parameters vs. sky brightness with one-sigma error bars in both filters (F475X in blue, F110W in red) for the NGC 777 artificial star test. Variations of the slope α with sky brightness are modest, while the 50% cutoff magnitude becomes brighter along with the sky background. *Left:* m_0 vs. sky counts. *Right:* α vs. sky counts.

Each image in this work has a bright elliptical galaxy at the centre and background light fading away toward the edges of the image. Because of that layout, completeness values are lower in the centre of the image and higher at the edges, and a study of completeness as a function of sky brightness would need to track that increase. To that end, I divided the space around each case study galaxy into twelve zones, each capturing a different level of sky brightness. Zones are numbered from 1 to 12, with 1 being closest to the centre of the galaxy and 12 being farthest. Because elliptical galaxy surface brightness falls off exponentially (e.g. Kormendy et al. 2009), I constructed the inner zones to be narrower than the outer zones. Each quartile of zones is 20 pixels wider in radius (for NGC 1016) or semiminor axis (for NGC 57 and NGC 777) than the previous quartile on the F110W image.

NGC 1016 is roughly spherical, so that galaxy’s zones are simple and circular. In that

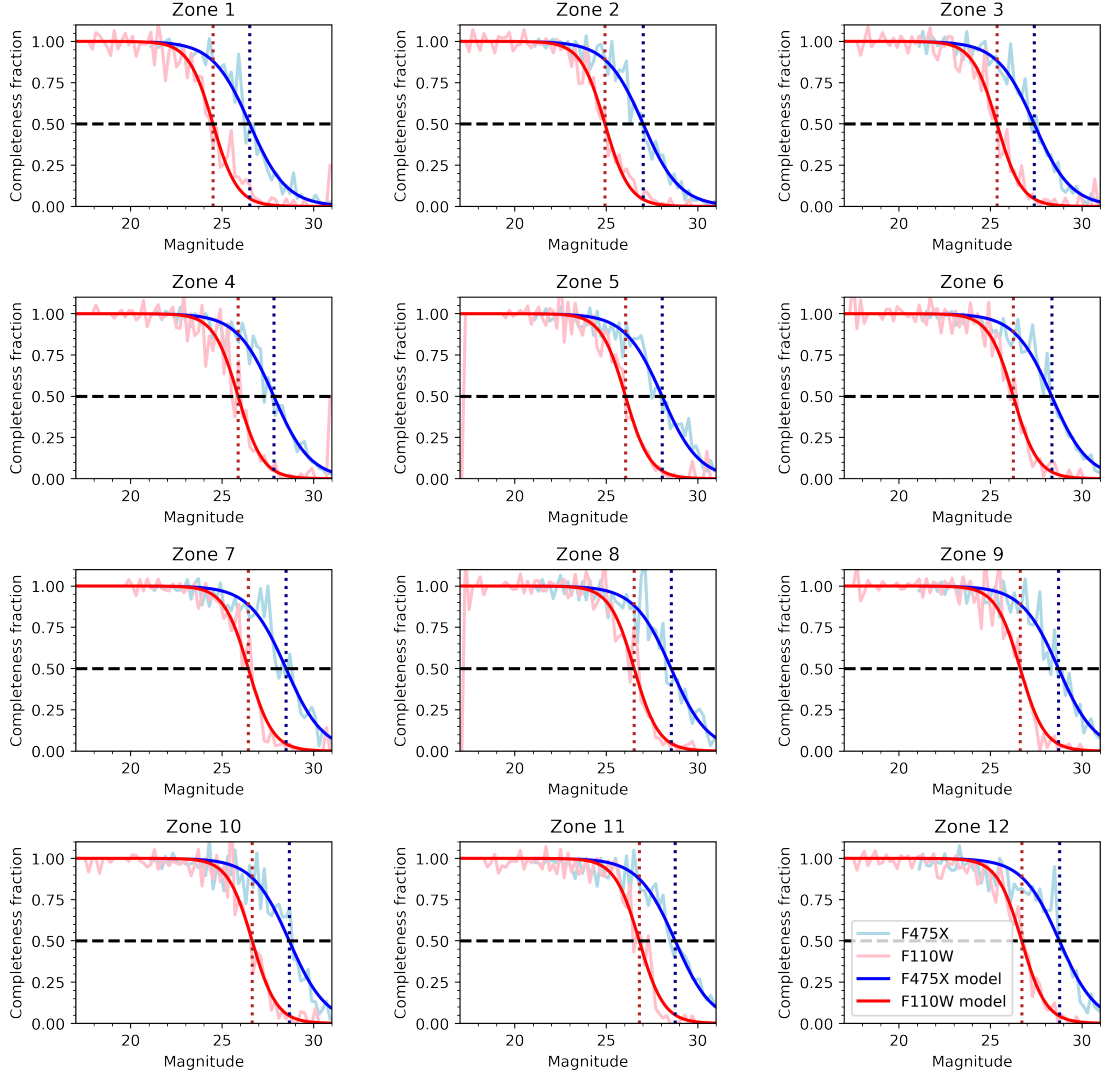


FIGURE 3.13: Completeness curves and fitted sigmoids for the twelve zones around NGC 777. Zone 1 is closest to the centre of the galaxy with the brightest sky background, and zone 12 is farthest with the lowest sky background. As the sky brightness drops, so does the 50% cutoff magnitude. Vertical dotted lines show the position of m_0 for both filters.

case, I sorted artificial stars into zones by their distance from the centre of the galaxy. NGC 57 and NGC 777 are elliptical, so their zone geometry is more complicated. Zones for those galaxies are ellipses of constant eccentricity, and instead of making a simple distance calculation, I checked whether an artificial star fell between the two elliptical edges of a given zone by using Equation 3.2, the equation for an ellipse:

$$1 = \frac{x^2}{a^2} + \frac{y^2}{b^2} \quad (3.2)$$

The elliptical zones for NGC 777 can be seen in Figure 3.11.

In order to obtain adequate recovered star counts in each zone, I increased the star count from 5000 to 42,000. DOLPHOT places artificial stars randomly and uniformly across the image, so I also redistributed them so that each zone contained 3500 stars placed randomly within that zone (I kept DOLPHOT’s magnitudes and colours and changed only the coordinates of the stars).

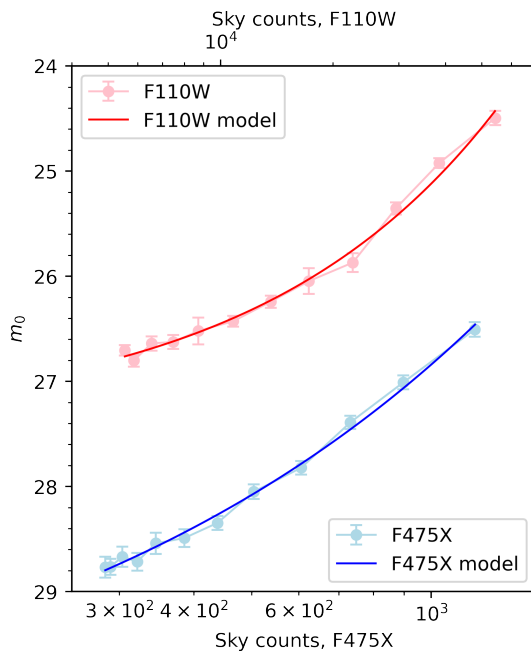


FIGURE 3.14: 50% cutoff magnitude vs. sky background counts in both filters for NGC 777. Both m_0 curves are fitted with an exponential function.

ues dependent more exclusively on background brightness. The results of the second sigmoid fitting can be seen in Figure 3.13, which illustrates how m_0 becomes fainter in the outer zones where sky brightness is lower.

Once I had sky brightness and m_0 for each filter in each zone, I could relate those two values to each other. I fitted two exponential curves, one for each filter, modelling m_0 as a function of sky brightness as seen in Figure 3.14. The exponential curves for both filters are very similar across all three case study galaxies (Figure 3.15) despite differences between those galaxies in distance and shape, demonstrating that the relation would hold across the entire sample of galaxies in this work. That generalization is possible because all fifteen galaxies were imaged with the same filters and the same exposure times, and so are all directly comparable. With this in mind, I calculated the mean exponential

After running photometry on the artificial stars, I determined which zone an output star fell into using the algorithms described above. I then repeated the process from the previous section to create completeness curves for each zone. At that point, the connection to sky brightness came into play; I fitted a sigmoid (Equation 3.1) to each completeness curve, finding m_0 and α values for both filters in each zone, and then plotted those values against each zone’s mean sky brightness, as seen in Figure 3.12. On the left side of Figure 3.12, m_0 becomes brighter with increasing sky counts, as expected (i.e. objects become too incomplete to use in analysis at brighter magnitudes), and on the right, any variations in α , the slope of the sigmoid, are modest in magnitude and of little importance compared to m_0 . After observing the behaviour of α , I found the mean α value in each filter, held it constant in the sigmoid fitting routine, and ran the routine again to get m_0 values

parameters for the three case study galaxies and adopted those as the parameters in my sample-wide relation between m_0 and sky brightness.

The relation shown in Figure 3.14 and Figure 3.15 determines the the 50% completeness cutoff magnitude m_0 for a given background brightness, but how is that information useful? In order to learn something about the completeness of my GCs, I had to take each GC backward through the processes described in this section and the previous one. In addition to magnitudes in each filter, DOLPHOT photometry provides a sky background count; I took those values for each object and used the exponential m_0 -sky relation to find m_0 in each object’s area of the image. Once I had m_0 for each object, I returned to the sigmoid curve (Equation 3.1); by inputting m_0 , the fixed alpha value, and each GC’s magnitude in both filters, I found a completeness fraction for each GC.

With completeness fractions for all of my objects, I made one more cut of the dataset for each galaxy, dropping GCs with a completeness fraction lower than 0.5. As mentioned earlier, it is not advisable to do statistics with GCs from areas of the image where less than half of the existing objects are being detected. For all fifteen galaxies, this removed a small handful of objects, less than 10% of the remaining sample.

Completeness fractions also allowed me to better estimate the true number of GCs in my images—by quantifying how partially DOLPHOT was counting the GC population, I could add the GCs it was missing back into the sample. Because completeness is the ratio of detected objects to existing objects, I could invert the completeness value f_b from F475X and f_r from F110W and multiply them by the number of detected objects, which in this case is 1:

$$N = \frac{1}{f_b} * \frac{1}{f_r} \tag{3.3}$$

Equation 3.3 gives the number of objects N that are actually there in a particular area of the image given that DOLPHOT has detected one. N is almost always slightly greater than 1 (the only instance in which it is exactly 1 is if the associated object has a completeness value of 1.0 in both filters, and because completeness values are never

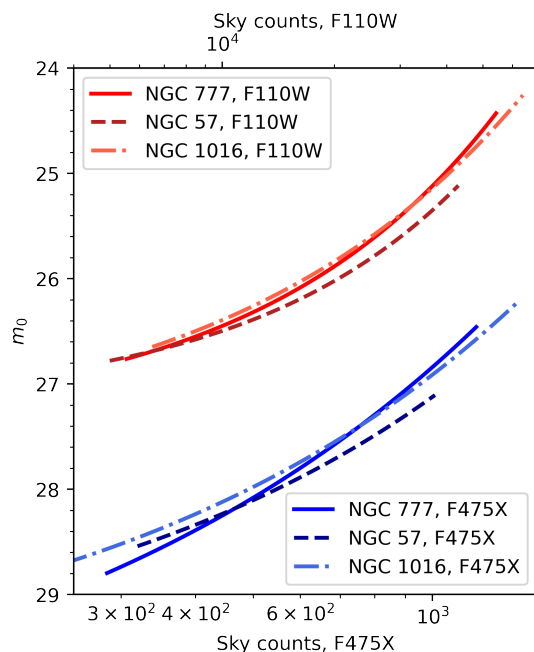


FIGURE 3.15: 50% cutoff magnitude vs. sky background counts for all three case study galaxies. The exponential curves are very similar between galaxies in both filters.

greater than 1.0, N is never less than 1), and totalling these completeness-corrected objects results in a non-integer number of GCs in each system. But this step is necessary to minimize the potential for observational bias (in case, for example, red GCs tend to be undercounted compared to blue ones, or vice versa), and it ensures that the proportions of GCs of various colours are as similar as possible to those found in the actual GC population. I made this completeness correction for all of the GCs that remained after the culling steps described in Chapter 2 and the cut of objects with a completeness value less than 0.5 and used the corrected GC counts for statistical analysis.

4.1 Colour vs. Metallicity for Optical-Infrared Colour Indices

I used the photometry described in Chapter 2 to create a colour distribution function (CDF)—a measure of how many GCs in a system are present at a given colour index value—for every galaxy in my sample. But while colour indices are easily observable, what we really want is the metallicity distribution function (MDF) for each GCS, a measure of how many GCs are present at a given $[M/H]$ value.

In an ideal world, metallicity measurements would come from spectroscopy—but spectroscopy is expensive and time-consuming, especially for faint extragalactic GCs. High-quality spectroscopic measurements for nearby extragalactic GCs have required major observing campaigns (e.g. [Villaume et al. 2019](#) and [Fahrion et al. 2020](#)), and those campaigns targeted objects within 25 Mpc. Spectroscopy for the sample of galaxies in this work, the nearest of which is 66 Mpc away, is completely out of the question without 30m-class telescopes.

Instead, this work relies on photometry and colour indices. In addition to being observationally efficient, colour in GCs is sensitive mainly to metallicity, which allows us to use GC colour as a proxy for metallicity. However, the relation between colour and metallicity is nonlinear across all wavelengths ([Brodie and Strader 2006](#), [Harris 2010](#)); GCs' integrated light comes primarily from giant stars, whose colour becomes redder at a faster rate as metallicity increases. This means that the red end of a CDF is compressed, and the blue end is broadened, upon conversion to an MDF. Because of this nonlinearity, the CDF must be converted to an MDF in order to get accurate proportions of metal-poor to metal-rich clusters.

Some colour indices, such as $(g - z)$ (e.g. Villaume et al. 2019 and Fahrion et al. 2020), equivalent to $(F475W - F850LP)$ in *HST* terms, have spectroscopic colour-to-metallicity calibrations. But because of the expense, time, and effort involved in creating such calibrations, other indices, including the $(F475X - F110W)$ index used in this work, do not. In order to move from CDFs to MDFs, therefore, I used a set of simulated GCs to jump from $(F475X - F110W)$ to $(F475W - F850LP)$, with an observed spectroscopic calibration, to metallicity.

4.2 Creating a Colour-Metallicity Conversion

To determine the colour-colour conversion, I created simulated GCs using the CMD 3.6 online tool from the Osservatorio Astronomico di Padova (http://stev.oapd.inaf.it/cgi-bin/cmd_3.6).

This tool makes use of PARSEC and COLIBRI stellar evolution tracks (Bressan et al. 2012 and Marigo et al. 2013, respectively), and allows the user to specify the physical properties of the stellar populations to be simulated. Users choose which versions of PARSEC and COLIBRI to employ, how to handle circumstellar dust, how to handle interstellar extinction, a prescription for post-main sequence long-period variability, an initial mass function (IMF), and age and metallicity ranges for the simulated stars.

For more information about these models, see Bressan et al. 2012, Chen et al. 2014, Chen et al. 2015, Tang et al. 2014, Marigo et al. 2017, Pastorelli et al. 2019, and Pastorelli et al. 2020.

The Padova tool also allows the user to choose a photometric system, and in stellar population mode, it produces a magnitude for each simulated star in each filter of the specified system. One of the options is the full set of *HST* WFC3 filters, which I used to compare my $(F475X - F110W)$ index to the spectroscopically calibrated $(F475W - F850LP)$ index.

For the other parameters, I specified an age of 12 Gyr, a typical age for a GC (Beasley 2020), and a single metallicity, and a total cluster mass of $10^5 M_{\odot}$, and left everything else in its default setting. The most metal-poor GC had a metallicity of $[M/H] = -2.2$, with subsequent GCs stepping up in metallicity by $[M/H] = 0.05$ until $[M/H] = -1.9$, and then by $[M/H] = 0.1$ until $[M/H] = 0.0$, with a return to the $[M/H] = 0.05$ steps from $[M/H] = 0.0$ to $[M/H] = 0.3$. The difference in metallicity between simulated GCs was smaller at the ends of the distribution so that those ends would be well-characterized during the curve fitting steps of the simulated GC procedure. This produced single-population GC simulations with a known metallicity; the output file contained stellar parameters and an absolute magnitude in every WFC3 filter for every simulated star. Using the F475X, F110W, F475W, and F850LP columns, I summed the light from each simulated star to obtain integrated GC magnitude in the relevant filters, and then calculated the two colour indices.

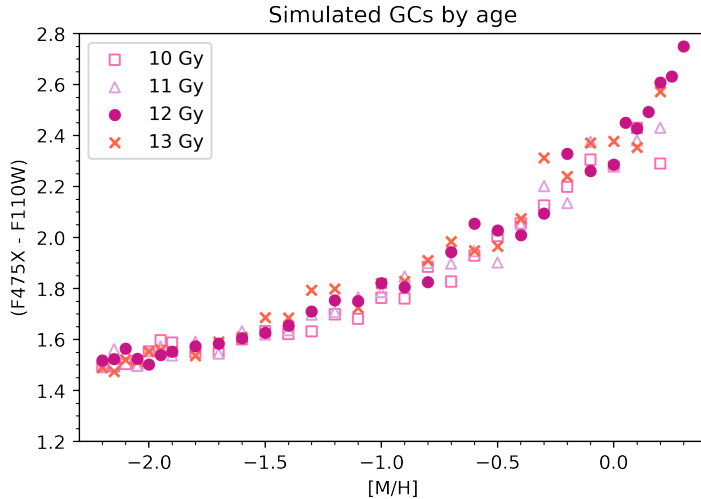


FIGURE 4.1: Colour vs. metallicity for simulated GCs of age 10 Gyr (open squares), 11 Gyr (open triangles), 12 Gyr (filled circles), and 13 Gyr (crosses). There is no discernible difference between GC age groups.

GCs can differ in age, so to avoid including age effects in my colour-colour-metallicity conversion, I also simulated GCs with ages of 10 Gyr, 11 Gyr, and 13 Gyr (all other parameters were the same as those for the 12-Gyr GCs). Figure 4.1 shows $(F475X - F110W)$ colour versus metallicity for all four age sets, with the 12-Gyr GCs as dark pink filled circles. There is no discernible difference between the different age sets, so I continued with the 12-Gyr GCs as my fiducial set of simulations.

The scatter in the simulated GC figures comes from the somewhat random determination (within IMF and other parameters) of the properties of individual stars. It is not expected to compromise curve fitting results.

After ensuring that the age of the GCs would not affect colour, I quantified the colour-metallicity relation. I fitted exponential, cubic, and quadratic curves to the 12-Gyr GC data, as seen in Figure 4.2. All three curves described the data well at mid-range

metallicities, but the quadratic curve tracked too low at high $[M/H]$ (colour index should redden faster at high metallicities) and swung back redward at low $[M/H]$ (despite colour becoming less sensitive as metallicity decreases, it should still move blueward). The exponential and cubic curves were comparable throughout the dataset; I continued with

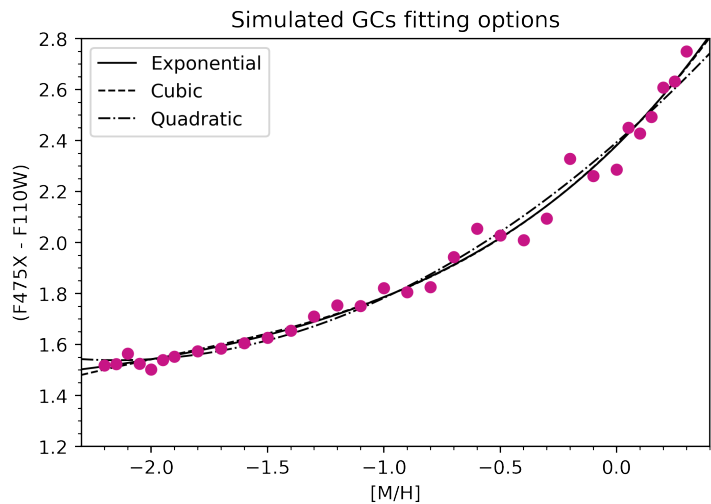


FIGURE 4.2: Exponential, cubic, and quadratic equations fitted to the set of 12-Gyr simulated GCs.

an exponential curve.

The simulated GC procedure used GCs with known metallicity and derived colour to find a colour-metallicity relation, but I needed the opposite—the photometry described in Chapter 2 produced objects with known colour, and I needed to derive metallicity. To get the metallicity-colour relation that such a derivation requires, I inverted the simulated GC data and again fitted an exponential curve. At this point in the process, I added a set of simulated GCs for the spectroscopically calibrated [F475W – F850LP] index. F475W and F850LP were filed under a different filter set in the Padova online tool, so they required a separate set of GCs; all parameters other than the filter set were identical to those for the (F475X – F110W) GCs, but because of the somewhat random nature of the simulations, the scatter is slightly different. Both sets of GCs can be seen in Figure 4.3 ((F475X – F110W) in pink, (F475W – F850LP) in blue), with the original colour-metallicity relation on the left, and the target metallicity-colour relation on the right. This figure illustrates the complexity of the ties between colour and metallicity—the two relations occupy different but overlapping areas in colour space and display different curvatures.

Because both sets of GCs came from the same program, I could compare them directly to get an internally consistent colour-colour relation, shown in Figure 4.4 with points coloured by metallicity. Again I fitted exponential, cubic, and quadratic equations to the relation, and again the exponential and cubic options described the data best at extreme colours. As before, I continued with the exponential fit:

$$(F475W - F850LP) = -5.000 \exp(-0.830 * (F475X - F110W)) + 2.785 \quad (4.1)$$

Equation 4.1 allowed me to input an observed (F475X – F110W) colour index and




FIGURE 4.3: *Left:* colour vs. metallicity for simulated GCs in (F475X – F110W) and (F475W – F850LP), with exponential fits. *Right:* metallicity vs. colour for simulated GCs in (F475X – F110W) and (F475W – F850LP), with exponential fits.

transform it to an $(F475W - F850LP)$ index. From there, I could use the inversion of the spectroscopic calibration shown in Equation 4.2, a quadratic fit to the Fahrion et al. 2020 spectroscopic data (Harris, private communication)

$$(F475W - F850LP) = 2.158 + 0.57081[Fe/H] + 0.10026[Fe/H]^2 \quad (4.2)$$

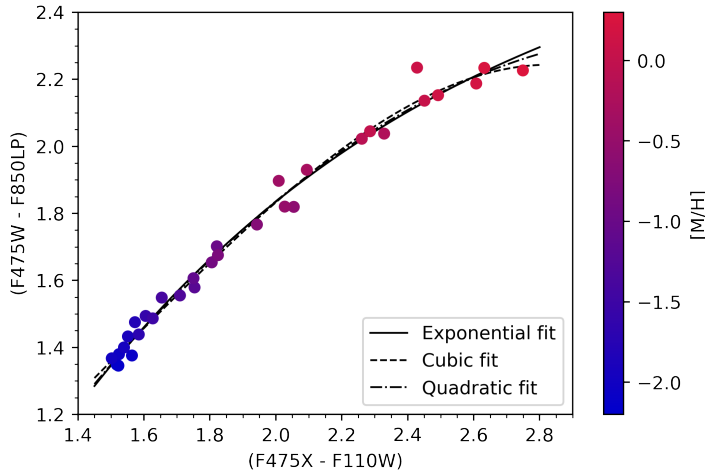


FIGURE 4.4: A colour-colour plot for the 12 Gyr GCs; $(F475W - F850LP)$ vs. $(F475X - F110W)$.

same shape, with a slight offset.

With the completeness correction outlined in Chapter 3 and the two-step conversion described in this chapter, I could derive useful and physically meaningful information from the photometric data for each galaxy in my sample of fifteen—an estimate of the true number of objects present in each GCS, and the metallicity of each object. The only item still needing to be addressed was the other aspect of this work: host galaxy environment.

to derive the associated metallicity.

Note that Equation 4.2 involves iron content $[Fe/H]$, but the Padova online tool parameter is general metal content $[M/H]$. The two colour-metallicity relations are typically about the same shape, with $[M/H]$ offset toward higher metallicity values than $[Fe/H]$. Figure 4.5 compares Equation 4.2, in purple, to the fitted exponential curve from the simulated GCs, in blue. As expected, the curves are roughly the




FIGURE 4.5: A comparison between the $[Fe/H]$ -based spectroscopic calibration and the $[M/H]$ -based simulated GCs.

CHAPTER 5

QUANTIFYING THE METALLICITY DISTRIBUTIONS

5.1 Applying Completeness Corrections and Metallicity Conversions

The completeness calculation described in Chapter 3 Equation 3.1 allowed for further cleanup of the GCS data in this work. As mentioned in that chapter, objects with completeness values less than 0.5 in any filter should not be used for analysis; the initial culling measures outlined in Chapter 2 removed most insufficiently complete objects from the dataset, and adding completeness itself as a culling criterion (i.e. rejecting objects with at least one completeness value < 0.5) removed the rest. For all fifteen galaxies in this work, the completeness cut removed only a handful of objects—NGC 777, for example, shown in Figure 5.1, lost 31 objects to this cut.

For the entire set of fifteen post-completeness cut CMDs, see Figures 5.2 (F475X vs. colour) and 5.3 (F110W vs. colour).

In order to better visualize the distribution of colours in each GCS, the colours seen on the x-axis of the Figure 5.1 plots can be replotted as a histogram. I sorted each GC into bins 0.1 magnitudes wide to create Figure 5.4. Although the presence of red and blue subpopulations can be seen in the CMDs of Figure 5.1 (look for the drop in object density at bright magnitudes near $(F475X - F110W) \approx 2$), it is clearer in Figure 5.4. The blue population appears as a peak centred at $(F475X - F110W) \approx 1.8$, and the red at $(F475X - F110W) \approx 2.3$.

At this point in the analysis, it was also possible to use the completeness correction described in Chapter 3 for its main purpose—to estimate how many objects are present for every object detected in the images for this work. Figure 5.4 shows two histograms; the shorter histogram, in light pink and labelled "uncorrected," tallies the number of GCs detected by DOLPHOT, while the taller histogram, in dark pink and labelled

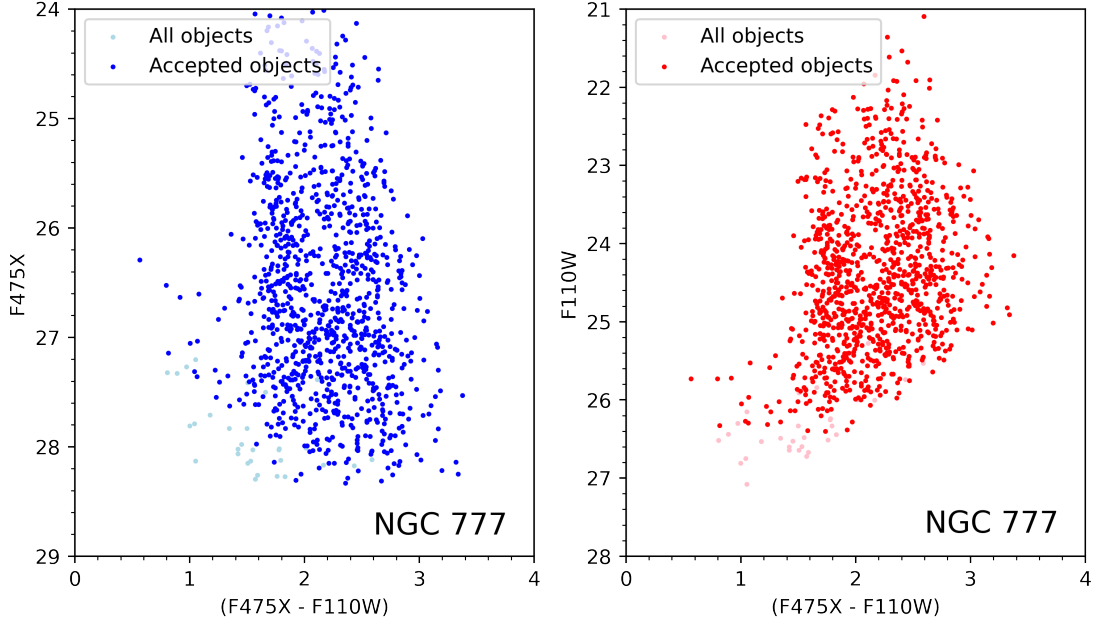


FIGURE 5.1: CMDs for NGC 777. All objects in the plots have made it through the initial cuts described in Chapter 2, and the light-coloured points were dropped because of low completeness values. *Left*: F475X CMD. *Right*: F110W CMD.

"corrected," shows the same data after treatment with Equation 3.3, which takes into account the completeness fractions of each object in both F475X and F110W. Equation 3.3 adjusts the initial count for each object (i.e. 1) up, approximating the number of objects actually present in the galaxy for each one observed. The adjusted histogram preserves the shape of the original histogram while providing a better estimate of object counts in each bin.

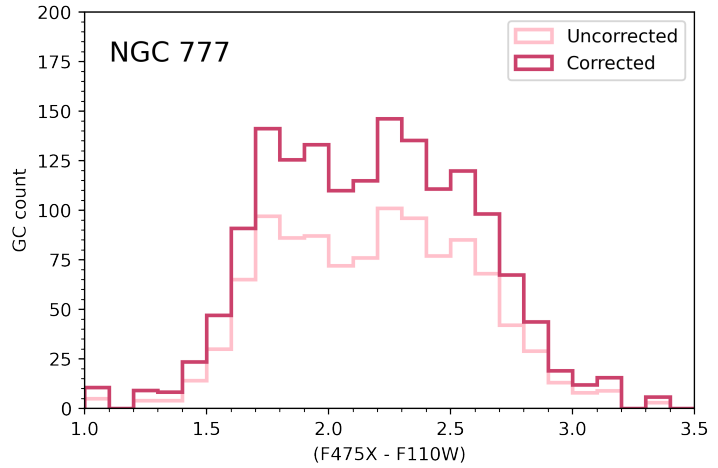


FIGURE 5.4: The colour histogram for NGC 777, with and without completeness corrections.

After de-reddening by subtracting extinction in each filter (see Schlafly and Finkbeiner 2011), colour values can be transformed into metallicity values, using Equation 4.2. As

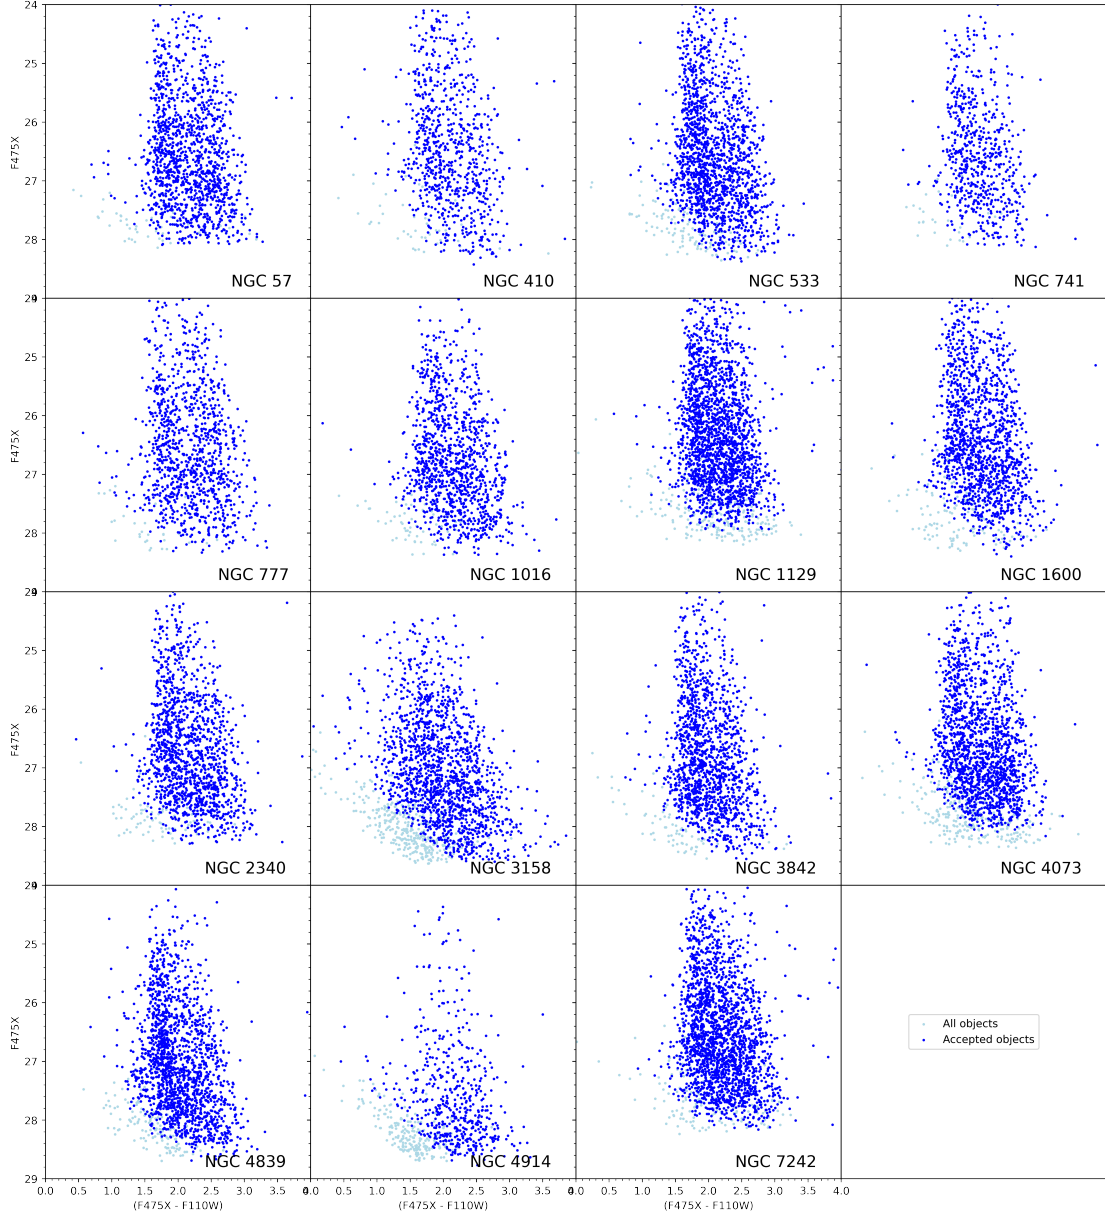


FIGURE 5.2: Post-completeness cut CMDs; F475X vs. colour.

explained in Chapter 4, the relation between GC colour and metallicity is nonlinear—and that nonlinearity is illustrated starkly in Figure 5.5. The left side of Figure 5.5 shows the colour histogram from Figure 5.4, and in a direct comparison, the right side shows the metallicity histogram from the same GCS. After applying Equation 4.2 to the DOLPHOT photometry, I sorted the resulting metallicity data into bins 0.15 dex wide. As colour in GCs reddens, its sensitivity to metallicity increases, leading to a significant difference in shape between the colour and metallicity histograms—the red end of the

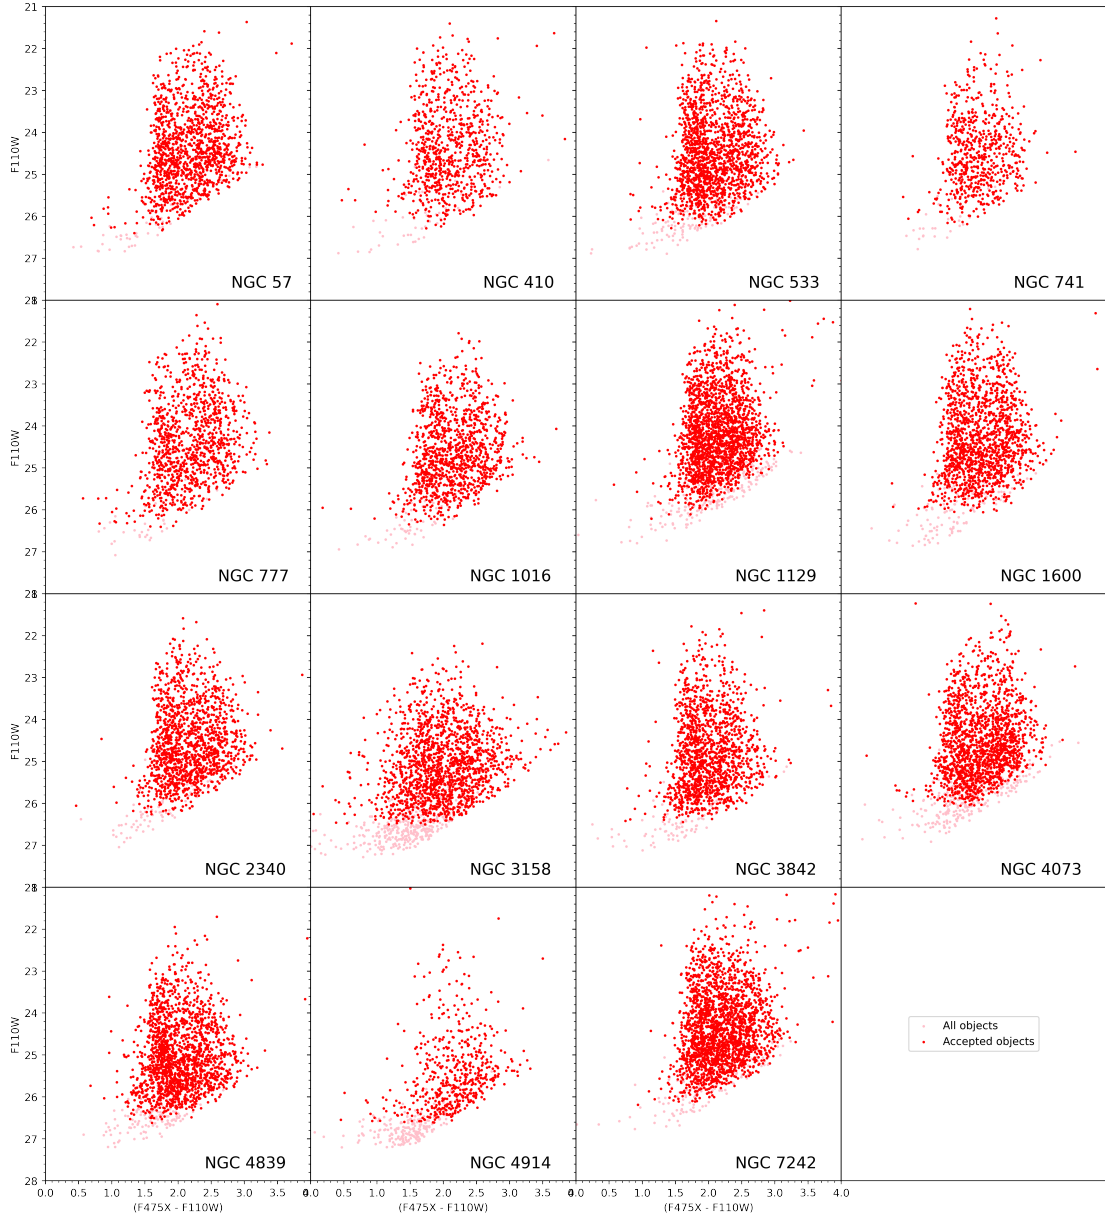


FIGURE 5.3: Post-completeness cut CMDs; F110W vs. colour.

metallicity histogram drops off much more sharply compared to the colour histogram, and the blue end has been stretched out into a long tail. It is still possible, however, to see the two peaks in the metallicity histogram; what was the blue peak in the colour histogram has become a bend in the blue tail at $[M/H] \approx -1.2$ dex in the metallicity histogram.




FIGURE 5.5: Colour and metallicity histograms for the GCS of NGC 777, with and without completeness corrections. The conversion from colour to metallicity sharpens the red peak and stretches the blue peak into a long tail. *Left*: colour histogram. *Right*: metallicity histogram.

5.2 Double Gaussian Fitting

The next step after acquiring metallicity distributions was to quantify them. As explained in Chapter 1, large GCSs tend to contain subpopulations of red GCs and blue GCs, so any quantitative model of the metallicity distributions of those GCSs must allow for multiple peaks. With that in mind, I focused on multiple Gaussian curve fitting for this work.

While testing a double Gaussian fit on un-completeness-corrected metallicity data, I used a Gaussian mixture modelling (GMM) program from Muratov and Gnedin 2010. This program tests for uni-, bi-, or multimodality in a distribution, fits multiple Gaussian curves to the distribution, and calculates a value quantifying the separation of peaks and other statistical metrics. GMM provided an efficient method for testing the validity of a double Gaussian model for the GCSs in this work; the GMM goodness of fit metrics and visual inspection favoured a double Gaussian model over a single Gaussian model, and showed negligible difference between bimodal and multimodal models.

For more information on this version of GMM, see the appendix of Muratov and Gnedin 2010.

GMM was not an option for the completeness-corrected metallicity data—it expects integer GC counts in each histogram bin, which is no longer the case after a completeness correction—so I fitted those histograms with a simple bespoke python routine. I used the `curve_fit` function from SciPy’s `optimize` package to fit a double Gaussian

$$N = A_{blue} * \exp \frac{-(x - \mu_{blue})^2}{2 * \sigma_{blue}^2} + A_{red} * \exp \frac{-(x - \mu_{red})^2}{2 * \sigma_{red}^2} \quad (5.1)$$

to each completeness-corrected metallicity histogram. To predict the number of GCs N at a given metallicity, Equation 5.1 simply stacks two Gaussian curves—one for the blue peak, with amplitude A_{blue} , mean μ_{blue} , and width σ_{blue} ; and one for the red peak, with amplitude A_{red} , mean μ_{red} , and width σ_{red} . This fitting produces a model like that shown in Figure 5.6—the bimodal fit follows the completeness-corrected metallicity distribution and shows significant overlap between the red and blue peaks.

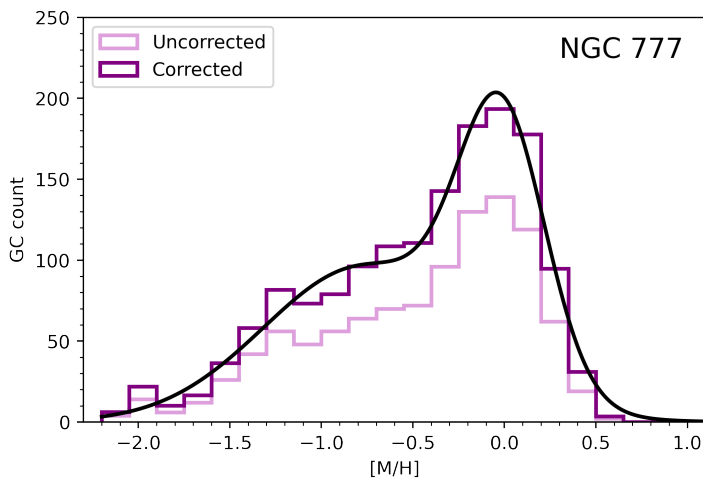


FIGURE 5.6: The metallicity histogram for NGC 777 fitted with a double Gaussian curve. The original histogram is plotted in light purple, the completeness-corrected histogram in dark purple, and the double Gaussian fit in black.

found the absolute magnitude of the faintest GC remaining in its dataset using the distance modulus

$$m - M = \log_{10} \frac{d}{10 \text{ pc}} \quad (5.2)$$

where m is an object’s apparent magnitude, M is its absolute magnitude, and d is its distance in parsecs. I then used Equation 5.2 to calculate absolute magnitudes for objects in the other GCSs, and discarded all those fainter than the faintest GC in NGC 4839. Because NGC 4839 is the farthest galaxy in the sample, its faint GCs have the faintest apparent magnitudes and are hardest for DOLPHOT to detect and measure; absolute magnitudes observable in a nearer galaxy may not be observable in NGC 4839, so it was necessary to use that galaxy for this cutting criterion. By imposing an absolute magnitude floor on all fifteen GCSs in this work, I ensured that the same portion of the GC luminosity function would be accounted for in each galaxy.

Second, I looked at NGC 1600, the nearest galaxy in the sample at $D \sim 66$ Mpc, determined the distance d in pixels from the centre of the galaxy to the nearest F110W

Although the exposure times are the same for each filter across all fifteen galaxies in this work, and although the galaxies lie within a narrow range of stellar masses, they are not uniformly distant; furthermore, the field of view of *HST*’s WFC3 camera is too small to capture the entire GCS of each galaxy. Therefore, in order to compare the same physical portion of each GCS, I made two final cuts to the dataset. First, I looked at the GCS of NGC 4839, the most distant galaxy in the sample at ~ 109 Mpc, and

image edge (F110W has a smaller field of view than F475X), and converted it to kiloparsecs:

$$d = D \left(\frac{1000\text{kpc}}{\text{Mpc}} \right) * r_{\text{arcsec}} \left(\frac{1^\circ}{3600''} \right) \left(\frac{\pi}{180^\circ} \right) \quad (5.3)$$

For NGC 1600, this yielded a physical distance of ~ 18.2 kpc; I then adopted that distance as a radial cutoff, found its equivalent in F110W pixels for the remaining fourteen galaxies in the sample using the inverse of Equation 5.3, and discarded all objects outside of that radius relative to the centre of each galaxy.




FIGURE 5.7: The remaining GCs in the sample for NGC 777 after initial culling, the completeness cut, the magnitude cut, and the radial cut.

GCSs of large elliptical galaxies can exhibit a metallicity gradient, with proportions of blue clusters higher in the outskirts of the galaxy than near the centre (Beasley 2020), so it was important to ensure that the GCS data came from the same physical portion of each GCS. The results of this cut can be seen in Figure 5.7 for NGC 777, which shows the remaining objects plotted over the F110W image of the galaxy; at ~ 68 Mpc, NGC 777 is not much more distant than NGC 1600, but it is still possible to see a small gap between the cutoff radius and the nearest image edge at the top of the figure. Like the magnitude floor, setting a cutoff radius based on a galaxy other than NGC 1600 (i.e. a more distant galaxy) would have led to some GCs within that radius being

unobservable within closer galaxies.

Having made GCS inclusion parameters as uniform as possible across galaxies, I remade the metallicity histograms and refitted the double Gaussian curves. The new, post-cut double Gaussian for NGC 777 appears in Figure 5.8; note the difference in scale between the full-image Figure 5.6 and the magnitude- and radius-cut Figure 5.8. Although the GC count in each bin has changed, the shapes of the pre- and post-cut metallicity histograms are approximately the same. As described above, I fitted a double Gaussian to the new metallicity histogram using SciPy’s `optimize` package and `curve_fit` function (see the black line in Figure 5.8). The fitting parameters, and their one-sigma errors calculated with the covariance matrix produced by `curve_fit`, can be seen in Table 5.1. For the full table with all fifteen galaxies, see Appendix C.

The values in Table 5.1 are typical of the GCSs in this work; the red peaks tend to be higher and narrower than the blue peaks, and they appear around Solar metallicity, while the blue peaks range between $[M/H] \sim -0.4$ and ~ -1.0 .

TABLE 5.1: Fitted double Gaussian amplitudes, peak positions, and widths for NGC 777, with one-sigma uncertainties.

Galaxy	Blue amplitude	Blue mean	Blue sigma	Red amplitude	Red mean	Red sigma
NGC 777	50.8 ± 3.3	-0.72 ± 0.11	0.52 ± 0.08	122.5 ± 11.4	0.00 ± 0.01	0.23 ± 0.02

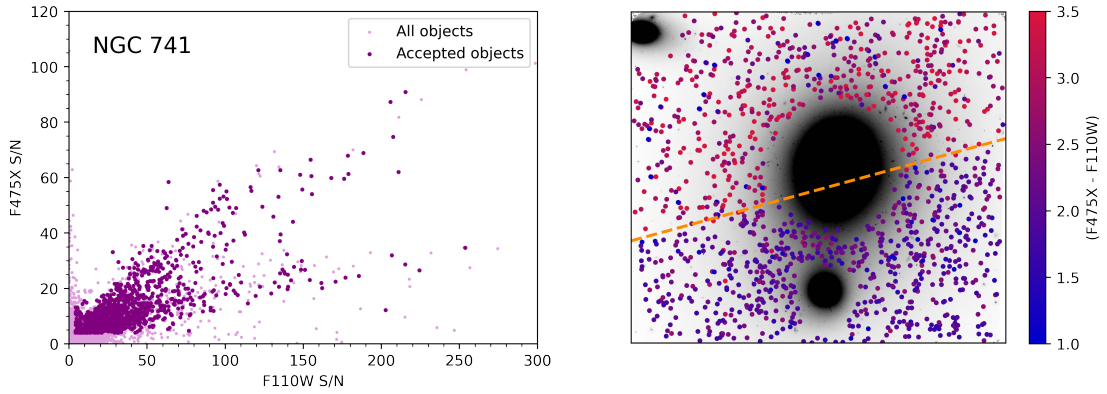


FIGURE 5.9: A problem with the NGC 741 observations in F475X. *Left:* S/N plot with two distinct populations, one with normal F475X S/N and another with lower F475X S/N. *Right:* image overlay with objects coloured by colour index. There is a distinct shift to redder colours from one F475X CCD chip to the other, denoted by the dashed orange line.

It should be noted before any full-sample analysis that there was a problem with the observations of NGC 741. One of the two F475X CCDs did not produce adequate results for this galaxy. Figure 5.9 illustrates the problem—on the left is a signal-to-noise plot which, unlike its normal counterpart from NGC 777 shown in Chapter 2, shows two distinct S/N distributions, and on the right is an overlay of objects on the F110W image with points coloured by colour index. Because




FIGURE 5.8: The final metallicity histogram for NGC 777 fitted with a double Gaussian curve, after the magnitude and radial cuts. Notice the scale compared to Figure 5.6.

of the abnormally low S/N values and the sharp contrast in colour at the junction of the two chips, I excluded all objects from the underperforming chip for NGC 741, which left $\sim 40\%$ of the GCs visible in both filters (the bluer points in the bottom of the right side of Figure 5.9). Data such as GC counts are therefore smaller for this galaxy than they are for others in the sample of fifteen, although proportional quantities such as red and blue GC fractions should not be affected. No amount of DOLPHOT parameter adjustment made a difference for NGC 741; the F475X images seem to have come off the telescope with the problem already present.

The double Gaussian parameters (amplitude, mean, width) for each peak, as well as full-sample metrics such as GCS mean metallicity, red and blue GC fractions, and separation between the two peaks (see Chapter 6 for more information on the latter two) provide a quantitative description of the metallicity distributions in this work. Further analysis of the data combined these metrics with a quantitative description of the environment of each host galaxy to focus on the original question for this research—is there a connection between GCS properties and host galaxy environment?

CHAPTER 6

ANALYSIS (OR THE CHAPTER OF A THOUSAND PLOTS)

6.1 Host Galaxy Environment Data

As mentioned in Chapter 2, the fifteen galaxies in this work came from the MASSIVE survey (Ma et al. 2014). MASSIVE targeted ~ 100 bright elliptical galaxies within about 110 Mpc. The survey imposed an absolute K-band magnitude limit of $M_K < -25.3$, which is effectively a stellar mass cutoff—the MASSIVE galaxies all have stellar masses of $M_* \gtrsim 10^{11.5} M_\odot$.

GCS properties are closely correlated with host galaxy stellar mass (Harris et al. 2013), but because the MASSIVE galaxies in this work have such similar stellar masses (see Table 6.1), those effects flatten out and we are able to see second order effects, such as those of environment. The fifteen galaxies in this work are also present in the Crook et al. 2007 group catalogues, which sort galaxies into groups based on relative distance and velocity. Crook and collaborators created a low-density-contrast (LDC) catalogue and a high-density-contrast (HDC) catalogue with respectively larger, less stringent and smaller, more stringent maximum relative distance and velocity between grouped galaxies. The LDC catalogue is perhaps more complete, while the HDC catalogue is less likely to include associations that are not actually gravitationally bound. Thirteen of the fifteen galaxies in this work appear in both catalogues; the remaining two, NGC 57 and NGC 4914, only appear in the LDC catalogue.

Both the MASSIVE survey and the group catalogues provided key galaxy and galaxy environment data for this work. MASSIVE (Ma et al. 2014, Veale et al. 2017, data in Table 6.1) measured the galaxies' stellar mass, necessary for ensuring that GCS-stellar mass relations would not drown out any signal from GCS-environment relations, and

TABLE 6.1: Data from the MAS-SIVE survey (Ma et al. 2014, Veale et al. 2017): absolute extinction-corrected K-band magnitudes and stellar mass.

Galaxy	M_K	M_* ($\log M_\odot$)
NGC 57	-25.75	11.79
NGC 410	-25.90	11.86
NGC 533	-26.05	11.92
NGC 741	-26.06	11.93
NGC 777	-25.94	11.87
NGC 1016	-26.33	12.05
NGC 1129	-26.14	11.96
NGC 1600	-25.99	11.90
NGC 2340	-25.90	11.86
NGC 3158	-26.28	12.02
NGC 3842	-25.91	11.86
NGC 4073	-26.33	12.05
NGC 4839	-25.85	11.83
NGC 4914	-25.72	11.78
NGC 7242	-26.34	12.05

the group catalogues (Crook et al. 2007) collected information on each galaxy’s neighbourhood. The group data (Table 6.2) include number of galaxies in the group, group virial mass, group virial radius, and group velocity dispersion, all important descriptors of the area around the galaxies in this work.

Missing from the Crook et al. 2007 catalogues is a measure of density—of the proximity of each galaxy to its neighbours. To fill that need, I made n th nearest neighbour measurements for each of the fifteen galaxies in this work. An n th nearest neighbour measurement is often made in projection and is essentially a surface density; it is a measurement of a circular area with radius defined as the distance between the central galaxy and that galaxy’s n th nearest neighbour. These measurements can be made with any integer number n , with different values of n conferring different benefits—higher values of n tend to be better for dense clusters, while lower values of n are better for groups, where an n too large would require leaving the group and would lead to an inaccurate picture of group density. In the literature, $n = 5$ is common and is considered a

good compromise value. In this work, I made two n th nearest neighbour measurements, one with $n = 5$, the standard compromise value, and the other with $n = 2$. Two of the fifteen galaxies in this work have fewer than 5 neighbours, with one only having 2; all galaxies therefore have an $n = 2$ measurement, and all except NGC 57 and NGC 4914 also have an $n = 5$ measurement.

To calculate n th nearest neighbour values, I used the coordinates from Crook et al. 2007, which has a limiting magnitude of $K = 11.25$, and the distances listed in Chapter 2 Table 2.2. I adopted the distance of the galaxy in this work for all galaxies in the corresponding group, for an n th nearest neighbour distance in projection. I then found the angular separation of the relevant galaxy and its fellow group members, converted those to physical distance, identified the 2nd and 5th nearest neighbours, and used those distances D_n to calculate surface densities Σ_n :

$$\Sigma_n = \frac{n}{\pi D_n^2} \tag{6.1}$$

The densities found with Equation 6.1 can be seen in Table 6.3 and Figure 6.1. The right panel of Figure 6.1 compares the $n = 2$ and $n = 5$ nearest neighbour calculations, with the black line indicating a one-to-one relation. In general, $n = 2$ produces higher




FIGURE 6.1: *Left*: completeness-corrected number of GCs vs. original GC counts. *Centre*: HDC group counts vs. LDC group counts, from Crook et al. 2007. *Right*: Σ_5 vs. Σ_2 , with the dashed line denoting a one-to-one relation. All positive correlations are expected, as each panel compares two different measurements of the same variable.

TABLE 6.2: Data from the Crook et al. 2007 group catalogue: group number, members in the group, total group virial mass, group virial radius, and group velocity dispersion. Data for NGC 57 and NGC 4914 from the LDC catalogue; all other data from the HDC catalogue.

Galaxy	Group number	Group members	Group virial mass (log M_\odot)	Group virial radius (Mpc)	Group velocity dispersion (km/s)
NGC 57	L11	4	14.092	2.85	198.8
NGC 410	H53	29	14.286	1.30	368.8
NGC 533	H74	3	13.769	1.13	217.9
NGC 741	H102	5	13.704	1.12	203.2
NGC 777	H109	7	13.434	0.89	167.2
NGC 1016	H159	8	13.940	0.92	294.1
NGC 1129	H185	33	14.730	1.44	583.5
NGC 1600	H294	16	14.193	0.96	384.2
NGC 2340	H426	18	14.086	1.71	254.8
NGC 3158	H570	6	13.019	0.75	112.9
NGC 3842	H672	42	14.701	1.22	614.5
NGC 4073	H692	10	14.024	1.07	300.3
NGC 4839	H745	49	14.880	2.00	588.0
NGC 4914	L929	3	12.943	1.42	75.1
NGC 7242	H1186	15	14.062	0.70	387.0

surface density values than $n = 5$, although as a set either value of n captures the overall trend—i.e. galaxies with a higher Σ_2 tend to also have a higher Σ_5 . Both values of n will appear in the analysis, with the caveat that $n = 5$ captures only thirteen of the fifteen galaxies in this work.

In addition to the comparison of the $n = 2$ and $n = 5$ nearest neighbour calculations,

I made several other test comparisons of environmental variables to ensure that there was no danger of environment-GCS relations being washed out by stronger correlations. Two of those tests appear alongside the nearest neighbour comparison in Figure 6.1. On the left is a comparison of completeness-corrected GC counts to the original pre-correction groups of GCs; the two counting methods display a tight linear correlation, as expected. There is also an (expected) correlation between the LDC and HDC group member galaxy counts, seen in the centre panel.

Figure 6.2 shows a battery of distance checks; from left to right and top to bottom, the figure shows LDC count, HDC count, completeness-corrected GCs, original GCs, group virial mass, group virial radius, group velocity dispersion, galaxy stellar mass, GCS mean metallicity, GCS red and blue modes, GCS red and blue sigmas, GCS red and blue fractions, two measures of the separation between peaks, Σ_2 , and Σ_5 versus distance. The only two of those variables that might be expected to be correlated with distance are the GC counts (first row, third and fourth columns), because of the radial distance cutoff described in Chapter 5—as a result of that cutoff, farther galaxies have a smaller area in their images in which GCs can be found, leading to lower GC counts. None of the other variables shows a correlation with distance.

Likewise, Figure 6.3 shows the same dependent variables checked against galaxy stellar mass (with the exception of row two column four, which shows distance versus galaxy stellar mass). This was a crucial precautionary measure, as GCS properties can be strongly correlated with stellar mass (see Chapter 1 for details). The most important items in Figure 6.3 are those involving GCS properties—all of row three, plus row four columns one and two. Those panels show no relation between fitted double Gaussian parameters and other GCS properties and galaxy stellar mass. The galaxy selection criteria described in Chapter 2 successfully controlled for stellar mass and its correlations with GCS properties, allowing for investigation of second-order variables such as environment.

Finally, Figure 6.4 comprises galaxy environmental relations and again checks for expected correlations and non-correlations. Column one shows correlations between

TABLE 6.3: 2nd and 5th nearest neighbour measurements, calculated with coordinates from Crook et al. 2007 and distances from Chapter 2. NGC 57 and NGC 4914 have fewer than five neighbours, so they only have a Σ_2 measurement.

Galaxy	Σ_2 (Mpc ⁻²)	Σ_5 (Mpc ⁻²)
NGC 57	0.2453	N/A
NGC 410	62.41	5.000
NGC 533	6.416	0.6980
NGC 741	1.924	0.7893
NGC 777	1.891	0.4923
NGC 1016	10.69	11.38
NGC 1129	475.5	23.74
NGC 1600	267.7	7.348
NGC 2340	22.56	0.4929
NGC 3158	16.00	0.8733
NGC 3842	100.1	95.05
NGC 4073	7.943	7.631
NGC 4839	43.77	1.500
NGC 4914	0.3473	N/A
NGC 7242	98.56	56.66

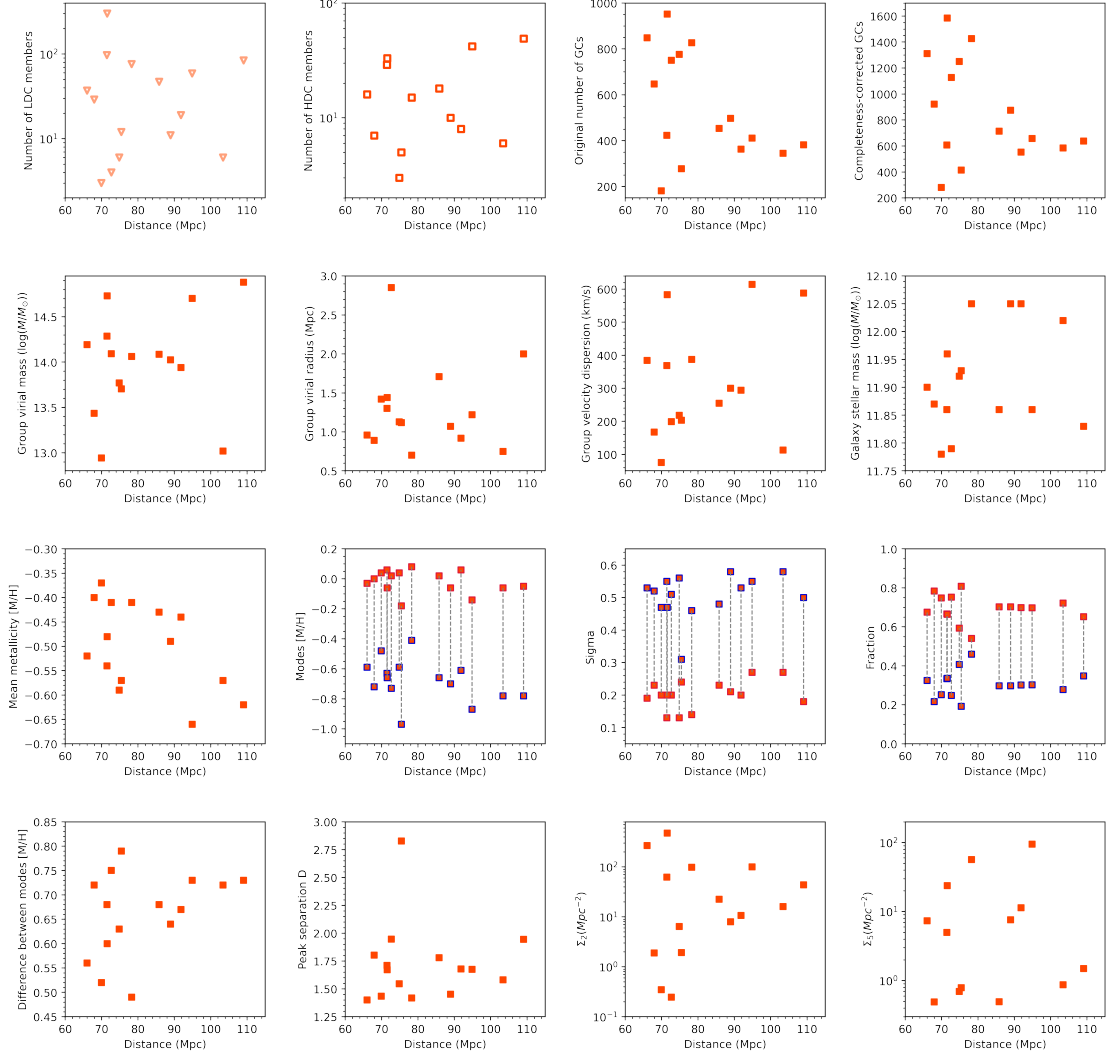


FIGURE 6.2: Tests to ensure that no variables are correlated with distance. LDC group member counts are marked with light-coloured open triangles (first row, first column), HDC group member counts are marked with dark-coloured open squares (first row, second column), blue double Gaussian peak fitting parameters are marked with a blue outline (third row), and red peak fitting parameters are marked with a red outline (also third row). Environmental data from Crook et al. 2007; stellar mass data from Veale et al. 2017.

number of group members and group virial mass, velocity dispersion, and neighbourhood density Σ_2 (all expected), and a non-correlation between virial radius and number of group members (also expected). The same is true of the remaining columns—all environmental variables show positive correlations with each other, with the exception of group virial radius. That behaviour is expected; groups with more members tend to be

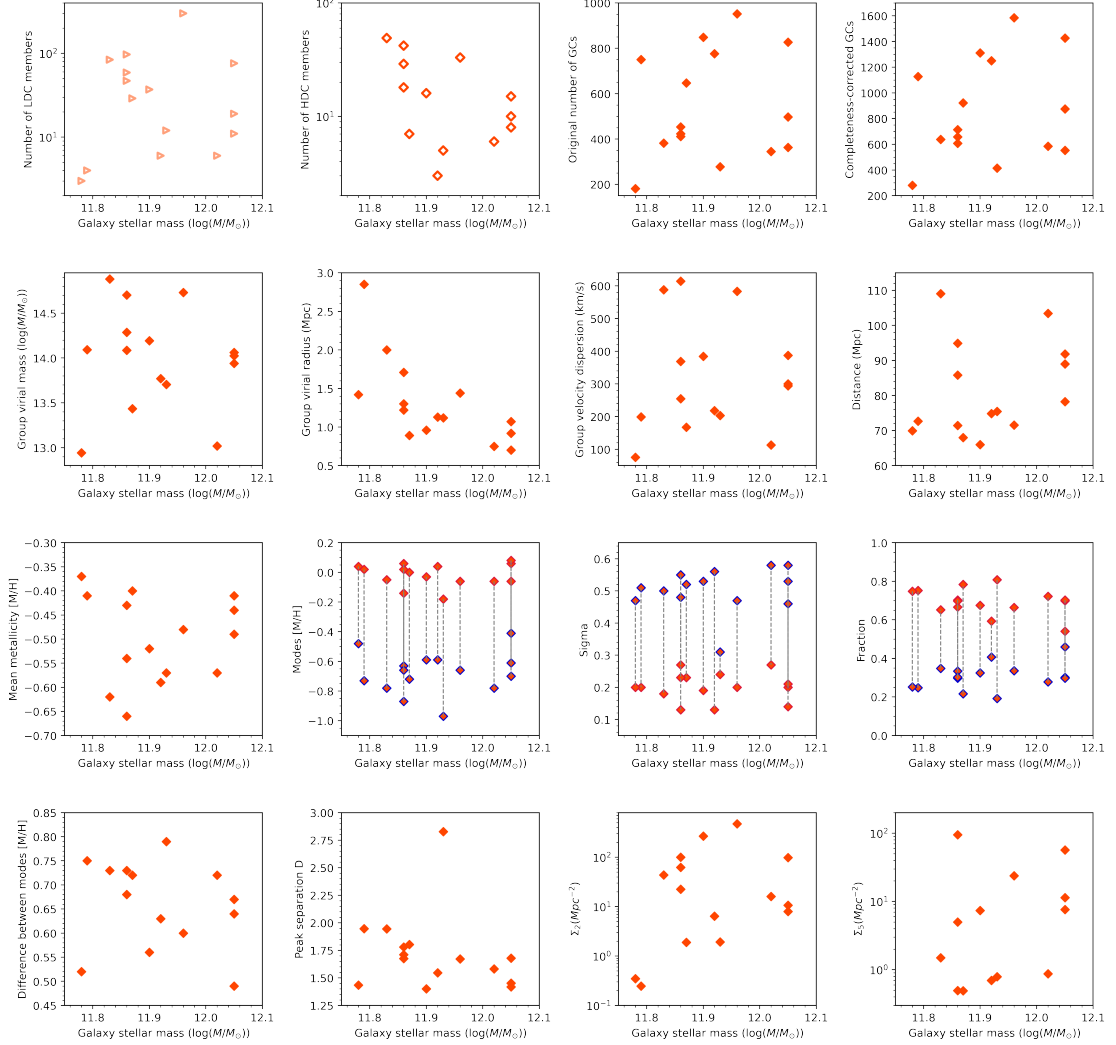


FIGURE 6.3: Tests to ensure that no variables are correlated with galaxy stellar mass. LDC group member counts are marked with light-coloured open triangles (first row, first column), HDC group member counts are marked with dark-coloured open diamonds (first row, second column), blue double Gaussian peak fitting parameters are marked with a blue outline (third row), and red peak fitting parameters are marked with a red outline (also third row). Environmental data from Crook et al. 2007; stellar mass data from Veale et al. 2017.

denser and more massive, with higher velocity dispersions.

With the assurance that no first-order correlations would drown out the signal from potential GCS-environment relations, I could start making comparisons between the Crook et al. 2007 and n th-nearest neighbour data and my GCS data. The following section will go into detail on the fitted double Gaussian parameters discussed in Chapter




FIGURE 6.4: Tests to ensure that environmental variables show the expected correlations and non-correlations. LDC group member counts are marked with light-coloured open triangles (first column) and HDC group member counts are marked with dark-coloured open circles (also first column). Every relation not involving virial radius is expected to show a positive correlation. Data from Crook et al. 2007.

5, and make comparisons between those and the environmental data discussed in this chapter (i.e. number of group members, group virial mass, group virial radius, group velocity dispersion, and the two nearest neighbour metrics).

6.2 GCS-Environment Comparisons

Throughout this work, we have been looking to NGC 777 as an example of each data analysis step. NGC 777’s GCS can be seen once again in Figure 6.5 (row 2, column 2), along with the fitted metallicity histograms from the other fourteen galaxies in the sample for this work. The GCSs have the same overall form, with a compact metal-rich (red) peak and an extended metal-poor (blue) tail, but with visually obvious differences in shape, which the fitted double Gaussian parameters quantify.

The red peaks are quite uniform across all fifteen galaxies in this work, having a mode around $[M/H] \sim 0.00$ and a characteristic width between $\sigma_r \sim 0.10$ and $\sigma_r \sim 0.30$ (recall Equation 5.1). Most differences between GCSs arise with the blue peak, which can lie anywhere from ~ 0.5 dex to ~ 0.8 dex away from the red peak—see Figure 6.5 and compare NGC 2340 to NGC 4839, for example. The blue peak can also vary widely in amplitude relative to the red peak—see Figure 6.5 and compare NGC 57 and NGC 7242. The relative sizes and shapes of the two peaks (characterized by the double Gaussian parameters listed in Table 6.4) affect what portion of the GC population is classified as blue and metal-poor versus red and metal-rich. Table 6.5 compiles GC counts (total, red, and blue), plus the blue fraction for each GCS. The blue fraction, f_b , always between 0 and 1, is defined as the ratio of blue GCs to total GCs; conversely, the red fraction is defined as $f_r = 1 - f_b$.

TABLE 6.4: Parameters from double Gaussian fitting, with one-sigma errors: blue amplitude, mean, and width; and red amplitude, mean, and width.

Galaxy	Blue amplitude	Red amplitude	Blue mode	Red mode	Blue sigma	Red sigma
NGC 57	68.6 ± 4.6	161.9 ± 14.2	-0.73 ± 0.10	0.02 ± 0.01	0.51 ± 0.08	0.20 ± 0.02
NGC 410	48.3 ± 5.0	65.2 ± 12.4	-0.63 ± 0.11	0.06 ± 0.02	0.55 ± 0.09	0.13 ± 0.03
NGC 533	106.1 ± 6.2	103.8 ± 14.9	-0.59 ± 0.06	0.04 ± 0.02	0.56 ± 0.05	0.13 ± 0.02
NGC 741	32.4 ± 4.8	57.2 ± 5.9	-0.97 ± 0.09	-0.18 ± 0.05	0.31 ± 0.09	0.24 ± 0.04
NGC 777	50.8 ± 3.3	122.5 ± 11.4	-0.72 ± 0.11	0.00 ± 0.01	0.52 ± 0.08	0.23 ± 0.02
NGC 1016	38.9 ± 3.8	59.5 ± 11.4	-0.61 ± 0.14	0.06 ± 0.02	0.53 ± 0.10	0.20 ± 0.04
NGC 1129	128.6 ± 9.3	161.0 ± 29.1	-0.66 ± 0.10	-0.06 ± 0.02	0.47 ± 0.07	0.20 ± 0.03
NGC 1600	97.4 ± 7.6	124.4 ± 16.5	-0.59 ± 0.08	-0.03 ± 0.02	0.53 ± 0.05	0.19 ± 0.03
NGC 2340	52.3 ± 6.2	74.3 ± 21.2	-0.66 ± 0.19	0.02 ± 0.03	0.48 ± 0.13	0.23 ± 0.05
NGC 3158	31.0 ± 4.2	47.9 ± 12.1	-0.78 ± 0.21	-0.06 ± 0.02	0.58 ± 0.13	0.27 ± 0.05
NGC 3842	42.9 ± 8.0	54.6 ± 25.2	-0.87 ± 0.31	-0.14 ± 0.05	0.55 ± 0.19	0.27 ± 0.08
NGC 4073	47.5 ± 6.2	114.5 ± 14.1	-0.70 ± 0.16	-0.06 ± 0.02	0.58 ± 0.10	0.21 ± 0.03
NGC 4839	49.9 ± 2.7	64.0 ± 7.8	-0.78 ± 0.07	-0.05 ± 0.02	0.50 ± 0.06	0.18 ± 0.02
NGC 4914	18.3 ± 4.0	38.8 ± 9.3	-0.48 ± 0.22	0.04 ± 0.02	0.47 ± 0.12	0.20 ± 0.04
NGC 7242	153.9 ± 7.5	90.8 ± 16.2	-0.41 ± 0.04	0.08 ± 0.02	0.46 ± 0.03	0.14 ± 0.03

In addition to the six double Gaussian parameters, I used the fifteen metallicity histograms in Figure 6.5 to calculate each GCS’s mean metallicity, the difference between the two double Gaussian modes (μ_b and μ_r in Equation 5.1), and another measure of peak separation:




FIGURE 6.5: Fitted metallicity histograms for all fifteen galaxies in this work, with a constant y scale. In general, farther galaxies tend to have fewer GCs, with three exceptions—NGC 741, for which only about half of the image was usable, and NGC 410 and NGC 4914, which seem to have fewer GCs than the other galaxies.

$$D = \frac{|\mu_b - \mu_r|}{\sqrt{\frac{\sigma_b^2 + \sigma_r^2}{2}}} \quad (6.2)$$

As with Equation 5.1, μ_b and μ_r are the blue and red modes, and σ_b and σ_r are the blue and red widths, respectively. Unlike a simple difference between modes, Equation 6.2 takes into account the width of each peak; a double Gaussian with wider peaks will have a smaller D value than one with the same modes and narrower peaks.

In general, very few combinations of GCS and environmental variables showed any sign of a correlation. Figure 6.6, showing completeness-corrected GC totals versus environmental parameters, is one of two exceptions. Visually, there seems to be a weak but consistent correlation between the number of GCs and the number of galaxies per

TABLE 6.5: Counts of total GCs (both as counted originally and with the completeness correction) and red and blue GC subpopulations (with the completeness correction), and GC blue fraction and peak separation D (see Equation 6.2.)

Galaxy	Total GCs (original)	Total GCs (complete)	Red GCs	Blue GCs	Blue fraction	D
NGC 57	750	1127.8	848.4	279.4	0.248	1.947
NGC 410	423	607.2	404.1	203.4	0.334	1.711
NGC 533	776	1250.3	740.9	509.3	0.407	1.545
NGC 741	278	414.4	335.0	79.4	0.192	2.828
NGC 777	647	921.9	722.6	199.4	0.216	1.803
NGC 1016	363	552.9	386.8	166.2	0.301	1.678
NGC 1129	952	1585.0	1053.3	531.7	0.336	1.671
NGC 1600	848	1310.7	884.2	426.5	0.325	1.400
NGC 2340	453	713.2	501.2	212.0	0.297	1.780
NGC 3158	345	584.7	421.0	163.7	0.278	1.581
NGC 3842	411	658.1	459.6	198.5	0.302	1.675
NGC 4073	497	874.7	615.2	259.6	0.297	1.451
NGC 4839	382	638.0	416.1	221.9	0.348	1.944
NGC 4914	181	280.9	210.0	70.9	0.252	1.434
NGC 7242	827	1426.9	772.6	654.3	0.459	1.419

group (leftmost panel, with triangles denoting LDC group counts and circles HDC group counts), the group virial mass (second panel from the left), and the n th-nearest neighbour surface density (rightmost panel, with triangles denoting the $n = 2$ metric and circles the $n = 5$ metric). For the group virial radius and velocity dispersion, however, there is no obvious relation.

Most GCS property metrics do not show a relationship with any of the environmental




FIGURE 6.6: Number of GCs in each GCS vs. environmental parameters. From left to right: number of galaxies in the group (LDC count in light purple triangles, HDC count in dark purple circles), group virial mass, group virial radius, group velocity dispersion, and n th-nearest neighbour measurement ($n = 2$ in medium purple triangles, $n = 5$ in dark purple circles).

variables in this work, including mean GCS metallicity, in Figure 6.7. The weak correlations seen in Figure 6.6 are not present in Figure 6.7—there is no indication of any ties between mean GCS metallicity and environment.

The same is true of both peak separation metrics—the difference between modes (i.e. $|\mu_b - \mu_r|$), shown in Figure 6.8, and D (Equation 6.2), shown in Figure 6.9. Neither GCS metric appears to be correlated with any of the environmental variables.




FIGURE 6.7: Mean GCS metallicity vs. environmental parameters. From left to right: number of galaxies in the group (LDC count in light purple triangles, HDC count in dark purple circles), group virial mass, group virial radius, group velocity dispersion, and n th-nearest neighbour measurement ($n = 2$ in medium purple triangles, $n = 5$ in dark purple circles).

(The outlier point in Figure 6.9 with a peak separation of $D \sim 2.8$ comes from NGC 741, the galaxy for which we were only able to keep photometry from half of the image—recall Figure 5.9.)

The story is the same for the double Gaussian parameters from Equation 5.1 and Table 6.4—i.e. the parameters associated with each individual peak in the metallicity distribution. Figure 6.10 shows double Gaussian amplitude pairs plotted against the same five environmental variables appearing in the previous four figures, with group virial




FIGURE 6.8: Difference between GCS blue and red mode vs. environmental parameters. From left to right: number of galaxies in the group (LDC count in light purple triangles, HDC count in dark purple circles), group virial mass, group virial radius, group velocity dispersion, and n th-nearest neighbour measurement ($n = 2$ in medium purple triangles, $n = 5$ in dark purple circles).

mass, virial radius, and velocity dispersion appearing in the top figure, and number of group members and n th-nearest neighbour surface densities in the figure on the bottom left. Red and blue amplitudes (shown as red and blue points, respectively) from the same system are connected with a thin dashed line in the top plot, and environmental measurements corresponding to the same amplitude are connected in the bottom plot; additionally, because of the slight correlation between the number of GCs per system and




FIGURE 6.9: Peak separation D vs. environmental parameters. From left to right: number of galaxies in the group (LDC count in light purple triangles, HDC count in dark purple circles), group virial mass, group virial radius, group velocity dispersion, and n th-nearest neighbour measurement ($n = 2$ in medium purple triangles, $n = 5$ in dark purple circles).

a few environmental variables, I normalized the amplitudes by GC count for each GCS. With the normalization, there is no indication of any relation between subpopulation peak amplitude and environmental metrics.

Looking at the mode of each peak (Figure 6.11) produces the same result; there is no relation between position of either of the peaks in metallicity space and host galaxy environment. Outside of its primary purpose, Figure 6.11 provides verification of the qualitative observations from Figure 6.5—the red, metal-rich peaks as a group are more

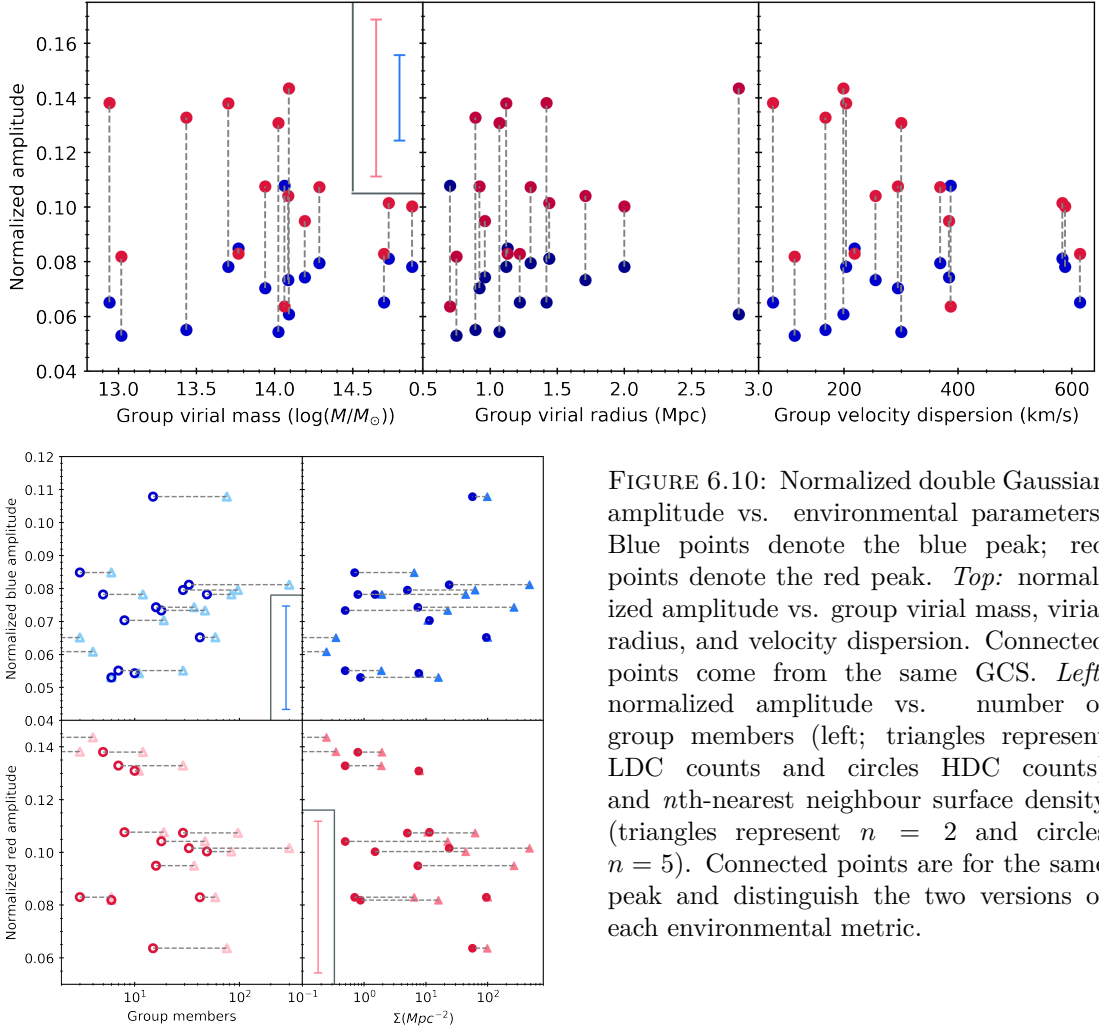


FIGURE 6.10: Normalized double Gaussian amplitude vs. environmental parameters. Blue points denote the blue peak; red points denote the red peak. *Top*: normalized amplitude vs. group virial mass, virial radius, and velocity dispersion. Connected points come from the same GCS. *Left*: normalized amplitude vs. number of group members (left; triangles represent LDC counts and circles HDC counts) and n th-nearest neighbour surface density (triangles represent $n = 2$ and circles $n = 5$). Connected points are for the same peak and distinguish the two versions of each environmental metric.

uniform than their blue, metal-poor counterparts, which display more variation in the collection of histograms and more scatter in metallicity space here. It should be noted that both peaks are ~ 0.5 dex more metal-rich than expected for a massive elliptical galaxy GCS; I will address this in more detail in Chapter 7.

As with the other double Gaussian parameters, σ_b and σ_r , the blue and red peak widths, respectively, are not correlated with any environmental variables. Figure 6.12 also confirms visual observations of a sharper (narrower) red peak and a longer (wider) blue tail. NGC 741 is once again visible among the σ_b values in particular, at $\sigma_b \sim 0.3$, probably because of the loss of half of the image.

The other exception to the trend of non-correlations is the blue fraction f_b , seen in Figure 6.13. Like the completeness-corrected GC count, f_b shows weak but consistent correlations with some of the environmental variables; galaxy group virial radius (middle

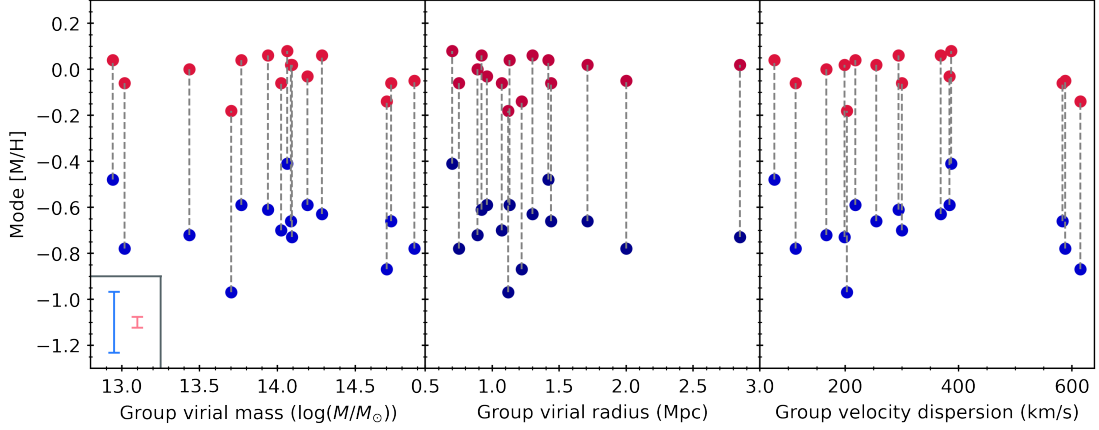



FIGURE 6.11: Double Gaussian mode vs. environmental parameters. Blue points denote the blue peak; red points denote the red peak. *Top*: mode vs. group virial mass, virial radius, and velocity dispersion. Connected points come from the same GCS. *Right*: amplitude vs. number of group members (left; triangles represent LDC counts and circles HDC counts) and n th-nearest neighbour surface density (triangles represent $n = 2$ and circles $n = 5$). Connected points are for the same peak and distinguish the two versions of each environmental metric.



panel) is the only environmental metric without a hint of a relationship with f_b , and the strongest correlation appears to be with n th-nearest neighbour surface density (rightmost panel), for both $n = 2$ and $n = 5$ (triangles and circles, respectively).

In order to test the strength of the correlations, I fitted simple linear models to the set of comparisons in Figure 6.6 and Figure 6.13 and calculated the corresponding R^2 values. The results are plotted as solid lines with one-sigma error shaded regions in Figure 6.14 and Figure 6.15. The number of GCs and the GCS blue fraction exhibit similar behaviour for each environmental variable—the relationship with galaxy group virial radius is almost flat, and all other correlations are positive, with n th-nearest neighbour being the strongest and the others being somewhat weaker. ("Strong" is relative here; see Figures 6.14 and 6.15 for R^2 values.) For the number of completeness-corrected GCs $N_{GCs,comp}$ versus Σ_2 , the linear fitting produced Equation 6.3 with $R^2 = 0.19$, and for the blue fraction f_b versus Σ_2 , the fitting produced Equation 6.4 with $R^2 = 0.36$.

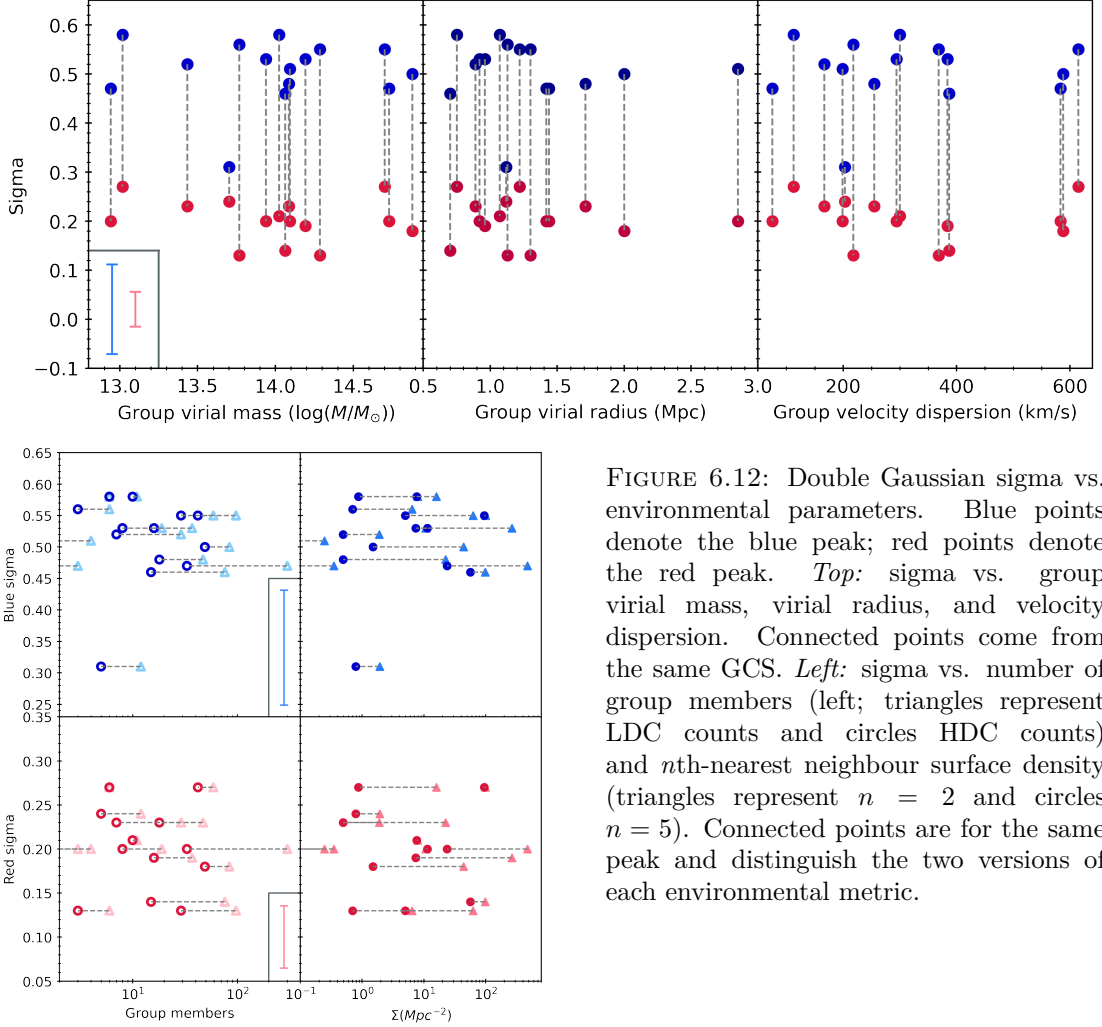


FIGURE 6.12: Double Gaussian sigma vs. environmental parameters. Blue points denote the blue peak; red points denote the red peak. *Top*: sigma vs. group virial mass, virial radius, and velocity dispersion. Connected points come from the same GCS. *Left*: sigma vs. number of group members (left; triangles represent LDC counts and circles HDC counts) and n th-nearest neighbour surface density (triangles represent $n = 2$ and circles $n = 5$). Connected points are for the same peak and distinguish the two versions of each environmental metric.

$$\langle N_{GCs,comp} \rangle = 175 * \log(\Sigma_2) + 661 \quad (6.3)$$

$$\langle f_b \rangle = 0.042 * \log(\Sigma_2) + 0.258 \quad (6.4)$$

The lines shown in Figures 6.14 and 6.15 were fitted using Σ_2 ; fitting the same linear model with Σ_5 produces similar results with slightly larger one-sigma errors and a smaller R^2 value, as seen in Figure 6.16. The overlays in the rightmost panels of Figure 6.16 compare the two fits directly; for both $N_{GCs,comp}$ and f_b , the zeropoint for $n = 5$ is larger, so the fitted line appears above the line for $n = 2$ in the plots.

Regardless of which value of n is used in fitting, the n th-nearest neighbour correlation, strongest among those for the environmental variables in this work, is weak in absolute




FIGURE 6.13: GCS blue fraction vs. environmental parameters. From left to right: number of galaxies in the group (LDC count in light blue triangles, HDC count in dark blue circles), group virial mass, group virial radius, group velocity dispersion, and n th-nearest neighbour measurement ($n = 2$ in medium blue triangles, $n = 5$ in dark blue circles).

terms. The f_b - Σ_2 relation in particular, however, is visually consistent (see the bottom left panel of Figure 6.16) and warrants more investigation.




FIGURE 6.14: Number of GCs in each GCS vs. environmental parameters (see Figure 6.6), with a linear model fitted. The GC count has a weak but consistent correlation with number of galaxies in the group and with n th-nearest neighbour surface density. The group virial mass panel also appears to have a weak but consistent correlation, although the uncertainty associated with that fit is large in comparison to the other fits. The purple line shows the linear model, and the shaded region shows one-sigma uncertainty.




FIGURE 6.15: GCS blue fraction vs. environmental parameters (see Figure 6.13), with a linear model fitted. The blue fraction has a weak but consistent correlation with number of galaxies in the group and with n th-nearest neighbour surface density. The blue line shows the linear model, and the shaded region shows one-sigma uncertainty.

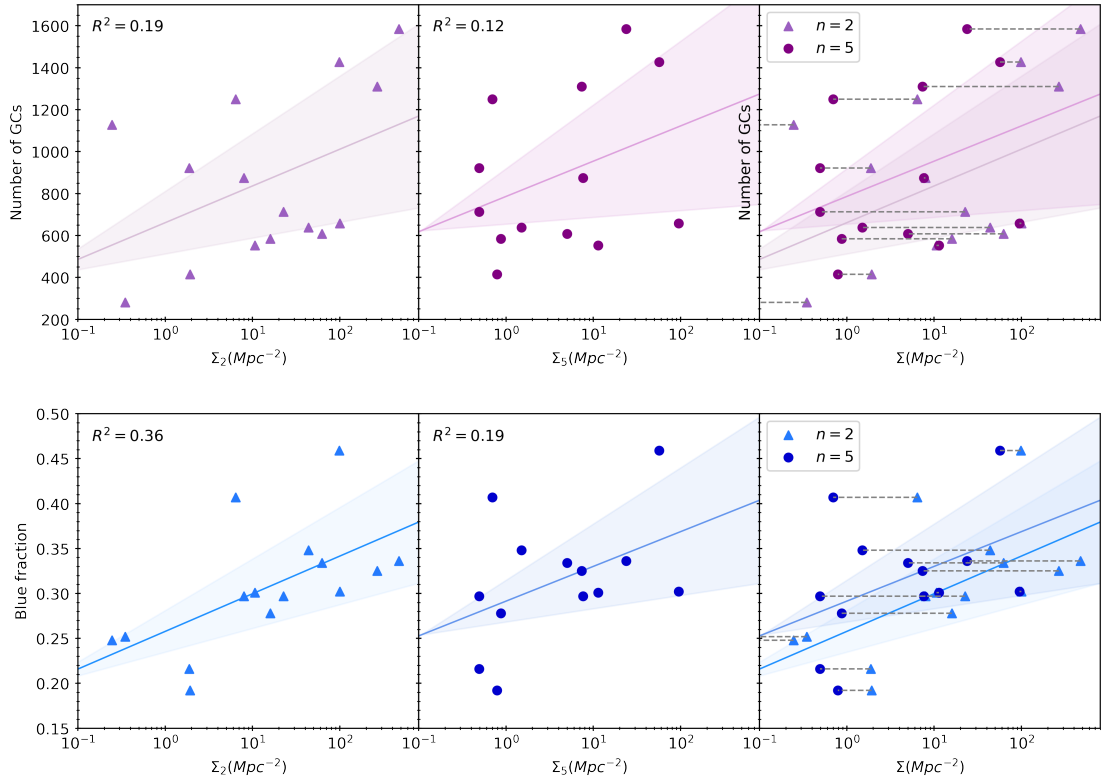


FIGURE 6.16: A comparison of the linear fits for Σ_2 and Σ_5 . Solid lines are the linear fits, shaded areas are the associated one-sigma errors, and dotted lines connect the $n = 2$ and $n = 5$ measurements for each GCS. *Top*: number of GCs vs. n th-nearest neighbour surface density. *Bottom*: blue fraction vs. n th-nearest neighbour surface density. From left to right: $n = 2$, $n = 5$, and both values of n overlaid. In both of the rightmost panels, the fit for $n = 5$ has a larger zeropoint and appears above that for $n = 2$.

7.1 GCS Properties by Host Galaxy Environment?

This study aimed to investigate the GCS property-galaxy environment parameter space, which had not been explored, and to look for any existing GCS-environment correlations. In this work, I performed photometry on *HST* images of fifteen BCGs of approximately the same mass, obtaining a (F475X – F110W) colour index for each GC detected. I developed and applied a method for correcting for completeness based on both object magnitude and background brightness level, and then used Padova stellar models to convert my colour index to the (F475W – F850LP) index, which has a direct spectroscopic metallicity calibration. I then used the metallicity values to determine the MDF for each galaxy and fit it with a double Gaussian to model the blue, metal-poor and red, metal-rich subpopulations. Finally, I compared GCS properties such as mean metallicity, blue fraction, and the fitted double Gaussian parameters to galaxy environment parameters such as number of galaxy group members, group virial mass, and n th-nearest neighbour surface density.

For almost all combinations of GCS and environmental variables, I found no correlation. The two exceptions to this pattern are the completeness-corrected number of GCs per system $N_{GCs,comp}$ and the blue fraction f_b , which are both weakly but consistently positively correlated with n th-nearest neighbour surface density Σ_n for both $n = 2$ and $n = 5$.

7.2 Conclusions

When we began this work, we naively expected to see a correlation between host galaxy environmental metrics such as number of galaxy group members, group virial mass, and group velocity dispersion, for which higher values correspond to denser galaxy groups,

and a higher ratio of blue GCs to red GCs in each galaxy’s GCS, indicated by e.g. a high blue fraction, low mean metallicity, high blue peak amplitude, etc. Rich galaxy groups have more metal-poor dwarf galaxies with metal-poor GCs per BCG available to accrete, so we expected BCGs in rich groups to have accreted a higher proportion of blue, metal-poor GCs. In reality, only two GCS variables (number of GCs and blue fraction) and one environmental variable (n th-nearest neighbour surface density) appear to be correlated. We find weak but consistent positive correlations between both GCS variables and Σ_n , with the strongest being f_b versus Σ_2 . These correlations match our initial expectations, although they are weaker than we anticipated.

There is more than one potential interpretation of the results in this work. One possibility is that environment affects GCS properties within a smaller sphere of influence than we expected. Another possibility, though, is that GCS metallicity properties depend on environment in a more complex way than we are able to tease apart in this work. We might also consider time—environmentally driven differences may emerge as BCGs evolve beyond the present day.

7.3 Future Work

In the short term, the procedure described in Chapter 3 could be refined. Although the completeness correction outlined in this work is more precise than imposing a blanket cutoff magnitude m_0 across an entire image without accounting for background brightness, it still relies on a twelve-zone binning and fitting procedure, which approximates sky background levels. For a more exact completeness correction, we plan on adapting the statistical procedures explained in Eadie et al. 2022 and applying logistic regression to the question of completeness. Eadie et al. 2022 used logistic regression to model the probability of whether a galaxy had at least one GC as a function of the galaxy’s mass; like the presence or lack of GCs, the detection or non-detection of an object is a binary (i.e. yes-or-no) variable, and the transition region between almost no objects being detected and almost all objects being detected can be modelled with a sigmoid curve. In Equation 3.1, the only variable in the exponential is the difference between object magnitude m and cutoff magnitude m_0 —but background brightness can be added to the exponential in order to account for it as well (see Equation 6 in Eadie et al. 2022 for a generalized form). Work on this refinement has already begun.

In the longer term, this work would benefit from an improved colour-to-metallicity conversion procedure. It is unrealistic to expect spectroscopic calibration for every photometric colour index (spectroscopy of point-source GCs is very observationally expensive), but by making direct photometric comparisons in several commonly used filters for a single dataset of GCs, we can derive conversions from typical *HST* colour indices to the $(g - z)$ index (or $F475W - F850LP$) in *HST* terms), which already has a spectroscopic calibration. Having those observationally grounded colour-to-colour conversions would allow us to more confidently set the zeropoints of GC colour-metallicity conversions. Our *HST* Cycle 30 proposal has been accepted, and we will image NGC 4874, the Coma Cluster BCG, in the ACS filters F435W, F555W, F606W, and F850LP, and in the WFC3




FIGURE 7.1: Figure 6 from Choksi and Gnedin 2019 (*Monthly Notices of the Royal Astronomical Society*), showing GC accreted fraction vs. stellar mass for blue (blue line) and red (red line) GCs. As galaxy mass increases, so does the fraction of both blue and red GCs that came from accreted satellites. Blue GCs become majority-accreted first, but eventually even red GCs, which require a massive host halo to form, come mostly from accreted satellites.

filters F438W and F475X. Additionally, the Hubble Legacy Archive includes images of this galaxy in F475W, F814W, F110W, and F160W with exposure times comparable to those suggested in the proposal. With all of this material, we will create new metallicity calibrations for all of the commonly used colour indices.

Following up on this work, we would like to conduct a more thorough review of galaxy growth simulations, with a focus on varied galaxy neighbourhoods and the precise contents and whereabouts of any accreted matter; this would give us a theoretical look

at the validity of the idea of galaxies in rich clusters having a higher proportion of metal-poor dwarfs to accrete. We should also consider studying the outer haloes of BCGs, rather than restricting our dataset to the middle halo as we did here—it is possible that accreted material, and subsequently differences in accreted material, resides predominantly in the outskirts of these galaxies.

Finally, farther afield, the analysis techniques in this work can be used in other projects. My colleagues and I would like to fill in the blue fraction-galaxy mass parameter space with all the data available for lower-mass systems, in order to track the correlation for Milky Way- and lower-mass galaxies. Blue fraction-galaxy mass is a three-part relation, with low-mass dwarf galaxies forming mostly blue GCs, the red fraction growing with galaxy mass up past Milky Way mass, blue fraction increasing again as host galaxies bring in significant numbers of blue GCs through merging with dwarfs, and finally a second increase in red fraction at very high mass where the merging galaxies are massive enough to form red GCs (see Figure 7.1). The techniques used in this work, along with the upcoming colour-colour calibrations, can be applied to lower-mass galaxies in a blue fraction-galaxy mass project, and to any work where analysis of the GCS MDF is needed.

APPENDIX A

SAMPLE CALCKY SCRIPT

calcky	icv411010_drz.chip1	15	35	-128	2.25	2.00
calcky	icv411vuq_flt.chip1	15	35	-128	2.25	2.00
calcky	icv411vvq_flt.chip1	15	35	-128	2.25	2.00
calcky	icv411vxq_flt.chip1	15	35	-128	2.25	2.00
calcky	icv411vzq_flt.chip1	15	35	-128	2.25	2.00
calcky	icv411w1q_flt.chip1	15	35	-128	2.25	2.00
calcky	idfr05bbq_flc.chip1	15	35	-128	2.25	2.00
calcky	idrf05bbq_flc.chip2	15	35	-128	2.25	2.00
calcky	idfr05beq_flc.chip1	15	35	-128	2.25	2.00
calcky	idfr05beq_flc.chip2	15	35	-128	2.25	2.00
calcky	idfr05biq_flc.chip1	15	35	-128	2.25	2.00
calcky	idfr05biq_flc.chip2	15	35	-128	2.25	2.00

APPENDIX B

SAMPLE DOLPHOT PARAMETER FILE

```
Nimg = 11          #number of images (int)
img0_file = icv411010_drz.chip1
img1_file = icv411vuqflt.chip1
img0_file = icv411vvqflt.chip1
img0_file = icv411vxqflt.chip1
img0_file = icv411vzqflt.chip1
img0_file = icv411w1qflt.chip1
img0_file = idfr05bbqflc.chip1
img0_file = idfr05beqflc.chip1
img0_file = idfr05biqflc.chip1
img0_file = idfr05bbqflc.chip2
img0_file = idfr05beqflc.chip2
img0_file = idfr05biqflc.chip2
#
# The following parameters can be specified for individual images (img1_...)
# or applied to all images (img_...)
img_shift = 0 0    #shift relative to reference
img_xform = 1 0 0  #scale, distortion, and rotation
```

```
img_PSFa = 3 0 0 0 0 0          #PSF XX term (flt)
img_PSFb = 3 0 0 0 0 0          #PSF YY term (flt)
img_PSFc = 0 0 0 0 0 0          #PSF XY term (flt)
img_RAper = 5.0                 #photometry aperture size (flt)
img_RChi = -1                   #Aperture for determining centroiding (flt); if <=0 use
RAper
img_RSky = 10 15                #radii defining sky annulus (flt>=RAper+0.5)
img_RPSF = 10                   #PSF size (int>0)
img_aprad = 10                  #radius for aperture correction
img_apsky = 25 30               #sky annulus for aperture correction
#
# The following photometers affect the finding and measurement of stars
photsec =                        #section: group, chip, (X,Y)0, (X,Y)1
RCentroid = 1                   #centroid box size (int>0)
SigFind = 3.0                   #sigma detection threshold (flt)
SigFindMult = 0.85              #Multiple for quick-and-dirty photometry (flt>0)
SigFinal = 3.5                  #sigma output threshold (flt)
MaxIT = 10                      #maximum iterations (int>0)
FPSF = Lorentz                  #PSF function (str/Gauss,Lorentz,Lorentz2,G + L)
PSFPhot = 1                     #photometry type (int/0=aper,1=psf,2=wtd-psf)
PSFPhotIt = 1                   #number of iterations in PSF-fitting photometry (int>=0)
FitSky = 2                      #fit sky? (int/0=no,1=yes,2=small,3=with-phot)
SkipSky = 1                     #spacing for sky measurement (int>0)
SkySig = 2.25                   #sigma clipping for sky (flt>=1)
NegSky = 1                      #allow negative sky values? (0=no,1=yes)
NoiseMult = 0.10                #noise multiple in imgadd (flt)
FSat = 0.999                    #fraction of saturate limit (flt)
Zero = 25.0                     #zeropoint for 1 DN/s (flt)
PosStep = 0.25                  #search step for position iterations (flt)
dPosMax = 3.0                   #maximum single-step in position iterations (flt)
```

```
RCombine = 1.5      #minimum separation for two stars for cleaning (flt)
SigPSF = 5.0       #min S/N for psf parameter fits (flt)
PSFStep = 0.25     #stepsize for PSF
MinS = 1.0         #minimum FWHM for good star (flt)
MaxS = 9.0         #maximum FWHM for good star (flt)
MaxE = 0.5         #maximum ellipticity for good star (flt)
#
# Settings to enable/disable features
UseWCS = 1         #use WCS info in alignment (int 0=no, 1=shift/rotate/scale,
2=full)
Align = 3          #align images? (int 0=no,1=const,2=lin,3=cube)
AlignIter = 2      #number of iterations on alignment? (int>0)
AlignTol = 0       #number of pixels to search in preliminary alignment
(flt>=0)
AlignStep = 1      #stepsize for preliminary alignment search (flt>0)
AlignOnly = 0      #exit after alignment
Rotate = 1         #allow cross terms in alignment? (int 0=no,1=yes)
SubResRef = 1      #subpixel resolution for reference image (int>0)
SecondPass = 1     #second pass finding stars (int 0=no,1=yes)
SearchMode = 1     #algorithm for astrometry (0=max SNR.chi, 1=max SNR
Force1 = 0         #force type 1/2 (stars)? (int 0=no,1=yes)
EPSF = 1           #allow elliptical PSFs in parameter fits (int 0=no,1=yes)
PSFsol = 1         #Analytic PSF solution (int -1=none, 0=con, 1=lin, 2=quad)
PSFres = 1         #make PSF residual image? (int 0=no,1=yes)
psfstars =         #Coordinates of PSF stars
psfoff = 0.0       #coordinate offset (PSF system - dolphot system)
ApCor = 1          #find/make aperture corrections? (int 0=no,1=yes)
SubPixel = 1       #subpixel PSF calculation (int>0)
FakeStars =        #file with fake star input data
FakeOut =          #file with fake star output data (default=phot.fake)
```

```
FakeMatch = 3.0      #maximum separation between input and recovered star
                      (flt>0)
FakePSF = 2.0        #assumed PSF FWHM for fake star matching
FakeStarPSF = 1      #use PSF residuals in fake star tests? (int 0=no,1=yes)
RandomFake = 1       #apply Poisson noise to fake stars? (int 0=no,1=yes)
FakePad = 0          #minimum distance of fake star from any chip edge to
                      be used
UsePhot              #if defined, use alignment, PSF, and aperture corr from
                      photometry
DiagPlotType =       #format to generate diagnostic plots (PNG, GIF, PS)
xytfile =            #position file for warmstart (str)
xytspf =              #reference PSF for image subtraction
VerboseData = 0      #to write all displayed numbers to a .data file
#
# Flags for HST modes
ForceSameMag = 0     #force same count rate in images with same filter? (int
                      0=no, 1=yes)
FlagMask = 4         #photometry quality flags to reject when combining magnitudes
CombineChi = 0       #combined magnitude weights uses chi? (int 0=no, 1=yes)
WFPC2useCTE = 1     #apply CTE corrections on WFPC2 data? (int 0=no, 1=yes)
ACSuseCTE = 0        #apply CTE corrections on ACS data? (int 0=no, 1=yes)
WFC3useCTE = 0       #apply CTE corrections on WFC3 data? (int 0=no, 1=yes)
ACSpstfType = 0     #use Anderson PSF cores? (int 0=no, 1=yes)
WFC3UVISpstfType = 0 #use Anderson PSF cores? (int 0=no, 1=yes)
WFC3IRpstfType = 0  #use Anderson PSF cores? (int 0=no, 1=yes)
InterpPSFlib = 1     #interpolate PSF library spatially
#
# Other flags not recommended for most users

```


APPENDIX C

TABLE OF FITTED DOUBLE GAUSSIAN PARAMETERS

TABLE A3.1: Fitted double Gaussian amplitudes, peak positions, and widths for all fifteen galaxies, with one-sigma uncertainties.

Galaxy	Blue amplitude	Blue mean	Blue sigma	Red amplitude	Red mean	Red sigma
NGC 57	68.6 ± 4.6	-0.73 ± 0.10	0.51 ± 0.08	161.9 ± 14.2	0.02 ± 0.01	0.20 ± 0.02
NGC 410	48.3 ± 5.0	-0.63 ± 0.11	0.55 ± 0.09	65.2 ± 12.4	0.06 ± 0.02	0.13 ± 0.03
NGC 533	106.1 ± 6.2	-0.59 ± 0.06	0.56 ± 0.05	103.8 ± 14.9	0.04 ± 0.02	0.13 ± 0.02
NGC 741	32.4 ± 4.8	-0.97 ± 0.09	0.31 ± 0.09	57.2 ± 5.9	-0.18 ± 0.05	0.24 ± 0.04
NGC 777	50.8 ± 3.3	-0.72 ± 0.11	0.52 ± 0.08	122.5 ± 11.4	0.00 ± 0.01	0.23 ± 0.02
NGC 1016	38.9 ± 3.8	-0.61 ± 0.14	0.53 ± 0.10	59.5 ± 11.4	0.00 ± 0.01	0.20 ± 0.04
NGC 1129	128.6 ± 9.3	-0.66 ± 0.10	0.47 ± 0.07	161.0 ± 29.1	-0.06 ± 0.02	0.20 ± 0.03
NGC 1600	97.4 ± 7.6	-0.59 ± 0.08	0.53 ± 0.05	124.4 ± 16.5	-0.03 ± 0.02	0.19 ± 0.03
NGC 2340	52.3 ± 6.2	-0.66 ± 0.19	0.48 ± 0.13	74.3 ± 21.2	0.02 ± 0.03	0.23 ± 0.05
NGC 3158	31.0 ± 4.2	-0.78 ± 0.21	0.58 ± 0.13	47.9 ± 12.1	-0.06 ± 0.02	0.27 ± 0.05
NGC 3842	42.9 ± 8.0	-0.87 ± 0.31	0.55 ± 0.19	54.6 ± 25.3	-0.14 ± 0.05	0.27 ± 0.08
NGC 4073	47.5 ± 6.2	-0.70 ± 0.16	0.58 ± 0.10	114.5 ± 14.1	-0.06 ± 0.02	0.21 ± 0.03
NGC 4839	49.9 ± 2.7	-0.78 ± 0.07	0.50 ± 0.06	64.0 ± 7.8	-0.05 ± 0.02	0.18 ± 0.02
NGC 4914	18.3 ± 4.0	-0.48 ± 0.22	0.47 ± 0.12	38.8 ± 9.3	0.04 ± 0.02	0.20 ± 0.04
NGC 7242	153.9 ± 7.5	-0.41 ± 0.04	0.46 ± 0.03	90.8 ± 16.2	0.08 ± 0.02	0.14 ± 0.03

BIBLIOGRAPHY

- El-Badry, K., Quataert, E., Weisz, D. R., Choksi, N., and Boylan-Kolchin, M. (Feb. 2019). The formation and hierarchical assembly of globular cluster populations. *Monthly Notices of the Royal Astronomical Society* 482(4), 4528–4552.
- Bastian, N., Pfeffer, J., Kruijssen, J. M. D., Crain, R. A., Trujillo-Gomez, S., and Reina-Campos, M. (Oct. 2020). The globular cluster system mass-halo mass relation in the E-MOSAICS simulations. *Monthly Notices of the Royal Astronomical Society* 498(1), 1050–1061.
- Beasley, M. A. (June 2020). Globular Cluster Systems and Galaxy Formation. In: *Reviews in Frontiers of Modern Astrophysics; From Space Debris to Cosmology*. Ed. by P. Kabáth, D. Jones, and M. Skarka. Springer International Publishing, 245–277.
- Beasley, M. A., Trujillo, I., Leaman, R., and Montes, M. (Mar. 2018). A single population of red globular clusters around the massive compact galaxy NGC 1277. *Nature* 555(7697), 483–486.
- Blakeslee, J. P., Tonry, J. L., and Metzger, M. R. (Aug. 1997). Globular Clusters in 19 Northern Abell Clusters. *Astronomical Journal* 114, 482–506.
- Bressan, A., Marigo, P., Girardi, L., Salasnich, B., Dal Cero, C., Rubele, S., and Nanni, A. (Nov. 2012). PARSEC: stellar tracks and isochrones with the PAdova and TRieste Stellar Evolution Code. *Monthly Notices of the Royal Astronomical Society* 427(1), 127–145.
- Brodie, J. P. and Strader, J. (Sept. 2006). Extragalactic Globular Clusters and Galaxy Formation. *Annual Review of Astronomy & Astrophysics* 44(1), 193–267.
- Chen, Y., Bressan, A., Girardi, L., Marigo, P., Kong, X., and Lanza, A. (Sept. 2015). PARSEC evolutionary tracks of massive stars up to $350 M_{\odot}$ at metallicities $0.0001 \leq Z \leq 0.04$. *Monthly Notices of the Royal Astronomical Society* 452(1), 1068–1080.
- Chen, Y., Girardi, L., Bressan, A., Marigo, P., Barbieri, M., and Kong, X. (Nov. 2014). Improving PARSEC models for very low mass stars. *Monthly Notices of the Royal Astronomical Society* 444(3), 2525–2543.

Bibliography

- Cho, J., Sharples, R. M., Blakeslee, J. P., Zepf, S. E., Kundu, A., Kim, H.-S., and Yoon, S.-J. (June 2012). Globular cluster systems of early-type galaxies in low-density environments. *Monthly Notices of the Royal Astronomical Society* 422(4), 3591–3610.
- Choksi, N. and Gnedin, O. Y. (Oct. 2019). Origins of scaling relations of globular cluster systems. *Monthly Notices of the Royal Astronomical Society* 488(4), 5409–5419.
- Choksi, N., Gnedin, O. Y., and Li, H. (Oct. 2018). Formation of globular cluster systems: from dwarf galaxies to giants. *Monthly Notices of the Royal Astronomical Society* 480(2), 2343–2356.
- Crook, A. C., Huchra, J. P., Martimbeau, N., Masters, K. L., Jarrett, T., and Macri, L. M. (Feb. 2007). Groups of Galaxies in the Two Micron All Sky Redshift Survey. *The Astrophysical Journal* 655(2), 790–813.
- De Bórtoli, B. J., Caso, J. P., Ennis, A. I., and Bassino, L. P. (Mar. 2022). Scaling relations for globular cluster systems in early-type galaxies - II. Is there an environmental dependence? *Monthly Notices of the Royal Astronomical Society* 510(4), 5725–5742.
- de Vaucouleurs, G., de Vaucouleurs, A., Corwin Jr., H. G., Buta, R. J., Paturel, G., and Fouque, P. (1991). *Third reference catalogue of bright galaxies, version 3.9*.
- Dolphin, A. (Jan. 2013). *DOLPHOT User’s Guide*.
- Dolphin, A. E. (Oct. 2000). WFPC2 Stellar Photometry with HSTPHOT. *The Publications of the Astronomical Society of the Pacific* 112(776), 1383–1396.
- Eadie, G. M., Harris, W. E., and Springford, A. (Feb. 2022). Clearing the Hurdle: The Mass of Globular Cluster Systems as a Function of Host Galaxy Mass. *The Astrophysical Journal* 926(2).
- Ennis, A. I., Bassino, L. P., Caso, J. P., and De Bórtoli, B. J. (Sept. 2019). Early-type galaxies in low-density environments: NGC 6876 explored through its globular cluster system. *Monthly Notices of the Royal Astronomical Society* 488(1), 770–781.
- Fahrion, K., Lyubenova, M., Hilker, M., van de Ven, G., Falcón-Barroso, J., Leaman, R., Martín-Navarro, I., Bittner, A., Coccato, L., Corsini, E. M., Gadotti, D. A., Iodice, E., McDermid, R. M., Pinna, F., Sarzi, M., Viaene, S., de Zeeuw, P. T., and Zhu, L. (May 2020). The Fornax 3D project: Non-linear colour-metallicity relation of globular clusters. *Astronomy & Astrophysics* 637.
- Goullaud, C. F., Jensen, J. B., Blakeslee, J. P., Ma, C.-P., Green, J. E., and Thomas, J. (Mar. 2018). The MASSIVE Survey. IX. Photometric Analysis of 35 High-mass Early-type Galaxies with HST WFC3/IR. *The Astrophysical Journal* 856(1).
- Harris, W. E. (Jan. 2010). Massive star clusters in galaxies. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 368, 889–906.
- Harris, W. E. (1998). Aperture and PSF Measurement. In: *Star Clusters*. Springer, 386–388.
- Harris, W. E., Blakeslee, J. P., Whitmore, B. C., Gnedin, O. Y., Geisler, D., and Rothberg, B. (Jan. 2016). Globular Cluster Systems in Brightest Cluster Galaxies. II. NGC 6166. *The Astrophysical Journal* 817(1).
- Harris, W. E., Ciccone, S. M., Eadie, G. M., Gnedin, O. Y., Geisler, D., Rothberg, B., and Bailin, J. (Jan. 2017). Globular Cluster Systems in Brightest Cluster Galaxies. III: Beyond Bimodality. *The Astrophysical Journal* 835(1).

Bibliography

- Harris, W. E., Harris, G. L., and Alessi, M. (Aug. 2013). A Catalog of Globular Cluster Systems: What Determines the Size of a Galaxy’s Globular Cluster Population? *The Astrophysical Journal* 772(2).
- Horta, D., Hughes, M. E., Pfeffer, J. L., Basian, N., Kruijssen, J. M. D., Reina-Campos, M., and Crain, R. A. (Jan. 2021). Linking globular cluster formation at low and high redshift through the age-metallicity relation in E-MOSAICS. *Monthly Notices of the Royal Astronomical Society* 500(4), 4768–4778.
- Hudson, M. J., Harris, G. L., and Harris, W. E. (May 2014). Dark Matter Halos in Galaxies and Globular Cluster Populations. *The Astrophysical Journal Letters* 787(1).
- Kormendy, J., Fisher, D. B., Cornell, M. E., and Bender, R. (May 2009). Structure and Formation of Elliptical and Spheroidal Galaxies. *The Astrophysical Journal Supplement* 182(1), 216–309.
- Kruijssen, J. M. D., Pfeffer, J. L., Crain, R. A., and Bastian, N. (July 2019). The E-MOSAICS project: tracing galaxy formation and assembly with the age-metallicity distribution of globular clusters. *Monthly Notices of the Royal Astronomical Society* 486(3), 3134–3179.
- Larsen, S. S., Romanowsky, A. J., Brodie, J. P., and Wasserman, A. (Nov. 2020). An extremely metal-deficient globular cluster in the Andromeda Galaxy. *Science* 370(6519), 970–973.
- Ma, C.-P., Greene, J. E., McConnell, N., Janish, R., Blakeslee, J. P., Thomas, J., and Murphy, J. D. (Nov. 2014). The MASSIVE Survey. I. A Volume-limited Integral-field Spectroscopic Study of the Most Massive Early-type Galaxies within 108 Mpc. *The Astrophysical Journal* 795(2).
- Marigo, P., Bressan, A., Nanni, A., Girardi, L., and Pumo, M. L. (Sept. 2013). Evolution of thermally pulsing asymptotic giant branch stars - I. The COLIBRI code. *Monthly Notices of the Royal Astronomical Society* 434(1), 488–526.
- Marigo, P., Girardi, L., Bressan, A., Rosenfield, P., Aringer, B., Chen, Y., Dussin, M., Nanni, A., Pastorelli, G., Rodrigues, T. S., Trabucchi, M., Bladh, S., Dalcanton, J., Groenewegen, M. A. T., Montalbán, J., and Wood, P. R. (Jan. 2017). A New Generation of PARSEC-COLIBRI Stellar Isochrones Including the TP-AGB Phase. *The Astrophysical Journal* 835(1).
- Muratov, A. L. and Gnedin, O. Y. (Aug. 2010). Modelling the Metallicity Distribution of Globular Clusters. *The Astrophysical Journal* 718(2), 1266–1288.
- Pastorelli, G., Marigo, P., Girardi, L., Aringer, B., Chen, Y., Rubele, S., Trabucchi, M., Bladh, S., Boyer, M. L., Bressan, A., Dalcanton, J. J., Groenewegen, M. A. T., Lebzelter, T., Mowlavi, N., Chubb, K. L., Cioni, M.-R. L., de Grijs, R., Ivanov, V. D., Nanni, A., van Loon, J. T., and Zaggia, S. (Nov. 2020). Constraining the thermally pulsing asymptotic giant branch phase with resolved stellar populations in the Large Magellanic Cloud. *Monthly Notices of the Royal Astronomical Society* 498(3), 3283–3301.
- Pastorelli, G., Marigo, P., Girardi, L., Chen, Y., Rubele, S., Trabucchi, M., Aringer, B., Bladh, S., Bressan, A., Montalbán, J., Boyer, M. L., Dalcanton, J. J., Eriksson, K., Groenewegen, M. A. T., Höfner, S., Lebzelter, T., Nanni, A., Rosenfield, P., Wood,

Bibliography

- P. R., and Cioni, M.-R. L. (June 2019). Constraining the thermally pulsing asymptotic giant branch phase with resolved stellar populations in the Small Magellanic Cloud. *Monthly Notices of the Royal Astronomical Society* 485(4), 5666–5692.
- Pfeffer, J., Kruijssen, J. M. D., Crain, R. A., and Bastian, N. (Apr. 2018). The E-MOSAICS project: simulating the formation and co-evolution of galaxies and their star cluster populations. *Monthly Notices of the Royal Astronomical Society* 475(4), 4309–4346.
- Pillepich, A., Nelson, D., Hernquist, L., Springel, V., Pakmor, R., Torrey, P., Weinberger, R., Genel, S., Naiman, J. P., Marinacci, F., and Vogelsberger, M. (Mar. 2018). First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies. *Monthly Notices of the Royal Astronomical Society* 475(1), 648–675.
- Salinas, R., Alabi, A., Richtler, T., and Lane, R. R. (May 2015). Isolated ellipticals and their globular cluster systems. III. NGC 2271, NGC 2865, NGC 3962, NGC 4240, and IC 4889. *Astronomy Astrophysics* 577(A59).
- Schlaafy, E. F. and Finkbeiner, D. P. (Aug. 2011). Measuring Reddening with Sloan Digital Sky Survey Stellar Spectra and Recalibrating SFD. *The Astrophysical Journal* 737(2).
- Spitler, L. R. and Forbes, D. A. (Jan. 2009). A new method for estimating dark matter halo masses using globular cluster systems. *Monthly Notices of the Royal Astronomical Society: Letters* 392(1), L1–L5.
- Stetson, P. B. (Mar. 1987). DAOPHOT: A Computer Program for Crowded-Field Stellar Photometry. *The Publications of the Astronomical Society of the Pacific* 99, 191.
- Tang, J., Bressan, A., Rosenfield, P., Slemer, A., Marigo, P., Girardi, L., and Bianchi, L. (Dec. 2014). New PARSEC evolutionary tracks of massive stars at low metallicity: testing canonical stellar evolution in nearby star-forming dwarf galaxies. *Monthly Notices of the Royal Astronomical Society* 445(4), 4287–4305.
- Veale, M., Ma, C.-P., Greene, J. E., Thomas, J., Blakeslee, J. P., McConnell, N., Walsh, J. L., and Ito, J. (Oct. 2017). The MASSIVE Survey - VII. The relationship of angular momentum, stellar mass and environment of early-type galaxies. *Monthly Notices of the Royal Astronomical Society* 471(2), 1428–1445.
- Villaume, A., Romanowski, A. J., Brodie, J., and Strader, J. (July 2019). New Constraints on Early-type Galaxy Assembly from Spectroscopic Metallicities of Globular Clusters in M87. *The Astrophysical Journal* 879(1).