

Computational Cosmology as seen through a Telescope:
Observational Properties of Simulated Galaxies

by

Connor Bottrell
B.Sc., University of Victoria, 2014

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Physics and Astronomy

© Connor Bottrell, 2016
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by
photocopying or other means, without the permission of the author.

Computational Cosmology as seen through a Telescope:
Observational Properties of Simulated Galaxies

by

Connor Bottrell
B.Sc., University of Victoria, 2014

Supervisory Committee

Dr. Luc Simard, Co-Supervisor
(Department of Physics and Astronomy)

Dr. Sara L. Ellison, Co-Supervisor
(Department of Physics and Astronomy)

Supervisory Committee

Dr. Luc Simard, Co-Supervisor
(Department of Physics and Astronomy)

Dr. Sara L. Ellison, Co-Supervisor
(Department of Physics and Astronomy)

ABSTRACT

The current generation of cosmological hydrodynamical simulations offer new levels of fidelity in galaxy formation and evolution that can be benchmarked against observations. However, it is crucial that the comparison between the simulation products and observations is performed on level-ground. Using mock-observations of galaxies from hydrodynamical simulations with observational realism, an image-based comparison is enabled between the simulations and modern galaxy surveys that allow galaxy properties to be derived consistently. A new methodology is presented that provides an unprecedentedly comprehensive suite of observational realism to synthetic images of galaxies from simulations and performs detailed decomposition of their morphological structures. The crux of the methodology is that the same procedure for image-based surface-brightness decompositions of galaxy structures is employed for the simulations and observations – facilitating a fair and unbiased comparison of galaxy properties. The methodology is piloted on galaxies from the Illustris simulation and is designed to enable comparison with galaxies from the Sloan Digital Sky Survey (SDSS). The biases from observational realism on the decomposition results for the simulated galaxies are characterized in detail using several controlled experiments. Then, the decompositions are used in a comparison of the galaxy size-luminosity and bulge-to-total vs. total stellar mass relations. The comparisons show that galaxies from Illustris contain too many discs and too few bulges at low masses $\log M_{\star}/M_{\odot} \lesssim 11$ relative to the SDSS. A comparison of the photometric and kinematic bulge-to-total is also enabled by the methodology. The comparison suggests that photometry tends to systematically under-estimate the bulge fractions relative to

the kinematics – allowing no discernible connection to be made between photo-bulges and kinematic bulges. Several possibilities that may be driving the bulge deficit in Illustris’ low-mass galaxies are discussed, though particle resolution is argued to be the main culprit. The methodology that is presented in this thesis has broad applications to comparisons between computational and observational galaxy astronomy and stands to provide a wealth of feedback between each community.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	v
List of Tables	viii
List of Figures	ix
Acknowledgements	xi
1 Galaxy Surveys and Simulations: Setting the Stage for a New Comparison	1
1.1 Introduction	1
1.2 Observing Galaxies	5
1.2.1 Modern Galaxy Survey Astronomy	5
1.2.2 Galaxy Photometry in the SDSS	6
1.3 Quantitative Morphologies	13
1.3.1 Non-Parametric Morphologies	14
1.3.2 Parametric Morphologies	15
1.4 Numerical Simulations	23
1.4.1 N-body Simulations	23
1.4.2 Semi-Analytic Models	24
1.4.3 Hydrodynamical Simulations	25
1.4.4 Mock Observations	26
2 Galaxies in the Illustris simulation as seen by the Sloan Digital Sky Survey - I: Bulge+disc decompositions, methods, and biases.	32
2.1 Introduction	34

2.2	Simulated Galaxies	37
2.2.1	Illustris Simulation	37
2.2.2	Synthetic Galaxy Images	38
2.2.3	Image Realism	40
2.3	Quantitative Morphologies	44
2.3.1	Deblending	45
2.3.2	Parametric Surface Brightness Profile Fits	48
2.3.3	Simulated Galaxy Population Samples and Catalogs	48
2.4	Characterization of Biases	59
2.4.1	The Propagation of Stellar Light from Unresolved Stellar Populations	59
2.4.2	Camera Angle	65
2.4.3	Environment and Crowding	68
2.4.4	Internal Segmentation by Artificially Discrete Substructure	75
2.4.5	Systematics of Size and Flux Estimates in the Full Illustris Galaxy Population	83
3	Comparison with Real Galaxies	90
3.1	The Galaxy Size-Luminosity Relation	91
3.1.1	Selection of a Comparison Sample	94
3.1.2	Galaxy Size-Luminosity Relations of SDSS and Illustris	95
3.1.3	Impact of Bulge and Disc Morphologies	99
3.2	The Photo-Bulge Deficit	108
3.2.1	Morphological dependence on stellar mass	108
3.2.2	Comparison with kinematic b/t	113
4	Summary	118
4.1	Observational Realism and Morphological Decompositions	118
4.2	Characterization of Biases	119
4.3	Comparison with Real Galaxies	122
4.4	The Photo-Bulge Defecit	123
4.5	Future Directions	124
4.5.1	Disc and Bulge Structural Properties	124
4.5.2	Photometric Estimates of Stellar Masses	125
4.5.3	Redshift Dependence on Structural Estimates	125

4.5.4	Zooming-in on Illustris	125
4.5.5	Application to other Simulations	126
	Bibliography	128

List of Tables

Table 2.1 Decomposition Catalogs	57
--	----

List of Figures

Figure 1.1	Hubble Morphological Classification Scheme	3
Figure 1.2	Comparison of Galaxy Redshift Surveys	7
Figure 1.3	Bias on flux estimates from the SDSS DR7 PHOTO pipeline sky measurements	10
Figure 1.4	Allocation of light in crowded photometry - SDSS vs. SExtractor	11
Figure 1.5	Completeness of SDSS Galaxy Samples	12
Figure 1.6	Mosaic of GIM2D Fits to Visually Classified Morphologies . . .	18
Figure 1.7	Colour-magnitude diagrams from 2-D decompositions of galaxy pairs	21
Figure 1.8	Evolution of a present-day elliptical and spiral in a hydrody- namical simulation	26
Figure 1.9	Blue galaxies in the Illustris Simulation	27
Figure 1.10	Red Galaxies in the Illustris Simulation	28
Figure 1.11	Illustris Galaxy Morphologies and Mocks	30
Figure 2.1	Addition of observational realism	43
Figure 2.2	Camera angles and FoF projection	47
Figure 2.3	RIGS 1-20	49
Figure 2.4	RIGS 21-40	50
Figure 2.5	RIGS 41-60	51
Figure 2.6	RIGS 61-80	52
Figure 2.7	RIGS 81-100	53
Figure 2.8	Mosaic of fits to simulated galaxies with realism	58
Figure 2.9	Magnitude bias from the choice of SLD scheme	60
Figure 2.10	HLR bias from the choice of SLD scheme	61
Figure 2.11	(b/t) bias from the choice of SLD scheme	62
Figure 2.12	Statistical bias from projection	66
Figure 2.13	Statistical bias from placement	70

Figure 2.14	Systematic biases on magnitude and HLR from placement . . .	74
Figure 2.15	Examples showing low sensitivity of parameter estimates to crowding	77
Figure 2.16	Examples showing high sensitivity of parameter estimates to internal segmentation	79
Figure 2.17	Examples of internal segmentation in diffuse galaxies	80
Figure 2.18	Median <i>bflags</i> for each galaxy in the ASKA catalog	82
Figure 2.19	Recovery of flux in the DISTINCT catalog	85
Figure 2.20	Accuracy of model HLR estimates in the DISTINCT catalog . . .	88
Figure 3.1	Size-luminosity relations of the SDSS and Illustris	96
Figure 3.2	Size-luminosity relations of the SDSS and Illustris synthetic images	98
Figure 3.3	Size-luminosity relations of late- and early-type galaxies . . .	101
Figure 3.4	Size-luminosity relations of pure discs and ellipticals	104
Figure 3.5	Bulge-to-total fractions in the SDSS and Illustris	105
Figure 3.6	Size-luminosity relations matched by mass and morphology . .	107
Figure 3.7	Un-matched (<i>b/t</i>) in SDSS and Illustris	110
Figure 3.8	Photometric vs. kinematic bulge-to-total fractions	115

ACKNOWLEDGEMENTS

I would like to thank:

Luc Simard, for taking me on as your student. Your shared encouragement, knowledge, and support through tough troughs were always welcomed.

Sara Ellison, for integrating me into your group and providing me with the resources and connections necessary to perform research that I enjoyed thoroughly. Your involvement, inclusiveness, and support were invaluable.

Paul Torrey, for your advise, direction, and resources throughout my research without which I would have been stuck with plain-old *real* galaxies.

Don and Shelly, for the pride and love you have unconditionally shown me.

The Boys at Westshore RFC, for keeping me active and “sane.”

Maëva Perez, la lumière de ma vie, you are the best and I love you.

Chapter 1

Galaxy Surveys and Simulations: Setting the Stage for a New Comparison

1.1 Introduction

The building blocks of the universe are galaxies. But instead of blocks with similar shape and equal size, imagine that each block is a piece of a jigsaw puzzle. Collectively, galaxies form vast structures that are seen and derived from large surveys (Shane & Wirtanen 1954; Groth & Peebles 1977; Lilje & Efstathiou 1988; reviews by Abell 1965; Oort 1983; Chincarini & Vettolani 1987; Rood 1988; Bahcall 1988). Individually, they demonstrate remarkable diversity in their morphologies (reviews by Sandage 1975; Buta 2011 and references therein). However, the morphologies of galaxies are not determined by an industrial press. They are the products of chaotic and varied formation and merger histories, environmental influences, star-formation histories, dynamical instabilities, among other astrophysical processes. Morphologies of galaxies are tied to the physics that govern the formation and evolution of structure in the universe. Therefore, any theory for the physics of galaxy formation and evolution must be reconciled with the observed properties of galaxies and reproduce their diverse morphologies. Visual and quantitative classifications of galaxy morphology facilitate the comparison of theory with observation (e.g. Abadi et al. 2003a; Scannapieco et al. 2010; Brooks et al. 2011; Christensen et al. 2014; Snyder et al. 2015; reviews by van den Bergh 1998; Buta 2011 and references therein).

Direct observational signatures of galaxies are accessible from their baryonic constituents. The chemical composition of the atomic and molecular gas within galaxies can be determined from spectroscopy (reviews by Kitchin 1987; York 1982 and references therein). Transitions of electrons between energy states occur at specific energies that correspond to the properties of their parent atoms or molecules. A transition from lower energy states to higher ones are facilitated by the absorption of electromagnetic radiation with the precise energies corresponding to the energy differences in the transitions. Inversely, higher to lower transitions emit electromagnetic radiation with these precise energies. The excitation and relaxation of atomic electrons at specific energies manifest themselves as absorption and emission lines in the spectra of galaxies. The lines are used to measure the relative velocities of galaxies (e.g. Hubble 1929). Electromagnetic radiation is blue-shifted (perceived increase in energy) and red-shifted (perceived decrease in energy) according whether the velocity of an emission source is directed toward and away from an observer, respectively. The gas dynamics can be determined by obtaining spectra that are spatially resolved across a galaxy. Spatially resolved spectra of individual galaxies are obtained using long-slit spectroscopy oriented along the primary projected axis of a galaxy (e.g. Reyes et al. 2011; Kauffmann et al. 2015) or an Integral Field Unit – often composed of arrays of tightly bundled spectroscopic fibres (e.g. Kuzio de Naray 2008). The formation and evolution of the gaseous components determine the capacity of a galaxy to form stars – which provide independent means for direct observational analysis of morphology.

Information about the spatial distribution and structures of the stellar components is obtained from the projection of stellar light in galaxies. Visual inspection of the photometry reveals that a range of morphological structures exist among populations of galaxies (e.g. Hubble Deep Field: Menanteau et al. 2006, Galaxy Zoo (GZ): Fortson et al. 2012; Willett et al. 2013). Stellar discs are revealed among inclined galaxies (e.g. Reynolds 1916). The disc typically contains the largest concentration of the interstellar matter and commonly contains the spiral arms if the galaxy is a spiral (Hubble 1926, 1936; de Vaucouleurs 1959a, review by Buta 2011). Analyses of light profiles of discs indicate that they empirically follow an exponential surface brightness distribution that eventually truncates at large radii in the absence of perturbations from other structures (de Vaucouleurs, 1959b; Freeman, 1970; van der Kruit & Searle, 1982; Freeman, 2007). Many disc-shaped galaxies contain a luminous stellar component at their centres, called a bulge, that often has a more centrally concentrated surface-brightness profile than the disc (de Vaucouleurs, 1953; Kormendy, 1977; Ko-

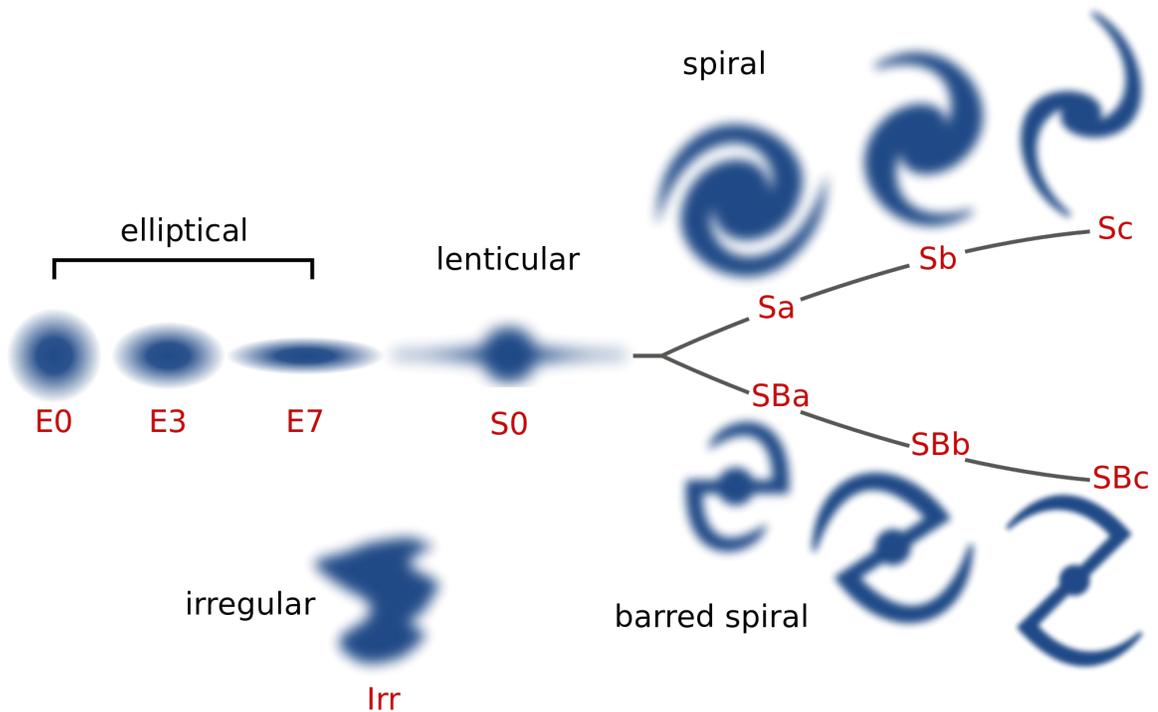


Figure 1.1: Hubble Morphological Classification Scheme with cartoon depictions of each class. Galaxies with elliptical morphologies are on the left. Spiral galaxies with and without discernible bar structures decrease in winding tightness from left to right. At the apex of spirals and ellipticals is the S0 lenticular class. The classes often correlate phenomenologically with star-formation histories, current star-formation rates, colours, merger histories, environmental effects, and other astrophysical phenomena.

rmendy & Djorgovski, 1989). The fraction of stellar light from the bulge relative to the disc can vary significantly. The shape of the surface brightness distribution in a bulge is related to its formation mechanism (Athanasoula, 2005). Classical bulges are thought to form from the hierarchical assembly of disc galaxies in the same way that elliptical galaxies are formed (review by Wyse et al. 1997 and references therein). If so, classical bulges have the same light profiles as ellipticals. However, alternative mechanisms for the formation of a bulge are required to describe variants in bulge structure (e.g. Kormendy 1985; Gadotti 2009). Disc instabilities can lead to the formation of a bar-like structure. Bars facilitate secular movement of gas within the disc toward the centre of the galaxy and provides the raw material to build bulge-like structures known as pseudo-bulges (reviews by Kormendy & Kennicutt 2004; Kormendy & Fisher 2008).

Observation-based classification of the various galaxy morphologies is central to galaxy formation and evolution theory – which places each class of objects into a global context. The most prevailing visual classification system was proposed by Hubble (1926, 1936) and later revisited and extended by de Vaucouleurs (1959a) and Sandage (1961). Figure 1.1 shows a cartoon of the Hubble classification scheme. Classifications are separated broadly into ellipticals and spirals. Each broad class is separated by slightly more detailed properties. Ellipticals are classified according to the ellipticity of their isophotes, the apparent absence of a stellar disc, and the smooth decline in its surface brightness distribution. Spirals are separated by the presence of a bar, central concentration, and by the tightness of the spiral arm winding. The success of the scheme is its simplicity. The vast majority of the galaxies that were observed could be broadly classified according to the scheme (Sandage, 1975). The trapping of visual classification schemes is that they are interpretive and generally non-quantitative. Separation of stellar light components, such as from the bulge and disc, and measuring their structural properties is challenging without modelling these components quantitatively. Furthermore, visual classification is time-consuming for surveys containing thousands of galaxies. The advent of modern galaxy survey astronomy has exacerbated the trappings of visual classification by collecting data volumes containing photometry for millions of galaxies (e.g. 2dF Galaxy Redshift Survey: Colless 1999, Sloan Digital Sky Survey: Stoughton et al. 2002).

Large-scale multi-filter imaging and spectroscopic redshift surveys such as the 2dF Galaxy Redshift Survey (2dFGRS) (Colless, 1999; Colless et al., 2001) and the Sloan Digital Sky Survey (SDSS) (York et al., 2000; Eisenstein et al., 2011) have generated substantial observational data volumes. Such surveys facilitate comparisons between theory and observation across cosmic time. Galaxy samples of the sizes provided by these surveys are collectively an enormous test-bed against which models in galaxy formation and evolution theory must be benchmarked (Somerville & Davé, 2015). The most direct way to create a link between theory with observations is with numerical hydrodynamical simulations that explicitly model the co-evolution of dark and baryonic matter (e.g. GASOLINE: Wadsley et al. 2004; FIRE: Hopkins et al. 2014). In particular, cosmological hydrodynamical simulations of large volumes facilitate statistically relevant comparisons of galaxy morphologies (e.g. EAGLE: Schaye et al. 2015; Illustris: Vogelsberger et al. 2014b). The crux of the interpretive power of these simulations is in whether they are fundamentally able to reproduce the observed morphological properties and diversity of galaxies seen throughout the

observable universe.

A fair quantitative comparison of galaxy properties requires translation between the camps of simulations and observations. Either the observed properties must be mapped into the space of physical parameters or the simulations must be placed into an observational context. The latter approach provides the most direct mapping (Conroy et al., 2010; Hayward et al., 2013a,b). Mock observations of simulated galaxies allow the same analysis tools that are used in large galaxy surveys to be used on the simulation products. Image-based comparisons between simulations and observations have only recently become accessible in ways that meaningful quantitative and statistical statements can be made about the similarities and differences between simulated and observed galaxy populations. The gap between the two can now be bridged in response to the physical fidelity, resolution, and scales achieved by the current generation of cosmological hydrodynamical simulations and the acquisition of vast observational constraints used to validate them.

1.2 Observing Galaxies

1.2.1 Modern Galaxy Survey Astronomy

Galaxy surveys were once confined to small regions of the sky and limited to statistically biased samples of objects (review by van den Bergh 1998). The two principle factors limited the relevance of early galaxy samples to models describing global formation and evolution: (1) insufficient observational volumes to achieve representative statistics of galaxy properties (cosmic variance); (2) Single-band observations that limit the interpretive power in studying individual galaxies. Modern surveys, on the other hand, are designed to uniformly probe large fractions of the sky with deep imaging and spectroscopy to obtain precise statistics for the distribution of galaxy morphologies and other properties throughout the history of the universe (e.g. Falco et al. 1999; Stoughton et al. 2002; Colless 1999; Davis et al. 2003; Le Fèvre et al. 2005; Lilly et al. 2007; Huchra et al. 2012). Figure 1.2 compares the covering area and density of spectra in several modern surveys. The number of photometric sources significantly exceed the number of sources with spectra in a typical survey. Modern surveys are often either volume limited (every galaxy within a specified luminosity range is expected to be observed within some fixed radial volume of the observable universe, e.g. MASSIVE: Ma et al. 2014) or surface-brightness limited (all

sources are included that can be identified with the surface brightness limits of the survey, e.g. SDSS: Stoughton et al. 2002; 2dFGRS: Colless 1999). Large volume and surface-brightness limited surveys that comprise of statistically meaningful populations of galaxies at several epochs provide a fundamental basis for theory of galaxy astrophysics. Any theoretical picture for galaxy formation and evolution must ultimately be benchmarked against the observed statistics of galaxy masses, luminosities, colours, kinematics, morphologies, and other properties.

While surveys that will achieve deeper imaging and cover larger volumes are currently in operation and on the horizon (e.g. GAMA: Driver et al. 2009; LSST: Ivezić et al. 2008), galaxy redshift surveys such as 2dF, CfA, DEEP, and the SDSS pioneered large optical and near-IR surveys. The SDSS imaged over 25% of the sky (now just over 35%) using a dedicated 2.5m wide-angle optical telescope at the Apache Point Observatory in New Mexico, United States. The SDSS uses the *ugriz* photometric system – which is comprised of five filters that cover the optical spectrum. The survey has surface-brightness limited photometry of millions of galaxies and spectra of $\sim 700,000$ (Abazajian et al., 2009). Using a source’s spectrum, its recession velocity can be determined using known features such as atomic emission lines. The velocities of galaxies are used to determine their distances using measurements and theory for the expansion history of the universe (e.g. Supernova Type Ia (SNeIa): Riess et al. 1998; Perlmutter 2003, Cosmic Microwave Background (CMB) Radiation Mapping: Bennett et al. 2013; Planck Collaboration et al. 2015, Distance Measures in Cosmology: Hogg 1999). The velocity is often called the redshift, z , when the relative velocity of a source is dominated by the universe’s expansion. The general shape of a source’s spectrum also shows whether the source is a single star within our own Milky Way, a distant galaxy, or a compact component of a distant galaxy. Such systems are more difficult to distinguish with photometry alone. Sources without spectra undergo a rigorous photometric classification analysis with the aim of separating extended sources (such as galaxies) from point sources (such as individual stars) out to limits of the survey’s detection thresholds (e.g. SDSS PHOTO pipeline: (Stoughton et al., 2002)).

1.2.2 Galaxy Photometry in the SDSS

The scope, detailed technical information about the SDSS, its biases and limitations, and the wealth of literature using its data products have made the SDSS a standard

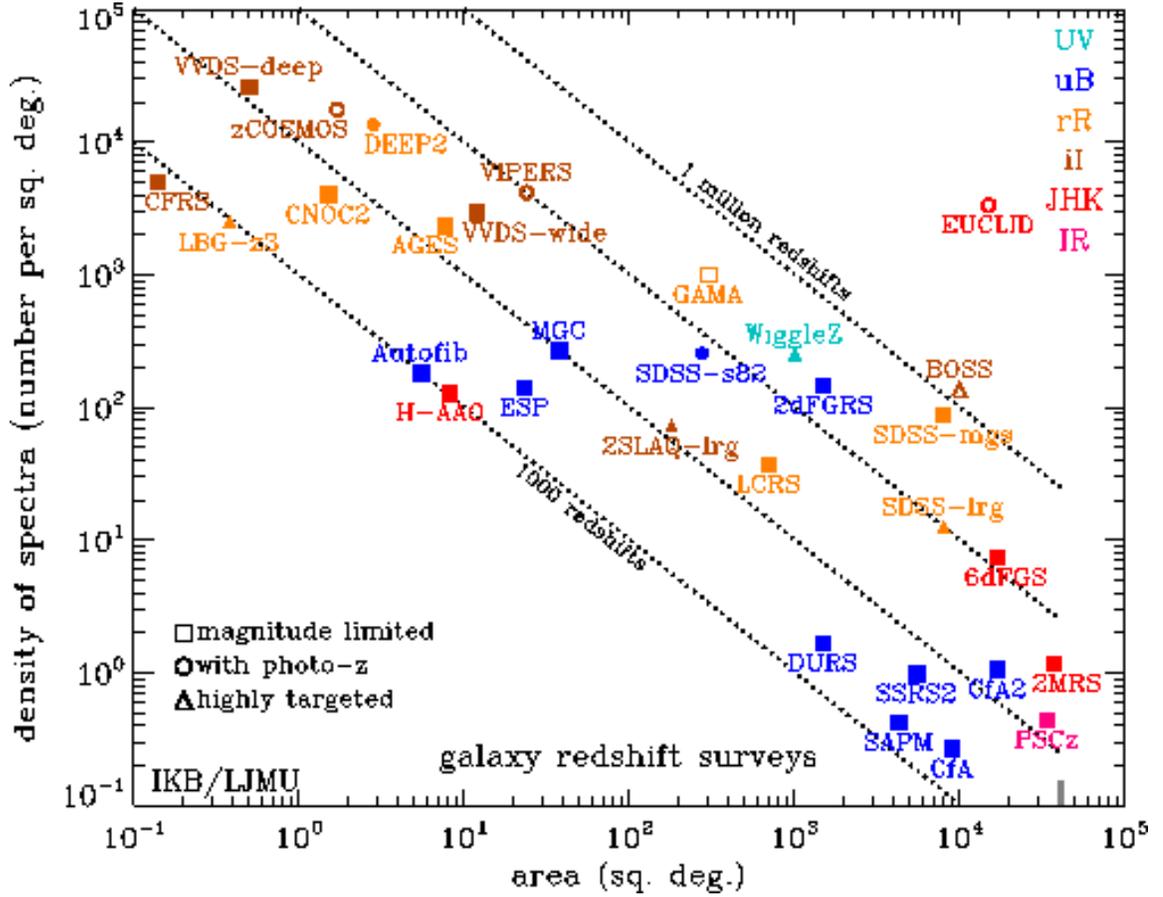


Figure 1.2: Comparison of past, current and future galaxy redshift surveys according to the covering area over which they extend and by the density of spectra per unit area. Surveys are coloured according to their selection wavelength shown by the key on the top right. Black dotted lines correspond to constant total numbers of spectra with 10^3 , 10^4 , 10^5 , and 10^6 shown. Surveys are categorized according to whether they are predominantly magnitude limited (squares), select galaxies using colour cuts for photometric redshifts (circles), or highly targeted surveys (triangles). Note that the photometric components of many of these surveys often consist of significantly larger numbers of galaxies. (Figure adapted from Baldry et al. 2010)

in modern galaxy surveys. Several biases affect galaxy photometry with ground-based telescopes like the SDSS such as atmospheric conditions and sky brightness. Understanding the biases is crucial to obtaining any information about the physical properties of galaxies from image data.

The point spread function (PSF) is a bias that limits the spatial resolution accessible in observations. The PSF arises from telescope optics, and, in the case of ground-based telescopes, from the blurring effect of atmospheric turbulence. Fortunately, the large field of view of the SDSS means that there are usually many spatially distributed stars that meet the selection criterion for being used to estimate the PSF – which is important because the size of the PSF can vary by up to 14% from one edge of a field to another (Gunn et al., 2006). The selected stars in each image are used to reconstruct the PSF at each pixel as described by Lupton et al. (2001) for the SDSS imaging pipeline, PHOTO. Accuracy of the PSF and its accessibility at each spatial location across the SDSS is crucial to image-based comparisons of galaxy morphologies – as this bias may manifest itself differently from galaxy to galaxy across the survey.

The SDSS sky brightness from scattered light is non-uniform on arc-minute scales (Stoughton et al., 2002). The contribution from the sky affects the ability to distinguish faint sources and to separate the wings of extended surface brightness profiles from the background light. The SDSS processing pipeline determines a clipped median sky within every 100'' square on a grid with 50'' spacing. The medians are linearly interpolated to each pixel on the grid and subtracted. This estimation of the background has been reported to be systematically biased in the presence of bright extended galaxies (r -band magnitudes $m_r < 16$) by up to $\Delta m_r \approx 0.8$. The bias suppresses the detection of faint sources in the vicinity of bright objects (Adelman-McCarthy et al., 2006; Mandelbaum et al., 2005) and leads to underestimates of the flux and structural scale sizes of bright galaxies (Lauer et al., 2007; Bernardi et al., 2007; Lisker et al., 2007). The biases associated with sky brightness are more rigorously quantified in the Seventh Data Release (DR7) of the SDSS (Abazajian et al., 2009) where its affect is shown to be less severe, but still significant. The reduced severity is shown in Figure 1.3, which compares the model magnitudes in the r -band recovered by PHOTO from analytic galaxy surface-brightness profile models that are inserted into SDSS skies to quantify the bias from sky measurement in PHOTO. Nevertheless, highly accurate estimations of galaxy flux and scales are not achieved with PHOTO alone – but require slightly more sophisticated methods of measuring skies

to facilitate measurement of light where the signal-to-noise (S/N) is poor (such as in the wings of extended sources).

In addition to spatial resolution and sky background effects, galaxy photometry in the SDSS is affected by crowding. Crowding occurs when the projected light profiles of sources on the sky overlap. Crowding causes to systematic errors in the morphological estimates of galaxies. The sizes of these errors depend sensitively on the algorithm that is used to allocate light for overlapping light profiles (Simard et al., 2011; Patton et al., 2011). Figure 1.4 compares the allocation of light for a crowded galaxy using two unique deblending schemes. The deblending mask in the top right panel is from Source Extractor (Bertin & Arnouts, 1996), where each pixel is allocated to a single source. The SDSS PHOTO deblending (Lupton et al., 2002) is shown in the bottom right panel. In the PHOTO deblending, the flux of an individual pixel may be allocated to more than one sources in proportion to its probability of belonging to each source. The relative fractions are determined by employing simple surface-brightness profile models on the photometry. The SExtractor scheme would be more strongly biased than the SDSS scheme if the flux of each source is computed by summing the the allocated fluxes directly. However, the allocation of each pixel’s flux to a single source has been shown to achieve accurate and precise morphological estimates with the aid of surface-brightness profile fitting strategies (Simard et al., 2011).

With a handle on the biases, the scale and public availability of the SDSS make it ideal for assembling large samples of galaxies with photometry and spectra even for studies with very specific selection criterion (e.g. studies of galaxy pairs in the SDSS: Ellison et al. 2008, 2010; Patton et al. 2011; Ellison et al. 2011; Scudder et al. 2012; Patton et al. 2013; Ellison et al. 2013). However, acquisition of galaxy photometry in SDSS does not uniformly sample the true density of galaxies because SDSS galaxy sample is magnitude-limited. Intrinsically bright galaxies are overrepresented in such a survey. The completeness of a galaxy survey measures the fraction of galaxies that it is able to detect as a function of magnitude. The completeness of the SDSS is determined by comparing detections in regions of the sky that are also scanned by telescopes that have superior surface brightness limits (completeness is also determined from analytic model recovery analyses, e.g. Annis et al. 2014). Figure 1.5 shows the completeness of the SDSS through comparison with the COMBO-17 survey S11 field (Wolf et al., 2003). COMBO-17 uses a Wide Field Imager at the MPG/ESO 2.2-m telescope on La Silla, Chile with photometry of one square degree of sky in 17 filters. While the camera exposure time for the SDSS is 53.9s with median PSF resolution

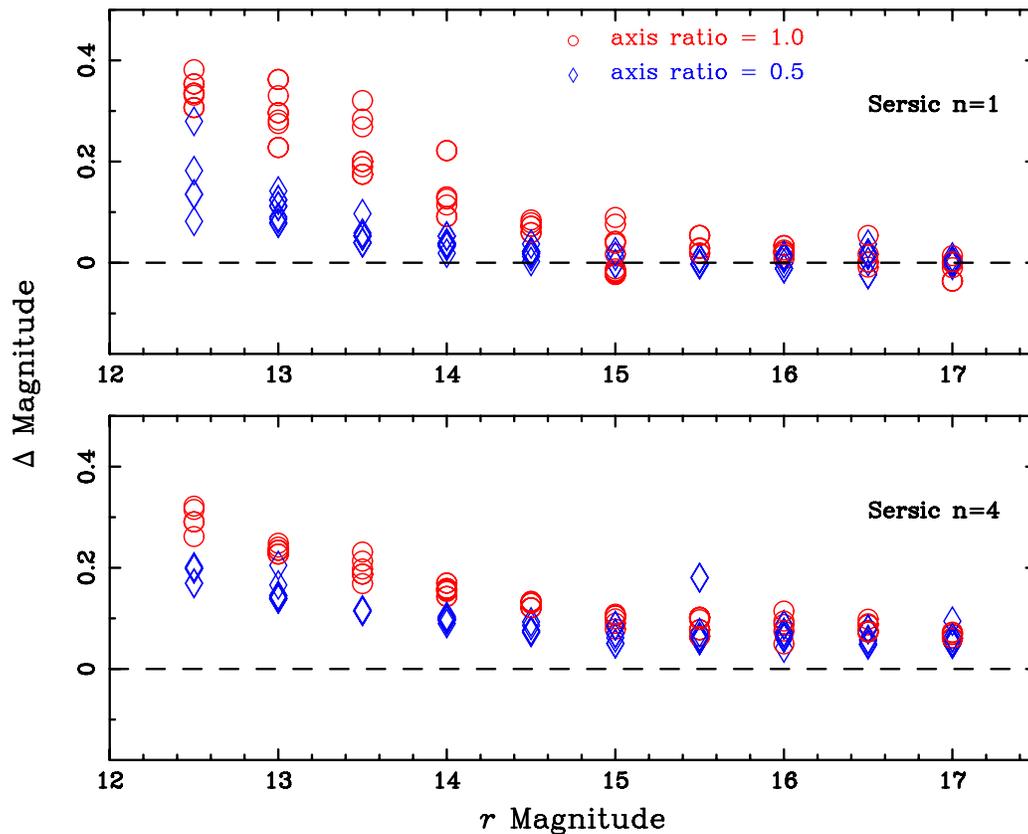


Figure 1.3: Bias on flux estimates for analytic surface brightness profile models of galaxies placed into SDSS skies. The analytic surface brightness profile models have *a priori* known parametric form and are parametrized using sérsic functions with sérsic indices $n = 1$ (used to model discs; upper panel) and $n = 4$ (ellipticals; lower panel) with axis ratios of either 1 (red) or 0.5 (blue). The analytic models were inserted into random high-latitude fields of the SDSS and run through PHOTO. The effect of the sky estimation from PHOTO is evaluated by comparing the PHOTO model magnitude estimates for the inserted model galaxies to their known properties, Δm_r . The bias is significant at the bright end with a maximum $\Delta m_r \approx 0.3$. However, the bias is less than the $\Delta m_r \approx 0.8$ reported in other works. All estimates are systematically fainter by 0.05 magnitudes as a result of truncation of the light profiles of PHOTO models beyond seven scale lengths. (Figure credit: Abazajian et al. 2009)

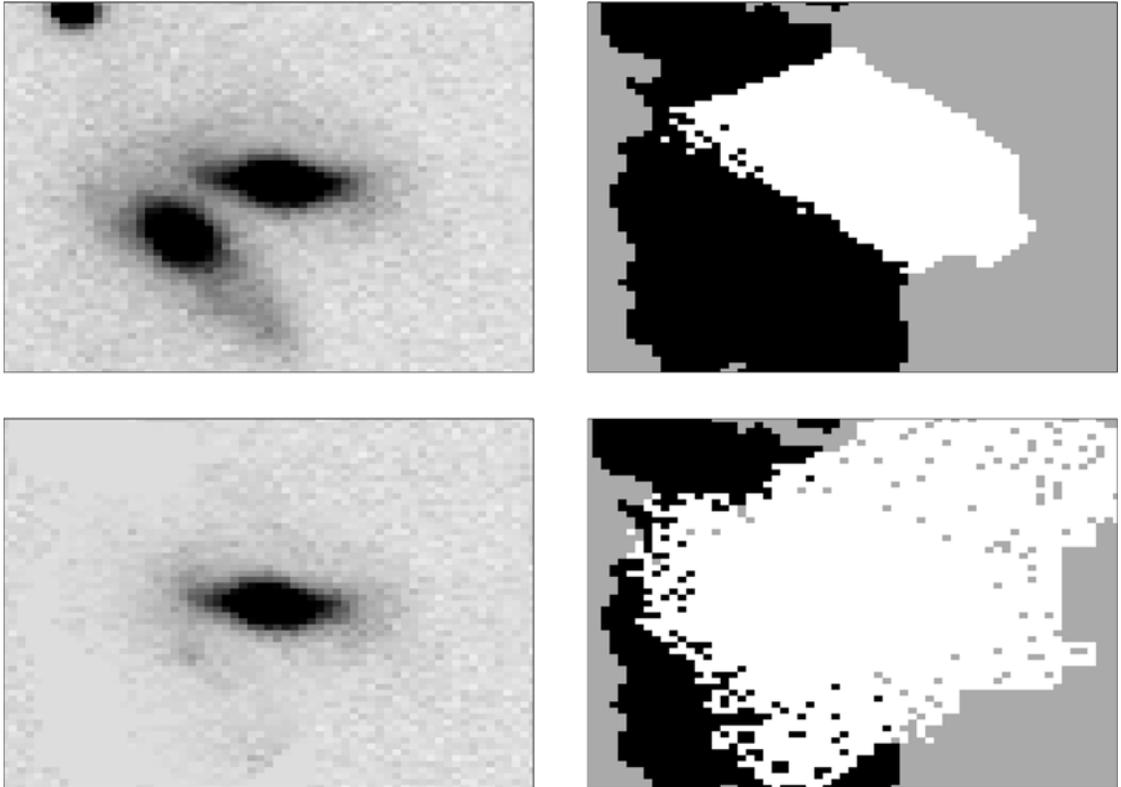


Figure 1.4: Science and deblending mask images of galaxies in a close pair. The top left panel shows the science image from the SDSS centred on the galaxy for which the flux and structure is being measured. The top right panel shows the allocation of pixels to each galaxy using the SExtractor sky determination and deblending scheme. Each pixel is allocated directly to an individual source. The bottom right panel shows the deblending of the sources using the SDSS PHOTO pipeline. Each pixel can allocate a fraction of its flux to different sources in the SDSS based on models. Significant overlap of a source's light profile over the profile of a neighbouring source is possible in this scheme. The bottom left panel shows the object image given by the SDSS deblending after masking of the light from external sources. (Figure credit: Simard et al. 2011)

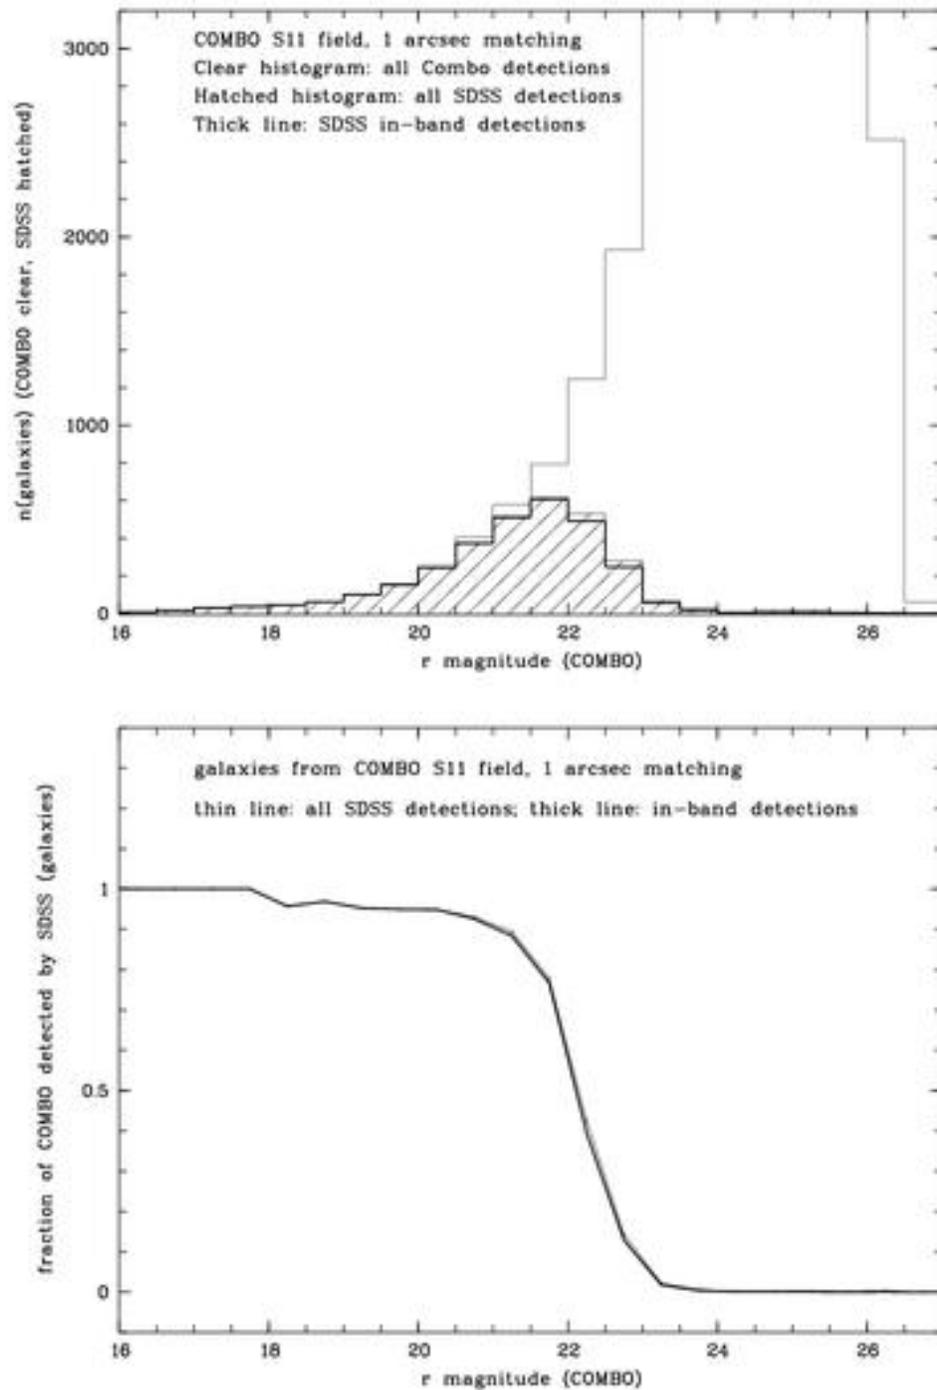


Figure 1.5: Comparison of SDSS galaxy detections and COMBO-17 in the S11 field overlap in the r -band. Top panel: COMBO-17 (clear histogram) demonstrates superior surface brightness limits than SDSS (hatched histogram) with significantly greater numbers of galaxy detections as a function of magnitude for fainter sources. Bottom panel: The completeness of SDSS as a function of apparent r -band magnitude computed from the division of statistics in the top panel. Figure taken from: <http://classic.sdss.org/dr7/products/general/completeness.html>

of 1.4 arcsec in the r -band, COMBO uses exposures of ~ 20 ks in the r -band with seeing below 0.8 arcsec and a smaller pixel scale. The superior sampling within the much smaller sky footprint enables high completeness with fainter sources. Figure 1.5 shows the distribution of galaxies by magnitude in each survey and the completeness of the SDSS measured by comparing the distributions. The SDSS galaxy sample is roughly complete for galaxies with $m_r < 18$ and less than 50% complete by $m_r = 22$. This is supported by recovery analysis of analytic galaxy models of known magnitudes that are inserted into Stripe 82 of the SDSS DR7 (Annis et al., 2014), which showed that DR7 is complete to $m_r \approx 19.0$ and has 50% completeness by $m_r \approx 21.3$. However, magnitude completeness does not imply volume completeness.

Even with the limitations described in this section, the SDSS is still considerably advantaged among modern galaxy surveys. Large, statistically representative samples of galaxies can be obtained readily in the form of direct imaging and spectrum or properties derived from the PHOTO pipeline and other post-processing routines. However, quantitative estimates of galaxy fluxes and morphologies from photometry must be carefully considered in light of these biases. The techniques used to achieve more accurate, precise, and unbiased measurements of galaxy properties are the topic of the next section.

1.3 Quantitative Morphologies

Visual classification schemes are successful in generating statistics of the distribution of galaxies within broad morphological classes. However, they do not provide measurements of key structural parameters that describe the physical components of galaxies. Furthermore, visual classification of large samples of galaxies is tedious and hard to reproduce. Major visual classification projects such as the Galaxy Zoo (Lintott et al., 2008) have helped to mitigate these problems in analyses of galaxies from large surveys such as the SDSS using vote-weighted morphological classification by the public (“citizen science”). Such projects have been instrumental in providing large samples of galaxy merger pairs (e.g. Darg et al. 2010), galaxies with bars (e.g. Masters et al. 2011), and various other large samples based on the classifications. Large visual classified samples are used to determine the connection between morphological classes and global properties such as colour (Tully et al., 1982; Bender & Moellenhoff, 1987; Dressler, 1984; Strateva et al., 2001; Bamford et al., 2009), mass (Calvi et al., 2012; Willett et al., 2015; Mortlock et al., 2013; Deng et al., 2013), lumi-

osity (Nakamura et al., 2003; Tempel, 2011; He et al., 2013), environment (Dressler, 1980; Bamford et al., 2009; Deng et al., 2013; Casteels et al., 2013), and star-formation (Deng et al., 2013; Schawinski et al., 2014; Willett et al., 2015). However, quantitative morphology tools are necessary to obtain more detailed information about the structural properties of galaxies within the samples created in visual classification schemes.

Quantitative morphological analyses fall into two main categories: non-parametric morphologies (Morgan, 1962; Abraham et al., 1994; Schade et al., 1995; Rix & Zaritsky, 1995; Abraham et al., 1996; Conselice, 1997; Bershady et al., 2000; Conselice et al., 2000b,a; Conselice, 2003; Takamiya, 1999; Papovich et al., 2003, 2005; Abraham et al., 2003; Lotz et al., 2004; Law et al., 2007; Freeman et al., 2013) and parametric morphologies (de Vaucouleurs, 1948, 1953; Sérsic, 1963; Kormendy & Djorgovski, 1989; Allen et al., 2006; Kormendy et al., 2009; Peng et al., 2010; Simard et al., 2011). Both are used in attempts to link the spatial distribution of stellar light (which traces the stars) to the evolutionary histories of galaxies.

1.3.1 Non-Parametric Morphologies

Non-parametric morphologies make no assumptions about the underlying functional form of galaxy surface-brightness distributions. The most common methods for measuring galaxy structure using non-parametric morphologies is with the Concentration-Asymmetry-Clumpiness (*CAS*) system (Conselice, 2003) and *Gini*- M_{20} (Lotz et al., 2004). The parameters are designed to capture the major features that signal for the underlying structure in galaxies (review by Conselice 2014). Many of these morphological quantities correlate with other parameters such as colour and tidally-induced disruptions in the stellar light distributions of galaxies. The principle advantage over parametric approaches is that computing these properties does not depend on the adequacy of optimization software or the accuracy of model components that are intrinsic to parametric morphologies. Furthermore, the biases that arise from choosing models with symmetric isophotes (invariant ellipticity with radial distance from the centre of a galaxy, for example) in parametric models do not affect non-parametric estimates. Recent work has shown the successes and limitations of non-parametric estimates to high-redshift galaxy morphologies and its general implications for low-surface brightness photometry (Thompson et al., 2015). Modern non-parametric morphologies demonstrate robustness to degradation of the S/N in galaxy images (Bershady et al.,

2000; Lotz et al., 2004; Graham et al., 2005).

The assumption underlying shapes for the surface brightness distributions of galaxies is often considered to be the greatest drawback of parametric morphologies relative to non-parametric statistics (Conselice 2003, 2014). Parametric models can be limited by the use of constant ellipticity isophotes which bias descriptions of galaxies with isophotes that change in ellipticity with radius from the galactic center. Because non-parametric morphologies do not assume underlying shapes of the surface-brightness profiles, they are free from the biases that arise in parametric fitting of galaxies with strong isophotal twists, for example. As with parametric morphologies, the choice of which pixels belong to a galaxy in crowded fields and the signal-to-noise (S/N) of the galaxy in the sky can significantly affect non-parametric statistics – making comparisons between galaxy morphologies in the low- and high-redshift universe challenging. Several modern non-parametric morphologies demonstrate robustness to degradation of the S/N in galaxy images (Bershady et al., 2000; Lotz et al., 2004; Graham et al., 2005). However, this robustness does not extend to all non-parametric statistics.

1.3.2 Parametric Morphologies

The principle motivation to use parametric morphologies is to attempt to decompose and describe the structural components of galaxies. While the decompositions require assumptions about the underlying surface-brightness profiles of the galactic components, they are advantageous with respect to non-parametric statistics in that the structural components are separable. Parameterizations of the structural components enable construction of global scaling relations and statistics for specific galactic components and their properties.

The most direct approach measuring the physical properties of the stellar components is by fitting analytic surface-brightness profile models to their projections in the sky. The technique was first employed by de Vaucouleurs (1948) who found that a power-law relation of the form $\exp(-r^{1/4})$ was successful in modelling the radial surface brightness profiles of elliptical galaxies. de Vaucouleurs (1959b) showed that this light profile is extendable to the dynamically hot bulges of galaxies, while the dynamically cold, flattened, and rotationally supported stellar discs were successfully modelled by a pure exponential $\exp(-r)$. These light-profiles are generalized by the sérsic profile (Sérsic, 1963). In 1-D, the sérsic profile is:

$$\Sigma(R) = \Sigma_e + \frac{2.5b_n}{\ln(10)} [(R/R_e)^{1/n} - 1] \quad (1.1)$$

where r is the projected distance from the galaxy's centre, $\Sigma(r)$ is the surface brightness at r , $\Sigma(r_e) = \Sigma_e$ is the surface brightness at the effective radius (which contains half of the total light from the model), and b_n is a constant that is defined in terms of the sérsic index n , which controls the shape of the light profile.

The empirical surface-brightness profiles of bulges and ellipticals and bulges ($n = 4$) and discs ($n = 1$) are distinct. The distinction between the profiles allows both components to be modelled simultaneously in either one or two dimensions. Parametric techniques that directly analyze the two-dimensional (2-D) images of galaxies are advantageous in that they make full use of the spatial information content that is accessible from photometry. This advantage is lost in 1-D techniques which either trace the light along individual axes or must azimuthally average image data about independently determined galaxy centroid coordinates. Furthermore, the goodness-of-fit in 2-D can be determined visually from the image residual (model subtracted from image) and statistically such as with the normalized χ^2_ν -statistic:

$$\chi^2_\nu = \frac{1}{N_{\text{dof}}} \sum_{x=1}^{n_x} \sum_{y=1}^{n_y} \frac{(f_{\text{image}}(x, y) - f_{\text{model}}(x, y))^2}{\sigma^2(x, y)} \quad (1.2)$$

N_{dof} is the number of degrees of freedom allowed by the model and image and is generally dominated by the number of pixels, $f(x, y)$ is the flux at a given pixel, and $\sigma^2(x, y)$ is the expected combined variance from noise contributions at each pixel.

The use of perfectly ellipsoidal models in 2-D parametric analysis is often cast negatively in comparisons with 1-D and non-parametric techniques because galaxies are not generally perfect ellipses in projection. However, Peng et al. (2010) argue that the benefits of the 2-D approach outweigh the drawbacks. Fundamentally, 2-D parametric techniques do not reduce spatial information content which often helps to break multiple-component profile degeneracies. The 2-D approach is more capable in handling situations where the bulge, disc, and other stellar components have semi-major axes that are not aligned. 1-D is unable to discern misalignment of the major axes of the stellar components without inherently utilizing 2-D information in some way (such as by performing the 1-D analysis along several axes). Ultimately, 2-

D parametric surface-brightness decompositions preserve the original pixel weights from the image data that are lost in procedures that are used to convert to 1-D.

The current generation of two-dimensional parametric quantitative morphology tools (GIM2D: Simard (1998); Simard et al. (2002), GALFIT: Peng et al. (2002, 2010), BUDDA: de Souza et al. (2004), MEGAMORPH: Vika et al. (2013), and others) encompass a broad range of model complexities, optimization algorithms, and weighting methods. GALFIT and MEGAMORPH allow arbitrary amounts of complexity in models including winding by Fourier and bending modes, hyperbolic rotation functions, and a wide variety of profiles. This is very useful to perform accurate and visually satisfying fits on the level of individual galaxies. However, the level of inspection and number of model components that may be necessary to generate an accurate model can become as tedious as visual classification and ultimately reduce the meaningfulness of any individual model component.

Each new layer of complexity or additional component in a model increases the total number of free parameters and stretches the amount of information that can be extracted consistently for large populations of galaxies. The bulge+disc model is often adopted as a trade-off between meaningful decompositions and a reasonable number of fitting parameters. It is noteworthy that while additional components (such as bars, spiral arms, and nuclei) are not modelled explicitly, higher-order complexity can be derived from the best-fitting bulge+disc models. Barred galaxies are effectively identified in histograms of the difference in position angle between the bulge and disc components, $\Delta\phi \equiv |\phi_b - \phi_d|$. Compact centrally bright sources (such as seen in quasar hosts) drive down the disc component and leave distinct residuals that can be identified using automated techniques. Automated analysis of the model residuals yields asymmetries that demonstrate ties with recent or on-going mergers. Furthermore, independently employing $(n_{\text{bulge}} = 4) + (n_{\text{disc}} = 1)$ fits, $(n_{\text{bulge}} = \text{free}) + (n_{\text{disc}} = 1)$ fits, and $n = \text{free}$ pure sérsic fits make it possible to determine the stability of each model component. Figure 2.8 shows a set of visually classified galaxies in the SDSS and their quantitative morphological decompositions using two different models.

Parametric surface brightness profile fits can be performed using various optimization techniques. Pixels that are allocated to the galaxy manually or using a deblending algorithm are used as weights in the fitting – where each pixel’s weight is often proportional to its relative intensity. Each iteration of testing a parametric model against a surface brightness distribution consists of subtraction of the model from the data and statistical evaluation of the residual. One common evaluation method is

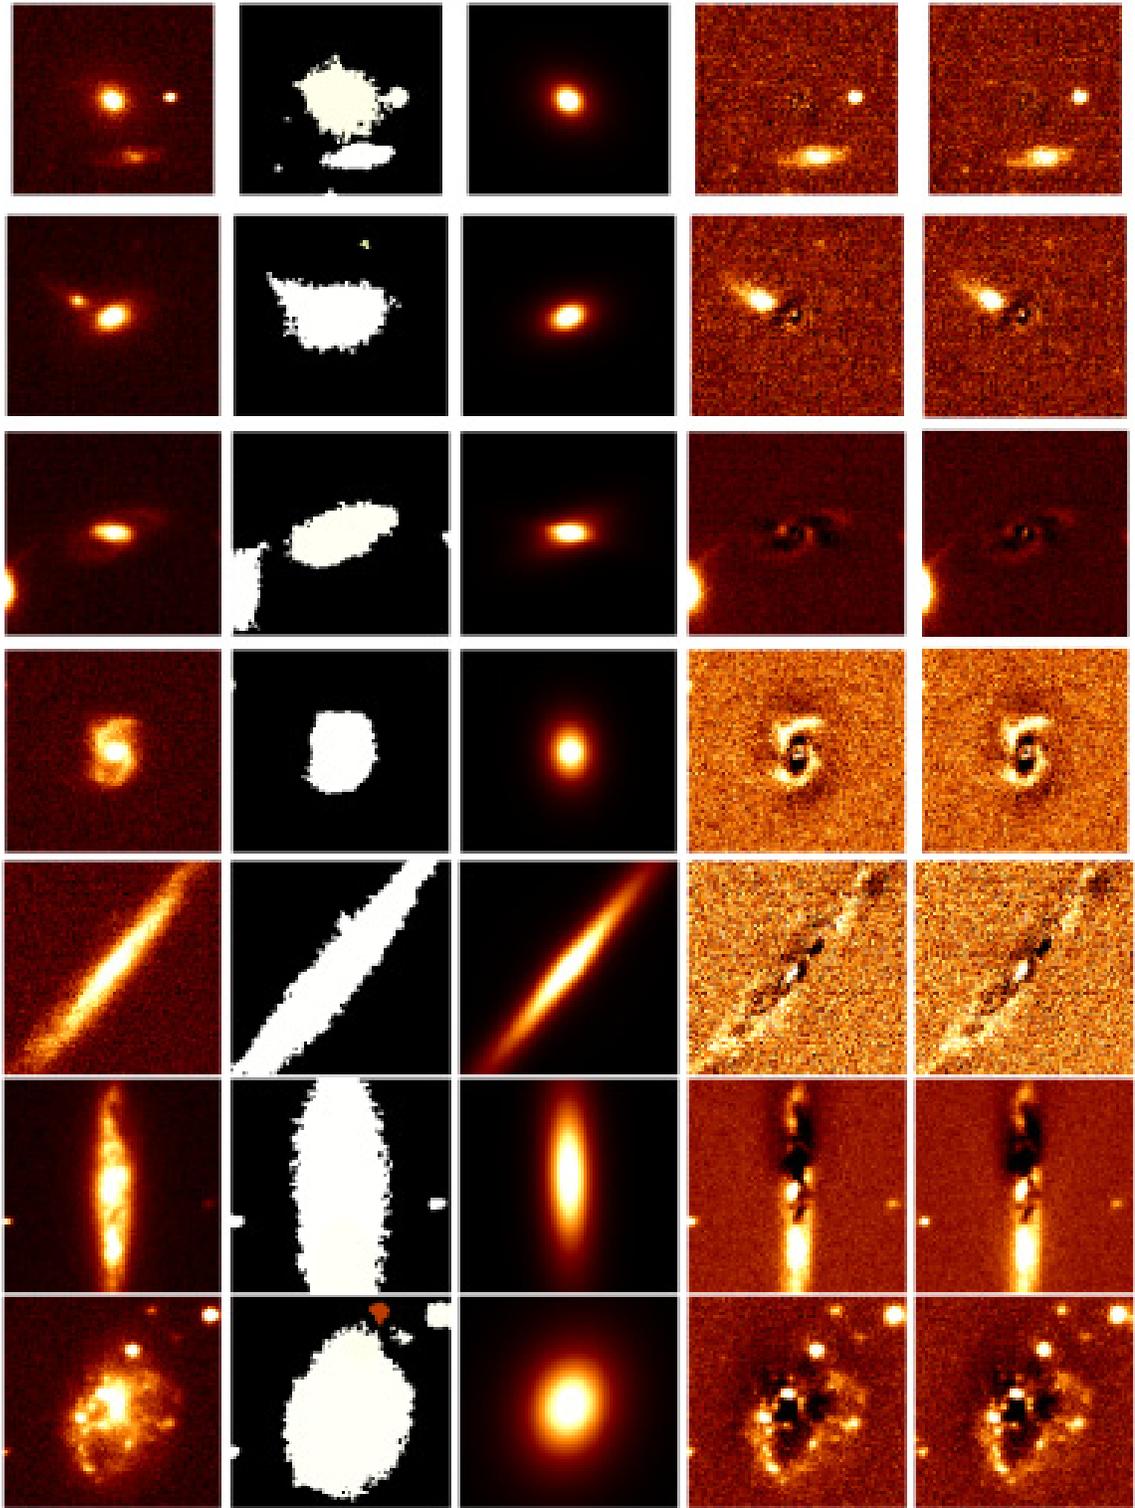


Figure 1.6: A mosaic of GIM2D fits to visually classified morphologies in the SDSS. Galaxies are classified as E, S0, Sa, Sb, Sc, Sd, and Irr from top to bottom. The SDSS r -band science image, SExtractor segmentation image (mask), best-fitting bulge+disc model between $(n_{\text{bulge}} = 4) + (n_{\text{disc}} = 1)$ and $(n_{\text{bulge}} = \text{free}) + (n_{\text{disc}} = 1)$, the $(n_{\text{bulge}} = 4) + (n_{\text{disc}} = 1)$ model residual, and the $(n_{\text{bulge}} = \text{free}) + (n_{\text{disc}} = 1)$ model residual from left to right. Figure taken from Tasca & White (2011).

minimization of the normalized χ^2_{ν} -statistic from Equation 1.2 (Levenberg-Marquardt minimization) that is employed by GALFIT among others. Using the gradients of χ^2_{ν} with respect to changes in the model parameters, an iterative search is performed for local minima in $\chi^2_{\nu}(\alpha_1, \dots, \alpha_N)$ where $\alpha_1 \rightarrow \alpha_N$ are the set of parameters in the model being optimized to fit the data. The search ceases when the change in χ^2_{ν} does not vary significantly over a set number of iterations – after which the optimization is said to be converged.

The local minimum in $\chi^2_{\nu}(\alpha_1, \dots, \alpha_N)$ -space that is obtained from iterative optimization is ideally the global minimum. However, algorithms that do not allow parameter steps that increase χ^2_{ν} invite the possibility of converging within a local minimum for certain initial sets of parameters. Markov-Chain Monte-Carlo (MCMC) techniques such the Metropolis-Hastings algorithm used by GIM2D mitigate this problem by assigning each parameter set that is sampled a posterior probability $P(\alpha_1, \dots, \alpha_N|D)$ for the model parameters given the data. The likelihood estimates that determine $P(\alpha_1, \dots, \alpha_N|D)$ are usually based on the value of $\chi^2_{\nu}(\alpha_1, \dots, \alpha_N)$ or some other statistic. In the Metropolis-Hastings MCMC code iterations, a candidate set of parameters is always sampled if the probability at the candidate position is higher than the current position, i.e. $P(\alpha_1, \dots, \alpha_N|D) > P(\alpha_{1,o}, \dots, \alpha_{N,o}|D)$. The difference between MCMC and direct minimization is that MCMC has a finite probability $P_{\text{step}} = P(\alpha_1, \dots, \alpha_N|D)/P(\alpha_{1,o}, \dots, \alpha_{N,o}|D)$ of stepping to lower probability positions than the current position. This enables escape from local minima and convergence on the global minimum with sufficient sampling. Furthermore, every sample is stored in memory which allows accurate estimates of the uncertainties on each parameter by marginalizing over every other parameter.

Parametric codes such as GIM2D and MEGAMORPH are designed to perform simultaneous multi-band fitting of galaxy images – for which the best-fitting model estimates are generally more stable (Simard et al., 2002; Häußler et al., 2013). Constraints across the bands are often placed on model parameters such as bulge and disc sizes, position angles of each component, and sérsic index, and centroid position that require them to be the same within a specified tolerance. In this way, each waveband will weight every other towards a common optimization of models. Simultaneous fits have N_{bands} –times the spatial sampling of an individual band which compensates for the increase in the number of free parameters. Simultaneous multi-band fitting of the bulge and disc components combined with individual fits in each band can be used to derive broad-band Spectral Energy Distributions (SEDs) for each stellar component

(e.g. Mendel et al. 2014). The SED of a source provides its relative intensity at each wavelength of light. Broadband SEDs are an averaging of the distribution of light produced in wavelengths covered by each band – producing a SED with resolution defined by the bandwidths and band intervals. The broadband SEDs can be used in tandem with surface-brightness decompositions to determine the relative colours and stellar masses of each component.

The success of any quantitative morphologies technique is sensitively tied to the method used to determine which pixels are allocated to sources (deblending). The result of deblending in an image is the construction of a segmentation map (see right panels of Figure 1.4). In general, the segmentation map delineates sources from the sky and each source from each other. While Lotz et al. (2004) provide a mechanism for eliminating the adverse effects of low S/N , crowding by neighbouring projected sources is a challenging bias to handle in non-parametric analyses. However, 2-D parametric morphologies are capable of obtaining accurate estimates for the properties of crowded galaxies with the right procedures. Simard et al. (2011) used the 2-D parametric surface brightness decomposition tool GIM2D to show the results from decompositions using several deblending strategies. Figure 1.7 shows 6 colour-magnitude diagrams using the difference in g - and r -band rest-frame magnitudes plotted against the rest-frame r -band magnitude, M_r . Galaxies are taken from the SDSS pair sample of Patton et al. (2011) which are selected for small projected separations and relative velocities of galaxy sources. The optimization of the bulge+disc models are performed in two ways: (SEP) Galaxy images in each band (g and r) are modelled separately – giving completely independent results; (SIM) g - and r -band models are optimized simultaneously and provide weight to each other.

The colour-magnitude diagrams in Figure 1.7 each contain two principal components: a tight red sequence (higher $g - r$) and blue cloud that is more broadly distributed. Galaxies generally form and evolve within the blue cloud until the star formation that is responsible for birthing young blue populations of stars is stifled by some mechanism – from whence they redden in colour as the blue stars die and passively evolve along the tighter red sequence. Comparison across each panel shows that the SDSS deblending and background measurement systematically offset the colours and magnitudes for a large number of galaxies. Relative to the SDSS Petrosian and Model magnitudes¹ the combination of SDSS deblending and single-band model

¹*Petrosian* magnitudes (Petrosian, 1976) are computed from aperture photometry. The Petrosian ratio \mathcal{R}_P is defined at the radius r from the object centroid where the local surface brightness within

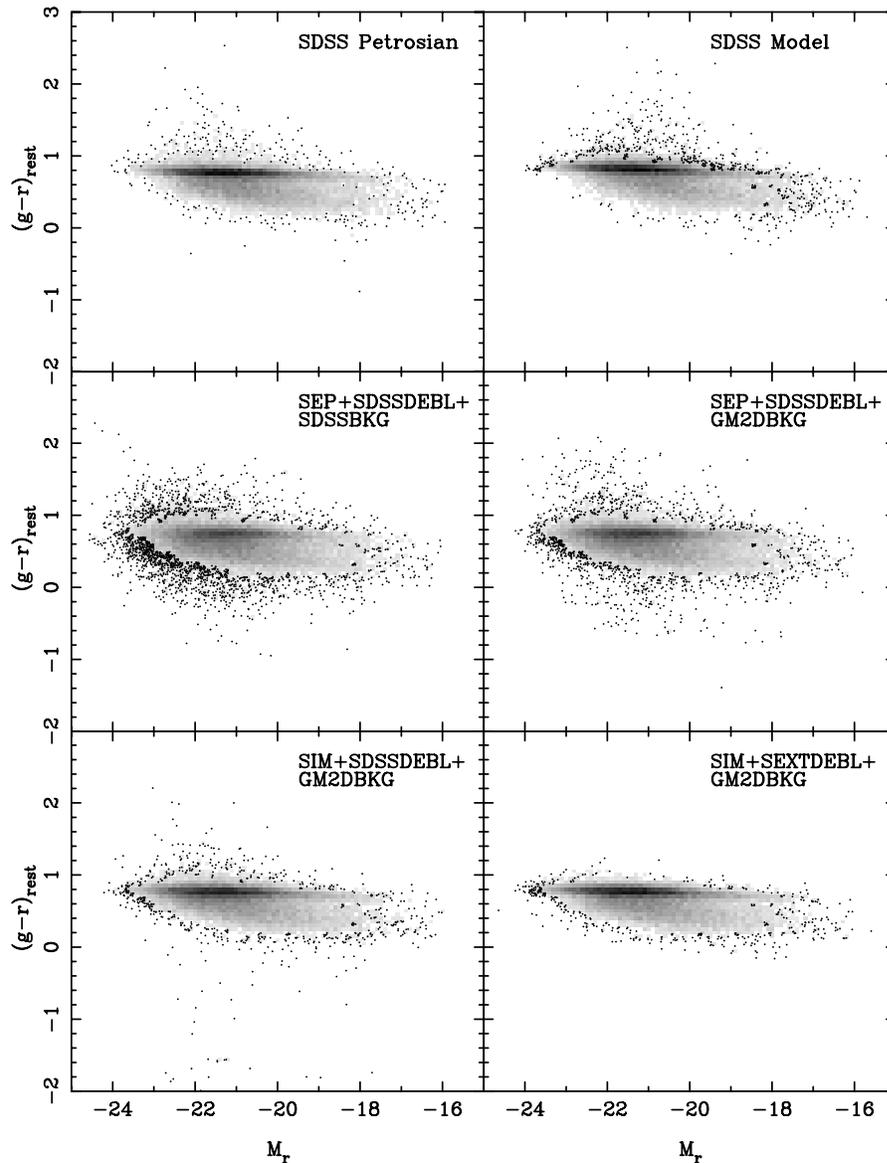


Figure 1.7: Colour-magnitude diagrams of galaxies in the galaxy pair sample of Patton et al. (2011). Top left: SDSS Petrosian. The Petrosian radius, r_P , is determined from circular aperture photometry as the radius at which surface brightness is equal to $0.2 \times$ the mean enclosed surface brightness $\mu(r_P) = 0.2\bar{\mu}(r < r_P)$. Top right: SDSS Model. Model magnitudes are computed from the best-fitting result from independent pure sérsic and de Vaucouleurs profile fitting. Both Petrosian and Model magnitude estimates use the SDSS deblending and sky measurement. Middle left: GIM2D magnitudes to galaxies with separate fitting in each band (SEP) using SDSS deblending (SDSSDEBL) and sky measurements (SDSSBKG). Middle right: SEP fitting with SDSSDEBL and GIM2D sky measurements (GM2BKG). Bottom left: GIM2D magnitudes from simultaneous multi-band fits (SIM) with SDSSDEBL and GM2DBKG. Bottom right: SIM fitting with SExtractor deblending and GM2DBKG. Grayscale is the same in each panel and bins containing only one point are replaced with scatter showing the position of the data. Figure taken from Simard et al. 2011.

optimizations produces erroneous scatter and a large fraction of outliers. However, the scatter is significantly reduced using simultaneous multi-band fitting with GIM2D and alternative deblending and sky estimation strategies. The simultaneous g - and r -band fits using SExtractor deblending and the GIM2D background estimation (bottom right panel) noticeably depopulates the outliers and improves the tightness and definition of the sequence and cloud relative to each other technique examined in the study. The scheme exhibited in the lower right panel is generally adopted the fiducial scheme in GIM2D analysis pipelines. The lack of red outliers in the fiducial scheme is particularly significant considering that these galaxies were ascribed as new population of unusually red pairs in the literature (see explanation by Patton et al. 2011). In particular, GIM2D has demonstrated the level of robustness for parameter estimates that is accessible with 2-D analyses in crowded images.

The quantitative parametrization of the galaxy components facilitates connections with the physical properties that are predicted from theory. However, accurate translation between the physical predictions from theory and the observationally derived properties from galaxies has limited the meaningfulness of these comparisons. Without realistic images of large numbers of galaxies from numerical simulations and theory, comparisons between observed and predicted galaxy populations are fundamentally indirect. Fortunately, recent advances in the scale and fidelity of numerical hydrodynamical simulations have made it possible to bridge the long-standing gap between analyses of galaxies using observational and numerical data.

a fixed annulus at r is a specified fraction of the mean surface brightness within r as described in Blanton et al. (2001) and Yasuda et al. (2001):

$$\mathcal{R}_P(r) \equiv \frac{\int_{0.8r}^{1.25r} dr' 2\pi r' I(r') / [\pi(1.25^2 - 0.8^2)r'^2]}{\int_0^r dr' 2\pi r' I(r') / (\pi r'^2)} \quad (1.3)$$

where $I(r)$ is the azimuthally averaged surface brightness profile. The Petrosian radius, r_P , is defined as the radius at which the Petrosian fraction $\mathcal{R}(r_P) = 0.2$ for the SDSS. The Petrosian magnitude is computed from the total flux contained within $2r_P$ of the galaxy centroid. Petrosian magnitudes are used to make consistent measurements for galaxy fluxes that do not vary strongly as a function of position and distance. *Model* magnitudes are computed from the best-fitting independent single-component exponential or de Vaucouleurs profile fits to the 2D surface brightness distribution of a galaxy.

1.4 Numerical Simulations

The limitation of the observable universe is that observers are confined to a static view of galaxy evolution. A dynamic view of the evolution of individual galaxies in the universe is required to make links between their observed properties and galaxy formation and evolution theory. Comparing populations of galaxies in different redshift ranges is the only way an observer can formally study their evolution. This is a very powerful technique for investigating changes to the statistical distributions of galaxy parameters across cosmic time. However, the ability to link individual high-redshift galaxies to local ones requires assumptions that often weaken their interpretive capacity.

The computational counterparts to observational galaxy surveys are numerical simulations. Simulations have evolved rapidly over the last few decades. The first applications of numerical simulations were used in approaches to the theory of stellar dynamics within small clusters of stars (von Hoerner, 1960; Toomre & Toomre, 1972). Simulations that model only gravitational interactions are known as N -body simulations and continue to be a staple of numerical astronomy because of their low computational overhead by modern standards (Springel et al., 2005; Teyssier et al., 2009; Kim et al., 2011; Angulo et al., 2012).

1.4.1 N -body Simulations

N -body simulations became particularly interesting following observations of the Andromeda Galaxy (M31) by Rubin & Ford (1970). Atomic hydrogen emission lines from the $3 \rightarrow 2$ electron transition ($H\alpha$) were identified in 67 locally bright regions containing young stellar populations (HII regions) and used to determine the rotational velocity of M31 out to large radii. They found that the rotation velocities were systematically larger than the rotation predicted from estimates of the stellar and gas mass distribution in the galaxy. Ostriker & Peebles (1973) used N -body simulations to show that the common morphologies of galaxies such as spirals and ellipticals were unstable configurations without support from a static, uniform distribution of mass that extended out to large radii and comprised a considerably larger total mass than the sum of the masses from the stellar particles. This “missing mass” could not be accounted for by the stellar and gaseous matter that is accessible via the electromagnetic spectrum. The inclusion of dark matter to galaxy formation theory brings the theoretical rotation velocities into agreement with observations. The

extended distributions of dark-matter in which galaxies are embedded are known as the galactic halo (or dark-matter halo) (White & Rees, 1978; Fall & Efstathiou, 1980; Frenk et al., 1985; Blumenthal et al., 1986; Dekel & Silk, 1986; Kauffmann & White, 1993). Current observational evidence for dark matter is indirect, but so far it is the most consistent model within the framework developed by observational and theoretical cosmology (e.g. Rubin & Ford 1970; Faber & Jackson 1976; Rubin et al. 1980; Aubourg et al. 1993; Moore 1994; Percival et al. 2007; Komatsu et al. 2009 or reviews by Refregier 2003; Bertone et al. 2005). Dark matter appears to be the dominant form of matter in the universe and provides the principle basis for structure formation (Peebles, 1980; Davis et al., 1982; Blumenthal et al., 1984; Davis et al., 1985; Gott et al., 1986; Bahcall, 1988). The strong candidacy of dark matter as a real and fundamental component of the mass-energy budget in the universe, its dominance over baryonic matter on large scales, and the tentative exclusivity of its interactions to gravity make it ideally suited for N -body simulations.

1.4.2 Semi-Analytic Models

Modern dark-matter only (DM-only) N -body simulations are used to make statistically representative predictions for dark matter structure, halo merger trees, the mass function of halos, etc. in large cosmological volumes (Springel et al., 2005; Teyssier et al., 2009; Kim et al., 2011; Angulo et al., 2012). Semi-analytic models (SAMs) are used to link dark halos with galactic components using the skeleton and growth-histories provided by N -body simulations (e.g. Kauffmann et al. 1993; Emsellem et al. 1994; Baugh et al. 1996; Somerville & Primack 1999; Cole et al. 2000; Hatton et al. 2003; Bower et al. 2006; Somerville et al. 2008; Guo et al. 2011). However, comparisons between the physical parameters predicted by the models and observationally derived properties are fundamentally indirect. Mock observations of galaxies produced in SAMs bridge this gap by converting the simulation products to catalogs of quantities that are familiar to observational galaxy astronomy (e.g. Overzier et al. (2013)). Using realistic images and spectra of galaxies, the predictions from theory and observational evidence are placed on level ground. Observational tools can be used to analyze the simulation and observation products consistently. Still, semi-analytic models are fundamentally limited by “recipe-based modelling” of the baryonic components.

Hydrodynamical simulations are designed to mitigate the constraints on SAMs by

including the baryons and their physics directly within the simulations. The complexity of physics that is necessary to accurately and explicitly model the baryons is the principle limit to the volume and resolution of these simulations due to the high computational expense relative to N -body simulations. However, current numerical hydrodynamical simulations that explicitly track the co-evolution of dark and baryonic matter comprise of sufficient scales and resolution that enable comparisons with observational galaxy surveys (Illustris: Vogelsberger et al. 2014b, EAGLE (Evolution and Assembly of GaLaxies in their Environments): Schaye et al. 2015).

1.4.3 Hydrodynamical Simulations

The advantage of numerical hydrodynamical simulations is that the properties of the simulation components (stellar, gaseous, or dark) are determined explicitly by their complete dynamical, chemical, and thermodynamic histories. This affords several advantages over more indirect results based the post-processing allocation of galaxies properties and morphologies to DM-only N -body halos. The properties of the stars and gas in hydrodynamical simulations can then be used in radiative transfer codes to probabilistically determine the projected spatial distribution of stellar light that would be observed with a camera in an arbitrary location (e.g. SUNRISE: Jonsson 2006). The images that are produced from the stellar particles within a simulation are called “synthetic stellar images” or “stellar mocks.” The transfer codes can be run for any snapshot of the evolution of the simulation. Therefore, mock images of individual galaxies throughout their formation histories can be used to directly tie the observed evolution of morphologies to their mass accretion, interaction, environmental, and star formation histories, for example.

Figure 1.8 shows an example of how star-formation rates, colours, and morphologies can be connected through inspection of their histories in the Illustris simulation. Figures 1.9 and 1.10 show colour-composite mock observations for galaxies in Illustris for samples of blue and red galaxies. The diversity of morphologies within the blue and red populations is a particular triumph of the current generation of cosmological hydrodynamical simulations. These simulations provide an unprecedented level of fidelity with respect to the observed morphologies of galaxies. Figure 1.11 shows a Hubble diagram composed of mock observations of Illustris galaxies and a comparison of a real and mock image using the Hubble Space Telescope (HST) Ultra-Deep Field (UDF). The stellar mocks enable image-based comparisons that show how well

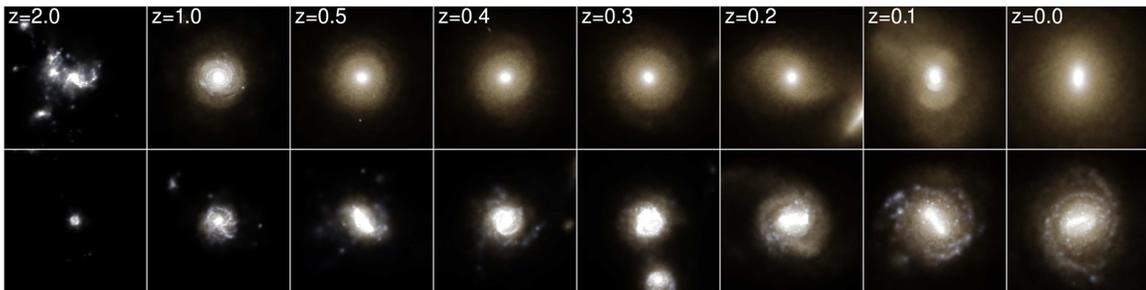


Figure 1.8: Mock observations of two selected galaxies from the Illustris Simulation from redshift, $z = 2$. The mocks are colour composites of the SDSS *gri* optical bands viewed from the rest-frame of the galaxy at each epoch. The top galaxy evolves to an elliptical morphology by $z = 0$ where the bottom has a spiral morphology with the visual presence of a bar. The true star-formation rates are known from the simulation. At high redshift, both galaxies have high on-going star formation. Star-formation has all but ceased by $z = 0$ for the galaxy in the top panel. The galaxy in the bottom panel has a high star formation rate at $z = 0$. The star-formation properties are correlated with the colours in the mock-observations – as they are in the real universe. Blue galaxies typically have higher specific star formation rates (rate of production of stars normalized by the total stellar mass of the galaxy) than red galaxies. Both of the selected galaxies experience recent mergers. The mergers lead to the formation of shell-like stellar density structures in the red galaxy and a central bar in the blue galaxy. Figure taken from Vogelsberger et al. (2014b).

simulated galaxies reproduce the properties of observed ones. These tests provide insights to the physical mechanisms that drive or suppress structure formation in galaxy formation and evolution theory.

1.4.4 Mock Observations

Realistic images from stellar mocks facilitate direct comparison of the simulation products with observations (e.g. Abadi et al. 2003a; Scannapieco et al. 2010; Brooks et al. 2011; Christensen et al. 2014; Snyder et al. 2015). Quantitative morphologies are ideally suited to this task because the same codes can be used for both real and mock observations without the interpretive biases that are intrinsic to visual classifications. It is important that the biases from the parametric model in the decomposition are separable from the differences in the fundamental parameters that describe the galaxies’ components. Therefore, decompositions that do not include an arbitrary number of components have potentially the most interpretive power – as they have the least sensitivity to biases from the choice of model. Pure sérsic and bulge+disc

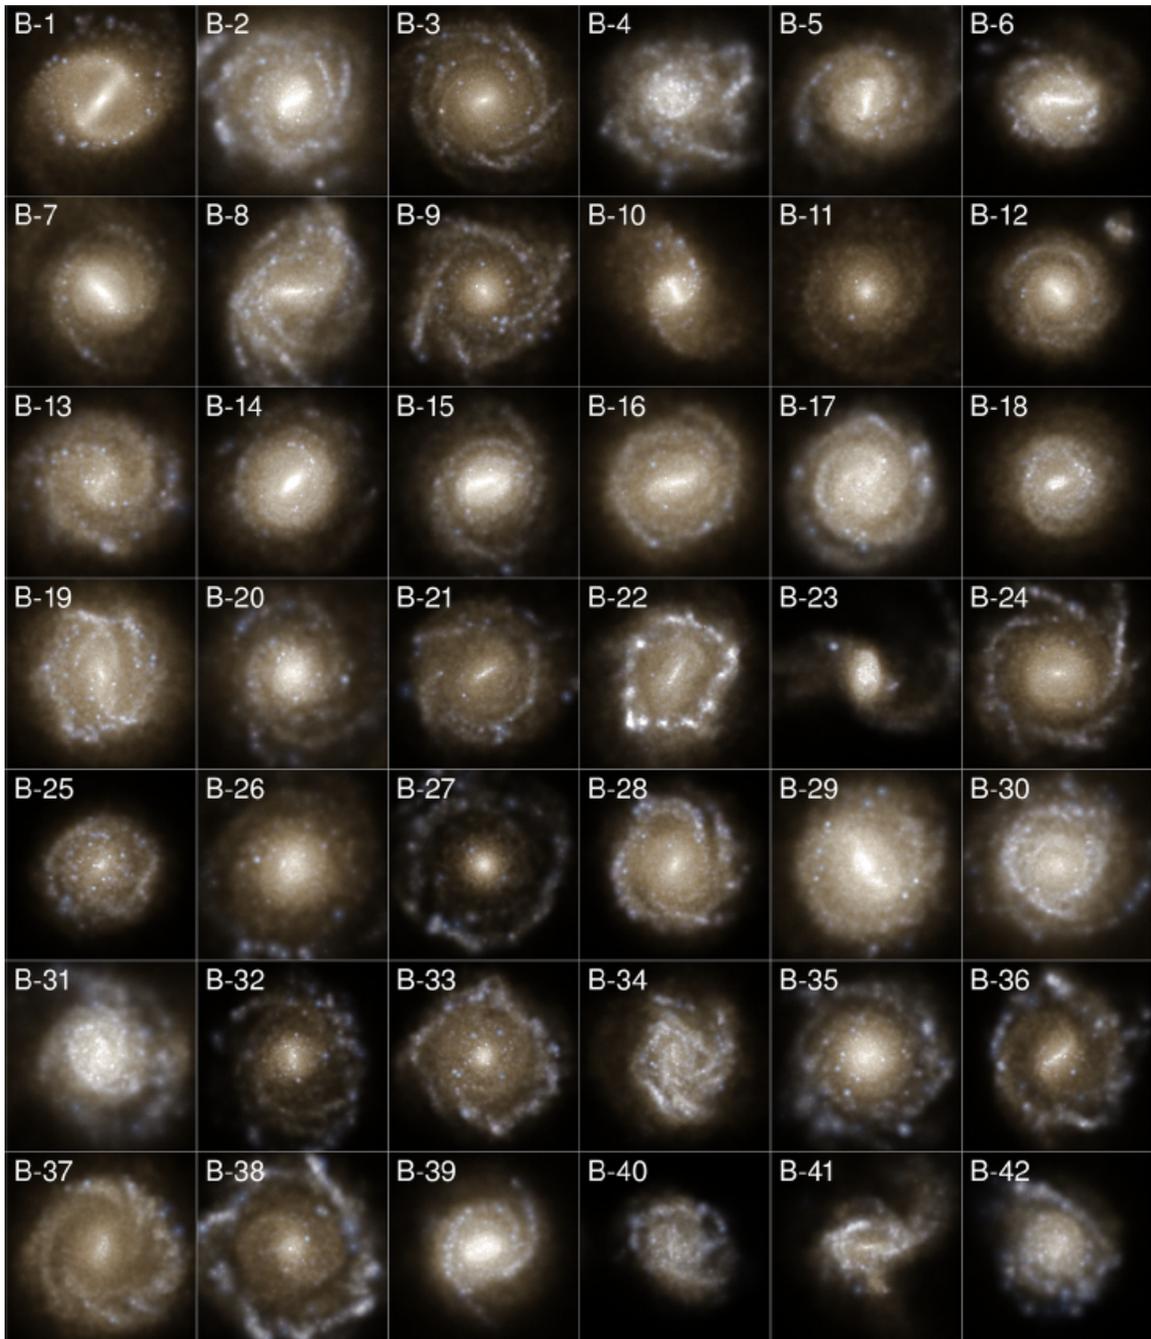


Figure 1.9: Mock observations of a sample of blue galaxies (shown face-on) from the Illustris Simulation. The mocks are colour composites of the SDSS *gri* optical bands. Most of the blue galaxies are disc-like. Tight blue clumps are regions of recent or on-going star-formation. Figure taken from Vogelsberger et al. (2014b).

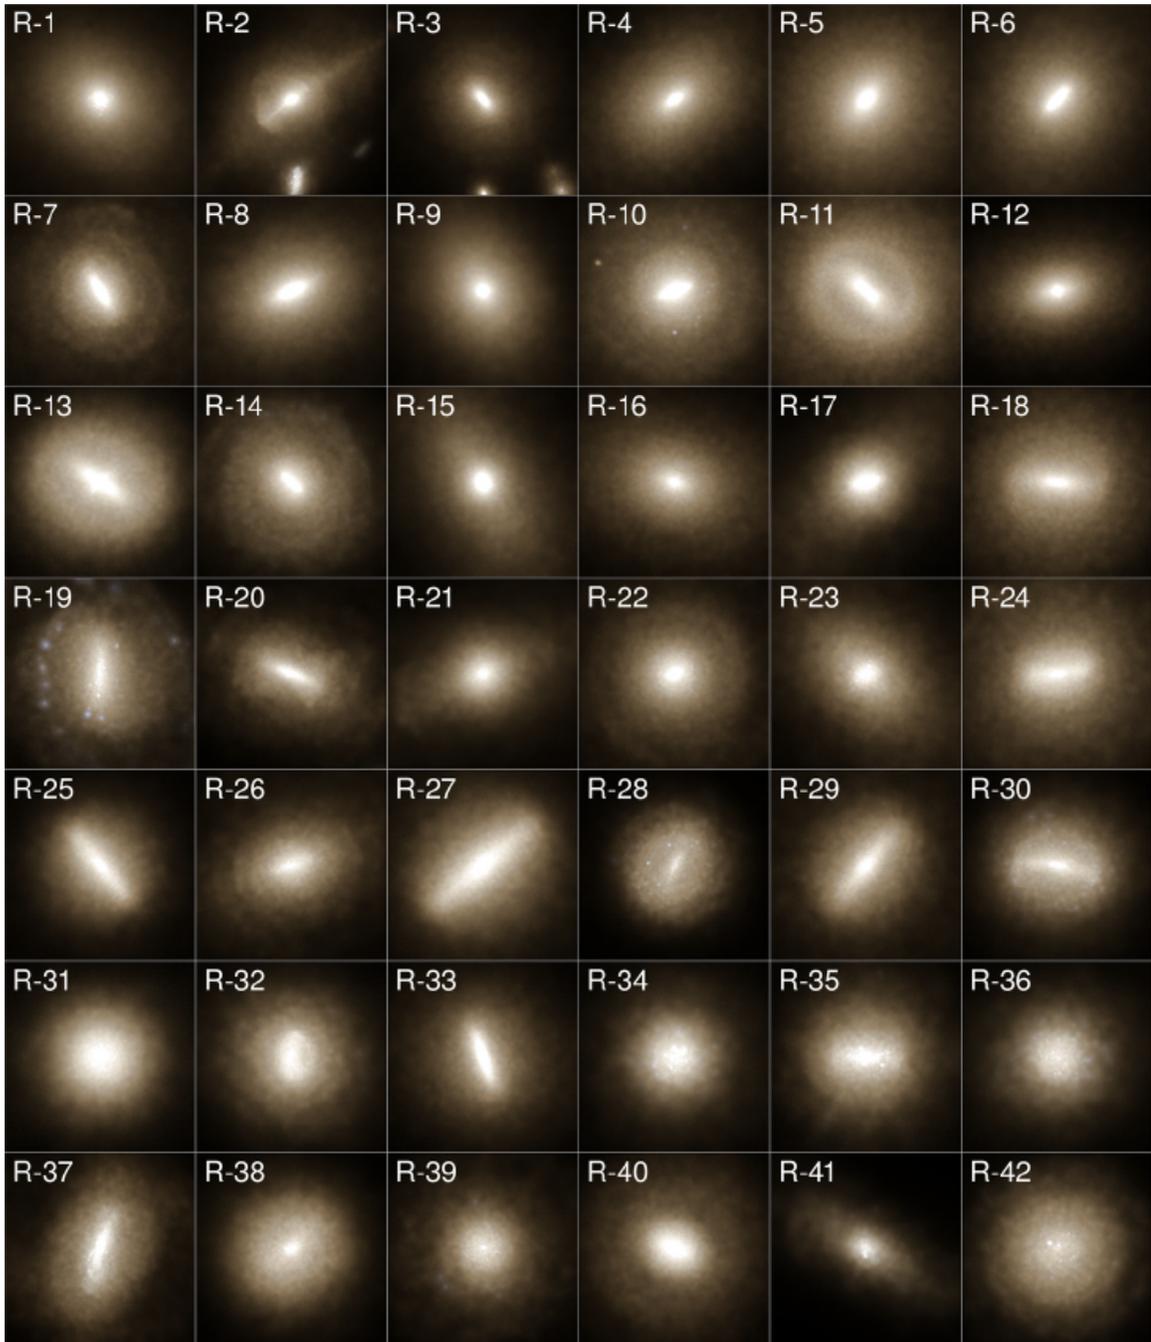


Figure 1.10: Mock observations of a sample of red galaxies from the Illustris Simulation. The mocks are colour composites of the SDSS *gri* optical bands. Most of the red galaxies are elliptical in morphology and have little on-going star-formation. Figure taken from Vogelsberger et al. (2014b).

surface brightness decompositions are the simplest models that maintain the flexibility required to link the parameters to the physical components of a galaxy. The sizes of galaxy populations produced by the current generation of hydrodynamical simulations also make automated fitting algorithms more attractive as long as they have low sensitivity to the initial guess and may escape from local minima in parameter space.

The crucial tenet of a fair comparison between the simulated and observed galaxies is ensuring that consistent and comprehensive observational realism is added to mocks. A synthetic image of a galaxy from a simulation contains no realism and no noise (apart from the small amount from the photon propagation scheme in the radiative transfer code). Therefore, the realism that burdens real galaxies must be added to the synthetic images either artificially or via insertion into real images. The dominant biases that need to be reproduced are summarized in §1.2.2. Principally, the realism that must be achieved is a statistically consistent distribution of resolution, signal-to-noise, and crowding by external sources between real and simulated galaxy populations. The most direct way to add this realism is by insertion of the mock galaxy flux into real observational images.

This thesis details a new method for image-based comparisons of the morphological properties of galaxies from simulations and observations. Quantitative morphologies using the GIM2D software are employed on galaxies from the Illustris simulation with a comprehensive suite of observational realism considerations. The realism ensures that analyses and comparisons of simulated and observed galaxy populations are consistent and fair – which is central to the interpretive power of their results. Synthetic images of galaxies from the Illustris simulation are inserted directly into SDSS images in a way that the biases from realism affect the derived morphological parameters in a statistically equivalent fashion. A detailed characterization of these biases is performed to quantitatively assess the effects of realism on derived galaxy properties. Once biases have been carefully characterized, it is then possible to disentangle them from real physical differences between Illustris and observations using well-known relations. The relation between size and luminosity and the relation between bulge-to-total fraction and total stellar mass are the two principle comparisons that are performed.

The automated and objective nature of the methodology facilitates consistent estimates of fundamental parameters that describe observed and simulated galaxy morphologies. Even though simplified models such as the ones used by GIM2D may

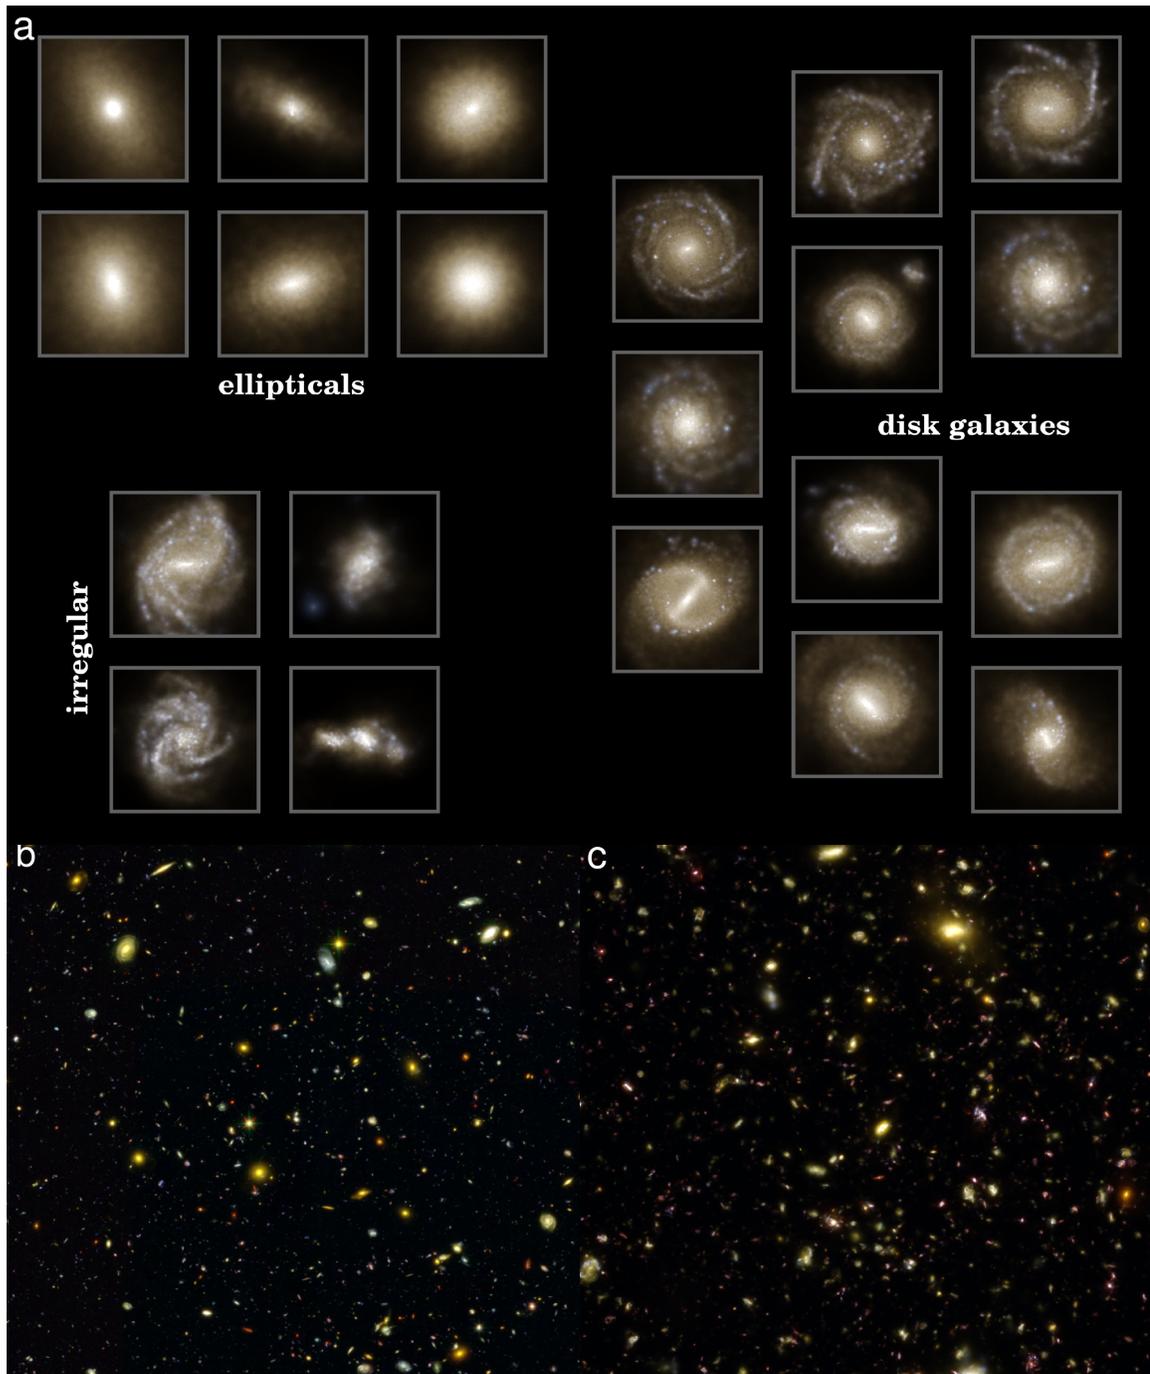


Figure 1.11: Galaxy morphologies and populations produced in the Illustris simulation. (a) Galaxy stellar mocks showing the diversity of morphologies produced by Illustris in the shape of a Hubble diagram. Stellar mocks are colour composites of the *gri* optical bands. (b) HST Ultra-Deep Field image (2.8 arcmins a side) using colour composites of the BZH bands. (c) HST UDF mock observation of the Illustris galaxy population created by partitioning the Illustris volume to create a light-cone. Figure taken from Vogelsberger et al. (2014a).

not fully reproduce the surface brightness distributions of more complex structures, this limitation should have the same impact on measured mock and real galaxy image properties – provided that the same models are used for both. Inversely, the absence, stifling, or over-production of morphological structures in the simulations can also be identified through comparisons of the decomposition results to observations. The discrepancies and similarities between the structural parameters derived from photometry provide insights that stand to advance the development of even higher-fidelity numerical techniques and models that enter the simulations. Image-based analyses of real and simulated galaxies enables a new comparison between theory and observations – a conjunction that stands to provide a wealth of feedback to the communities of computational and observational astronomy alike.

Chapter 2

Galaxies in the Illustris simulation as seen by the Sloan Digital Sky Survey - I: Bulge+disc decompositions, methods, and biases.

Connor Bottrell,¹ Paul Torrey,² Luc Simard,³ and Sara L. Ellison¹

¹University of Victoria, Department of Physics and Astronomy, Victoria, British Columbia, V8P 1A1, Canada

²Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

³National Research Council of Canada, Herzberg Institute of Astrophysics, 5071 West Saanich Road, Victoria, British Columbia, V9E 2E7, Canada

Abstract

The current generation of large-scale hydrodynamical cosmological simulations are valuable instruments for understanding the formation and evolution of galaxies. However, the crux in the interpretive power for such simulations is fundamentally in whether they are able to reproduce the observed properties of galaxies. Furthermore, it is essential that comparisons are fair – which is enabled by deriving the properties of simulated and real galaxy populations on level-ground. We present a new, image-based method for comparing the structural properties of galaxies produced in hydrodynamical simulations to observed galaxies. First, we apply an unprecedentedly rigorous suite of observational realism to synthetic images of simulated galaxies that statistically matches the biases for image-based analyses of real galaxies in the Sloan Digital Sky Survey. We then perform 2-D bulge+disc decompositions of simulated galaxies using the GIM2D pipeline of Simard et al. (2011). We apply our methodology to the $z = 0$ synthetic image catalog of galaxies from the Illustris simulation (Torrey et al., 2015). In this chapter, we detail our realism suite, methods, and decomposition catalogs. We quantitatively characterize observational biases on our decomposition results to enable meaningful insights in comparisons with real galaxies. We find that our decomposition results are largely robust to observational biases affecting decompositions of real galaxies. However, we identify several diffuse galaxies in Illustris are prone to internal segmentation in the deblending (roughly 30% of the full population) – leading to systematically reduced flux estimates by up to a factor of 6, smaller half-light radii by up to a factor of ~ 2 , and erroneous estimates for the bulge-to-total fraction. Nonetheless, our methodology and bulge+disc decomposition results enable unbiased comparisons with observations. The adaptive and scalable nature of our methods to comparisons using various observing instruments and simulations facilitate important conjunctions between observational and computational galaxy astronomy that stand to provide a wealth of feedback to each community alike.

2.1 Introduction

A range of tools have been developed in the past several decades to model the formation and evolution of galaxies with the overarching goal of reproducing the observed properties of galaxies and their populations in nearby and distant epochs of the universe (see Somerville & Davé 2015 for a recent review). Validation of the models is fundamentally determined by the accuracy of the model products with respect to observational constraints (e.g. Abadi et al. 2003a; Brooks et al. 2011; Agertz et al. 2011; Guedes et al. 2011; Christensen et al. 2014; Agertz & Kravtsov 2015; Furlong et al. 2015; Lange et al. in prep). The growth in observational data from modern observational large galaxy redshift surveys such as the SDSS (Eisenstein et al., 2011), 2dF (Folkes et al., 1999), DEEP2 (Newman et al., 2013), CANDELS (Grogin et al., 2011) and that is expected from forthcoming projects such as the LSST (project description in LSST Science Collaboration et al. 2009) provides an increasingly precise and complete test-bed against which the model products must be benchmarked.

The most direct way of deriving observable properties of galaxies from theoretical predictions is to numerically track the coevolution of dark and baryonic matter in hydrodynamical cosmological simulations (e.g. Katz et al. 1992, 1996; Weinberg et al. 1997; Murali et al. 2002; Springel & Hernquist 2003; Kereš et al. 2005; Ocvirk et al. 2008; Crain et al. 2009; Croft et al. 2009; Schaye et al. 2010; Oppenheimer et al. 2010; Vogelsberger et al. 2012). While limitations on spatial and mass resolution require certain prescription models that operate on the *sub-grid* level (see §3 of Somerville & Davé 2015 and references therein) that are often similar to those used in semi-analytic models (SAMs) of galaxy formation, the explicit inclusion of hydrodynamics in the formation of structure and galactic dynamics facilitates significantly broader comparisons with observations. In particular, tracking the dynamics of both the dark and baryonic matter self-consistently to small spatial scales allow predictions to be made about the internal structure of galaxies including the distribution of gas (Kereš et al., 2012; Torrey et al., 2012) and the structures formed from the stellar components (Abadi et al., 2003a,b; Governato et al., 2004; Agertz et al., 2011; Sales et al., 2012; Marinacci et al., 2012). Numerical hydrodynamical simulations place galaxy formation and evolution theory in a complete evolution-based cosmological context and are therefore an important tool for interpreting the observed properties of galaxies. In particular, the latest generation of high-fidelity multi-scale cosmological hydrodynamical simulations (e.g. Illustris: Vogelsberger et al. 2014b; EAGLE:

Schaye et al. 2015; FIRE: Hopkins et al. 2014; APOSTLE: Sawala et al. 2016; BAHAMAS: McCarthy et al. 2016) are designed specifically to allow straight-forward comparisons to observational constraints on the scales of individual galaxies and full galaxy populations.

Explicit tracking of the birth mass, chemical evolution, mass loss, ages and motions of stellar populations within galaxies in hydrodynamical simulations allows the assignment of a full spectrum to each stellar population at every time-step. With this information, radiative transfer codes can be used to project stellar light from populations of stars within these galaxies and create synthetic images that provide realistic representations of what a simulated galaxy would look like if it had been observed with a real photometric or spectroscopic instrument (and at any stage in its evolution) (Jonsson 2006; Li et al. 2008; Baes et al. 2011; Robitaille 2011; see Steinacker et al. 2013 for a recent review). The tools for converting hydrodynamical simulations into synthetic images bridge the long-standing gap between theory and observations to the point that fully consistent analysis of the data products of each camp are possible. While data products of SAMs can also be formatted in ways that are accessible to conventional observational analysis software (e.g. the Millennium Run Observatory (MRObs): Overzier et al. 2013), hydrodynamical simulations provide a more explicit, evolution-based framework for comparisons with observations. However, the crucial requirement of a fair comparison between the properties derived from observed galaxies and simulations is that the same biases affect the interpretation of their results (Conroy et al., 2010; Hayward et al., 2012, 2013a,b). Furthermore, the apparent disparity between structural estimates, such as bulge- and disc-to-total ratios (e.g. Scannapieco et al. 2010), derived from photometric decomposition analyses of synthetic images and estimates from the orbital properties obtained directly from the simulations also demonstrates that consistency in methodology is fundamental to the interpretative power of such comparisons.

In perhaps the first statistically relevant, image-based comparison between the properties of galaxies from a hydrodynamical simulation and observations, Snyder et al. (2015) (S15) obtained non-parametric optical morphology estimates of synthetic images of galaxies from the Illustris simulation (Torrey et al., 2015) (T15) and demonstrated the diversity of bulge, disc, and composite morphologies produced therein. Two crucial elements of this work must be noted: (1) the application of observational realism by inserting the synthetic images directly into real fields; (2) consistent methods of computing the morphological estimates. While the realism

provided by S15 was incomplete (which we correct for by including the statistical biases from crowding, sky brightness, and PSF resolution in our observational realism suite), they showed that their non-parametric morphology estimates for simulated galaxies roughly occupied the same space as real galaxies – an unprecedented and statistically meaningful success in the structural comparison of galaxies from hydrodynamical simulations and the real universe. While several puzzles are highlighted by the study, the morphological similarities shown between simulated and observed galaxies in their proof-of-concept analysis motivates a complete and detailed study of the structural morphologies of galaxies produced in simulations that is realized in a fully observational context.

In this first paper in our series, we describe our methodology and data products that facilitate consistent, image-based comparisons between observationally accessible structural properties derived from real galaxies and galaxies evolved in hydrodynamical simulations. In particular, we derive structural estimates from SDSS renderings of the synthetic galaxy image catalog from the Illustris simulation (T15) that are consistent with SDSS analyses using the GIM2D parametric surface-brightness decomposition analysis tool. We apply a comprehensive suite of observational realism to the Illustris synthetic images to ensure that the same biases in resolution, signal-to-noise, and projected crowding affect both model and real galaxies. Furthermore, we provide a detailed characterization of these biases by conducting several “mini”-experiments that quantify the random and systematic errors associated with these biases and identify their correlations with morphology. We also investigate the biases associated with projected line of sight on the morphological parameters. Questioning the adequacy of the synthetic images themselves, we examine the biases associated with post-processing choices of how stellar light is propagated from full unresolved stellar populations and their effect on structural estimates. Having understood the limitations of parameter recovery, biases and uncertainties, in Bottrell et al. (in prep) we present a detailed comparison of the morphologies of galaxies in Illustris compared with those in the SDSS.

The paper is organized as follows. Section 2.2 provides a basic description of the simulation products and synthetic images and a detailed description of our observational realism suite. Section 2.3 describes the application of our 2D parametric quantitative morphologies to mock-observed galaxies and the resulting catalogs. The biases on structural parameters from realism are investigated and discussed in Section 2.4. We adopt cosmological parameters that are consistent with WMAP-9 measure-

ments in a Λ CDM cosmogony: $\Omega_m = 0.2726$; $\Omega_\Lambda = 0.7274$; $\Omega_b = 0.0456$; $\sigma_8 = 0.809$; $n_s = 0.963$; and $H_0 = 100h \text{ km s}^{-1}\text{Mpc}^{-1}$ where $h = 0.704$ (Hinshaw et al., 2013).

2.2 Simulated Galaxies

In this section, general descriptions are provided for the Illustris simulation and the production of synthetic stellar mock galaxy images as context for a full description of our methods. We refer the reader to the source literature found in Vogelsberger et al. (2014a) and Torrey et al. (2015) for more complete descriptions and detailed discussions of the properties of the Illustris simulation and stellar mocks, respectively, that are briefly summarized in Section 2.2.1 and Section 2.2.2. We provide a complete description of our image realism considerations in Section 2.2.3 and emphasize the continuity between the realism we provide and the biases affecting analyses of real galaxies in the SDSS (Simard et al., 2011) (S11).

2.2.1 Illustris Simulation

The Illustris simulation is a large-volume cosmological hydrodynamical simulation (Vogelsberger et al., 2014a,b; Genel et al., 2014) that employs a set of physics and feedback models that includes procedures for treating the sub-resolution ISM, gas cooling (including primordial and metal-line cooling), star-formation, stellar evolution and enrichment, kinetic stellar feedback from supernovae, black-hole seeding and merging procedures, and AGN feedback (radiative, radio-mode, and quasar-mode) among others. For details of the simulation models, we refer the reader to Vogelsberger et al. (2013, 2014b,a). The parameters of these physical models – specifically the free parameters in star-formation feedback models – are tuned to reproduce the stellar mass function and stellar mass-halo mass relation at $z = 0$ and global star-formation rate density statistics across cosmic time. Genel et al. (2014) show that this result is realized through comparison with observations as well as the non-tuned reproduction of the stellar mass and halo occupation functions to high redshift predicted in Torrey et al. (2014).

The diversity of the redshift $z = 0$ galaxy morphologies in Illustris is demonstrated by Snyder et al. (2015) (S15 who performed a non-parametric quantitative morphologies analysis on the synthetic images of galaxies with stellar masses $9.7 < \log M_*/M_\odot < 12.3$. The synthetic images in S15 are constructed using the

SUNRISE radiative transfer code (Jonsson, 2006; Jonsson et al., 2010) and have added realism in the form of SDSS skies and convolution with a Gaussian point-spread function (PSF) with resolution set by the average seeing for SDSS images for each optical band. S15 showed that Illustris contains a variety of morphological structures at $z = 0$ including spirals, ellipticals, and irregulars. In particular, the presence of rotationally-supported disc galaxies in coexistence with elliptical populations is an interesting result. It is not only relevant to the capacity with which the Illustris volume can be compared with the observable universe, but also in resolving the long-standing problem of general angular momentum deficit, high central concentrations, and unrealistic rotation curves in disc-formation experiments (Navarro & Steinmetz, 2000; Scannapieco et al., 2012). The implication is that the general failure of realistic disc-formation in previous simulations is not in conflict with galaxy formation/evolution theory or Λ CDM, but rather a symptom of the inadequacy of numerical algorithms to this task (Torrey et al., 2012).

Current large-volume hydrodynamical simulations such as Illustris and EAGLE in producing diverse galaxy populations have been shown to broadly reproduce fundamental relations and morphologies observed in the real universe (e.g. Snyder et al. 2015; Furlong et al. 2015; Crain et al. 2015; Trayford et al. 2015, 2016; Lange et al. in prep.). Their new levels of fidelity make them well-suited for experiments comparing the observationally accessible properties of simulated galaxies with those of real ones. A quantitative morphology analysis that employs bulge+disc surface brightness decompositions and will yield detailed structural information about the physical components is ideally fashioned for this task. The key to the interpretive power in relating the results with observations is the condition that the same realism that exists for real galaxies is provided to the simulation sample. We describe our considerations for creating realistic images that facilitate consistent comparisons with observations in the following sections.

2.2.2 Synthetic Galaxy Images

We employ the full synthetic galaxy image catalog of Torrey et al. (2015) (hereafter T15) which is comprised of 6891 galaxies that meet the stellar particle resolution standard of $N_{\star} = 10^4$ stellar particles. The particle resolution cut places a lower limit on the total stellar mass of galaxies in our sample at $\log M_{\star}/M_{\odot} \gtrsim 10$. Systems that do not satisfy this criterion generally have poorer spatial resolution. Consequently,

imaging and morphological analysis of simulated galaxies that do not satisfy this criterion are less meaningful for comparisons to real galaxy populations for these masses. All galaxies in our sample are taken from the redshift $z = 0$ snapshot of the simulation and their surface brightness distributions are artificially redshifted to $z = 0.05$. The corresponding luminosity distance and angular scale at this redshift, assuming the cosmological parameters stated at the end of Section 2.1, are $d_L = 221.3$ Mpc and 0.973 kpc/arcsecond.

Stellar Light & Surface Brightness Smoothing

The methods for producing synthetic SEDs and idealized simulated galaxy images are described in detail in T15. In short, stellar particles inherit an initial mass $M_\star \approx 1.3 \times 10^6 M_\odot$, time of birth, and a metallicity from the local ISM from whose constituents the particle is created. As a result of aging stellar populations, the present mass and metallicity of a particle at each time-step are subject to change. Using the STARBURST99 single-age stellar population synthesis models (Leitherer et al., 1999; Vázquez & Leitherer, 2005; Leitherer et al., 2010) an SED is assigned to each stellar particle based on the mass, metallicity, and age. In an effort to spatially distribute stellar light arising from stellar particles whose masses generally represent unresolved full stellar populations, T15 use adaptive smoothing of each star particle using a gaussian kernel with full-width at half maximum equal to the 16th nearest neighbour distance. The smoothing is intended to give simulated galaxy images substructures that are similar in appearance to real ones. However, T15 caution that neither this, nor any other adaptive or fixed-length smoothing prescription that they explore is any more or less valid when treating individual star particles as full unresolved stellar populations. While a more physical light assignment procedure may exist, we limit our exploration of methods for distributing stellar light only to the biases on specific parameters in the morphological decompositions.

Photon packets containing the integrated SEDs of each stellar particle are propagated probabilistically from the spherical volumes derived from the smoothing kernel. Monte Carlo propagation of the stellar light is facilitated with the radiative transfer code SUNRISE using 10^8 photon packets - which was confirmed to provide convergent surface brightness distributions in the synthetic images. Nonetheless, T15 warn that the detailed structure of low surface brightness features should be explored with caution.

The dust absorption, scattering, or emission functionalities of the SUNRISE code are not used in the propagation scheme. While dust effects are important considerations for the creation of realistic synthetic images, Jonsson et al. (2010) show that dust-inclusive radiative transfer results from SUNRISE are not expected to be converged for the mass and spatial resolution (\sim kpc) employed by the Illustris simulation. However, a simple empirical model for dust obscuration (Charlot & Fall, 2000) and nebular emission (Dopita et al., 2005, 2006a,b; Groves et al., 2008) are considered with respect to limitations of the spacial resolution of the Illustris simulation for these phenomena. The details of these models and their implementation for Illustris synthetic images are described in T15.

Creation of Synthetic Images

Photon packets are propagated into pinhole cameras with 256×256 pixel resolution providing each pixel with the integrated SED of the photon packets incident upon it - effectively creating a mock integral field unit (IFU) data cube. The field of view from each camera is $10 \text{ } rhm_{\star}$ where rhm_{\star} is the stellar half-mass radius measured from the gravitational potential minimum of a galaxy. Cameras are located on the arms of a tetrahedron whose centroid is located at the gravitational potential minimum. The camera locations are not aligned to the principle angular momentum vectors of each galaxy, but are aligned consistently within the simulation volume. This gives 4 camera angles with random orientation with respect to the galaxy of interest where the angles of each camera to the principle angular momentum component of each galaxy are still accessible *a priori*. The mock IFU can then be convolved with an arbitrary filter transmission function and resampled onto the desired angular resolution. The broadband image data products of T15 can be easily manipulated to create idealized galaxy images in 36 unique bands using the dedicated python module SUNPY.¹

2.2.3 Image Realism

Comprehensive image realism is applied to the synthetic images of the simulated galaxy sample from Illustris to enable meaningful and unbiased comparisons between real and simulated galaxy populations. First, we generate synthetic images of simulated galaxies in SDSS bands g and r using SUNPY (as described in Section 2.2.2)

¹<http://www.github.ptorrey/sunpy-master/>

to be consistent with the canonical simultaneous g - and r -band morphological decomposition analysis of 1.12 million real galaxies in the SDSS DR7 Legacy Survey (Simard et al., 2011). We then use the following procedure to create an unbiased Illustris simulated galaxy population in SDSS in g and r bands:

- (1) **SELECTION OF SDSS FIELDS AND PHOTOMETRIC QUANTITIES:** We randomly select a unique galaxy objID from SDSS DR7 Legacy photometric galaxy catalog. The SDSS atlas (prefix "fpAtlas"), PSF (prefix "psField"), and g - and r -band corrected images (prefix "fpC") for the `run`, `rerun`, `camcol` and `field` containing this galaxy are then obtained from the SDSS Data Archive Server (DAS). The corrected images are field images that have been reduced through bias-subtraction, flat-fielding, and purging of bright stars. The PSF image contains all of the necessary metadata to reconstruct the PSF in any band and location on the corrected image field. The Atlas image is only collected for the CCD gain that is contained in its image header. The photometric zero points, airmasses, extinction coefficients (`PhotoPrimary` table), and CCD gain (Atlas image) for the SDSS fields are collected to be used to convert the synthetic Illustris galaxy image fluxes from nanomaggies to counts.
- (2) **PHOTOMETRIC SEGMENTATION AND LOCATION ASSIGNMENT:** We employ `SOURCE EXTRACTOR` (`SExtractor`) (Bertin & Arnouts, 1996) to create a segmentation map of the SDSS r -band corrected image. We select a random location on the SDSS corrected image as the designated position of the centroid of the simulated galaxy images. While upholding our intention of recreating realistic crowding statistics of the simulated galaxy population, we restrict selection *of the centroid location* to pixels which have not been flagged as belonging to other objects identified in the segmentation map. Our choices of deblending parameters in `SExtractor` are the same as those used in S11.
- (3) **SDSS PSF IMAGE CONVOLUTION AND FLUX CONVERSION.:** We reconstruct the g - and r -band PSF images specific to the selected location in the corrected image using the `read_PSF` software utility from SDSS². We remove the soft-bias added by SDSS, normalize the PSF images, and convolve them with the noiseless synthetic galaxy images to provide realistic SDSS resolution. Using the photometric information obtained in (1) specific to the choice of SDSS cor-

²http://classic.sdss.org/dr7/products/images/read_psf.html

rected images, we convert the fluxes of the convolved synthetic images from nanomaggies (default data product units) to DN counts.

- (4) ADDITION OF SIGNAL SHOT NOISE TO SIMULATED GALAXY.: Photon shot noise is generated and added to the convolved synthetic galaxy images following Poisson statistics. The contribution of Poisson noise to the total variance in each pixel is expected to be small relative to the sky – which is the dominant noise term for photometry in the SDSS (Simard et al., 2011). In a preliminary analysis of galaxies at $z = 0$ we found that the inclusion/exclusion of Poisson noise did not affect parameter estimates. Nonetheless, Poisson noise is implemented in our analyses for continuity with extensions of our methodology to comparisons using other observational data in which Poisson noise may constitute an appreciable component of the noise in image photometry.
- (5) PLACEMENT OF SIMULATED GALAXY IMAGE INTO THE SDSS.: The simulated galaxy images whose fluxes, PSF resolution, and signal shot noise are matched to their designated location in the SDSS corrected image fields are then inserted into the SDSS corrected images with image centroids (which double as the location of the gravitational potential minimum of the galaxy) at the designated location from (2). This provides the bias contributions from the sky, crowding, and any other field-specific properties that, complete with the previous steps, finalize our image realism considerations.

Figure 2.1 is a visual demonstration of the steps described above. After (5), the same bulge+disk decomposition quantitative morphologies pipeline used in S11 can be used on the SDSS $g-$ and $r-$ band images containing the simulated galaxy. The only prior is the position of the simulated galaxy in the corrected image – which is consistent with analysis of real galaxies.

There are some important subtleties to our image realism suite. Applying steps (1) and (2) to the placement of simulated galaxies achieves effective matching to the statistics of crowding, resolution, and signal-to-noise within the real photometric galaxy catalog. In this scenario, fields containing higher concentrations of galaxies have a naturally higher probability of being selected for placement. While galaxies in isolation will still be represented, this method guarantees that any systematic biases in the recovered parameters from quantitative morphological analysis are (a) consistent with the real galaxy population and (b) statistically quantifiable through analysis of the same simulated galaxy across multiple placements and viewing angles.



Figure 2.1: Addition of Observational Realism to Idealized Simulated Galaxy Images. *Left panel:* Synthetic SDSS r -band image of Illustris subhaloID 312287 in a nearly face-on projection. *Middle panel:* convolution of the r -band synthetic image with the PSF corresponding to its insertion into the SDSS and addition of signal shot noise. *Right panel:* placement of PSF-convolved synthetic image flux into an SDSS corrected image. The logarithmic scale is identical in each image. The flux in each pixel in the right panel has been lifted by 5 counts so that the sky is visible using this scale.

2.3 Quantitative Morphologies

The image realism described in the last section puts the Illustris simulated galaxy population in a fully observational context that is well suited for quantitative comparisons with real galaxy populations. Comparisons of this kind will not only serve to identify the differences and similarities in the detailed structural parameter spaces of simulated and real galaxy populations, but also provide technical insight into the adequacy of the methods employed to create synthetic galaxy images from collections of particles that each represent unresolved stellar populations in cosmological simulations.

To facilitate our comparison, we use the same quantitative morphology analysis of S11 which performs 2D photometric surface brightness decomposition with parametric component models. Unlike the non-parametric analysis of S15, our parametric bulge+disc and single component decompositions will enable comparison of the detailed photometric properties of the structural components in observed and simulated galaxies. Furthermore, our comprehensive realism considerations place our morphological estimates on level ground with observations. While S15 employ the same quantitative analysis and methods for source-sky and source-source delineation in the simulated and observed galaxy populations, any biases that are *intrinsic* to the source-delineation or morphological analysis that correlate with resolution, signal-to-noise, and crowding will manifest themselves differently in the observations and simulations if the statistical distributions for these realisms are not the same. Our more complete treatment of observational biases may have important consequences on the comparison between models and observations.

In this section, we detail our quantitative morphologies analysis and emphasize its consistency with the analysis of 1.12 million galaxies in the SDSS by S11. We describe our methods for delineating photometric boundaries between sources and the sky and between closely projected sources. We then explain our choice of parametric models in the surface brightness decompositions and the structural parameters that are afforded by these choices. Finally, we describe the design of our catalogs – each of which have associated dedicated experiments aimed characterization of the biases that our realism considerations have on structural parameter estimates.

2.3.1 Deblending

The methods that are used to delineate object boundaries (deblending) have been shown to affect morphological parameters – particularly in crowded images (Simard et al., 2011). The standard SDSS PHOTO pipeline attempts to isolate the flux from an object to reconstruct the image of what the object would have looked like as if it was the only source in the image. Therefore, pixels that share flux from multiple sources are attributed to the area of each source with an associated fractional flux contribution based on the reconstruction modelling. However, deblending with the PHOTO pipeline has been shown to produce erroneous photometric and structural estimates such as the production of red-outliers and large scatter in the colour-magnitude diagrams of pair galaxies – which were previously and erroneously ascribed to a new population of extremely red galaxies in pair systems (Alonso et al., 2004; Perez et al., 2009; Darg et al., 2010). Inaccurate photometric estimates in pairs is an indication that the same inaccuracies are relevant in all objects with closely projected external sources. S11 showed that photometric and structural estimates derived from SExtractor³ deblending combined with GIM2D sky measurement and bulge+disc decompositions improved upon other schemes using several sensitive tests: the size-luminosity relation of discs, and the colour-magnitude diagrams and fiber colours of pairs. They showed that the deblending used in the PHOTO pipeline and associated magnitude and colour estimates was the source of the outliers and scatter in the colour-magnitude diagrams of pairs. S11 also demonstrated that deblending using SExtractor in tandem with parametric bulge+disc decompositions reduced the scatter and eliminated the outlier populations in the colour magnitude diagrams of pairs – leaving a tight red sequence and clearly separated blue cloud using SDSS pair catalog of Patton et al. (2011). We therefore employ the SExtractor source deblending procedure used by S11. We do not presume that the S11 scheme is optimal or unique in defining object-sky boundaries and separating objects whose constituent pixels may have shared

³SExtractor deblending uses a multi-threshold flux tree. Starting with the lowest isophotal threshold and moving up, troughs are identified that separate branches which meet the criterion of containing a specified fraction of the total flux. The minima of these troughs delineate the flux boundaries and each pixel is given a flag corresponding to a unique object – creating a segmentation image. No pixels are shared and therefore no object’s segmentation map area extends into the area associated with another object identified through this scheme. Although the fluxes measured directly from the pixels associated with an object would be systematically be underestimated in this scheme in the presence of close neighbours, S11 showed that the missing flux is recovered by fitting a surface brightness profile model – which integrates the flux of the model (whose form is determined only by the pixels flagged as belonging to the object of interest) out to large radii. In practice, fitting the model recovers the missing flux of SExtractor deblended objects in crowded environments.

contributions from other sources. However, we note that although any biases from the S11 deblending and sky estimation scheme may affect our morphologies, application of the same scheme to our mock-observed simulated galaxies ensures that the biases will be consistent.

Before application of our five-step realism suite described in Section 2.2.3, the synthetic images generated using SUNPY have no sky or noise other than a residual noise contribution from the Monte Carlo photon propagation scheme in SUNRISE. T15 point out that this noise manifests itself only as small fluctuations in pixel-to-pixel intensity at low surface brightnesses. However, since the synthetic images are constructed from the Friends-of-Friends halo finder in Illustris, there may be contributions from other stellar sources within the field of view that give rise to fluxes that do not truncate to zero at the synthetic image boundaries. When added to the SDSS corrected image after the other steps in Section 2.2.3, the non-zero flux from other sources in the FoF group result in boxlike flux boundaries between synthetic image and SDSS corrected image into which it is placed. Figure 2.2 shows an example where a FoF companion to the galaxy of interest (centred) is in the projected field-of-view (two bottom rows of panels). The undesirable effect is exemplified in the bottom row, where the companion is projected along the line of sight between the galaxy and the camera position.

Galaxy projections with unrealistic artifacts such as shown in Figure 2.2 are easily separable by comparing the total flux at an arbitrary camera angle with any other. Identification of a significant positive systematic bias in the total flux of a synthetic image of a particular galaxy subhaloID with respect to, for example, an estimate of the mode flux for all camera angles is an effective flagging scheme for these situations. Situations such as seen in the third row of panels in Figure 2.2 are not problematic for our quantitative morphologies pipeline due to our deblending scheme and valid comparisons between the morphological parameters to those from other camera angles and real galaxies can be made. They are not useful, however, in analyses that compare properties of the synthetic images such as total flux and photometric aperture half-light radius with those recovered from the fitting. The aim of our analysis is to decompose the surface brightness profiles of the primary galaxy – which contains the subhalo’s gravitational potential minimum. Therefore, parametric estimates of flux and size in our analysis will have an apparent systematic bias relative to similar quantitative properties of the synthetic images for lines of sight that include companions and satellites.

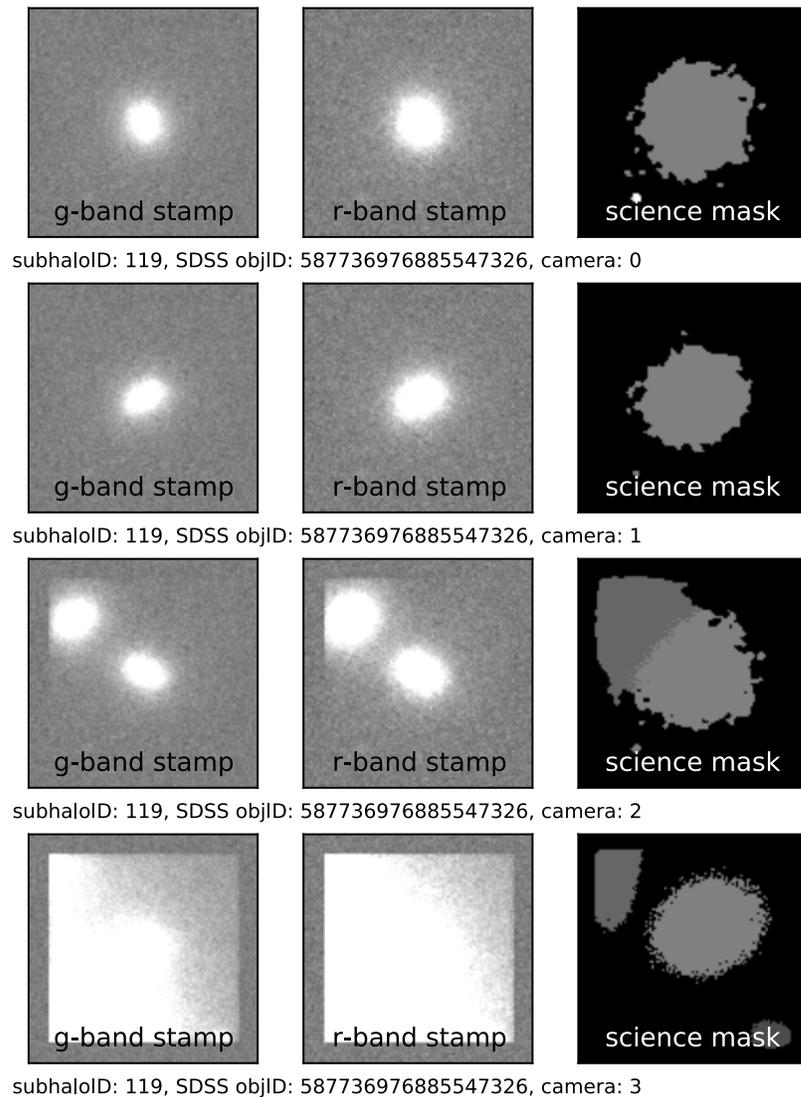


Figure 2.2: Four camera angle realizations of the same simulated galaxy with full observational realism. The galaxy is inserted into the same location in SDSS for each camera angle – separated by row in the Figure. The science r -band image stamp, g -band image stamp, and segmentation stamp are shown from left to right. The third row of panels indicate that the central galaxy has a companion that is not within the field of view for the first two rows. Note the delineation of flux for the companion at the boundary of inserted synthetic image. Only pixels for which the central galaxy has a dominant flux contribution are selected as data used in the fitting. In the fourth row, the companion’s orientation with respect to the central galaxy places it in the line of sight of the cameras field of view. The projection effect results in a large discontinuity between the flux at the boundary of the synthetic image and the rest of the SDSS sky and an erroneous science mask that does not properly delineate between the galaxy’s flux and and the visible flux from the companion.

2.3.2 Parametric Surface Brightness Profile Fits

We perform simultaneous parametric 2D surface brightness profile fitting of the g - and r -band image data using the GIM2D software package (Simard et al., 2002). GIM2D employs a Metropolis-Hastings Markov-Chain Monte Carlo algorithm for deriving the best-fitting models based on the image data. Only science image pixels identified to belong to the principle object in the mask are used in the likelihood calculations. We use two models to fit the images: a bulge+disc ($b + d$) model which are respectively represented by sérsic functions (de Vaucouleurs, 1953, 1959b; Freeman, 1970) with indices $n_b = 4$ and $n_d = 1$; and a pure sérsic (ps) model where the sérsic index is a free parameter but confined to the domain $0.5 \leq n_{ps} \leq 8.0$. The fourteen free parameters of the $b + d$ decompositions are the total fluxes in each band f_g, f_r , bulge-to-total ratios $(b/t)_g, (b/t)_r$, semi-major axis bulge effective radius r_e , bulge eccentricity e , bulge position angle (clockwise y-axis $\equiv 0$) ϕ_b , disk scale length r_d , disk inclination i , disk position angle ϕ_d , and centroid positions $(dx)_g, (dy)_g, (dx)_r$, and $(dy)_r$. The parameters for position angles of the bulge and disk, disk inclination, bulge ellipticity, bulge effective radius, and disk scale length have the added constraint that they must be equivalent in both bands. Furthermore, the centroid positions of the bulge and disk components of the model are constrained to be the same. Similarly, the ten free parameters for the ps model are f_g, f_r , sérsic index n_{ps} , semi-major axis disk effective radius r_e , disk eccentricity e , disk position angle ϕ_d , and centroid positions $(dx)_g, (dy)_g, (dx)_r$, and $(dy)_r$. The position angle of the profile, ellipticity, and effective radius are constrained to be the same in each band in the ps fits.

2.3.3 Simulated Galaxy Population Samples and Catalogs

Several experiments were conducted with increasing completeness to fully characterize and understand the complexities and biases that will affect estimates for the full Illustris galaxy population. Our samples and catalogs are described in this section. The decomposition catalogs are summarized in Table 2.1 at the end of this section.

Representative Illustris Galaxy (RIG) Sample

We began by selecting a small, but representative Illustris galaxy sample (RIGs) in the parameter space of stellar mass and half-stellar mass radius. The sample was

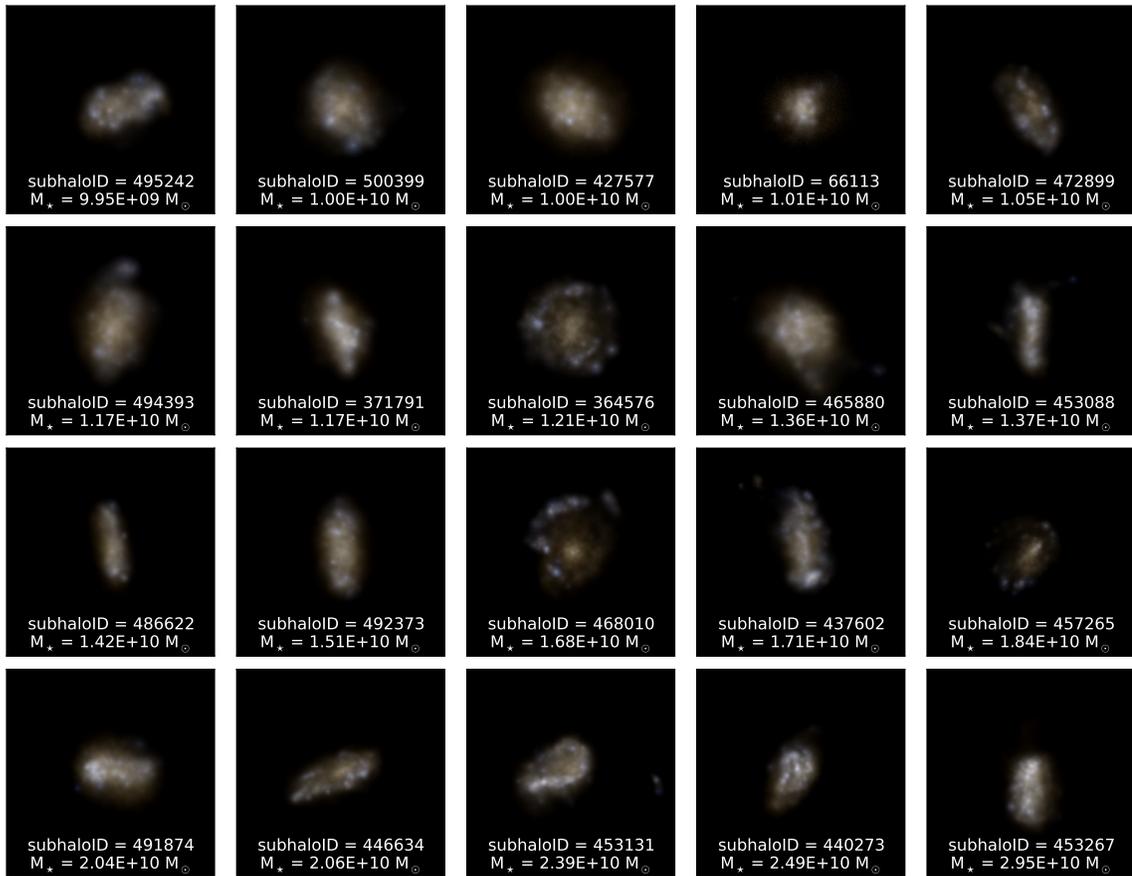


Figure 2.3: Representative Illustris Galaxy (RIG) Sample synthetic images ordered by total stellar mass. The synthetic images are composites of the SDSS *gri* colours.

selected by uniformly sampling the 2D plane of $M_* - rhm_*$ with selection of the nearest matching galaxy to 100 random placements. The tolerance of spacing between each point on the plane is then iteratively increased until uniform sampling of the plane is achieved. Finally, we perform a cull through visual identification of galaxies that are in mergers, show strong tidal disturbances, or have problematic projection effects such as in row 4 of Figure 2.2 that would strongly affect our “representative” morphologies. The full process is repeated on galaxies that survived the cull until 100 galaxies are found that uniformly span the $M_* - rhm_*$ plane and will not have morphologies that are strongly influenced by the systematic effects eliminated by performing the culling at each iteration. The finalized RIG sample contains a diversity of visual morphologies. Figures 2.5-2.7 show each RIG in the CAMERA 0 projection.

The representative sample is used to perform repeated analyses of galaxies in SDSS fields to obtain distributions of parameter estimates for each galaxy and each

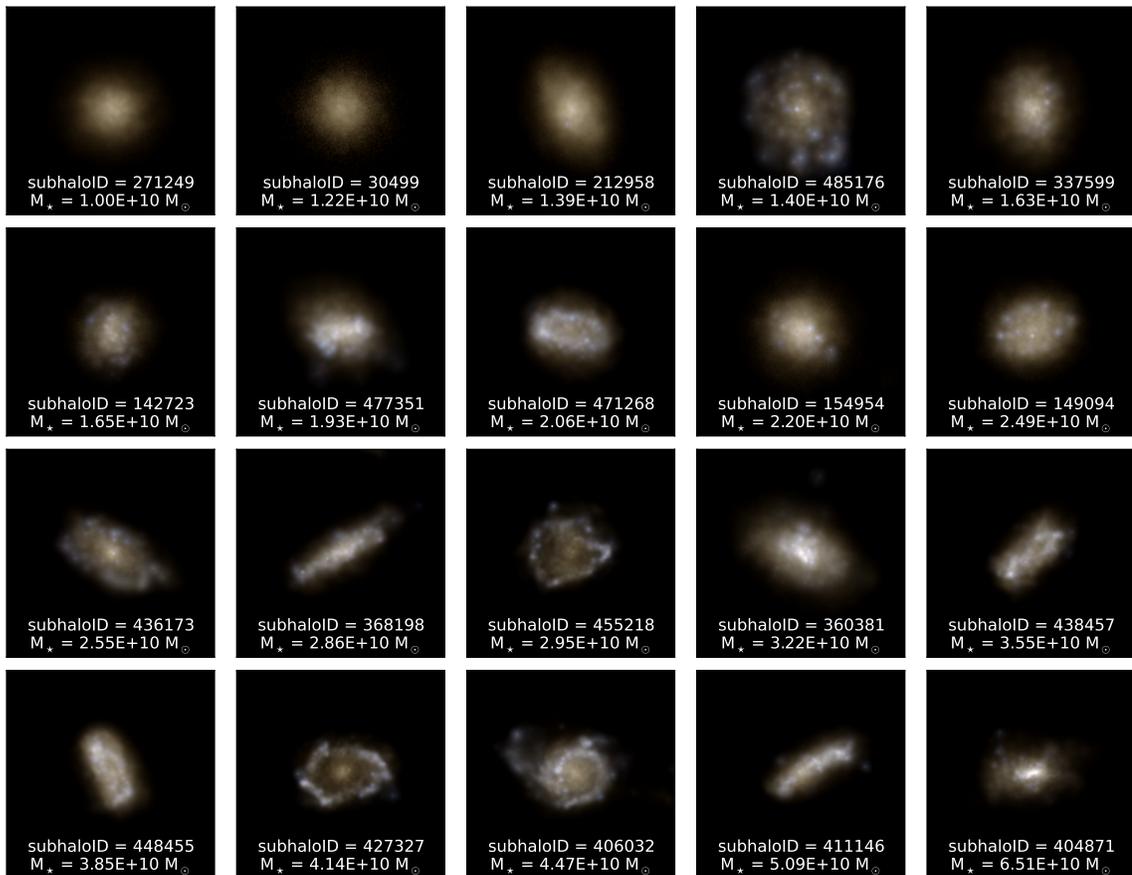


Figure 2.4: Representative Illustris Galaxy (RIG) Sample synthetic images ordered by total stellar mass. The synthetic images are composites of the SDSS *gri* colours.

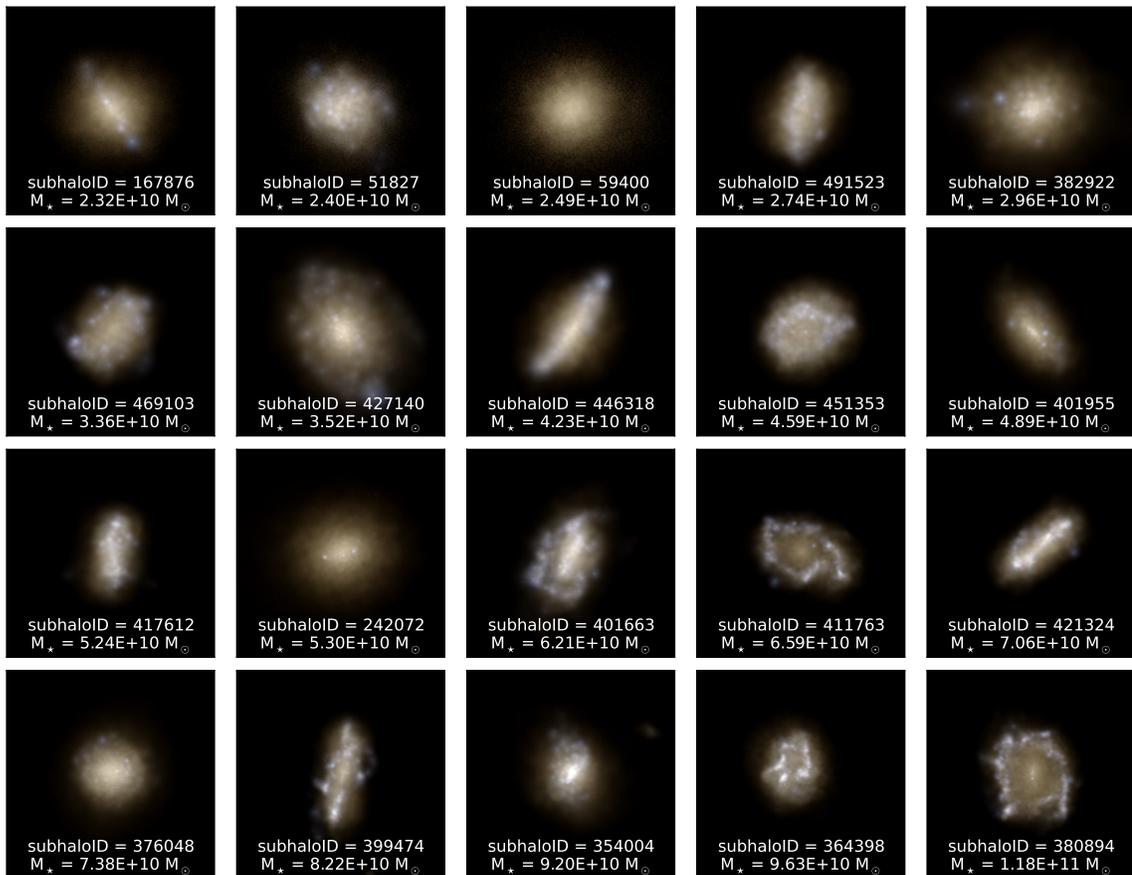


Figure 2.5: Representative Illustris Galaxy (RIG) Sample synthetic images ordered by total stellar mass. The synthetic images are composites of the SDSS *gri* colours.



Figure 2.6: Representative Illustris Galaxy (RIG) Sample synthetic images ordered by total stellar mass. The synthetic images are composites of the SDSS *gri* colours.

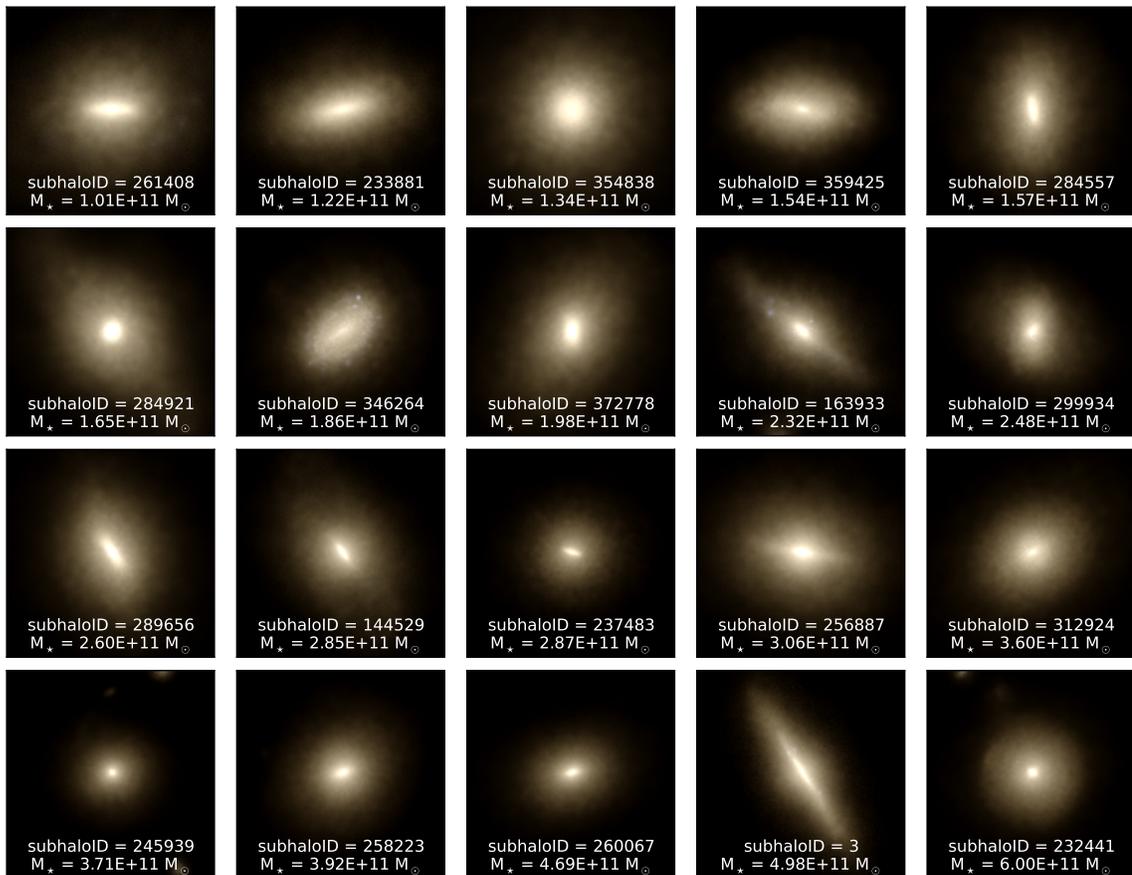


Figure 2.7: Representative Illustris Galaxy (RIG) Sample synthetic images ordered by total stellar mass. The synthetic images are composites of the SDSS *gri* colours.

camera angle – facilitating a quantitative analysis of the systematic and random biases from crowding and projection for the various morphologies. We further exploit the representative sample in an analysis of the biases that arise from the distribution of stellar light in SUNRISE. We construct several catalogs for our analysis of the representative sample that are organized to separate each potential bias. Our final catalog includes a single decomposition for all camera angles of every galaxy in the Illustris synthetic image catalog of T15. The characterization of our biases and errors in the analyses preceding the full realization of our analysis of the entire Illustris galaxy population ensures that we are able to make accurate statements about the degree of confidence in our structural and photometric parameter estimates for specific domains of parameter space – which is valuable considering our aim of comparing simulated galaxy populations to real ones.

ASKA catalog

An **All SKy Analysis** was performed on one hundred galaxies that uniformly sample the plane of stellar mass and half-mass radius in Illustris (RIGs). The RIGs are placed all over the SDSS sky with an average of ~ 100 unique fields for each galaxy following the steps in Section 2.2.3. All four camera angles of a galaxy are modelled in each placement – providing a distribution of ~ 100 unique sets of best-fitting morphological parameters for each camera angle. The full **ASKA** catalog constitutes $\sim 40,000$ decompositions of galaxies from the RIG sample. The distribution of best-fitting parameters for each camera angle of a galaxy samples the real statistics for crowding, signal-to-noise, and resolution that exists for SDSS galaxies as a result of our placement criteria and realism procedures. We use the **ASKA** catalog to quantify the random and systematic effects from the combined effects of crowding, signal-to-noise, and resolution by holding camera angle fixed. The distributions of best-fitting parameters for each camera angle facilitate the statistical quantification of the scatter and systematics from biases associated with placement.

In the following two sub-sections, we describe image catalogs in which the sky placement is held constant, in order to test issues associated with internal galaxy image generation and projection, rather than sky insertion, effects.

SMOOTHING catalog

The **SMOOTHING** catalog is constructed to explore the choices for distributing stellar light in creating the synthetic images (Torrey et al., 2015). Synthetic images are generated for the RIG sample using three alternative stellar light distribution (SLD) schemes to compliment the existing images constructed using the fiducial scheme. All SLD realizations of galaxies from the RIG sample are placed in the same uncrowded location in SDSS and fitting is performed for a single camera angle: **CAMERA 0**. The **SMOOTHING** catalog therefore contains 400 decompositions with 4 for each RIG (one for each SLD scheme). The biases on photometric and structural parameters from the SLD schemes may be evaluated with respect to the fiducial scheme and to each other. Four SLD schemes are explored in our analysis:

- (1) **FIDUCIAL SMOOTHING (fn16)**: Radiative transfer from each stellar particle within the FoF subhalo is projected probabilistically in **SUNRISE** from the commonly applied spherical volume cubic B-spline profile with a characteristic scale set by the distance to the 16th nearest neighbouring stellar particle. Adaptive smoothing allows for more concentrated distributions of stellar light while avoiding unrealistic compactness around largely isolated particles at large distances from the galactic centre.
- (2) **CONSTANT SMOOTHING (fc1kpc)**: Radiative transfer from each stellar particle within the FoF subhalo is projected probabilistically from a spherical volume with cubic B-spline profile with a constant characteristic scale set to 1kpc. Total stellar light is conserved with respect to the fiducial scheme, but compact surface brightness features with projected spatial distributions less than the characteristic scale are distributed more broadly.
- (3) **RESAMPLED ADAPTIVE SMOOTHING (rn16)**: Young, bright stellar populations associated with the $\sim 10^6$ solar mass simulation stellar particles result in artificially distinct and circular light artifacts in the synthetic images. Surface brightness features that are highly discrete relative to the local surface brightness of the rest of the galaxy may be delineated from the rest of the galaxy in the deblending – effectively removing their (often significant) contributions to the total flux. To mitigate this effect, we resample the light associated with these young, bright star particles into 100 particles (child particles). The child particles contain 1/100th the mass of the parent particle, are spatially distributed

within the parent particles original light kernel, and are assigned Gaussian age and metallicity distributions with one-sigma values of 10% of the original age and metallicity. The resulting stellar population flux is roughly conserved, but the sharp light profile edges are somewhat reduced. The stellar light from all child particles and remaining particles is projected probabilistically from a spherical volume cubic B-spline profile with a characteristic scale set by the distance to the 16th nearest neighbour, as in the fiducial scheme.

- (4) **RESAMPLED CONSTANT SMOOTHING (rc1kpc)**: Same resampling as in (3) but radiative transfer from each of the child and remaining particles is projected probabilistically from a spherical volume cubic B-spline profile with a constant characteristic scale set to 1kpc.

CAMERAS catalog

The **CAMERAS** catalog is constructed to evaluate the biases on parameter estimates from projection. Each galaxy in the RIG sample is placed in a single, uncrowded location in SDSS and fitting is performed for all four camera angles. This guarantees consistent and controlled environment, resolution, and sky in each decomposition and across analyses of each RIG. The **CAMERAS** catalog contains 400 decompositions with 4 for each RIG (one for each camera angle). The variation in best-fitting photometric and structural estimates in each RIG are obtained from the **CAMERAS** catalog.

DISTINCT catalog

The **DISTINCT** catalog contains decompositions for all galaxies in synthetic image catalog of T15. Each camera angle for a given galaxy is assigned probabilistically to a location in the SDSS following item (2) of Section 2.2.3. The catalog contains 27,564 decompositions from the 4 camera angles for each of 6891 galaxies from Illustris. Figure 2.8 shows examples of $b + d$ and ps decompositions of galaxies taken from the **DISTINCT** catalog. The catalog is designed to investigate the global observational properties of the full Illustris galaxy population. The **DISTINCT** catalog also forms the basis for comparisons between simulated galaxies and populations of real galaxies in SDSS. Galaxies in the **DISTINCT** catalog can be sampled to match the luminosity or stellar mass distributions of real galaxies around $z = 0.05$. We aim to extend our study of redshift $z \sim 0$ galaxies out to $z \sim 0.2$ using galaxies sampled from the snapshot history of Illustris with projected positions and redshifts matched to

Table 2.1: Decomposition catalogs. Each catalog contains as many bulge+disc decompositions N_{decomp} as pure sérsic decompositions. SLD denotes which stellar light distribution methods are used to create the synthetic images of each galaxy. Cameras denotes which camera angle projections of each galaxy are used in a catalog. Whether the insertion into the SDSS sky is fixed to a single location for all decompositions (Fixed) or follows the randomized procedure described in Section 2.2.3 (Random) is indicated by Insertion. N_{gal} is the number of galaxies used in the decompositions. (*) Decompositions of all camera angles of galaxies are performed in each insertion location.

Catalog	SLD	Cameras	Insertion	N_{gal}	N_{decomp}
ASKA	fn16	0-3	Random*	100	42319
SMOOTHING	all	0	Fixed	100	400
CAMERAS	fn16	0-3	Fixed	100	400
DISTINCT	fn16	0-3	Random	6891	27564

the SDSS spectroscopic galaxy sample. An analysis of this kind would provide an evolutionary context to image-based comparisons between real and simulated galaxy populations.

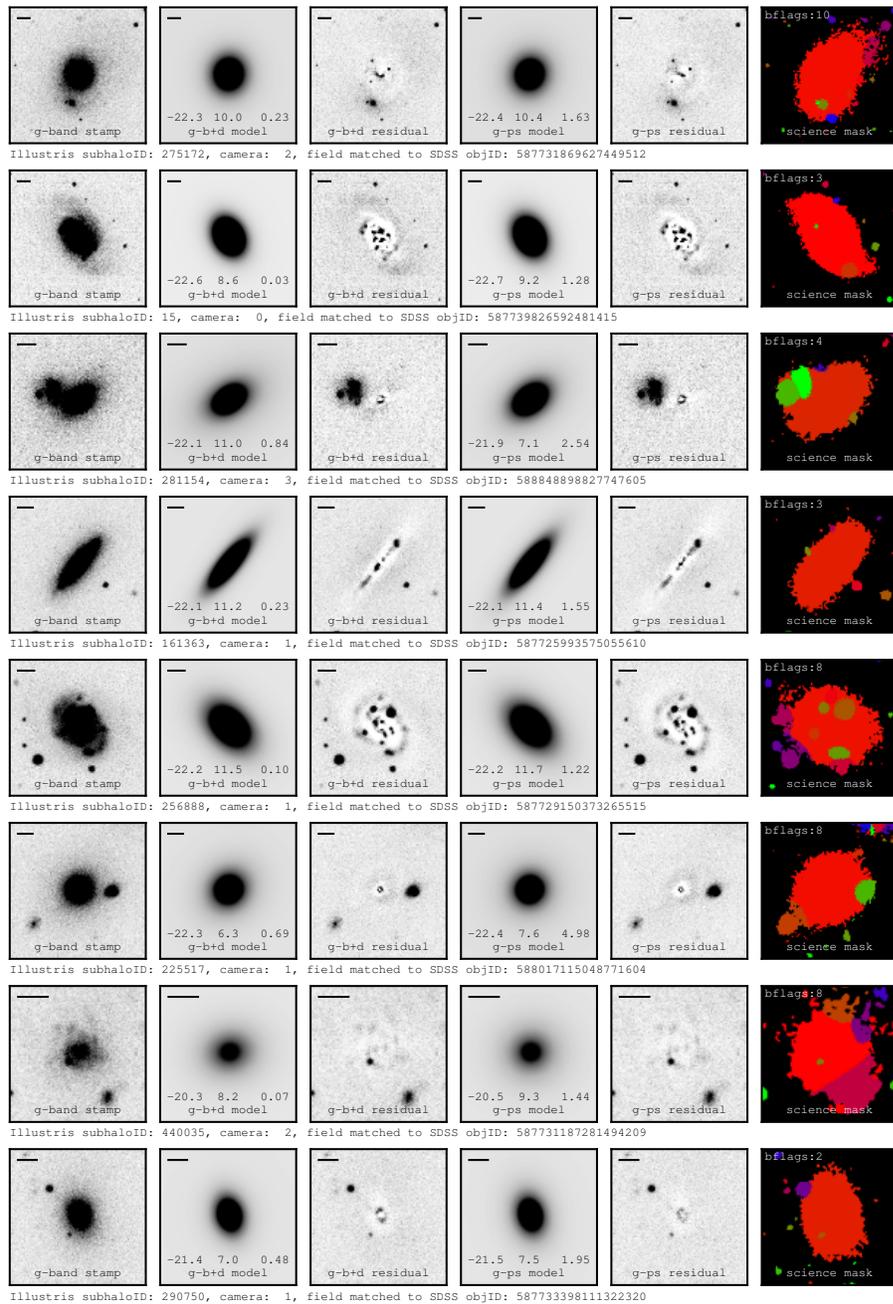


Figure 2.8: Mosaic of *g*-band images and best-fitting models for simulated galaxies with realism. *Left to Right*: (1) science cut-out (stamp) of simulated galaxy with realism; (2) *b* + *d* model absolute magnitude, circular half-light radius, and bulge-to-total ratio parameters listed from left to right; (3) *b* + *d* residual image (model subtracted from science cut-out); (4) *ps* model with absolute magnitude, circular half-light radius, and sérsic index parameters listed from left to right; (5) *ps* residual image; (6) source delineation mask (science mask). The science masks are colorized by flags associated with distinct objects. Red corresponds to the pixels that are used in the fitting. *bflags* is the number of sources that directly neighbour pixels used in the fitting. A scale in the top left of each panel denotes 10 kpc at $z=0.05$.

2.4 Characterization of Biases

In this section we explore estimates for total flux, circular half-light radii, and bulge-to-total light ratios in our decomposition catalogs. We use the derived values to disentangle sources of random and systematic bias on our parametric models. In particular, we explore the consequences of post-processing choices in how stellar light is distributed spatially and propagated from particles that embody full unresolved stellar populations in the simulation. We then assess the sensitivity of photometric estimates to the observational realism of projection and crowding, respectively. In each analysis of a potential bias, we take judicious precautions to control all other potential biases such that the variation in parameters will be sensitive exclusively to the bias under examination.

2.4.1 The Propagation of Stellar Light from Unresolved Stellar Populations

Stellar particles in the Illustris have a birth mass of $\sim 10^6 M_\odot$ – representing an unresolved stellar population. A method is required to convert the discrete distribution of simulation particles into a continuous light distribution in the synthetic images. In this section, we explore our fiducial scheme for the distribution of stellar light. While no particle-to-light conversion is formally any more correct than any other, the stellar light distribution (SLD) scheme may bias photometric and structural estimates. We employ four simple and general methods for translating the distribution of simulation particles into light to test the impact of this choice on the derived galaxy properties.

We produce synthetic images for each of the 100 galaxies in the RIG sample for a single projection using 4 unique SLD schemes to characterize the biases of our fiducial and alternative SLD schemes. The specifications for each scheme are detailed in Section 2.3.3. We visually selected a location in the SDSS that had low sky variance and no sources that could be identified by SExtractor within the nearby vicinity. We consistently performed decompositions for 400 galaxy images inserted in the same uncrowded location in the SDSS to eliminate biases that might arise from sensitivity to environment and projection. The decompositions collectively constitute the SMOOTHING catalog. Our control of placement and projection for decompositions in the SMOOTHING catalog ensures that the variations in parameter estimates for each galaxy is exclusively due to SLD schemes.

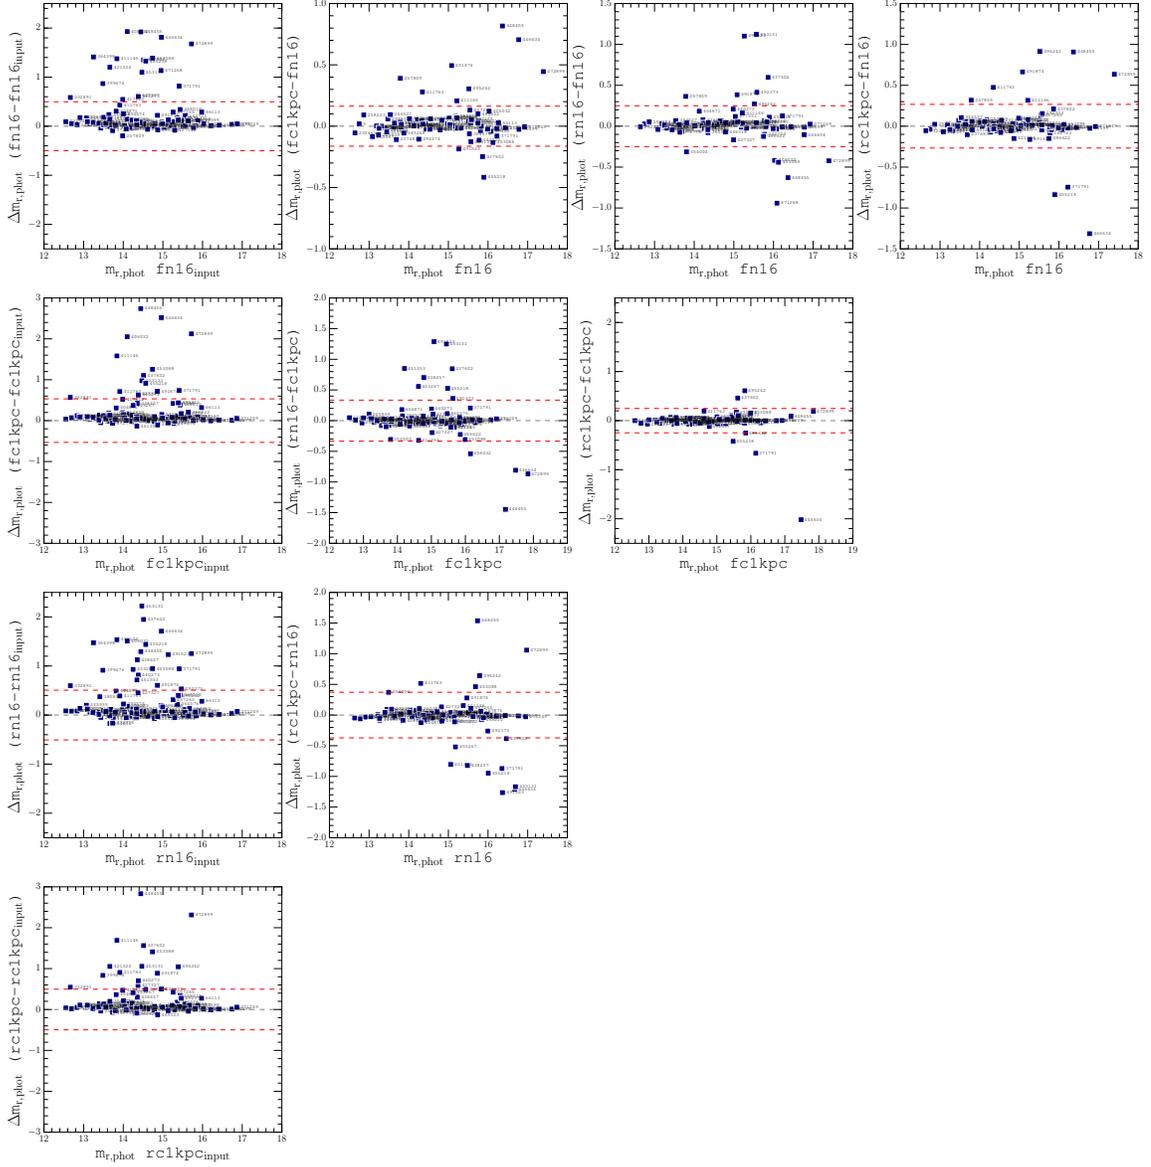


Figure 2.9: Variation and systematic biases in integrated magnitude for alternative SLD schemes. The first panel of each row shows the systematic bias on r -band integrated magnitude computed from the difference between the magnitude of the best-fitting $b + d$ model and the synthetic image for the specified SLD scheme. All other panels compare the magnitude estimates from each SLD scheme with every other. Next to each point is the corresponding galaxy subhaloID being compared. Equality between estimates is marked by the black dashed lines. Red dashed line show 1σ in each direction about zero.

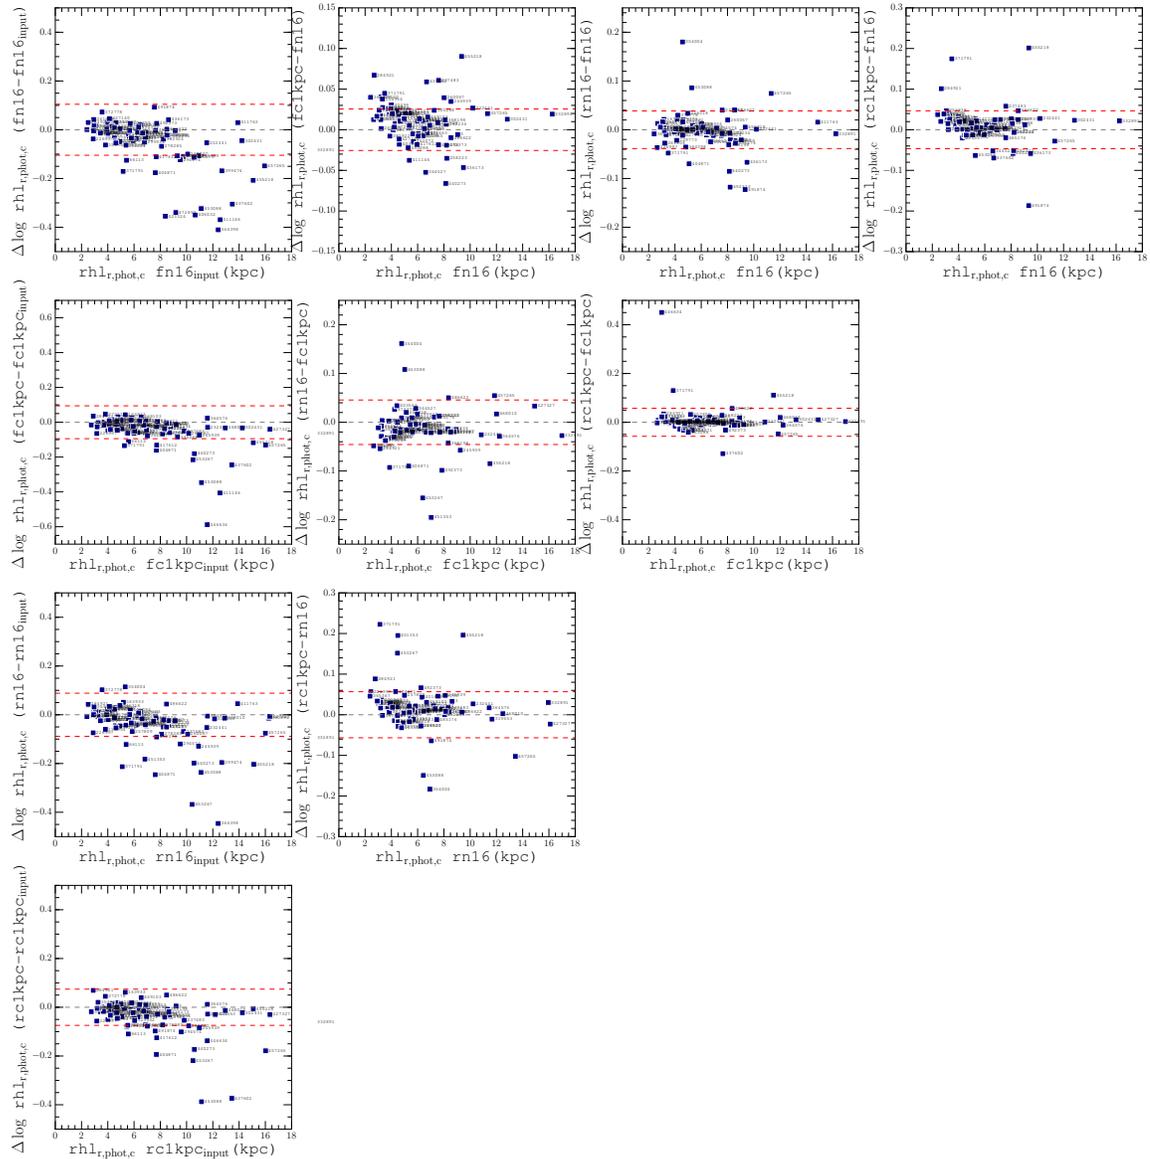


Figure 2.10: Variation and systematic biases in the circular aperture half-light radius for alternative SLD schemes. The first panel of each row shows the systematic bias on r -band half-light radius computed from the ratio between the half-light radius of the best-fitting $b+d$ model and the synthetic image for the specified SLD scheme. All other panels compare the half-light radius estimates from each SLD scheme with every other. Next to each point is the corresponding galaxy subhaloID being compared. Equality between estimates is marked by the black dashed lines. Red dashed line show 1σ in each direction about zero.

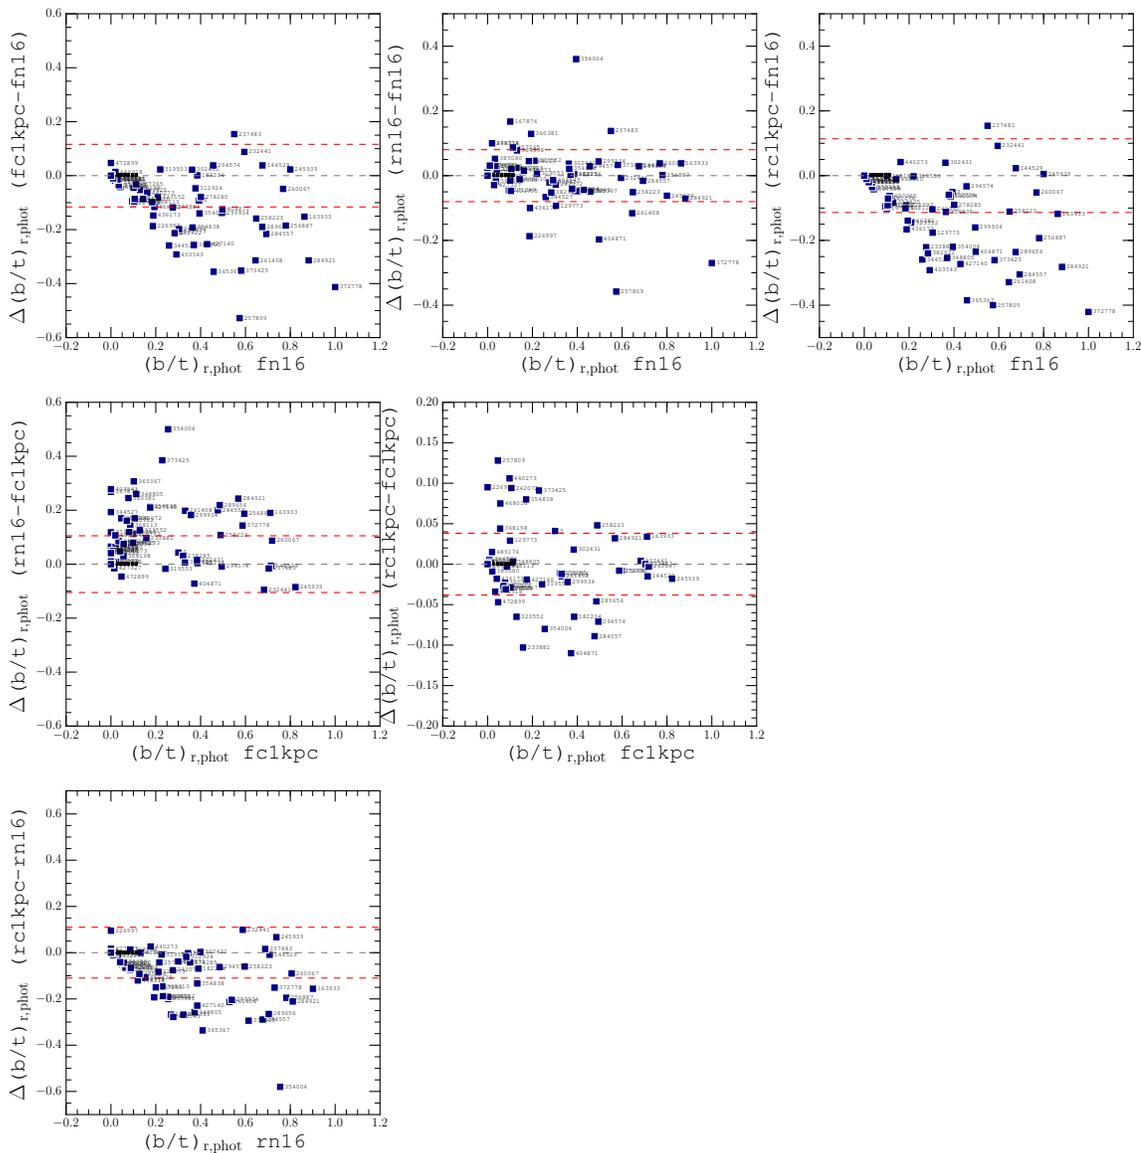


Figure 2.11: Variation in the photometric bulge-to-total light ratio for alternative SLD schemes. Each panels compares the r -band (b/t) from each SLD scheme with every other. Next to each point is the corresponding galaxy subhaloID being compared. Equality between estimates is marked by the black dashed lines. Red dashed line show 1σ in each direction about zero. Note that the axis limits are not the same in each panel.

The decompositions from the **SMOOTHING** catalog are used to quantify the biases from the SLD schemes on integrated magnitude, circular aperture half-light radius, and photometric (b/t). Integrated magnitude and half-light radius can both be computed from the synthetic images as well as the best-fitting ($b+d$) models for each galaxy and SLD scheme. The systematic biases on integrated magnitude and half-light radius are determined by comparing properties of the best-fitting models m_{b+d} and $rhl_{b+d,c}$ with the corresponding properties of the synthetic images m_{synth} and $rhl_{\text{synth},c}$ – which exclude realism. The integrated magnitude and circular half-light radii of the galaxies without realism can be computed from the total flux and aperture photometry of the synthetic images without any assumption about the form of their surface brightness profiles. Comparisons between the properties of galaxies in each SLD scheme are compared among themselves as shown in Figures 2.9 and 2.10.

Figure 2.9 compares estimates of integrated magnitude in the $b+d$ decompositions from the **SMOOTHING** catalog. Each panel compares the decomposition results for RIGs in two SLD schemes or the decomposition results from a specific SLD scheme with the magnitudes derived from the synthetic images. The systematic biases in model magnitudes for each SLD scheme with respect to their corresponding synthetic images are shown in the first panel of each row. $\Delta m_{r,\text{phot}}$ is the magnitude difference for the RIGs in each comparison. Red dashed lines denote 1σ above and below zero on the $\Delta m_{r,\text{phot}}$ axes. The majority of our integrated magnitude estimates for the representative sample of galaxies show excellent consistency with those computed directly from the synthetic images for all SLD schemes. Meanwhile, a handful of galaxies have larger integrated magnitude estimates than those computed from their synthetic images for all SLD schemes. No SLD scheme that we employ significantly reduces the number of galaxies with large systematic offsets. Visual inspection of the images and masks in each SLD scheme shows us that the outliers are strongly internally segmented due to the prevalence of substructure in their surface brightness distributions. However, each scheme provides unique variations in the segmentation map. Therefore, each best-fitting model to the surface brightness profile of the galaxy is determined from a unique set of science pixels and the best-fitting parameter estimates vary correspondingly. Such changes are reflected in the varying degree by which the outliers are offset in each SLD scheme.

Figure 2.9 also shows comparisons of integrated magnitude estimates for each galaxy between each SLD scheme. Panels in the first row show comparisons against the fiducial scheme. The majority of the magnitude estimates are consistent and are

correspondingly concentrated around the black dashed line – showing little sensitivity to SLD scheme. The outliers in each comparison appear randomly distributed about zero and are the same galaxies that demonstrated large positive systematic offsets in all SLD schemes. As expected from the panels showing systematic offsets, the random distribution of the outliers indicates that no SLD scheme alleviates the internal segmentation. Instead, the choice of SLD scheme simply provides unique changes to the segmentation maps.

The biases on photometric sizes of galaxies should be sensitive to the accuracy with which the flux is recovered. Figure 2.10 shows the systematic offsets and comparisons between SLD schemes for estimates of circular aperture half-light radius. The outliers in the panels showing the systematic offsets for each SLD scheme are the same galaxies that were systematic outliers in integrated magnitude. Again, no SLD scheme significantly reduces the population of outliers. The negative correlation between the offsets and half-light radii computed from the synthetic images hints that size estimates tend to be less robust for galaxies that are extended. However, several galaxies with large sizes computed from the synthetic images have consistent sizes estimated from the models. The galaxies with accurate sizes have high total stellar masses relative to those with inconsistent estimates from the models.

As in Figure 2.9, the compared estimates in each panel off of the first column of Figure 2.10 show a large population with highly consistent half-light radius estimates between the various SLD schemes – particularly in comparisons that both use adaptive or constant smoothing kernels. Interestingly, a systematic offset is present at small half-light radii in comparisons between the adaptive and constant smoothing SLD schemes. The half-light radii computed from the best-fitting models of galaxies with small intrinsic sizes are systematically larger in constant SLD schemes when compared to adaptive schemes. However, the choice of SLD scheme does not strongly affect estimates of half-light radius and full galaxy size estimates do not demonstrate a strong systematic offset based on the SLD method.

Finally, we investigate the variation in photometric (b/t) in comparisons between alternative SLD schemes in 2.11. We omit comparisons between the derived and intrinsic (b/t) values owing to ambiguities when defining the intrinsic (b/t). Such comparisons are saved for a follow up paper. (b/t) estimates show greater sensitivity to the choice of SLD scheme than integrated magnitude and half-light radius.

The comparison between (b/t) estimates that are both derived from decompositions of galaxies produced using adaptive schemes (`rn16` - `fn16`) shows that the

estimates of (b/t) for the majority of galaxies are consistent. However, several galaxies have discrepant (b/t) between SLD schemes, and these are different from the delinquent cases in the size and magnitude comparison. A similar outlier population exists in the two constant smoothing SLD schemes but with reduced scatter. However, there is a significant reduction of the (b/t) values derived from the constant schemes compared to adaptive schemes (`fc1kpc-fn16`, `rn16-fc1kpc`, `rc1kpc-fn16`, and `rc1kpc-rn16`). Many galaxies with large (b/t) using adaptive schemes to become virtually bulgeless in the constant schemes. One explanation for the reduction in (b/t) in constant schemes is that the light stellar particles near the centre of the bulge (on which the identification of the bulge relies sensitively) is smoothed too broadly relative to adaptive schemes that enable more spatially concentrated light profiles for tight clusters of stellar particles.

The systematic discrepancy in (b/t) estimates using alternative SLD schemes indicates that caution should be exercised in the choice of how stellar light is distributed to facilitate *realistic* propagation of photons from particles representing full unresolved stellar populations. While the total fluxes and sizes are generally robust for all smoothing types (apart from the handful of consistent outliers with strong internal segmentation), the choice of SLD scheme can clearly bias estimates of fundamental structural properties of galaxies.

2.4.2 Camera Angle

In the previous section, we investigated the biases on parameter estimates that arise from the choice of scheme for distributing stellar light. In this section we address biases introduced in the derived best-fitting parameters for galaxies from variations in the galaxy viewing angle. We employ the **CAMERAS** catalog described in Section 2.3.3 which contains a single decomposition of each camera angle projection of the RIGs. Synthetic images are constructed using fiducial SLD scheme `fn16`. Placement of each image is restricted to the same uncrowded location in the SDSS described in the previous section. Inserting each projection of each galaxy in the same location ensures that the resulting variation in parameters are in response to projection alone.

Figure 2.12 shows the sensitivity of bulge+disc model parameters to projection for the RIGs. The greyscale shows the distribution of the full T15 catalog of Illustris subhalos in the size-mass plane of stellar half-mass radius and total stellar mass. The locations of the coloured boxes indicate the sizes and masses of the RIGs. The

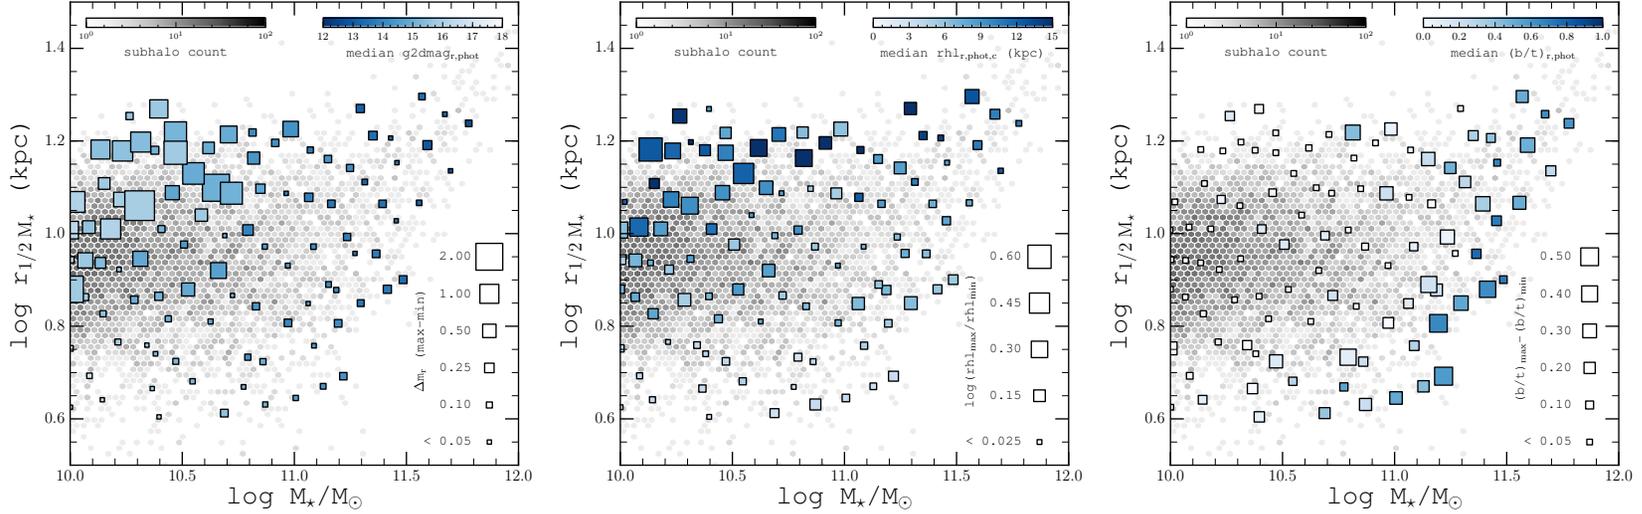


Figure 2.12: Characterization of random error from camera angle on r -band photometric magnitude (*left panel*), half-light radius (*middle panel*), and photometric bulge-to-total ratio (*right panel*). The $2D$ greyscale histogram in each panel shows the distribution of the full T15 catalog of Illustris subhalos in the plane of total stellar mass, M_* , and stellar half-mass radius, $r_{1/2M_*}$ with colours on a logarithmic scale and colourbar located at the top left of each panel. Coloured boxes are positioned in this size-mass plane according to the corresponding subhalo properties of each RIG. The colour of each box indicates the median of the best-fitting parameters for the 4 projections of each galaxy in a manually selected, uncrowded field. The corresponding colourbar is located at the top right of each panel. The size of each box is calculated from the full range in the recovered parameters. A key of random errors and sizes, along with the exact computation method, is located along bottom right of each panel.

size of each box indicates the variation of the model parameters from projection for each RIG. We quantify the variation using the maximum and minimum parameter estimates from the four decompositions of each RIG (one for each camera angle) due to our small number of estimates.⁴ The max-min variations are denoted by $\Delta_{\text{max-min}}$. The exact computation for the range in each parameter is given alongside the key at the bottom right of each panel.

The left panel of Figure 2.12 shows the variation in $b + d$ photometric magnitude from projection with respect to position on the size-mass plane. The magnitudes from the $b + d$ decompositions of each projection are consistent for the majority of the RIGs ($\Delta m_{r,\text{max-min}} \lesssim 0.1$ mag). However, a handful of galaxies with relatively large sizes and low total stellar mass ($\log(rhm_*/\text{kpc}) \gtrsim 0.8$ and $\log M_*/M_\odot \lesssim 10.8$) have large variations in magnitude. Galaxies with large variation, $\Delta m_{r,\text{max-min}} \gtrsim 0.5$, are identified as the same galaxies that showed significant systematically positive magnitude offsets for all SLD schemes in the previous section. The outliers in the SLD comparison were found to be caused by segmentation issues, so the same factor may also be the cause of the camera angle variations.

We confirmed that internal segmentation is the source of the large variations in magnitude between camera angles by visual inspection of the images and masks for our galaxy sample. Nearly all of the galaxies with large variations in apparent magnitude are extended discs (and some irregulars) – for which alternative camera angles allow the largest variation in projected surface brightness distribution. For example, it is less likely for substructure in edge-on discs to be deblended from the rest of the galaxy because the general condition for deblending is that the source must be locally discrete in surface brightness. Face-on discs are more prone to segmentation because substructure is more likely to dominate the local surface brightness distribution where that local flux is integrated along the thinnest axis of the disc. The deblending is strongly affected by changes in how these features are distributed with respect to the orientation of the galaxy. The largest variations occurred when a galaxy is strongly segmented in particular projections and not others – such as when a galaxy is face-on in one projection and edge-on in another, respectively. Though we only show random variation in Figure 2.12, we inspected the systematic offsets of our integrated magnitude estimates relative to the synthetic images. All magnitude estimates computed from the models are positively offset in all projections of diffuse

⁴The differences between maximum and minimum parameter estimates will be more akin to the 2σ than 1σ statistical estimates.

galaxies with large amounts of substructure. The middle panel of Figure 2.12 shows variations in half-light radius for the RIGs in the **CAMERAS** catalog. As expected, the variations in half-light radius roughly mirror the variations in the flux for each galaxy.

The right panel of Figure 2.12 shows the random errors on bulge-to-total light ratio from camera angle. Since many galaxies have estimated (b/t) values of ~ 0 , quoted errors are measured using linear differences. There is an apparent relationship between the variation in (b/t) and the median (b/t) of the four projections. Galaxies with intermediate median (b/t) estimates show the largest variations in (b/t) estimates in each projection while high and small (b/t) estimates have smaller variations. The implication is that (b/t) estimates can vary significantly with projection when significant bulge and disc components both exist in a galaxy, but vary weakly when one component dominates.

2.4.3 Environment and Crowding

In previous sections, we applied controlled experiments to investigate the variation in recovered parameters for specific observational biases. In each case, the same manually selected, uncrowded location in the SDSS with low sky variance was selected for analysis of each galaxy in the RIG sample. We took these measures to ensure that the variation in estimates for each galaxy in these experiments would not be further biased by variations from factors that depend on the location into which a galaxy is inserted into the SDSS. In this section, we remove our placement constraint in order to investigate the random and systematic variation in parameter estimates that are associated with location. For this experiment, we employ the decompositions from the **ASKA** catalog. The **ASKA** catalog uses synthetic images of galaxies from the RIG sample that are generated with the fiducial SLD scheme. Decompositions of each RIG were performed in approximately 100 unique locations in the SDSS and for all four camera angles⁵ at each location to build the **ASKA** catalog (see Section 2.3.3). We analyze the ~ 100 decompositions for each simulated galaxy and projection to explore the effects of external crowding.

As stated in Section 2.3.3, the galaxies in RIG sample have been visually inspected

⁵We state that we wish to remove any biases from projection, so the use of all 4 projections naively seems at odds with this goal. However, it is always possible to extract and compare results only for a specified projection of a galaxy – as we do. We also wish to perform a sanity check that any trends for the variation in parameters for galaxies holding each camera angle fixed, respectively, are qualitatively similar. Analysis in all four camera angles at each location facilitates such a check.

to ensure that there are no projection effects from other objects within the subhalo for any camera angle and that there are no *obvious* structural disturbances from merger activity or pair proximity. The absence of projection effects from other sources in the subhalo’s FoF group ensures that the statistics for crowding are uniform and the sources of crowding are external for all galaxies in all camera angles. Galaxies that are prone to internal segmentation are not excluded, since they represent a *bona fide* population of galaxies in Illustris. As stated in the description of the ASKA catalog in Section 2.3.3, we selected galaxies quasi-randomly in the plane of stellar mass and half-mass radius with the aim of characterizing the biases in photometric and structural parameter estimates on this fundamental size-mass parameter surface. The stellar mass and half-mass radius are easily computed as they are integrated properties of the particles belonging to a subhalo in the simulation. Furthermore, in the absence of dust, these properties ought to roughly scale to luminosity and size.

Our experimental results are comprised of $\sim 40,000$ fits from the 100 RIGs that are each placed in ~ 100 unique locations in the SDSS with analyses of all 4 camera angles at each location. We compute the random variation and median systematic offsets using the distribution of best-fitting $(b+d)$ model parameters for each camera angle of the RIGs. We employ the same metrics from previous sections: integrated magnitude, circular aperture half-light radius, and photometric (b/t) . The random variation in parameter estimates is computed using the 16th – 84th percentile range in the distribution of estimates for each galaxy. We also compute the systematic errors by comparing the median in the distribution of parameter estimates for each galaxy with the respective values computed from the synthetic images. The median systematic offset in r -band integrated magnitude, $\Delta m_{r,50\%}$, for example, is computed from the difference between the median integrated magnitude in the distribution of estimates from the models to the integrated magnitude of the synthetic image. Median systematic offsets in half-light radius $\Delta \text{rhl}_{r,50\%}$ are computed similarly – taking the ratio of the median half-light radius from the distribution of best-fitting models parameters and the half-light radius computed from photometry of the synthetic image.

Random Variation

Figure 2.13 shows the random variation in parameter estimates from crowding effects for the RIG sample. As in Figure 2.12, the greyscale shows the distribution of the full

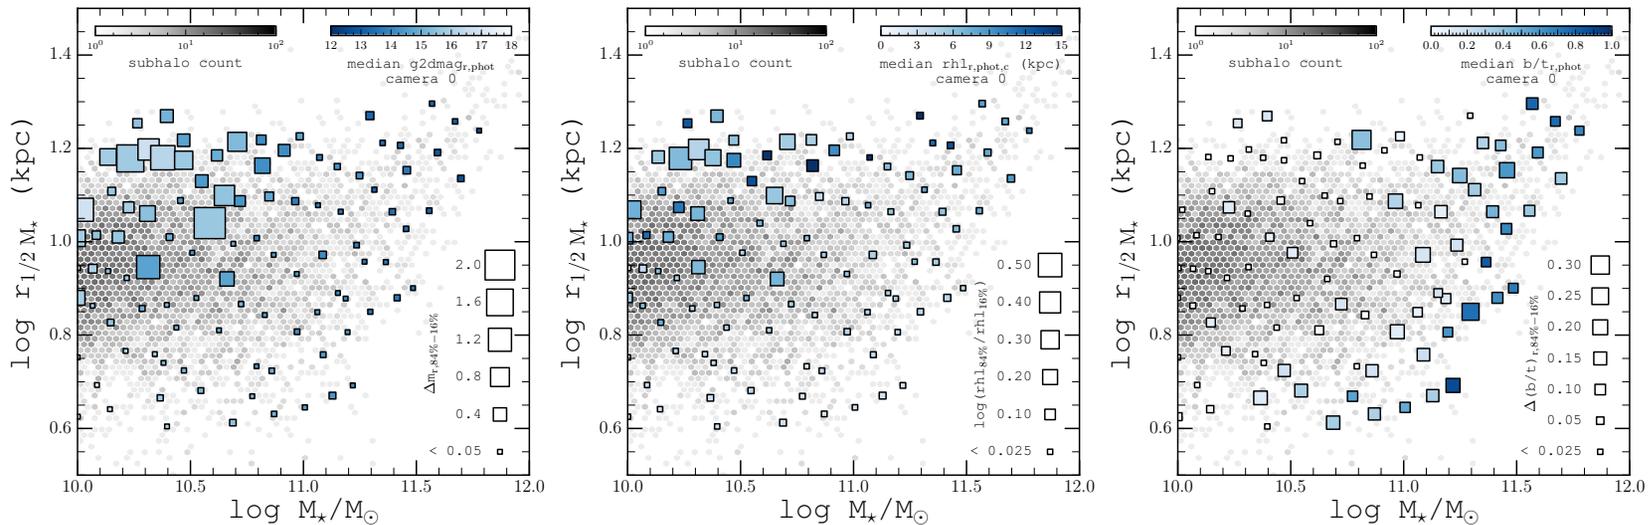


Figure 2.13: Characterization of random error from crowding on r -band photometric magnitude (*left panel*), half-light radius (*middle panel*), and photometric bulge-to-total ratio (*right panel*) where each galaxy from the ASKA catalog is projected onto its position in the plane of stellar half-mass radius, $r_{1/2M_*}$, and total stellar mass, M_* . The *2D greyscale histogram* in each panel shows the distribution of the full T15 catalog of Illustris subhalos in on this parameter surface with colours on a logarithmic scale and corresponding colourbar on the top left in each panel. Coloured boxes are positioned in this plane according to the properties of each RIG subhalo. The colour of each square indicates the median in the distribution of estimators obtained from placing each galaxy in an average of 100 randomly selected SDSS locations with fixed camera angle (CAMERA 0 in all cases here). The corresponding colourbar is located on the top right of each panel. The size of each square shows the random error that can be expected for a parameter estimate for an arbitrary SDSS location that is computed by comparing the 16th and 84th percentiles of the distribution of best-fitting estimates for each galaxy. A key of random errors and sizes, along with the exact computation method, is located in the bottom right of each panel.

T15 catalog of Illustris subhalos in the size-mass plane of stellar half-mass radius and total stellar mass. The coloured boxes are positioned corresponding to the location of each RIG according to their intrinsic total stellar masses and half-mass radii. From left to right, the panels show the random error that one can expect on an estimate of r -band integrated magnitude, rest-frame half-light radius in kpc, and photometric bulge-to-total fraction for an arbitrary placement in the SDSS. The colour of each square denotes the median of the distribution of parameter estimates for the galaxy. The size of each square is set by the 16th – 84th percentile range in the distribution of estimates for each galaxy. The exact computation for the comparison is given to the left of the key for random error and square size, located at the lower right of each panel. We show only the results for `camera 0` decompositions of each galaxy in the Figure because the results are qualitatively similar for all camera angles.

The left panel of Figure 2.13 shows the 16th - 84th percentile range in the integrated magnitude estimates derived from the GIM2D ($b + d$) models. We have verified that the $\pm 1\sigma$ interval about the mean assuming normally distributed parameter estimates provides very similar results. As with the camera angle tests, many RIGs with total stellar mass $\log M_\star/M_\odot \lesssim 11$ and half-mass radius $\log(rhm_\star/\text{kpc}) \gtrsim 0.8$ have an unusually high sensitivity to placement inferred from the range in their parameter estimates. Meanwhile, the remainder of our sample is largely robust to biases associated with location in the SDSS. Typical random errors, excluding galaxies in the upper left corner, are $\Delta m_{r,84\%-16\%} \lesssim 0.05$ mags which corresponds to flux variations of approximately $\pm 5\%$ in flux about the median. For the relatively low mass, highly extended galaxies at the upper left, random errors are as large as $\Delta m_{r,84\%-16\%} \approx 2$ mags, which roughly correspond to flux variations by factors of 6 about the quoted median.

The variation in estimates of circular aperture half-light radius shown in the middle panel of Figure 2.13 mirrors the variation in the integrated magnitudes, i.e galaxies that demonstrate weak sensitivity to placement also have small variation in their size estimates. Galaxies with large random variation on the magnitude estimates have correspondingly large variation in the size estimates – which are again confined to a particular region of the size-mass space. Precise size estimates, as with flux, for a given galaxy has two requirements: accurate evaluation of the sky level and proper delineation of boundaries between multiple sources. The fact that every galaxy in our sample statistically experiences the same variations in crowding and sky level indicates that the large sensitivity to placement for extended but relatively low mass

galaxies is driven by their morphologies. Note also that the median estimates for half-light radius in each galaxy, denoted by the colours in the boxes, do not demonstrate a clear correlation with half-mass radius among galaxies with large random variation. Further statement about the driving source of the large random errors requires inspection of the the images and *systematic* errors on these estimates. But the fact that the highly sensitive galaxies we see here is the same troublesome group with high sensitivity to camera angle and large systematic offsets in all SLD schemes gives a strong case for internal segmentation as the source of the problem.

The characterization of the random error on $(b/t)_r$ is shown in the right panel of Figure 2.13. A similar relationship between the variation in (b/t) and the median (b/t) that was seen in the camera angle analysis is also manifested here. Galaxies with intermediate median values of (b/t) have the largest random errors about these medians. This is expected because the separation of the bulge and disc components is the most challenging in decompositions of galaxies with simultaneously significant bulge and disc components. Furthermore, the significance of the bulge component is sensitive to the accuracy with which the location and light profile of the peak in the bulge surface brightness distribution can be inferred – which may be affected by placement specific biases such as crowding. Accurate modelling of the surface brightness profile at large radii where the S/N becomes small will also provide variation in both components’ flux estimates. Galaxies with median $(b/t)_r > 0.8$ and $(b/t)_r < 0.2$ have relatively small random errors. We also note that the variation in estimates for the low-mass, extended galaxies in the upper left of the distribution does not mirror the variations in size and magnitude estimates shown in the other panels, and all have median (b/t) estimates that are close to zero. In general, it appears that the variation in (b/t) that arises from biases associated with placement are not correlated with variations in integrated magnitude and half-light radius estimates.

Systematic Offsets

We have seen that the estimates of integrated magnitude and half-light radius are robust to sky position for all but a handful of extended low-mass discs and irregulars which show poor precision as a function of sky position and camera angle. We now inspect the median systematic offsets for integrated magnitude and half-light radius relative to the synthetic images to examine accuracy of our measurements. The systematics offsets will inform us on adequacy of our source deblending scheme in allo-

cating pixels to the fitting that facilitate consistent estimates of flux and photometric size in crowded environments and against real sky backgrounds. Furthermore, they will provide insight on the source of the high sensitivity to placement of integrated magnitude and half-light radius estimates seen among the diffuse discs and irregulars in our sample.

Figure 2.14 shows the systematic errors on our integrated magnitude estimates and half-light radii from the bulge+disc fitting of the RIG sample. The properties of each figure are similar to those described for Figure 2.13, with the difference that we now look at median statistics of each galaxy relative to the the respective properties of the synthetic images before realism is added. $\Delta m_{r,50\%}$ denotes the difference between the integrated magnitude computed from the total flux in the synthetic images and the median in the distribution on estimates for integrated magnitude from the $(b + d)$ decompositions (input- $b + d$).⁶ Similarly, $\Delta r_{hl,50\%}$ is the logarithm of the ratio between the median in the distribution of half-light radii over all placements and the half-light radius computed from the synthetic image. The colour of the borders on each square denote negative (*red*) and positive (*black*) systematics according to their computation, shown along the keys in each panel. The colour of each square shows the respective input integrated magnitudes and half-light radii from the synthetic images against which each median statistic is compared.

The left panel of Figure 2.14 shows the median systematic offsets for the distributions of best-fitting $b + d$ model magnitudes of the RIGs over all placements for a single camera angle. We show only the results for each galaxy in a single camera angle because we find that the general trends are the same when holding each other camera angle fixed. Systematic offsets in magnitude are small apart from the usual suspects and a few of the higher mass extended galaxies. Furthermore the systematics are randomly distributed about $\Delta m_{r,50\%} = 0$. The random distribution of systematic small systematic offsets indicates that the effects of crowding and other placement-sensitive biases do not systematically affect estimates of the flux. Crowding and other positional biases only provide scatter about $\Delta m_{r,50\%} = 0$ for the majority of our sample of galaxies and morphologies. The second panel shows the systematic offsets of half-light radius estimates. The half-light radius offsets mirror of the systematics on integrated magnitude – both for the majority of the sample and for the diffuse galaxies with large negative systematic offsets.

⁶The unconventional calculation of the difference using (input- $b + d$) rather than $(b + d$ -input) is solely for consistency in the colours of the figure.

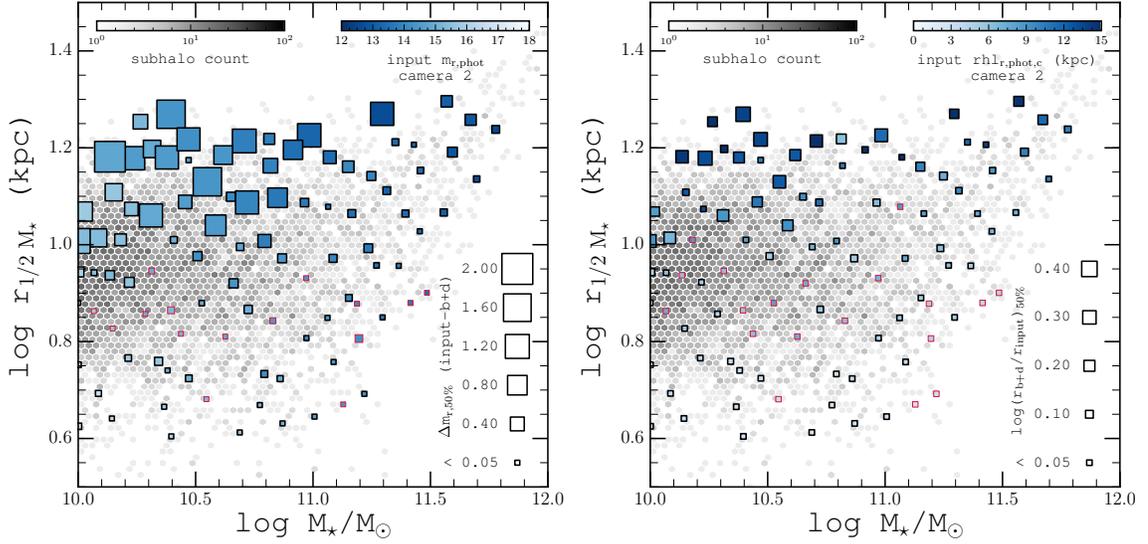


Figure 2.14: Characterization of systematic errors on best-fitting estimates of magnitude (*left panel*) and half-light radius in kpc (*right panel*). Figure properties are the same as for the first two panels of Figure 2.13 with the exceptions that the size of each coloured square denotes systematic error relative to the properties computed directly from the synthetic images and the border colour denotes whether the systematic indicated along the key is positive (*red*) or negative (*black*). Note that a negative systematic error in magnitude corresponds to a higher magnitude estimate from the model relative to the synthetic image and therefore a relative decrease in flux. Median systematic errors for the magnitudes in each galaxy are computed from difference between the integrated $b + d$ analytic model magnitudes (output) and the magnitude of the galaxy as calculated from the sum of the flux in the synthetic images before any realism is added (input). Similarly, output and input half-light radii are computed from circular aperture photometry of the best-fitting models and the synthetic images, respectively. Note that positive systematic errors are all very small and are randomly distributed amongst galaxies with small negative systematic errors – indicating that crowding in these galaxies provides weak variations about 1:1 flux recovery for these galaxies, but does not systematically drive the best-fitting model estimates in a particular direction.

In summary, our bulge+disc decomposition results for the majority of the RIGs show weak sensitivity to the the biases associated with placement – which principally include: crowding, sky background, and PSF resolution. However, a recurring handful of diffuse galaxies have large systematically negative median offsets in both integrated magnitude and half-light radius – consistent with the systematics seen in the smoothing analysis in Section 2.4.1. The corresponding random variations for these galaxies are also large and show strong sensitivity to both placement and camera angle (Section 2.4.2). The discrepancies that we have highlighted for these diffuse galaxies are the following: systematic under-estimations of flux and size over all environments; significantly greater random sensitivity to placement, projection, and SLD scheme than other galaxies on the size-mass plane; and that errors of this magnitude do not exist in our analyses of real galaxies. We explore the source that drives the unusually high sensitivity to observational biases in Illustris’ diffuse galaxies in the next section.

2.4.4 Internal Segmentation by Artificially Discrete Substructure

The morphological sub-population of diffuse galaxies in the RIG sample have consistently demonstrated high sensitivity to each observational bias. We found in Section 2.4.1 that most of these galaxies had large systematic errors that varied in magnitude with the choice of SLD scheme. We showed that small changes to the surface brightness distributions of the diffuse galaxies can change the measured half-light radius by up to a factor of 4. Small changes to the surface brightness distributions of these galaxies, such as those produced in alternative SLD schemes, are exacerbated by the observational biases of PSF resolution, S/N, projection, and crowding. In this section, we investigate the reasons that observational biases on the surface brightness distributions of the diffuse RIGs generate large random and systematic errors in their best-fitting model estimates.

We begin by comparing the effects of placement on the diffuse RIGs with an example taken from the rest of the population. Figure 2.15 first shows an example of the robustness of our structural and photometric estimates to crowding effects that is representative among the majority of the RIGs. Each row shows the decompositions for the same galaxy in a different location in the SDSS. We use the *bflags* parameter to quantify the segmentation of a galaxy from internal and external sources. As

outlined in Section 2.3.1, SExtractor allocates each pixel in an image to a source by assigning it a flag in the segmentation map. The pixels that have the same flag as the central pixel of the synthetic image (i.e. belong to the same source) are always used in the fitting. $bflags$ is computed as the number of uniquely flagged sources that are directly contiguous to the source pixels being used in the fitting. For example, in the middle row of Figure 2.15, the galaxy is located in a relatively uncrowded field with no external sources directly bordering the pixels used in the fitting. In this image incarnation, $bflags=0$. The top and bottom rows show the same galaxy with the same camera angle in locations where bright stars contaminate the line of sight and significantly crowd the galaxy. The delineation of boundaries between the galaxy and the external sources along with the corresponding $bflags$ number are shown in the corresponding science masks. The best-fitting $b + d$ parameter estimates listed in the second panel of each row show maximum variations on the order of a few percent in magnitude, half-light radius, and bulge-to-total fraction, respectively. The corresponding best-fitting ps results show remarkable consistency in integrated, half-light radius, and sérsic index. The quality of the residual images in each case aside, the residuals for the galaxy light profile are consistent in each case. The example in Figure 2.15 illustrates why the decomposition results for the majority of the RIGs show weak sensitivity to crowding – despite possibly significant variations in crowding by external sources. However, the best-fitting estimates for the diffuse galaxies in the RIG sample show high sensitivity to the observation biases from placement. Figure 2.16 shows an example of the variations that arise for the diffuse RIGs. As for Figure 2.15, each row of 2.16 shows the same galaxy in a different location in the SDSS. Galaxies such as the one shown in Figure 2.16 have intrinsically large amounts of segmentation, even in uncrowded fields; as shown in the image incarnation in the middle row along with the corresponding $bflags$ number in the science mask. The top and bottom rows of Figure 2.16 show that the observational biases associated with placement (crowding, sky background, and PSF resolution) drive significant variations in the segmentation maps of diffuse galaxies such that decomposition results are not consistent. In the top row of Figure 2.16, the presence of external sources in the form of two stars leads to a significantly higher estimate of the integrated absolute magnitude, $\Delta M_{r,(b+d)} = 0.3$ mag, relative to the uncrowded placement in the middle row. In the bottom row of Figure 2.16, the science mask is sufficiently segmented that a locally bright feature in the surface brightness distribution of the galaxy is identified as a distinct source. The collection of pixels that share the flag of the locally

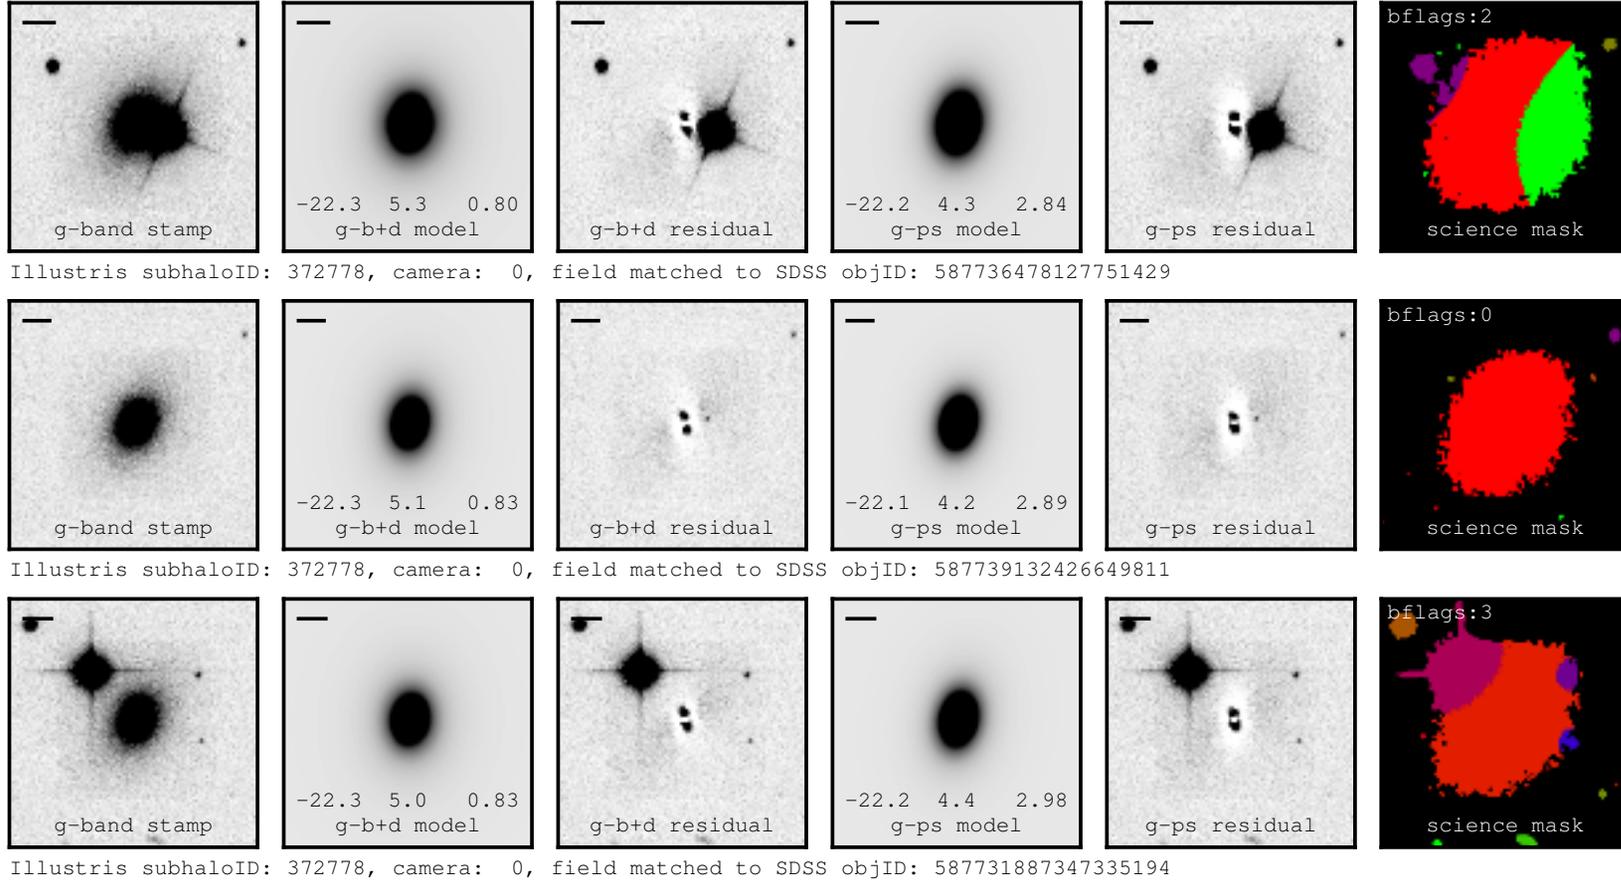


Figure 2.15: Examples of structural and photometric parameters varying weakly as a function of crowding. Each row of panels corresponds to the same camera angle of the same galaxy inserted in a unique SDSS field location. The middle row shows the galaxy in a relatively uncrowded environment. Note that this galaxy is not intrinsically prone to internal segmentation. The top and bottom rows show the galaxy where it happens to have been placed in environments where it is strongly crowded by bright field stars. The model parameters are nearly identical despite the significant variation in crowding effects and proximity of the additional source. A scale in the top left of each panel denotes 10 kpc at $z=0.05$.

distinct “source” includes the central pixel in the image (which is aligned with the gravitational potential minimum). Only pixels that have this flag (coloured yellow-green in the segmentation image) are used in the fitting. The segmentation leads to a factor of 5 reduction in the flux and a factor of 3 in the half-light radius compared to estimates where the galaxy is not so brutally shredded by segmentation (middle row of 2.16). In addition, each segmentation image (even in the uncrowded field) excludes a significant component of the galaxy’s flux that is bound in locally bright substructure. We inspected the fields into which the galaxies from the middle and bottom rows of Figure 2.16 were inserted to find no significant sources of crowding. Therefore, crowding, sky background variations, PSF resolution, or combinations thereof can all provide sufficient modification to the surface brightness distribution to cause additional *internal* segmentation of galaxies but only when a galaxy has locally discrete substructures in its surface brightness distribution that make it prone to internal segmentation.

The most extreme scenarios where we see differences as large as $\Delta m_{r,50\%} \approx 2$ mag in Figure 2.14 (a factor of 6.3 decrease in total flux) arise from two possible situations: (1) a large fraction of the galaxy’s light is bound in substructure components that are not included in the fitting; (2) discrete substructure is projected relative to the galaxy such that it is identified as the source on which to perform the fitting (which occurs when the substructure is discrete in a particular projection and overlaps with the image centre). In either scenario, the galaxy is not be identified as a single source, but is a contiguous collection of deblended features that each have a unique flag in the mask. Figure 2.17 shows two examples of galaxies whose science masks are shredded by segmentation from discrete substructure that contains appreciable fractions of the galaxies’ total fluxes. The galaxies in the top and bottom rows of Figure 2.17 have stellar masses $\log M_{\star}/M_{\odot} = 11.3$ and $\log M_{\star}/M_{\odot} = 10.3$, respectively. The supplementary synthetic colour images in right-most panel of each row show blue rings of containing knots of substructure in the form of young and highly luminous populations of stars that seemingly orbit at a fixed radius from the galactic centre. While it is *visually* apparent that both $b + d$ and ps models reproduce the fraction of the surface brightness distributions of the galaxies that excludes the rings, the residuals and masks demonstrate that a significant fraction of the total flux is lost because the bright substructure is masked out. For the galaxy in the first row of Figure 2.17, the median ratio of the fluxes determined by the $b + d$ decompositions to those of the synthetic images is $f_r/f_{r,\text{synth}} = 0.35$ across all environments, meaning

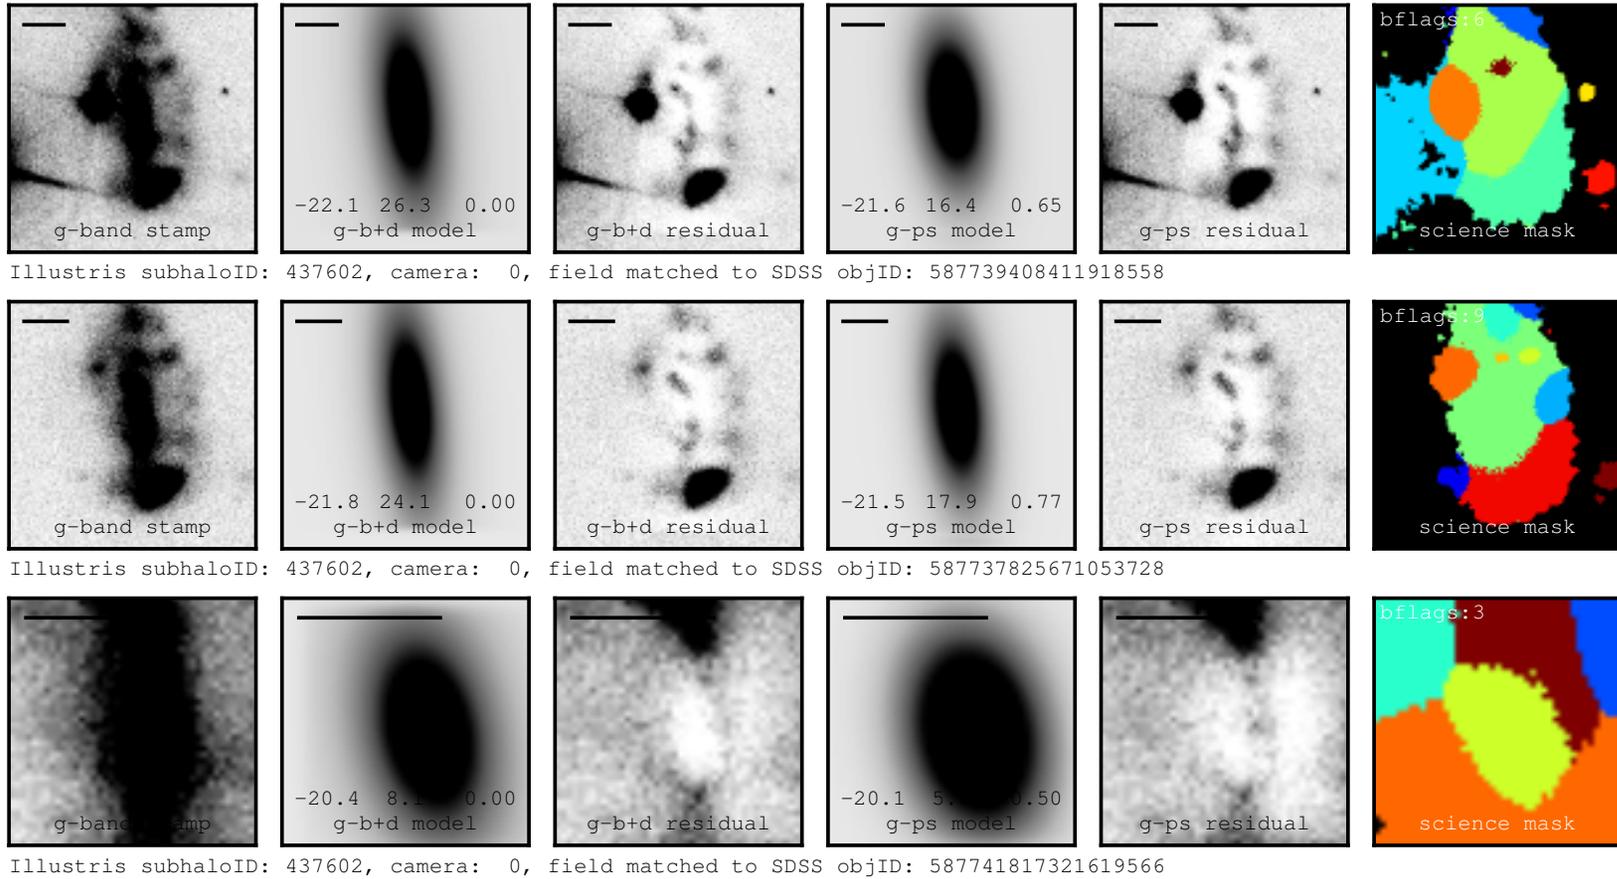
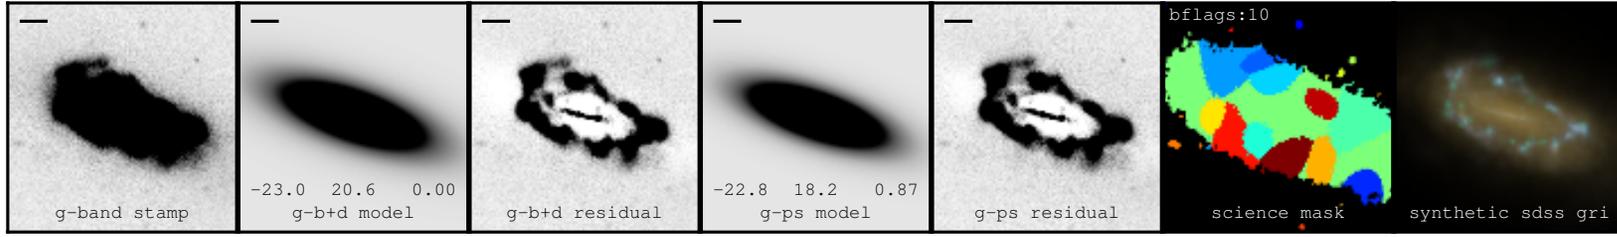
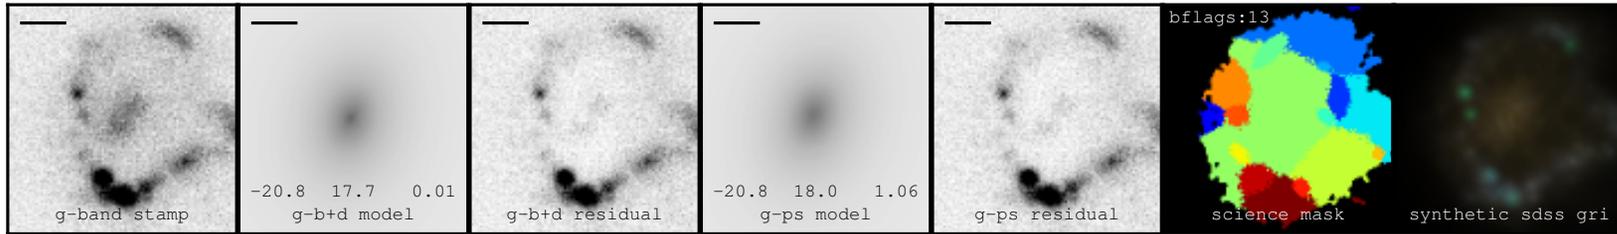


Figure 2.16: Examples of structural and photometric parameters varying strongly as a function of internal segmentation. Each row of panels corresponds to the same camera angle of the same galaxy inserted in a unique SDSS field location. Variations in the crowding by external sources and sky variance provide unique variations in the segmentation map for each placement. When there are intrinsically large amounts of internal segmentation due to high surface-brightness substructure within the galaxy, further change to the segmentation map through crowding can provide significant random error on the best-fitting estimates. Changes to the colourmap are made in the science masks of this mosaic to aid in visual distinction between zones in highly segmented systems. A scale in the top left of each panel denotes 10 kpc at $z=0.05$.



Illustris subhaloID: 332891, camera: 2, field matched to SDSS objID: 588017604159602881



Illustris subhaloID: 446634, camera: 2, field matched to SDSS objID: 587727225694126245

Figure 2.17: Examples of highly extended galaxies where luminous and localized substructure creates a highly segmented science mask for the galaxy, despite no significant presence of external crowding. The last panel in each row shows the synthetic SDSS *gri* image of each galaxy. These two galaxies (subhaloIDs 332891 in the *upper row* and 446634 in the *lower row*) have half mass radii of 18.6 kpc and 15.8 kpc and total stellar masses, $\log M_{\star}/M_{\odot}$, of 11.29 and 10.31, respectively. Their best-fitting model r-band half-light radii are 13.3 kpc and 15.0 kpc. The galaxy in the lower row is significantly more diffuse, with a total stellar mass that is an order of magnitude lower than the galaxy in the upper row but similar size when size is quantified by both mass and photometry. Bright green in the science mask denotes the pixels allocated to fitting. A scale in the top left of each panel denotes 10 kpc at $z=0.05$.

that roughly 2/3 of the total flux is in the substructure. The galaxy in the second row is 10 times less massive, but has roughly the same half-mass radius – making it significantly more diffuse. The median flux fraction for this galaxy is $f_r/f_{r,\text{synth}} = 0.56$ – so roughly half of the light is locked up in masked substructure. In each case, the substructure systematically represents a significant fraction of the total stellar light. While variation in sky and crowding will produce variations in what substructure is masked, the large negative systematic errors on flux and photometric size arise from the exclusion of bright substructure from the fitting.

The final piece of evidence that internal segmentation is driving the large random and systematic biases among diffuse galaxies is shown in the median *bflags* estimates over all placements and projections of each galaxy on the half-mass radius and total stellar mass plane, shown in Figure 2.18. The *bflags* parameter is not sensitive to whether the source of segmentation is internal from discrete substructure or external from crowding. However, our procedure for creating the ASKA catalog dictates that the placement of any particular galaxy is random. Therefore, the distribution of best-fitting model parameter estimates for each galaxy over all placements is affected by the same crowding *statistics* as any other and *bflags* should be roughly uniform for each galaxy. Although the *bflags* value should not depend on galaxy properties, Figure 2.18 shows that the diffuse galaxies have significantly greater median *bflags* estimates than the rest of the population. This indicates that the intrinsically large median *bflags* estimates are driven by internal segmentation. The fact that the galaxies with large amounts of internal segmentation seen here directly coincide with ones having high sensitivity to observational biases confirms that internal segmentation by discrete substructure is the source of the sensitivity.

The properties of the diffuse RIGs explain their high sensitivity to observational biases. Figure 2.14 showed that large median systematic offsets in integrated magnitude and half-light radius are seen for morphologically diffuse galaxies. The source of the systematic offsets is the deblending of locally discrete substructure that contain appreciable fractions of the flux. The sensitivity of the best-fitting model estimates to internal segmentation in diffuse galaxies is exacerbated by observational biases. The biases provide unique variations in the the segmentation that generate the large random variation and inconsistency between estimates in different environments, projections, and SLD schemes seen in Figures 2.13, 2.12, and 2.9, respectively. Galaxies for which (1) the majority of the galaxy’s total stellar light is contained in masked substructure or (2) a locally discrete feature of the galaxy is systematically identified

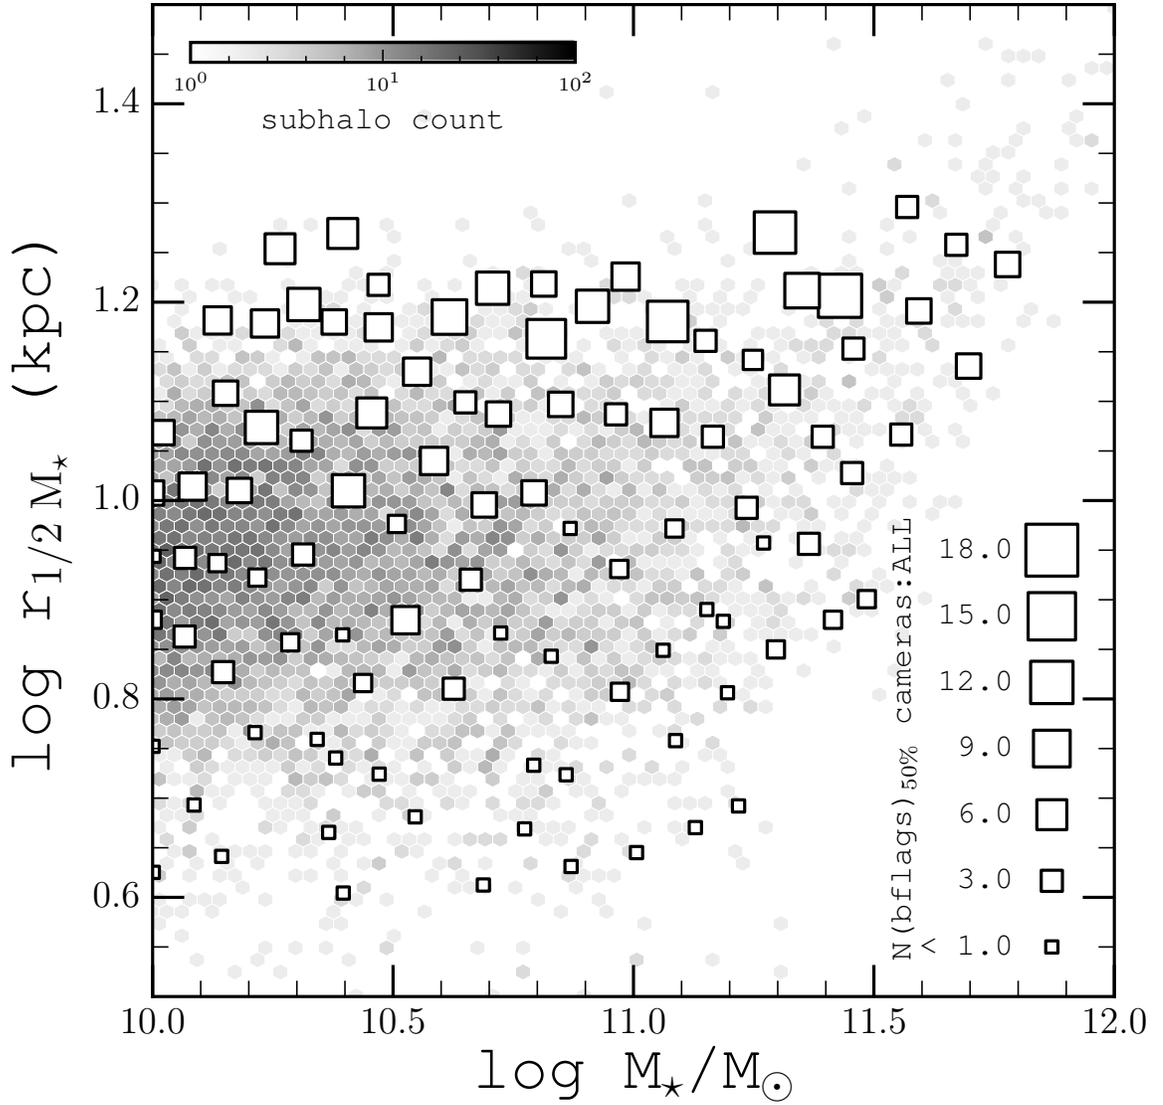


Figure 2.18: The median number of *bflags* across all placements and camera angles for each RIG. *bflags* is computed as the number of flagged sources that directly border pixels that are flagged as science pixels (pixels that are used to determine the likelihood of the models). It is a parameter that is sensitive to external crowding and internal segmentation. The locations of each square denotes the position of a RIG on the plane of total stellar mass and half-mass radius. The size of the square indicates the median *bflags* parameter over all placements and camera angles in the ASKA catalog. Despite uniform crowding statistics between galaxies, the extended and diffuse population of Illustris have intrinsically large amounts of segmentation.

as the source on which to perform fitting have the largest negative systematic offsets in flux and half-light radius. Conversely, galaxies with compact stellar light distributions are less likely to have their substructure deblended and therefore show little to no variation with observational biases and have best-fitting model properties that are consistent with those of the synthetic images.

In the next section, we investigate the systematic offsets on estimates of apparent integrated magnitude and circular aperture half-light radius for the full Illustris galaxy population to characterize the influence of these biases in comparisons with real galaxies.

2.4.5 Systematics of Size and Flux Estimates in the Full Illustris Galaxy Population

In previous sections, we have shown that our estimates of integrated magnitude and half-light radius are generally robust to post-processing biases in the form of SLD schemes, and the observational biases of projection, crowding, PSF resolution, and sky brightness with the exception of a population of diffuse galaxies in Illustris with intrinsically large amounts of discrete substructure. In this section, we investigate the systematics on total flux and half-light radius estimates in the full Illustris population using the DISTINCT catalog. The catalog contains the best-fitting model estimates from a single $b + d$ and ps decomposition of each camera angle for every galaxy in the Illustris synthetic image catalog of T15. Placement of each camera angle projection the galaxies is randomized following the procedure described in Section 2.2.3. Any systematic offset between the decomposition results and the derived properties of the synthetic images for a particular galaxy or morphology generally arises from the inadequacy or inflexibility of our quantitative morphology analysis in handling the particular surface brightness distribution of the galaxy in its environment. Our analysis will allow us to characterize the collective systematic biases of that arise from morphology and observational realism in the full Illustris population including the prevalence of bias from internal segmentation. For simplicity, in this section we will refer to every *insertion of a galaxy* simply as a galaxy because there are four decompositions of each Illustris galaxy in the catalog (one for each camera angle).

Recovery of the Integrated Flux

We inspect the systematics on r -band integrated flux, f_r , by comparing our estimates derived from the best-fitting $b + d$ models of galaxies with observational realism to the sum of the flux in the corresponding r -band synthetic images, $f_{r,\text{synth}}$ in Figure 2.19. The apparent integrated magnitude estimates, m_r and $m_{r,\text{synth}}$, used in previous sections are computed directly from the respective integrated fluxes, f_r and $f_{r,\text{synth}}$, of the model and synthetic images. Plotted in the lower panels of Figure 2.19 is the ratio of recovered to input fluxes against input apparent magnitudes. In the lower left panel of Figure, we show the distribution of flux recovery for the DISTINCT catalog – which includes a single fit for every galaxy in the synthetic image catalog of T15. In the lower right panel, we show the same plot for galaxies that have $bflags < 2$. The images in the upper panels correspond to numbered outliers in the lower right panel and will visually aid in characterizing these outliers.

The lower left panel of Figure 2.19 shows the distribution of systematic flux offsets for all galaxies from the DISTINCT catalog. The distribution demonstrates that the majority of flux estimates that are consistent with the fluxes computed from the synthetic images. However, roughly 30% of galaxies across the full range of magnitude have systematically negative flux offsets. In previous sections, we have demonstrated that the only bias that is compatible with these large systematic offsets in flux is internal segmentation by discrete substructure in relatively diffuse galaxies. In the lower right panel of Figure 2.19, we attempt to visualize the systematics on flux estimates in the absence of the population of internally segmented galaxies in the DISTINCT catalog by performing a $bflags < 2$ cut. We find that the cut largely eliminates the population of galaxies with systematically poor flux recovery – with only a few remaining outliers. However, the cut also significantly reduces the density of galaxies that reside around $\log(f_r/f_{r,\text{synth}}) = 0$ for the full magnitude range. The loss of density among galaxies with accurate flux recovery is not surprising, however, when we recall that genuine external crowding (which $bflags$ is also sensitive to) in the absence of internal segmentation provides no systematic bias on the integrated fluxes and sizes. Therefore, the cut is also eliminating galaxies with accurate flux recovery but experienced crowding by external sources in their placements.

We showed in Section 2.4.3 and 2.4.3 that our estimates of flux and size in placements with various degrees of crowding are consistent for the majority of the RIGs with no discernible systematic biases. In an effort to roughly quantify the number

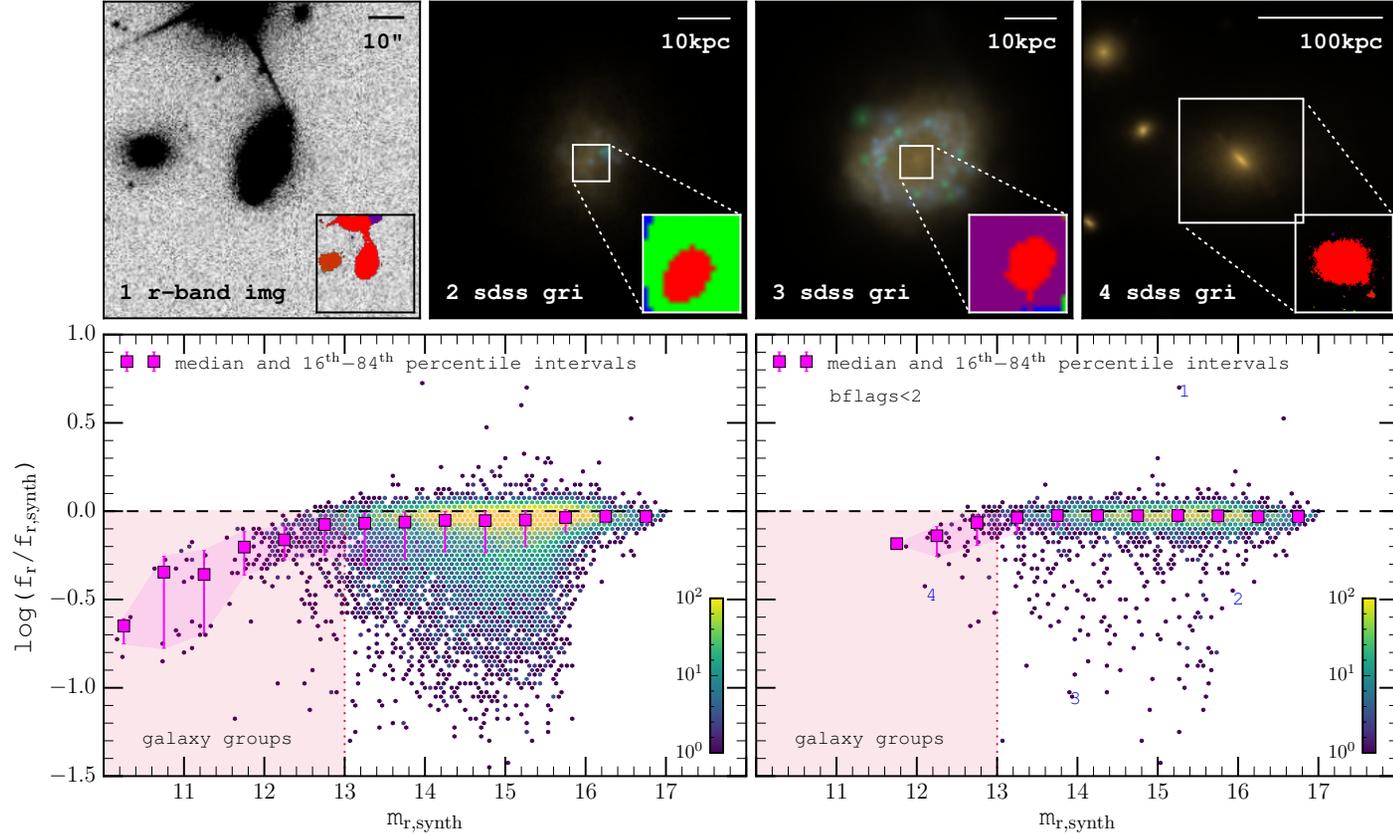


Figure 2.19: *Lower Panels:* The distributions of recovered, f_r , and synthetic image, $f_{r,\text{synth}}$, total r -band fluxes plotted against the apparent magnitudes derived from the sum of the flux in the synthetic images, $m_{r,\text{sim}}$ for the DISTINCT catalog. The lower left panel shows the distribution of flux offsets for the full DISTINCT catalog. The lower right panel shows only galaxies for which the number of flagged sources that directly border the pixels used in the fitting is $b\text{flags} < 2$. The *black dashed line* in each of the lower panels denotes a 1:1 relation between model and synthetic image fluxes. *Magenta symbols* show the 1-dimensional medians of recovered flux fraction with binning along $m_{r,\text{synth}}$. Errorbars denote the 16-84% percentile range about the median. The *salmon region* that is left of the *red dotted line* at $m_{r,\text{synth}} = 13$ denotes the domain of $m_{r,\text{synth}}$ computed from subhalos that predominantly contain galaxy groups. *Upper Panels:* characterization of outliers in the DISTINCT catalog where $b\text{flags} < 2$, numbered 1-4 in the lower right panel.

of galaxies with accurate flux estimates but that were culled because of the *bflags* parameter’s insensitivity to the source of segmentation (whether internally from discrete substructure or externally from crowding), we compare the number of galaxies in the range $-0.1 < \log(f_r/f_{r,\text{synth}}) < 0.1$ before and after the $bflags < 2$ cut. The total number of galaxies in the left and right panels of Figure 2.19 are $\sim 28,000$ and $\sim 7,700$, respectively, so 72% of galaxies in the full population do not satisfy the cut. We find that the number of galaxies that satisfy $-0.1 < \log(f_r/f_{r,\text{synth}}) < 0.1$ are 19,300 (70%) in the DISTINCT catalog (left panel) and 7,100 (92%) after the $bflags < 2$ cut (right panel). Therefore, approximately 12,200 galaxies have $bflags \geq 2$ due to genuine crowding effects or weak internal segmentation that does not strongly affect the overall surface brightness distribution of the galaxy. The contribution to the density of galaxies around $\log(f_r/f_{r,\text{synth}}) = 0$ provided by these galaxies are *accidentally* removed in the *bflags* cut as a result of the sensitivity of *bflags* to external crowding. As an example, note that all of the systems in Figure 2.8 would be excluded under the $bflags < 2$, yet in most cases have visually satisfying residuals and prove to have accurate flux estimates relative to their synthetic images (not shown in the Figures). Their *bflags* number is indicated in the top left corner of each science image mask. Taken together, our analysis indicates that roughly 8,000 (30%) galaxies in the DISTINCT catalog are affected by strong internal segmentation or other bias that has so far not been identified.

The upper panels of Figure 2.19 show four example galaxies that are significant outliers from $\log(f_r/f_{r,\text{synth}}) = 0$ that survive the $bflags < 2$ cut. Panel (1) shows the r-band science image of a galaxy that has been placed in the close vicinity of a bright star such that the galaxy’s surface brightness distribution overlaps partially with a diffraction spike. Inset in this panel is the associated science mask where red corresponds to the flag for pixels that are used in the fitting, other colours correspond to additional sources that do not, and sky pixels are transparent. In this case the galaxy’s flux is insufficient to warrant branching from the flux tree of the star as detailed in Section 2.3.1. Cases such as this are rare. We find that this scenario is common to each of the few large systematically positive outliers.

In panels (2) and (3) of Figure 2.19 we illustrate the scenarios that we found to be common amongst the systematically negative outliers (apart from those in the salmon region left of $m_{r,\text{synth}} = 13$). A discrete component at or near the centre of the galaxies’ surface brightness distribution has been separated from the larger structure due to internal segmentation. Galaxies such as these survive the cull $bflags < 2$ cull

because the pixels allocated for fitting are bordered by only a single non-sky source – but still have the undesirable properties of discrete substructure that lead to negative systematic bias in flux estimates. Shown are the SDSS *gri* false-colour images of these galaxies without any realism with a zoomed inset of the post-realism science image mask. In panel (2) a small young stellar population has been delineated from the rest of the galaxy but contains the pixel that is centred on the gravitational potential minimum of the galaxy. The substructure is therefore allocated as the sole source on which the fitting is performed, but contains only a fraction of the galaxy’s total flux.

Similarly, in panel (3) of Figure 2.19 we demonstrate that a discrete nuclear component that is embedded within a ring of star formation has been identified as the primary source. The ring-like morphologies such as this one are described in S15 with possible origins arising from choices of ISM and feedback models within Illustris.

We investigated the source of the negative systematic flux offsets for galaxies with $m_{r,\text{synth}} < 13$ shown in the salmon coloured region in each of the lower panels of Figure 2.19. The *gri* false-colour synthetic image of the galaxy subhalo in panel (4) shows that the galaxy that is decomposed in the fitting is nested at the gravitational potential minimum of a group containing several galaxies. Therefore, the apparent systematic error in the flux does not necessarily arise because of segmentation or problematic fitting. Instead, it is often a natural consequence of the total flux in the input image being the sum of the flux from all galaxies that belong to the subhalo. The fitting is performed only on the galaxy containing the minimum of the gravitational potential – which will contain a fraction of the total flux from the group. Only a handful of the groups on the bottom left panel satisfy the $bflags < 2$ criterion shown in the right panel, because there may be several projections for which the surface brightness distribution of a galaxy that belongs to a group is contaminated by other group members. The systematic descent with magnitude below $m_{r,\text{synth}} = 13$ is associated with groups that contain increasingly larger number of constituent galaxies and therefore increased total fluxes relative to the flux that can be obtained from the central. Such discrepancies should also be evident in estimates of the half-light radius of the best-fitting models and the synthetic images.

Recovery of the Half-Light Radius

In the last section, we characterized the systematics on estimates of the flux in comparisons between the integrated fluxes derived from the models and the synthetic

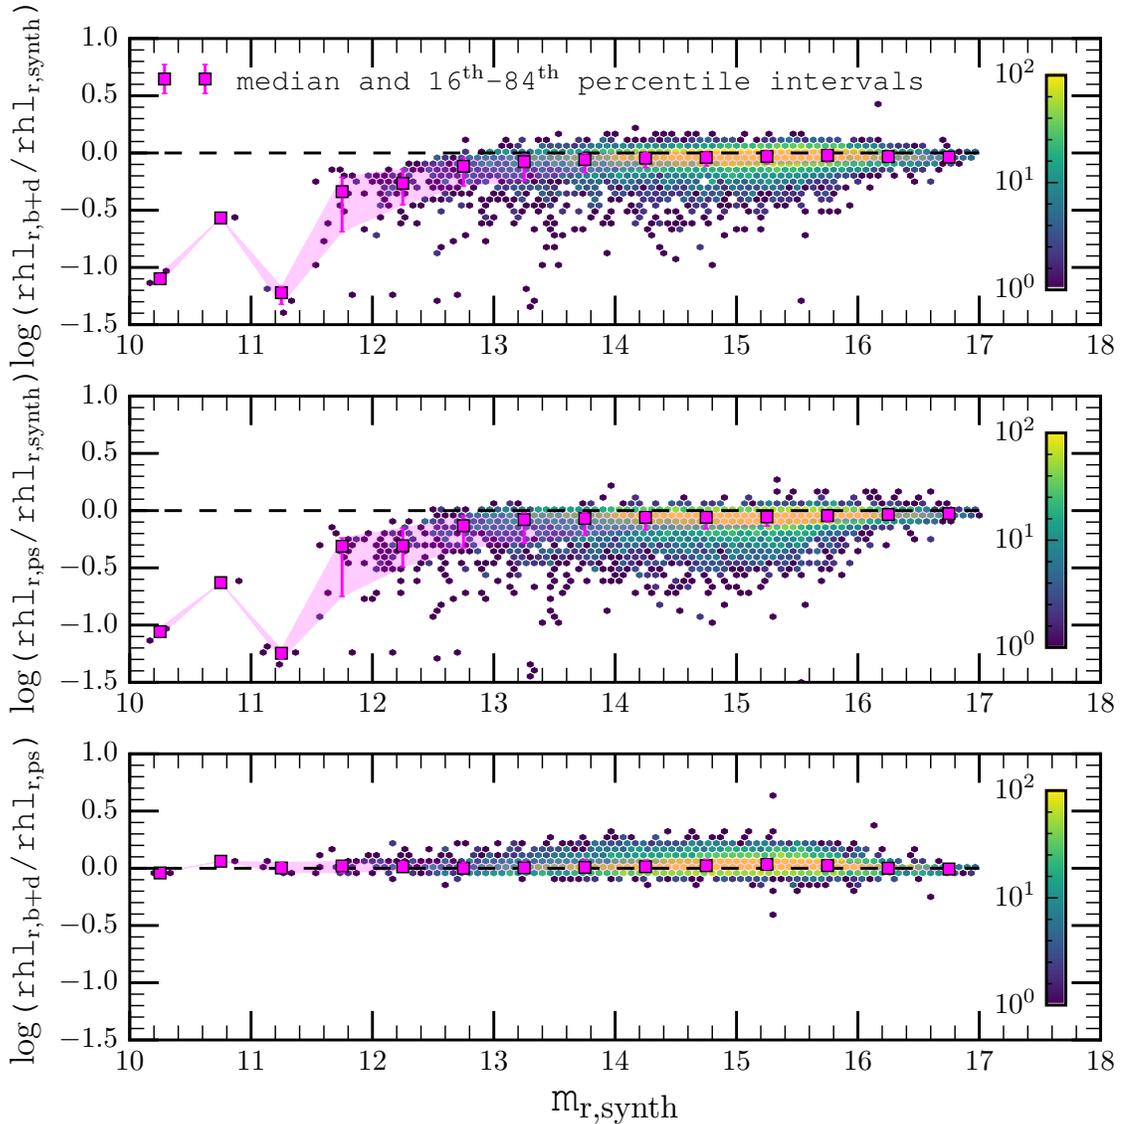


Figure 2.20: Comparisons of circular half-light radius estimates derived from aperture photometry on the models. *Top panel:* systematic offsets between the half-light radii computed from the $b + d$ models and the respective synthetic images without realism. *Middle panel:* systematic offsets between the half-light radii computed from the ps models and the respective synthetic images without realism. *Bottom panel:* comparison of the estimates of half-light radii computed from the best-fitting $b + d$ and ps models.

images for the full population of Illustris galaxies in the DISTINCT catalog. These systematics should be mirrored in the estimates of half-light radius due to the dependence of accurate half-light radius estimates on accurate flux recovery over the full surface brightness distribution of the galaxy. In this section, we perform comparisons of circular aperture half-light radii that are computed from the best-fitting models and synthetic images. Furthermore, we test the robustness of half-light radius estimates to the choice of surface brightness profile model by comparing the half-light radii derived from the best fitting $b + d$ (bulge: sérsic $n=4$, disc: sérsic $n=1$) models to those from the ps (sérsic n =free) models for each galaxy.

Figure 2.20 shows the photometrically derived size estimates in the DISTINCT catalog as a function of r -band input integrated apparent magnitude. In the top and middle panels, the r -band half-light radii that are computed from photometry of the $b + d$ and ps models, $rhl_{r,b+d}$ and $rhl_{r,ps}$, are compared with those of the raw synthetic images $rhl_{r,synth}$. In the bottom panel, the best-fitting $b + d$ and ps model estimates for each galaxy are compared directly with each other to check the consistency of estimates with the choice of model. The median and 16th – 84th percentile intervals are shown by the magenta points and filled region.

The top and middle panels of Figure 2.20 show the same systematic negative offsets at $m_{r,sim} < 13$ that reflect the downturn in the flux recovery for groups in the DISTINCT catalog. The majority of our estimates at $m_{r,synth} > 13$ are consistent with the derived properties of the galaxies from the synthetic images without observational realism. However, image segmentation leads to a population of galaxies with negative systematic errors in their size estimates. The similarity of the distributions indicates that results from the best fitting surface brightness profiles are affected similarly.

The bottom panel of Figure 2.20 shows the comparison of estimates of half-light radius from each surface brightness profile model directly. Any systematics that tend to affect estimates using one of the $b + d$ or ps models similarly affect the other. Half-light radius estimates from $b + d$ and ps models of galaxies with observational realism are generally consistent across the full magnitude range.

Chapter 3

Comparison with Real Galaxies

Interpretive power in a cosmological simulation depends critically on whether the simulation is fundamentally able to reproduce the observed properties of galaxies. In the previous chapter, a rigorous characterization of the biases in our parametric estimates was performed for simulated galaxies with observational realism. Several controlled experiments were used to quantify the biases on decomposition results from the choice of SLD scheme §2.4.1, projection §2.4.2, positional biases including crowding, resolution of the PSF, and S/N §2.4.3, and internal segmentation by discrete substructure §2.4.4. In particular, focus was placed on the systematic and statistical effects of these biases for fundamental estimators of size and flux – the half-light radius and integrated magnitude. Having thoroughly quantified the potential biases of the image processing for these key parameters, we can now make a comparison with observations.

The comprehensive suite of observational realism applied to galaxies in the **DISTINCT** catalog ensures that all of the observational biases associated with position manifest themselves in a statistically equivalent fashion with real galaxies from the SDSS. The foundation of our comparison is that the **GIM2D** quantitative morphologies pipeline that is employed in our analysis of the mock-observed Illustris galaxy population is identical to the pipeline previously used by S11 for 1.12 million real galaxies in the SDSS. Using the integrated magnitudes, half-light radii, and bulge-to-total fractions, for which the biases were examined in the previous chapter, a morphological comparison can be performed between the two populations. Two fundamental relations that may be examined using our luminosity, size, and morphology estimates are the size-luminosity relation and the bulge-to-total distribution as a function of stellar mass. In the following sections, we employ our decomposition results for galaxies in Illustris

in comparisons of the size-luminosity and (b/t)-stellar mass relations with galaxies from SDSS.

3.1 The Galaxy Size-Luminosity Relation

The observed relationship between size and luminosity is a crucial benchmark within the framework of hierarchical assembly of galaxies (Schade et al., 1996; Shen et al., 2003; Trujillo et al., 2004; Courteau et al., 2007; Graham & Worley, 2008; Brooks et al., 2011; Dutton et al., 2011; Baldry et al., 2012; Christensen et al., 2014; Furlong et al., 2015; Crain et al., 2015). The morphologies that are determined by photometric analyses are governed by the growth and evolution of stellar populations and their distribution within the physical components of galaxies. Reproducing the observed size-luminosity relation within populations of simulated galaxies requires comprehensive treatment of key physics including: stellar and gas kinematics; chemical enrichment of the inter-stellar medium (ISM); gas-cooling; star formation, feedback, and cessation; stellar population synthesis and evolution; and the influence of galaxy interactions and merging on these processes (Brooks et al., 2011; Christensen et al., 2014; Furlong et al., 2015; Crain et al., 2015). Similarities and differences between the size-luminosity relations of the simulation products and observations is informative on the successes and trappings of the models employed by the simulations.

The sizes of galaxies in hydrodynamical simulations have only recently demonstrated consistency with observations. In earlier generations of simulations, stellar masses and luminosities were too large and their sizes were too small (e.g. see comparison of various hydrodynamical codes by Scannapieco et al. 2012). The main culprit was the catastrophic transfer of angular momentum from the dense gas to the halo which is exacerbated if too much gas is turned into stars (Katz & Gunn, 1991; Navarro & White, 1994). The loss of angular momentum in the gas leads to overly compact galaxies with centrally concentrated and run-away star-formation. The problem can be overcome by the inclusion of energetic feedback in the models – which regulates the formation of stars and tends to expel the lowest angular momentum gas (e.g. Sales et al. 2010; Brook et al. 2011, 2012a). High-resolution hydrodynamical simulations that include realistic and efficient feedback have recently enabled the production of more reasonable disc galaxy sizes in small galaxy samples and for targeted mass ranges (Governato et al., 2004; Okamoto et al., 2005; Sales et al., 2010; Brook et al., 2012b; McCarthy et al., 2012a; Munshi et al., 2013; Aumer et al., 2013; Hopkins et al.,

2014; Marinacci et al., 2014). However, reproducing the distribution of sizes of full galaxy populations in cosmological simulations remains a challenge. The size-mass and size-luminosity relations of galaxies depend sensitively on a broad range of galaxy scaling relations including the stellar mass and luminosity functions, and $M_{\star} - M_{\text{halo}}$ relations in addition to models for energetic feedback – which must all be accurate to reproduce observed galaxy sizes (Scannapieco et al., 2012; Crain et al., 2015).

Galaxy sizes in statistically relevant samples from hydrodynamical simulations and their dependencies on the sub-grid physics models have been studied in several recent works (e.g. OWLS $z = 2$: Sales et al. 2010; OWLS $z = 0$: McCarthy et al. 2012b; GIMIC $z = 0$: Crain et al. 2009). The sub-grid feedback parameters generally have large uncertainties and are often calibrated to reproduce global scaling relations for galaxy populations at specific epochs (e.g. see Genel et al. 2014). In the Illustris simulation, the poorly constrained parameters for the efficiency of energetic feedback are calibrated to roughly reproduce the history of cosmic star-formation rate density and the $z = 0$ stellar mass function (Vogelsberger et al., 2013). However, the *evolution* of these relations are predictions of the simulation. Snyder et al. (2015) performed an image-based comparison using non-parametric morphologies derived from mock SDSS observations of galaxies from the Illustris simulation (Torrey et al., 2015) to show that galaxies in Illustris were roughly twice the size of observed galaxies for the same stellar masses at $z = 0$. In the EAGLE simulation (Schaye et al., 2015), the feedback efficiency parameters were calibrated to reproduce the galactic stellar mass function and the sizes of discs at $z = 0$ and the observed relation between stellar mass and black-hole mass (Crain et al., 2015). EAGLE has been shown to successfully reproduce the *evolution* of both disc *and* spheroid sizes out to $z = 2$ using inferred scalings between *physical* and photometric properties (Furlong et al., 2015). Comparison of the predictions from large-volume and high-fidelity hydrodynamical simulations such as Illustris and EAGLE to observations enables improved constraints on the sub-grid physics that govern galaxy sizes. However, it is important that such comparisons are made to be fair by deriving the properties of galaxies consistently in observations and simulations. Realistic mock observations facilitate translation of the physical information from the simulations to image-based photometric information that can be compared directly with observations.

Creating mock observations of simulated galaxies is the most direct way to consistently derive galaxy properties. Realistic images place analyses of simulated and observed galaxies on level-ground – as the same quantitative morphology tools can be

used to derive the photometric and structural properties of each. Mock observations of galaxies from hydrodynamical simulations have been successful in reproducing the observed trends for targeted morphologies – but have been limited to small samples of galaxies. Brooks et al. (2011) used high-resolution zoom-in hydrodynamical cosmological simulations and dust-inclusive radiative transfer to produce mock observations of a sample of 8 disc galaxies. Bulge+disc decompositions of the surface-brightness profiles were performed on B -band images to estimate the disc scale-length, r_d (or often, h , in the literature) and magnitudes of the bulge and disc components. The size-luminosity relation of the discs demonstrated remarkable agreement with observed trends at redshifts, $z = 0$ and $z = 1$ (van Zee, 2000; Graham & Worley, 2008; MacArthur et al., 2003, 2008; Miller et al., 2011). In particular, the $z = 0$ discs were consistent for a dynamic range in magnitude that is as large as for the observational constraints.

The size-luminosity relation for the bulges within discs has also been also been compared with observational constraints using mock photometry. Christensen et al. (2014) identified 2 galaxies with significant bulge components within the high-resolution disc galaxies simulated by Brooks et al. (2011). The properties of the bulge components were derived from H -band photometry using bulge+disc decompositions – consistently with photometric decompositions of observed bulges in late-type galaxies and in ellipticals (Fisher & Drory, 2010; Fisher et al., 2013). The properties of the bulges identified in the simulations were consistent with the size-luminosity relation derived from observations within a more narrow magnitude range than previously shown for the discs.

The principal limitations to previous studies aimed at comparing simulated and observed structural relations are at least one of the following factors: (1) the limitations on statistical relevance in comparisons to observations due to small simulated galaxy samples; (2) inconsistent derivations of simulated and observed galaxy properties; (3) incomplete observational realism that biases the distributions of derived properties of simulated galaxies in comparisons with observations. Each of these limitations can immediately be addressed using realistic mock-observations of galaxies in large-volume and high-fidelity cosmological hydrodynamical simulations. The current generation of large-volume hydrodynamical simulations contain appreciable populations of galaxies which can be used to compare the distributions of galaxies and their morphologies on the global size-luminosity relation. Mock observations of galaxies from these simulations (e.g. Torrey et al. 2015; Trayford et al. 2015) ensures that

their derived properties are statistically affected by the same observational biases on real galaxies.

The distribution of a population of galaxies on the size-luminosity plane is governed by the distribution of stellar populations within the physical components of galaxies. The size-luminosity relation is therefore well suited to examine the successes and inconsistencies in the structural morphologies of the simulation products using the methods from the previous chapter. In this section, we employ our estimates of half-light radius and integrated apparent magnitude for simulated galaxies with observational realism in an unprecedentedly fair comparison of the size-luminosity relations of galaxies in Illustris and the SDSS. The construction of comparison samples are described in Section 3.1.1 and our results are shown in Section 3.1.2.

3.1.1 Selection of a Comparison Sample

The catalog of 1.12 million quantitative morphologies of galaxies from the SDSS by S11 represents a reservoir from which we can draw populations of galaxies for comparisons to the simulated galaxies in the DISTINCT catalog. However, several important criteria must be met for the comparison to be fair. While the design of the DISTINCT catalog ensures statistical equivalence in the distributions of crowding, resolution, and sky, we recall that all of our galaxies are inserted into the SDSS at redshift $z = 0.05$. One observational bias that we have therefore not explored is the robustness of our parameter estimates with the surface brightness degradation at higher redshifts. A criterion of the SDSS control sample that is consequently necessary is that the control galaxies are confined to some thin redshift range around $z = 0.05$ so that any biases that arise from surface brightness increases or degradation do not enter into the comparison. Such a criterion for the redshift of a galaxy further requires that the estimate of the redshift is accurate – which requires the additional criterion that galaxies must have spectroscopically derived redshifts. We therefore impose the following criteria on the S11 catalog:

- (1) Galaxies are selected from the Spectroscopic Sample of the SDSS DR7 Legacy Survey ($\sim 700,000$ galaxies)
- (2) Galaxies are confined to the volume corresponding to the spectroscopic redshift range $0.04 < z < 0.06$ ($\sim 68,000$ galaxies)

Biases from surface brightness and volume incompleteness in the samples of simulated and real galaxies are removed by sampling the galaxies in the **DISTINCT** catalog to match the normalized stellar mass distribution of the SDSS over $0.04 < z < 0.06$ where $\log M_*/M_\odot > 10$. The latter criterion is imposed because it is the stellar mass lower limit of Illustris galaxies for which there are synthetic images (Torrey et al., 2015). The SDSS stellar masses are derived from combined surface brightness profile model estimates and SED template fitting by Mendel et al. (2014) and are used to perform the matching and cut in stellar mass. Galaxies are drawn with replacement from the 28,000 galaxies in the **DISTINCT** catalog using a Monte Carlo accept-reject scheme to match the stellar mass distribution of the SDSS sample.

3.1.2 Galaxy Size-Luminosity Relations of SDSS and Illustris

Figure 3.1 shows the distributions of Illustris (red) and SDSS (blue) in the plane of rest-frame r -band half-light radius (h^{-1} kpc) and absolute r -band magnitude from the bulge+disc decompositions. The luminosities in each sample span roughly 4 magnitudes at the 95% level (coloured dashed lines). However, the Illustris luminosities are brighter by roughly 0.7 magnitudes (2 times brighter) on average. Figure 3.1 also demonstrates a discrepancy in sizes between the distributions of Illustris and SDSS galaxies. Galaxies with high luminosities are systematically larger in Illustris than galaxies observed in the real universe by roughly +0.4 dex (or a factor of 2 larger, consistent with Snyder et al. 2015). But, the far more more puzzling discrepancy is the fundamental lack of correlation between size and luminosity for Illustris galaxies with respect to the SDSS for the same stellar masses. The slope of the global size-luminosity relation for galaxies in Illustris is significantly shallower than for galaxies in the SDSS – implying weak evolution of size with stellar mass in Illustris. In the following sections, we discuss several possible biases and their impacts on size-luminosity relation of Illustris.

Internal Segmentation

The large offset in magnitude between Illustris and the SDSS in Figure 3.1 occurs despite the known systematics from internal segmentation in an appreciable fraction of galaxies in the **DISTINCT** catalog. In the previous chapter, we showed that the **DISTINCT** catalog contains many galaxies that have systematically *positive* magnitude estimates from internal segmentation – corresponding to fainter fluxes by *up to a*

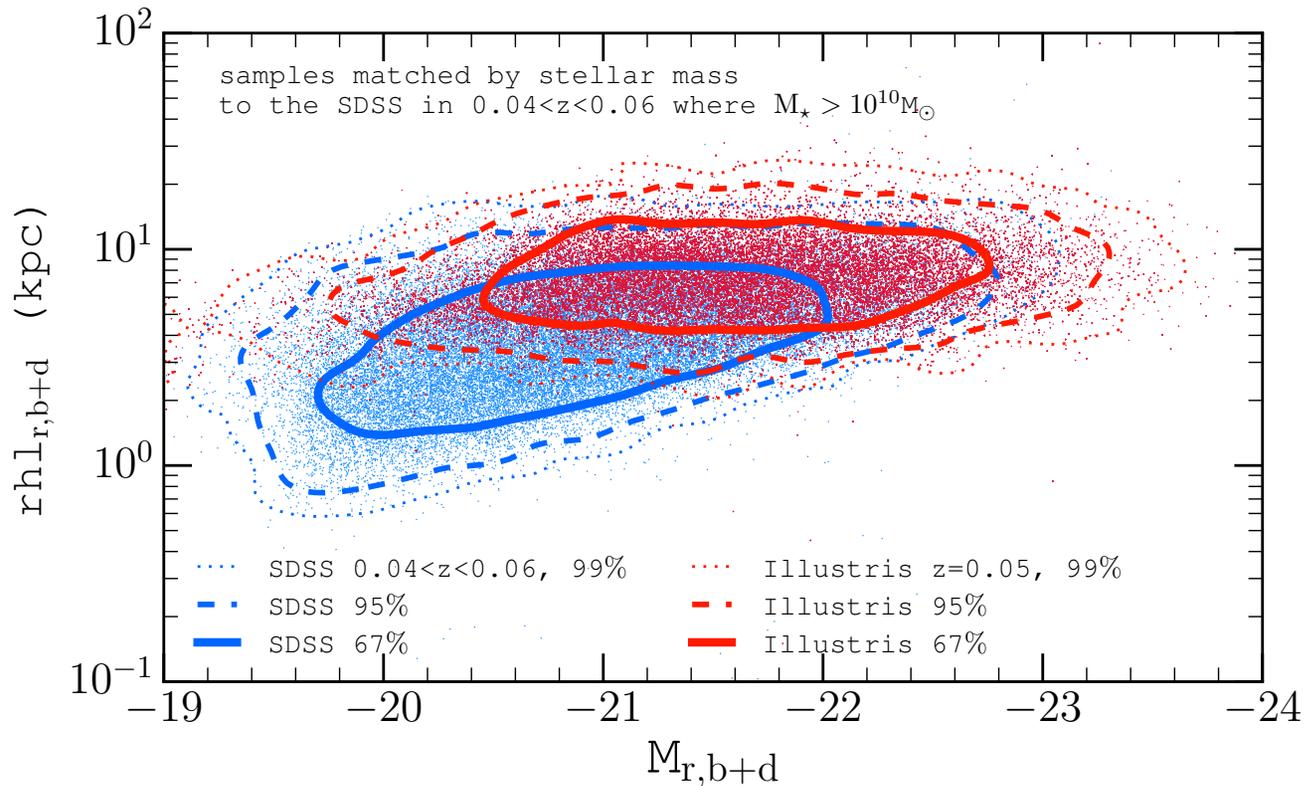


Figure 3.1: The size-luminosity relations of the SDSS (blue) and Illustris (red) in a stellar mass matched sample. The 99th, 95th, and 67th percentiles of the mass-matched distributions of Illustris and SDSS galaxies on the r-band $b + d$ absolute magnitude and circular half-light radius surface are shown in solid, dashed, and dotted lines, respectively. Stellar masses for the SDSS galaxies are taken from the catalog of Mendel et al. (2014). Illustris galaxies are sampled from the stellar mass distribution of SDSS within the $0.04 < z < 0.06$ range to minimize the biases from improvements or degradation of surface brightness on the distributions of best-fitting parameters in each population. Illustris and SDSS each span roughly 4 magnitudes in luminosity, but galaxies in Illustris are intrinsically brighter and larger. The scaling between size and luminosity is also shallower for galaxies from Illustris – leading to a stark contrast of size estimates at the low-luminosity end of the distributions.

factor of ~ 6 for roughly 30% of galaxies in the DISTINCT catalog. However, the offset in magnitude that is seen in Figure 3.1 for size-luminosity distributions of Illustris relative to the SDSS galaxies is *negative* – in conflict with the bias from internal segmentation. In fact, the systematically larger magnitude estimates from internally segmented galaxies seem to only give the impression that the 99% contours (coloured dotted lines) for the two samples occupy the same magnitudes at the low luminosity end. Figure 3.2 shows that plotting the luminosities and sizes derived directly from the synthetic images, $M_{r,\text{synth}}$ and $rhlc_{r,\text{synth}}$, removes these outliers from the Illustris distribution. However, replacing the decomposition results with the synthetic image properties only weakly shifts the 95% level of the Illustris size-luminosity distribution to lower magnitudes – but brings the width of each distribution at 99% into agreement. Furthermore, Figure 3.2 shows that replacing the decomposition results with the synthetic image properties affects only the sizes at higher magnitudes $M_{r,b+d} \gtrsim -21.5$. The increase in size at higher magnitudes occurs because intrinsically larger but diffuse galaxies stand to lose more of their shape through spurious internal segmentation if a particular piece of substructure is selected as the sole source for the fitting. Ultimately, replacing the decomposition results with the synthetic image properties in Figure 3.2 generates an Illustris size-luminosity relation that has a slightly greater slope but is shifted by an additional 0.2 magnitudes brighter with respect to the SDSS. The biases from internal segmentation are therefore insufficient to explain the discrepant offset in magnitude and difference in slope in the size-luminosity relations of Illustris and the SDSS.

Dust Obscuration

The discrepancy in Figure 3.1 between the magnitudes and sizes Illustris and the SDSS galaxies owes, in part, to the absence of dust in Illustris galaxies. The inability to resolve dust physics within the simulation makes it difficult to quantify its role quantitatively. However, absence of internal dust extinction within Illustris galaxies and its prevalence in real galaxies could serve to offset one relation from the other in magnitude and may affect the correlation between size and luminosity. Our comparison of the size-luminosity relations of Illustris and the SDSS galaxies is therefore complicated by the presence of dust in the real universe (Gadotti et al., 2010) – which is not distributed uniformly within galaxies (e.g. see McKinnon et al. (2016) and references therein). When dust is present in galaxies, disc scale-lengths from bulge+disc

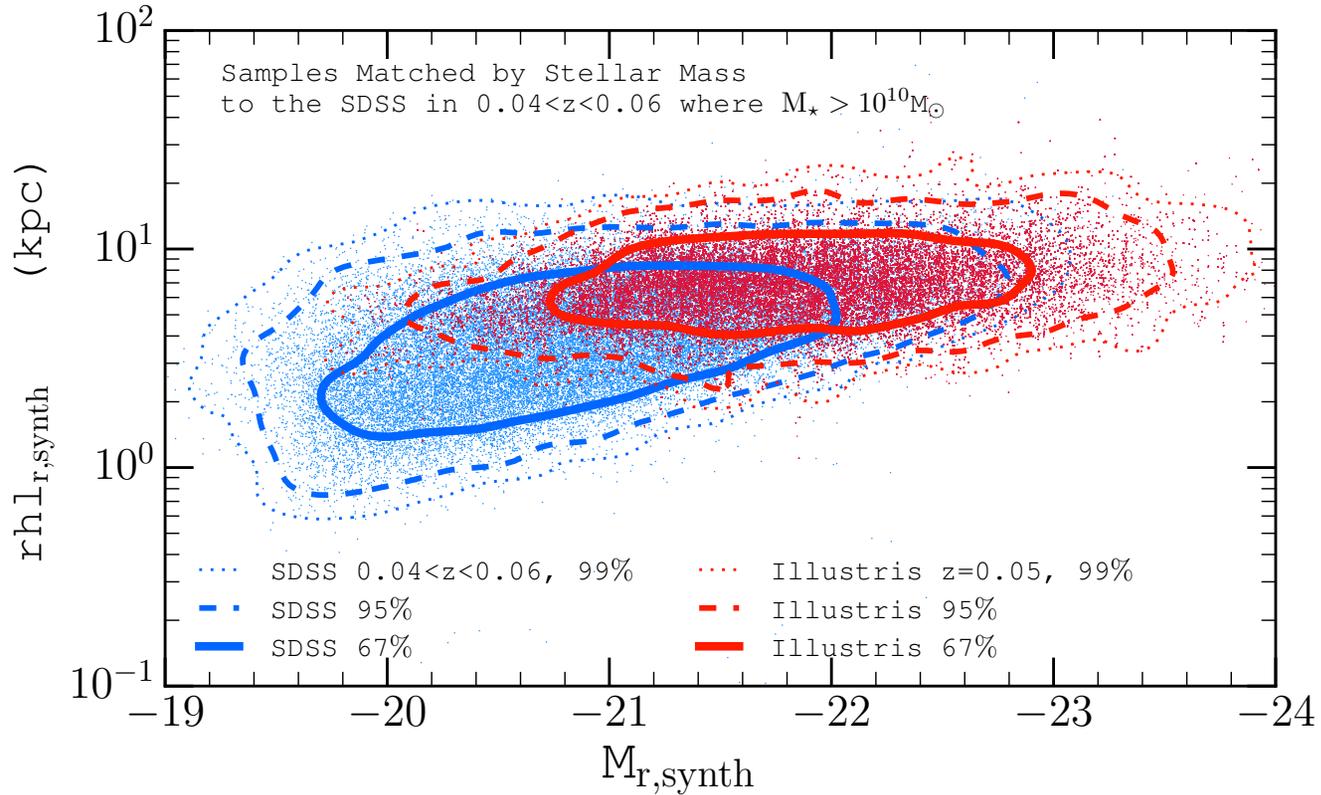


Figure 3.2: Same as Figure 3.1 but for the properties derived directly from the synthetic images without realism or decompositions. $M_{r,\text{synth}}$ and $rh_{l,r,\text{synth}}$ are computed directly from the flux in the synthetic images. The distribution for Illustris galaxies is shifted by roughly 0.2 magnitudes brighter than in Figure 3.1 due to the removal of bias from internal segmentation. The removal of internal segmentation effects leads to galaxy sizes that are larger at the high-luminosity end of the distribution – creating a slightly clearer correlation between size and luminosity.

decompositions are systematically over-estimated and this systematic is exacerbated by inclination (with edge-on discs biased most strongly) (Gadotti et al., 2010). Meanwhile, the sérsic indices and effective radii of bulges and spheroids are systematically underestimated (Gadotti et al., 2010). In populations of discs and spheroids, the effects of dust may therefore serve to increase the scatter *and* modify the scaling between size and luminosity. However, the differences in size at the low-luminosity end of Figure 3.1 cannot be reconciled by dust since, at low luminosities, galaxies in the SDSS are systematically smaller than in Illustris for the same luminosities – whereas dust in real galaxies is predicted to over-estimate disc sizes (Gadotti et al., 2010). For dust to cause this shift would first require a significantly larger population of low-luminosity spheroids whose sizes would be systematically underestimated due to dust. Therefore, while dust may partially explain the offset in luminosities between Illustris and the SDSS, a difference in morphologies between the two populations is necessary to cause the changes to the scaling between size and luminosity from dust. Ultimately, the absence of dust in the simulation and radiative transfer code used to produce the synthetic images presents a fundamental limitation in the realism of the mock observations. However, our choice to not employ a dust model in the radiative transfer is motivated by the uncertainties involved in the distribution of dust in galaxies. Further analysis of dust effects for our quantitative morphologies would yield interesting results that could be compared with the dust-less models.

3.1.3 Impact of Bulge and Disc Morphologies

In the previous sections, we argued that the discrepancy in the scaling between size and luminosity in our comparison of the SDSS and Illustris is not consistent with biases from internal segmentation or dust effects. Alternatively, we proposed that morphological differences between the populations of the real and simulated galaxies may be the source of the different scalings between size and luminosity seen in Figure 3.1. The presence and growth of a stellar bulge component in galaxy morphologies is strongly linked to many key processes of galaxy formation theory. The photometric bulge-to-total fractions obtained in the structural decompositions of each galaxy are targeted at estimating the prominence and properties of the bulge. Given the additional importance of the bulginess of galaxies to various scaling relations including size-luminosity, the bulge-to-total fractions are well-suited to identify the morphological differences between Illustris and the SDSS. In this section, the ob-

served size-luminosity relations of late-type (right side of S0 on the Hubble Diagram, disc-like) and early-type (left of S0, spheroid- or bulge-like) are examined to provide context for a morphological comparison using the photometric bulge-to-total fraction and total stellar mass.

Morphological dependence of the size-luminosity relation

The size-luminosity relations of bulge- and disc-dominated galaxies differ considerably in observations of real galaxies (Courteau et al., 2007; Blanton & Moustakas, 2009). Figure 3.3 shows the relationships between $g - r$ colour, sérsic index, $r - band$ half-light radius (r_{50} , here), and absolute $r - band$ magnitude for a large sample of SDSS galaxies split by visual morphological classifications stored in the NASA Extragalactic Database (NED) (catalogs from: Corwin et al. 1994; Nair & Abraham 2010; and many others). The underlying distribution in all panels shows the distribution of 77,153 SDSS galaxies with $z < 0.05$ from DR6 (Adelman-McCarthy et al., 2008) and an update to the low-redshift sample (Blanton et al., 2005). Galaxies in the left panel are late-type disc galaxies and are further split into Sd, Sc, Sb, and Sa in order of increasing strength of the spheroidal bulge component (see Figure 1.1). Galaxies in the right panel are early-types – consisting of dwarf/compact ellipticals, spheroids, ellipticals, and central dominant galaxies found near the centres of rich galaxy clusters.

The left panel of Figure 3.3 shows that late-types which contain visually prominent spheroidal components (Sa, for example) have a steeper relationship between size and luminosity, span a larger range of sérsic indices in pure sérsic decompositions, and are generally redder than discs with less significant bulge components. The right panel of Figure 3.3 shows that the size-luminosity relation of early-types is collectively tighter than any individual class of late-types. The slope of the size-luminosity relation for early-types is also considerably steeper than for later-types, with the exception of Sa galaxies which are classified for containing significant bulge components which will strongly weight their quantitative morphologies.

Courteau et al. (2007) and Fisher & Drory (2010) quantitatively confirmed the changes to the size-luminosity relations among classes of late-types and early-types, respectively, using bulge+disc decompositions. Fisher & Drory (2010) also performed a comparison of the size-luminosity relations for classical and pseudo-bulges – finding that the size-luminosity relation of classical bulges is the same as for ellipticals. Pseudo-bulges, which are often classified by sérsic index, $n \lesssim 2$, have a steeper slope

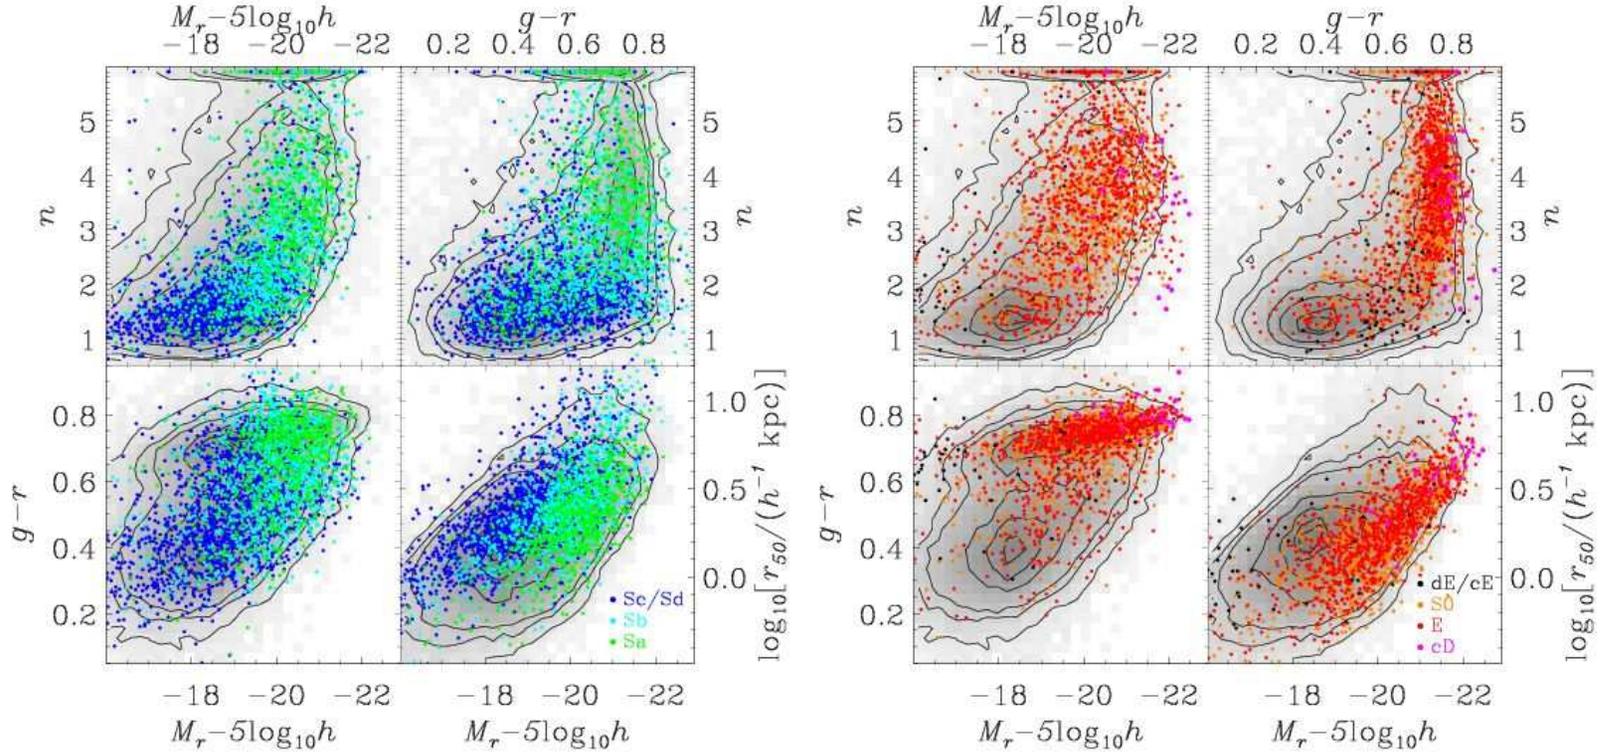


Figure 3.3: Size-luminosity relations of late- and early-type galaxies using data from the SDSS DR6 (Adelman-McCarthy et al., 2008) and an update to the low-redshift sample (Blanton et al., 2005). Relations between sérsic index, half-light radius, $g-r$ colour and absolute magnitude are shown for late-type (left) and early-type (right) galaxies. The contours and background histogram in each panel show the distribution of the full sample. Scatter plots show specific morphologies that are visually classified using the Hubble morphological classification scheme and taken from various catalogs stored in NED (including: Corwin et al. 1994; Nair & Abraham 2010; and many others). Figure taken from Blanton & Moustakas (2009).

than classical bulges and ellipticals on the size-luminosity relation but significantly greater scatter (Gadotti, 2009). The discrepancy between pseudo-bulges and classical bulges is expected as pseudo-bulges are believed to have a different formation mechanism and to be structurally different from classical bulges (e.g. Kormendy & Kennicutt 2004, and references therein). Classification of bulges into pseudo-bulges and classical bulges using sérsic index is generally imperfect, but can be considered as an approximation (Fisher & Drory, 2008). Many bulges with $n < 2$ follow the tight size-luminosity relation for classical bulges and, conversely, some bulges that are offset from this relation have sérsic indices that are consistent with the distribution of classical bulges (Gadotti, 2009).

Our analysis does not discern between classical and pseudo-bulges – as our bulge+disc decompositions use a bulge component with fixed sérsic index. However, our analysis of the simulated and observed galaxy populations is internally consistent. If pseudo-bulges in the simulations are equally represented and structurally similar to observed pseudo-bulges, then their effect on the distribution of structural parameters will be the same. Therefore, while the structural estimates for pseudo-bulge properties may be inaccurate in our bulge+disc decompositions, any discrepancies between the distributions of bulge properties between the simulations and observations will be sourced by true structural differences of these components.

Bulge and disc morphologies in Illustris and the SDSS

Morphological differences between our samples of observed and simulated galaxies are accessible through the bulge-to-total structural parameter from the decompositions results. For example, the different trends observed in the size-luminosity relations of discs and spheroids shown in the previous section are easily reproduced using (b/t) in our observational sample. Figure 3.4 shows the size-luminosity relations of *pure* discs and spheroids in our observational SDSS sample, selected within $0.04 < z < 0.06$ and identified using bulge-to-total ratios $(b/t)_r < 0.05$ and $(b/t)_r > 0.95$, respectively. Discs are generally larger than bulges and samples that contain galaxies with a dominant disc component are decoupled from samples of galaxies containing dominant bulges. The disc relation is also shallower and has more scatter than for the bulge-dominated galaxies, as expected from the literature. While our choice of sérsic indices for each component assumes that the bulges are more compact than discs for the same magnitudes, this bias is consistent in the Illustris and SDSS decompo-

sitions. Therefore, given the difference in the size-luminosity relations of bulge and disc morphologies, we can examine the morphological differences between the SDSS and Illustris using the bulge-to-total fraction and assess their impact on the galaxy size-luminosity relations in Figure 3.1.

Given that morphology is clearly critical in driving the normalization, slope, and scatter of the size luminosity relation, it is germane to compare the (b/t) distributions of the SDSS and Illustris samples. Figure 3.5 shows the distribution of r – *band* photometric (b/t) as a function of total stellar mass in the SDSS (left panel) and Illustris (right panel). The samples are the same as for Figure 3.1 which were matched in stellar mass. Figure 3.5 shows that the SDSS has a diversity of morphological populations including bulge-dominated, disc-dominated, and a large number of composite structure galaxies within this mass distribution. The diversity is not shared by Illustris – which is bereft of photometric bulges. Only for stellar masses $\log M_*/M_\odot \gtrsim 10.6$ do any galaxies with significant bulge components become more common – indicating a stronger correlation between bulge-fraction and total stellar mass within Illustris than exists in the observations. The right marginal for the Illustris distribution shows that over 72% of galaxies in Illustris have $(b/t)_r < 0.05$, a fraction which rapidly declines for larger photometric bulge fractions.

The startling result from Figure 3.5 has an obvious implication for the discrepancy between the size-luminosity relations of Illustris and the SDSS. Illustris contains a much larger disc population than the SDSS for the sample matched by total stellar mass to the distribution of spectroscopic SDSS galaxies in $0.04 < z < 0.06$. Figures 3.3 and 3.4 showed that discs which do not contain significant bulge components are elevated in the size-luminosity relation and have a shallower slope than early-types. The SDSS distribution in Figure 3.1 is analogous to the background distributions from Figure 3.3. The galaxy size-luminosity relation for the SDSS is broadened vertically at low luminosities and has bent contours because it contains populations of discs, bulges, and composite systems. The bulges in SDSS weigh the distribution to more compact half-light radii, as shown in 3.4. However, Illustris is bereft of bulges at low stellar masses (luminosities). Therefore, there is no downward weight from bulge-dominated galaxies at the low-luminosity end of Illustris to bring the slope and scatter into agreement with the SDSS.

The impact of the morphological differences between the SDSS and Illustris on the size-luminosity relation can be determined by matching samples in both total stellar mass and bulge-to-total ratio. If the bias to low (b/t) morphologies in Illustris

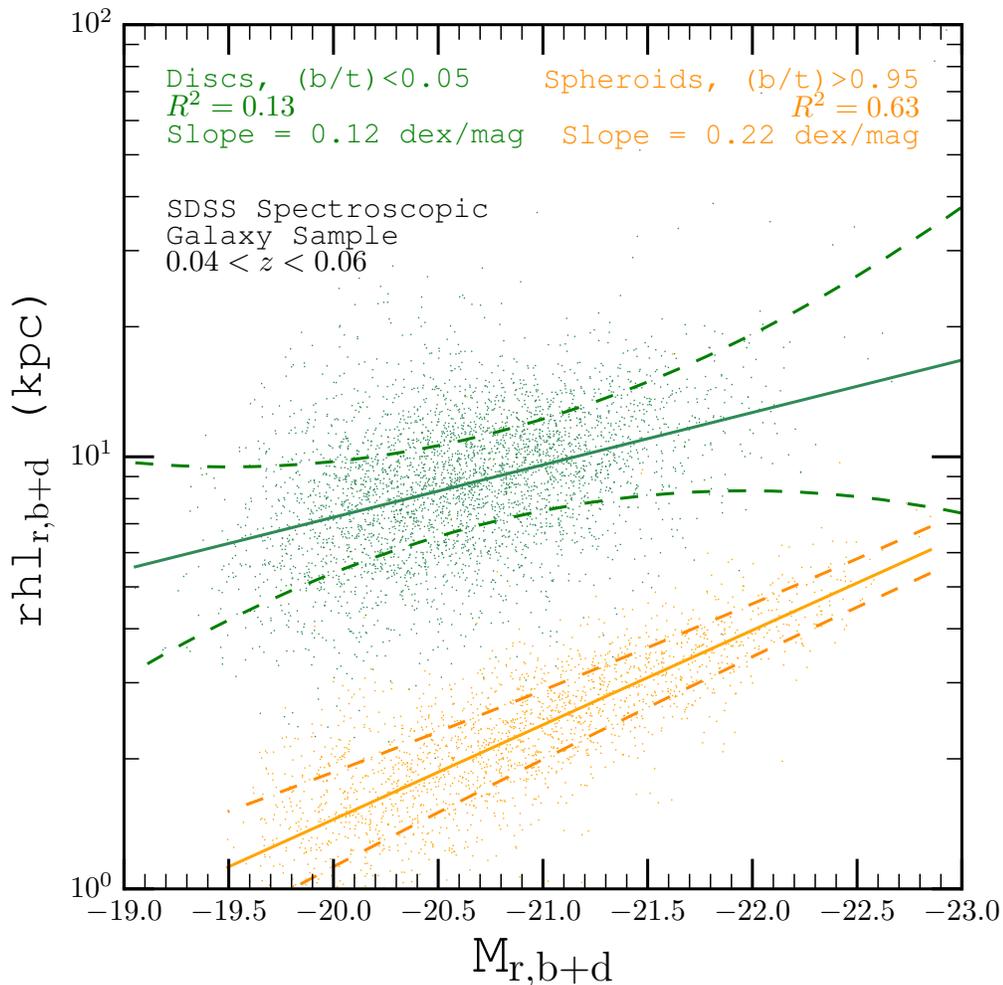


Figure 3.4: Size-luminosity relations of *pure* discs and ellipticals from our observational comparison sample. Galaxies are selected as described in §3.1.1, taken from the spectroscopic galaxy sample of the SDSS DR7 within $0.04 < z < 0.06$. Galaxies with bulge-to-total ratios $(b/t)_r < 0.05$ (green) and $(b/t)_r > 0.95$ (orange) are taken to approximate the size-luminosity relations of *pure* discs and ellipticals, respectively, using our estimates of magnitude and half-light radius. Solid lines indicate the best-fitting linear fit to each population. The slope of the size-luminosity relation for the ellipticals is roughly double that for discs and the distribution is considerably tighter, as seen in Figure 3.3. Dashed lines show the errors on the best-fitting relations computed by fitting a second order polynomial to the residuals of the best-fitting model. R^2 is the coefficient of determination for the linear regression which compares the variance of the residuals to the raw variance. Larger R^2 indicates smaller scatter about the best-fitting model.

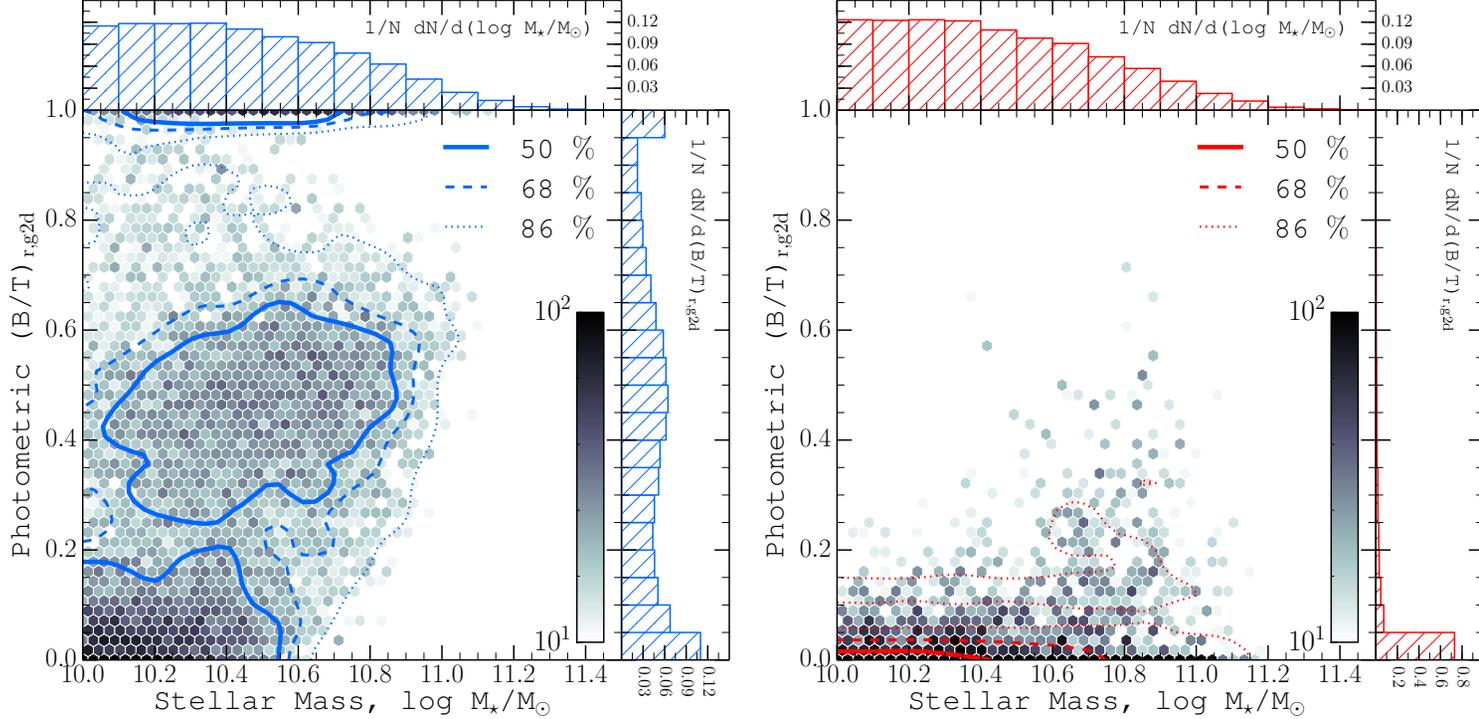


Figure 3.5: The distributions of galaxy bulge-to-total fractions in mass-matched galaxy samples of the SDSS (left panel) and Illustris (right panel) as a function of total stellar mass. The galaxies that form these distributions are the same ones that form the size-luminosity relations shown in Figure 3.1. Stellar masses for SDSS galaxies are taken from Mendel et al. (2014) and are the same masses used in the matching of the samples. The upper marginals show the stellar mass distributions in each sample and demonstrate their consistency. The right marginals show the distributions of photometric bulge-to-total fractions, $(b/t)_r$. The central panels show the 2-dimensional distribution of total stellar mass and photometric b/t on a logarithmic colourmap shown at along the lower right of the main panels. Coloured lines show the 50% (solid), 68% (dashed), and 86% (dotted) contours of the distribution. The observations from the SDSS contains diverse morphologies classified by b/t including galaxies dominated by both discs and bulges as well as a population of composite systems. Illustris galaxies are bereft of bulges in the mass-matched sample.

is responsible for the discrepancy in the size-luminosity relations of Illustris and the SDSS, then matching the SDSS morphologies (which are more diverse) *and* stellar masses to the Illustris galaxies from Figure 3.1 should bring the size-luminosity relations into better agreement.¹ Figure 3.6 shows the size-luminosity relations for the SDSS and Illustris where matching is performed for both stellar mass and bulge-to-total fraction. At low luminosities, the large discrepancy in galaxy sizes from Figure 3.1 is largely removed. The improved agreement at low-luminosities is consistent with a deficit in bulges in Illustris relative to galaxies in the SDSS – as seen in Figure 3.5. Figure 3.4 showed that the size-luminosity relations of bulges and discs are most discrepant at low-luminosity. Matching samples by morphology largely removes the bulge-dominated systems in the SDSS and produces a size-luminosity relation that is dominated by disc morphologies. With less weight towards more compact sizes from morphologies with high (b/t), the SDSS size-luminosity relation in Figure 3.6 is significantly more shallow than in Figure 3.1. The scaling between size and luminosity Figure 3.6 for Illustris and SDSS galaxies is brought into greater agreement by the morphology matching. Galaxies are also slightly larger on average – bringing the sizes of galaxies in each sample into better agreement.

Our results show that Illustris is dominated by disc morphologies for stellar masses matched to the SDSS controls. The dominance of disc morphologies at these masses is not consistent with the observations. The morphological differences between the simulated and observed galaxy populations generate discrepancy in their size-luminosity relations. The deficit of bulges in Illustris yields a size-luminosity relation that is offset to larger sizes and has a shallower slope than for galaxies in the SDSS. Matching SDSS galaxies to Illustris by (b/t) morphology, in addition to stellar mass, improves the agreement between the size-luminosity relations – but demonstrates a fundamental discrepancy in the morphologies of each sample. The remaining disagreement in the average sizes and luminosities of the mass-morphology matched samples is difficult to discuss meaningfully without the explicit inclusion of dust within the simulations or radiative transfer. Dust corrections in synthetic images of galaxies have recently been considered for galaxies in the EAGLE simulation (Trayford et al., 2015). Proper treatment and inclusion dust effects will yield greater consistency in the observational

¹Note that matching in the other direction does not work. Illustris contains too few galaxies with low (b/t) to be matched to the much larger population of bulges in the SDSS at these masses. In order to maintain the same stellar mass distributions as in our previous comparisons, we match galaxies one-to-one from the SDSS by stellar mass and (b/t) to the sample of Illustris galaxies from Figures 3.1 and 3.5 that are matched to the distribution of stellar masses in SDSS in $0.04 < z < 0.06$.

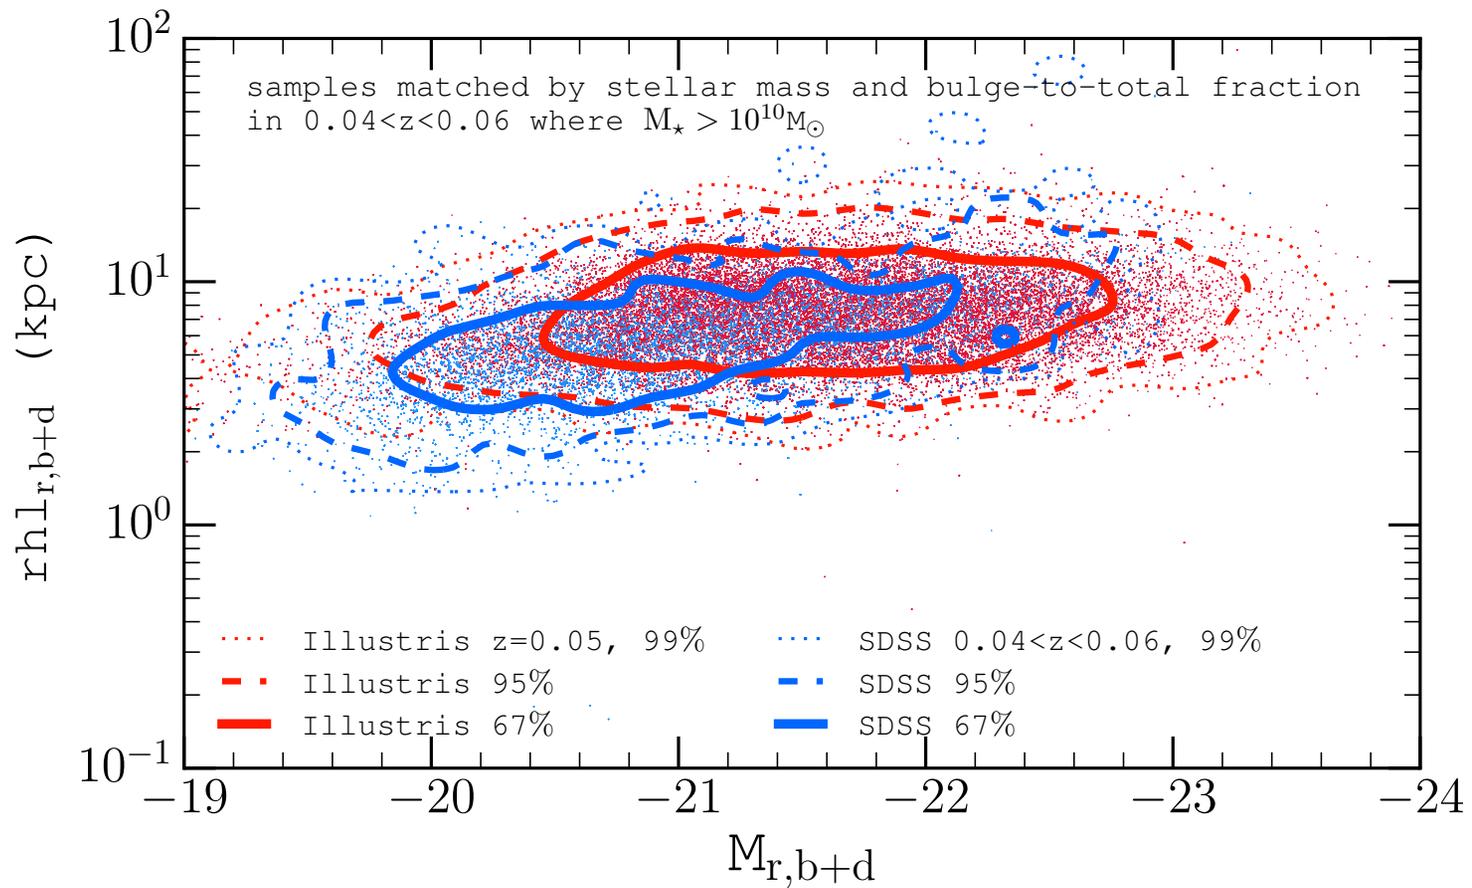


Figure 3.6: Same as Figure 3.1 but with Illustris and SDSS samples matched by stellar mass *and* (*b/t*) morphology.

realism of mock observations of galaxies.

3.2 The Photo-Bulge Deficit

The results from the previous section demand further investigation into the lack of bulges in Illustris. The lack of bulges in Illustris contrasts with historical problems in hydrodynamical simulations – in which the *physical parameters* of galaxies indicated that they were too bulge-dominated (Navarro & White, 1994; Katz & Gunn, 1991; Scannapieco et al., 2012). The question is whether (a) photometric bulges (photo-bulges) systematically do not exist in Illustris; (b) photo-bulges only do not strongly appear in the mass-matched sample that was used to examine the size-luminosity relation; or (c) they are just not well extracted by GIM2D. In the mass matched sample to $0.04 < z < 0.06$ of the SDSS, galaxies at the high mass end ($\log M_{\star}/M_{\odot} \gtrsim 11$) of the $z = 0$ stellar mass function of Illustris are not represented – despite the larger cosmic volume in the observed sample.² Examination of the bulge-to-total fractions at higher masses in Illustris and comparisons with observations will yield insights on the discrepancy between the morphological dependence on stellar mass in Illustris and the SDSS.

In this section, the relationships between morphological (b/t) fractions and stellar mass in the full populations of Illustris and SDSS are examined to provide insight on the deficit of bulges in Illustris seen in the previous section. We argue possible scenarios that cause the discrepancies with observations. A comparison of the kinematically derived stellar bulge-to-total fraction with the photometric fraction is also performed to examine the consistency between the information taken from the stellar orbits and stellar light.

3.2.1 Morphological dependence on stellar mass

Galaxies in Illustris contain bulges – albeit few at low stellar masses. Figure 3.7 shows the distribution of bulge-to-total fractions in the SDSS and Illustris (i.e. in the

²There are a number of reasons for the larger number of high-mass galaxies in Illustris relative to the SDSS volume to which we performed the matching. One possible reason is that Illustris slightly over-predicts the redshift $z = 0$ stellar mass function (SMF) at the high-mass end (Genel et al., 2014). Further biases could arise in the analysis of galaxies at the centres of rich clusters – which may provide discrepant stellar mass estimates with respect to the known masses from the simulations. A dedicated study of the systematics on photometric stellar mass estimates for all morphologies is required to fully understand the biases in the mass matching.

full DISTINCT catalog) – this time without any matching. The left panel of Figure 3.7 shows the distribution of (b/t) in the SDSS subdivided by stellar mass. The right panel shows the distribution for Illustris and includes an additional subdivision containing galaxies from the high-mass end of the Illustris stellar mass function – which is not represented in our comparisons of the size-luminosity and (b/t) vs. stellar mass relations. There are very few galaxies in the SDSS spectroscopic sample in $0.04 < z < 0.06$ with stellar masses $\log M_{\star}/M_{\odot} > 11.5$ and so no subdivision is allocated to these galaxies in the SDSS panel. A zoomed inset aids in the visualization of the Illustris distribution for $(b/t) > 0.05$ due to the prominence of $(b/t) < 0.05$ for the first two mass divisions.

Figure 3.7 shows that significant differences exist in the correlations between mass and photometric (b/t) for observed galaxies and in Illustris. The first mass division where $\log M_{\star}/M_{\odot} < 10.5$ (red) is a basic recovery of the distributions shown in Figure 3.5 – as a large fraction of galaxies in both the SDSS and Illustris fall within this mass window. The SDSS contains diverse morphological structures and Illustris contains few photometric bulges. Over 80% of Illustris galaxies are completely disc-dominated in the lowest mass bin. At slightly higher masses, $10.5 < \log M_{\star}/M_{\odot} < 11$ (green), the observed galaxies have fewer disc-dominated systems and more composites, but Illustris still contains $\sim 50\%$ discs and few composites or bulge-dominated systems. The distributions become more similar in the highest mass bin that can be related between the two populations (purple). Illustris contains an appreciable number galaxies with higher bulge fractions in $11 < \log M_{\star}/M_{\odot} < 11.5$ which is similar to the observations. The final mass division for Illustris (cyan) strongly resembles the SDSS distribution in $11 < \log M_{\star}/M_{\odot} < 11.5$.

Figure 3.7 demonstrates that although both simulated and real galaxies have a (b/t) -stellar mass dependence, there is a stronger correlation between photometric bulge-to-total ratio and total stellar mass in Illustris than in the observations from the SDSS. Galaxies in SDSS have diverse morphologies within each mass division – where Illustris is significantly more dichotomous when comparing high mass (b/t) morphologies to low mass. Both populations share the trend that bulges become more frequent at higher stellar masses, but the dependence is stronger in Illustris. We may discuss several possible biases in attempt to qualitatively reconcile the discrepancy between photometric bulge-to-total fractions in Illustris and the SDSS.

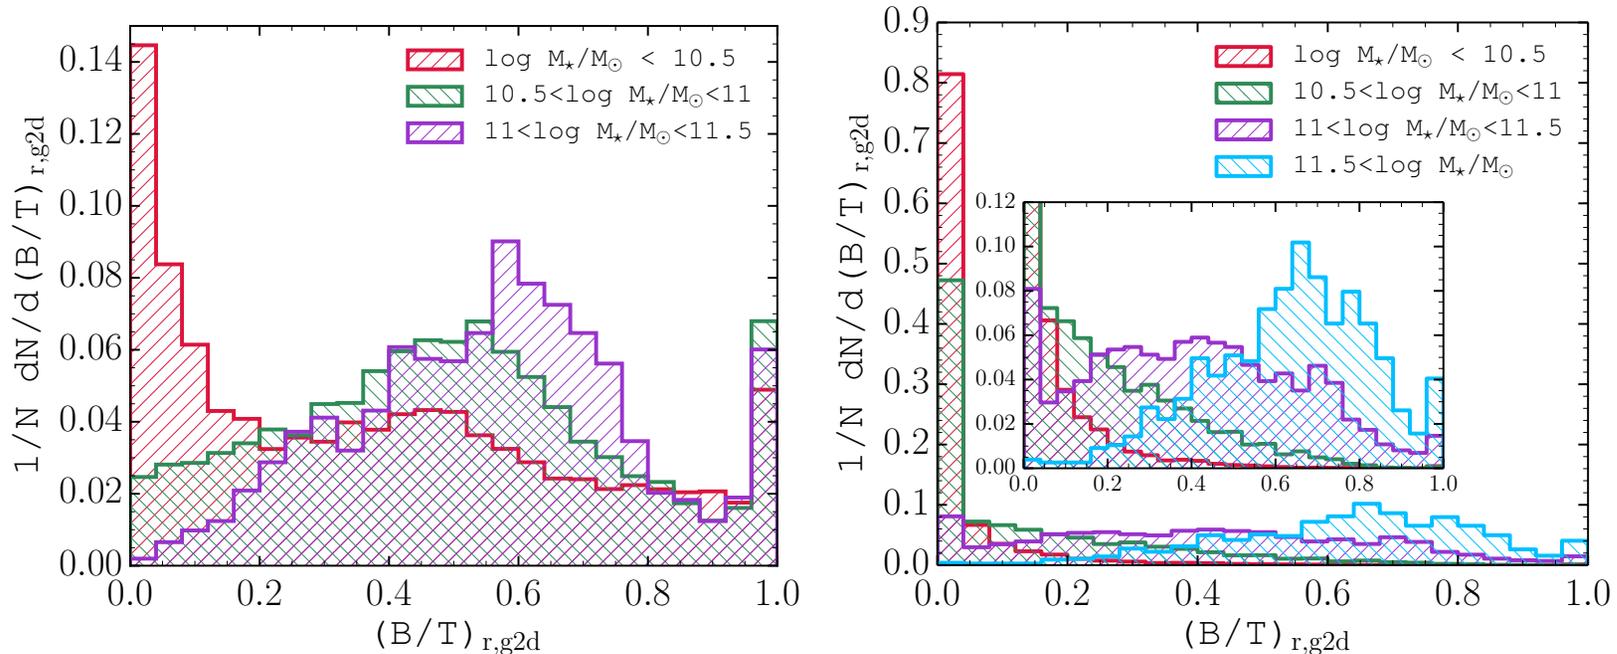


Figure 3.7: SDSS (left panel) and Illustris (right panel) r -band bulge-to-total fraction distributions in un-matched samples. SDSS galaxies are taken from the SDSS spectroscopic galaxy sample where $\log M_*/M_\odot > 10$ and $0.04 < z < 0.06$. Illustris galaxies are taken from the full DISTINCT catalog – which includes galaxies from the entire Illustris volume where $\log M_*/M_\odot > 10$. Each sample is subdivided by total stellar mass. A zoomed inset is provided for the Illustris distribution to aid in the inspection of the higher mass divisions. The SDSS sample does not contain an appreciable number of galaxies with $\log M_*/M_\odot > 11.5$ so this bin is excluded in the SDSS sample.

Impact of Dust Obscuration

A deficit of photo-bulges is not expected for galaxies in which there is no dust. The inclusion of a dust-model in the radiative transfer would serve only to under-estimate the presence of photo-bulges in Illustris by reducing the overall brightness of bulges and the pixels corresponding to the peak of the bulge surface brightness distribution (Graham & Worley, 2008; Gadotti et al., 2010). The de Vaucouleurs $n = 4$ model for the bulge is strongly dependent on the surface brightness from the inner 100 pc of a galaxies light profile. Attenuation of the light by dust at the centre of the bulge drives the free sérsic index and bulge half-light radius down in pure sérsic models (Gadotti et al., 2010). In fixed n bulge+disc decompositions the bulge model brightness is forced down to accommodate the decrease in flux from the central pixels and the exponential disc model is driven up. The combined effects lead to a reduction in bulge integrated magnitude and half-light radius, an increase in disc integrated magnitude and half-light radius, and a corresponding reduction in bulge-to-total fraction. The inclusion of dust tends to weaken the photo-bulge relative to the disc. Therefore, the exclusion of dust in the radiative transfer does not cause the photo-bulge deficit.

Stellar Particle Resolution

Accurate photometry for the inner region of the surface brightness distribution of a galaxy is essential for interpreting the bulge component. In Section 2.4.1, it was shown that the choice of stellar light distribution (SLD) scheme did not bias the global properties of the galaxy such as total integrated magnitude and half-light radius, but could strongly bias the structural parameter (b/t) (Figure 2.11). We showed that broader smoothing kernels (such as the constant 1 kpc kernel relative to the $N = 16$ nearest-neighbour kernel) artificially limit the resolution in the inner regions of galaxies. Broad, constant smoothing kernels reduced concentration of flux the from central regions of the galaxy and correspondingly reduced estimates of (b/t) systematically (to zero in many cases, even for galaxies with (b/t) as large as 0.6 in the fiducial scheme). While it is possible that the fiducial SLD scheme may be biasing (b/t) towards smaller bulge fractions in our decompositions, that leads to the notion that there is a “correct” SLD scheme for the particle mass resolution of Illustris (and the particle resolution of any hydrodynamical simulation). The choice of the N_{16} nearest-neighbour smoothing as the fiducial model was motivated by simplicity (Torrey et al., 2015). However, higher particle mass resolutions tend to reduce the

spatial separation between stellar particles in galaxies produced in hydrodynamical simulations and effectively increase the spatial resolution in adaptive SLD schemes. Figure 3.7 showed that the majority of high mass systems ($\log M_*/M_\odot \gtrsim 11$) in Illustris (that contain larger numbers of particles) have bulges and that low mass galaxies largely do not. The spatial distribution of particles determines the surface brightness distribution of a simulated galaxy. Larger numbers of particles reduce the smoothing radii in our fiducial SLD scheme and generally improve the spatial resolution of a galaxy surface brightness distribution. Improvements to the spatial resolution in the bulge surface brightness distribution (in particular to the inner 100 pc that are essential for discerning its profile from a disc) facilitate greater accuracy in modelling of the bulge component. Therefore, the strong mass dependence for the bulge-to-total fraction in Illustris seen in Figure 3.7 may arise from inadequate particle resolution for creating photo-bulges in synthetic images of low mass galaxies with smaller numbers of stellar particles. We have not yet investigated the co-dependence on stellar particle resolution (which tends to increase with stellar mass) in our SLD scheme analysis in Section 2.4.1. Ultimately, the stellar particle resolution that is accessible in Illustris may be reducing (b/t) measurements in our decompositions for galaxies with $\log M_*/M_\odot \lesssim 11$.

One way to test the particle resolution dependence on bulge fractions directly is to perform hydrodynamical zoom-in simulations of lower mass systems in Illustris with the same numerical techniques and simulation models (e.g. Sparre & Springel 2016). Comparison of the decomposition results from the high-resolution and low-resolution simulations would yield insight on the effects of particle resolution on (b/t) estimates from mock observations. A more extensive investigation of the biases on structural morphology from particle resolution and the simulation models that regulate the formation of structure may be performed by comparing our decomposition results with consistent decompositions of galaxies from other large hydrodynamical simulations such as EAGLE which has comparable mass resolution (Schaye et al., 2015; Crain et al., 2015). However, in such a comparison, the biases from differences in the simulation models on the morphological estimates would need to be carefully examined in order to assess whether particle resolution is the main culprit of the strong mass-dependence on (b/t) estimates in the mock observations.

3.2.2 Comparison with kinematic b/t

An interesting test that is feasible from our image-based decompositions of simulated galaxies is a comparison between the properties derived from kinematic and photometric information from the stellar particles. Comparisons of photometric and kinematic bulge fractions in galaxies have been performed previously in the literature without the large numbers or extensive realism considerations provided here. Scannapieco et al. (2010) used bulge+disc decompositions of mock observations of Milky Way-mass galaxies from hydrodynamical simulations with similar mass resolution to Illustris ($M_{\star} \sim 10^6$) to show that (b/t) is systematically lower from photometry relative to the kinematics. Obreja et al. (2016) reproduced this result in a sample of 18 cosmological zoom-in simulations of galaxies. Each galaxy from Obreja et al. (2016) had a photometric bulge-to-total ratio of $(b/t) \approx 0$ but kinematic ratios ~ 0.5 . However, the zoom-in simulations from Obreja et al. (2016) used adaptive particle mass resolution to each halo – making it difficult to ascertain the effects of particle mass resolution. Still, the implication of each study is that exponential structure of mock observed surface brightness profiles of simulated galaxies does not imply a cold rotationally-supported kinematic disc (i.e. low photometric b/t does not necessarily indicate the lack of a kinematic bulge). These results are complicated by (Christensen et al., 2014) who, while not investigating the differences between photometry and the kinematics, demonstrated that realistic mock-observed photometric bulges can be produced in high-resolution simulations that match well with the photometric properties of real bulges. So, while realistic photo-bulges may be produced, they may not necessarily couple well with the kinematics.

In the simulations, the angular momentum for each particle about the principle rotational axis in a galaxy can be derived easily using the particle velocities and locations relative to the galactic potential. Stars that belong to the bulges of galaxies tend to have a gaussian distribution of velocities (see Kormendy & Ho 2013 for a recent review) whereas stars in the disc have rotationally supported orbits and generally larger coherent angular momenta. Stars (and particles representing stars) can be approximately associated to their stellar components using their angular momentum information to estimate the kinematic bulge-to-total fraction. One measure of the kinematic (b/t) that is common in the literature (see Nelson et al. 2015; Genel et al.

2015) is:

$$(b/t)_{\text{kin}} = 1/N_{\star,\text{tot}} \quad 2 \times N_{\star}(J_z/J(E) < 0) \quad (3.1)$$

where $N_{\star,\text{tot}}$ is the total number of star particles belonging to the galaxy, J_z is an individual particle's component of angular momentum about the principle rotation axis (computed from the angular momenta of all stars within 10 half-mass radii), and $J(E)$ is maximum angular momentum of stellar particles ranked by binding energy ($U_{\text{gravity}} + \nu^2$) within 50 ranks of the particle in question. $N_{\star}(J_z/J(E) < 0)$ is an approximation to the number of stars whose motions are not coherent with the bulk rotation. Because the velocities in the bulge are expected to be normally distributed about $J_z/J(E) = 0$, symmetry provides that $2 \times N_{\star}(J_z/J(E) < 0)$ should approximate the number of stars in the bulge of a galaxy – which can be normalized by the total number of stars to get the kinematic bulge-to-total ratio.

The left panel of Figure 3.8 compares the kinematic and photometric estimates of (b/t) using the decompositions from the ASKA catalog. The RIGs are used to provide a sense of the uncertainties associated with photometric (b/t) by employing the distributions of decomposition results from all placements and camera angles for each galaxy. The vertical position of each point represents the median photometric (b/t) over all placements and camera angles for each galaxy. The error bars show the 95% range centred on the median of the distribution of estimates. The horizontal position for each system is the kinematic (b/t) derived from the stellar orbits. Note that none of the galaxies in this sample have kinematic bulge fractions less than $(b/t)_{\text{kin}} = 0.2$ – which creates immediate tension with our photometric decomposition results. A large number of galaxies with large kinematic bulge fractions have no photo-bulge. So whilst the kinematic (b/t) indicate than many galaxies are bulge dominated, the photometric results are completely disk dominated! Furthermore, no discernible correlation is seen between the kinematic and photometric bulge fractions. The results are consistent with previous findings that the photometric bulge-to-total fractions are systematically lower than the kinematic fractions (Scannapieco et al., 2010; Obreja et al., 2016).

Three RIGs in the left panel of Figure 3.8 are highlighted by star symbols and text showing their subhaloIDs. The highlighted RIGs were selected to enable visual inspection of galaxies with $(b/t)_{\text{phot}} > (b/t)_{\text{kin}}$ (upper right row), $(b/t)_{\text{phot}} \approx (b/t)_{\text{kin}}$ (middle right row), $(b/t)_{\text{phot}} < (b/t)_{\text{kin}}$ (bottom right row). The panels show *gri*

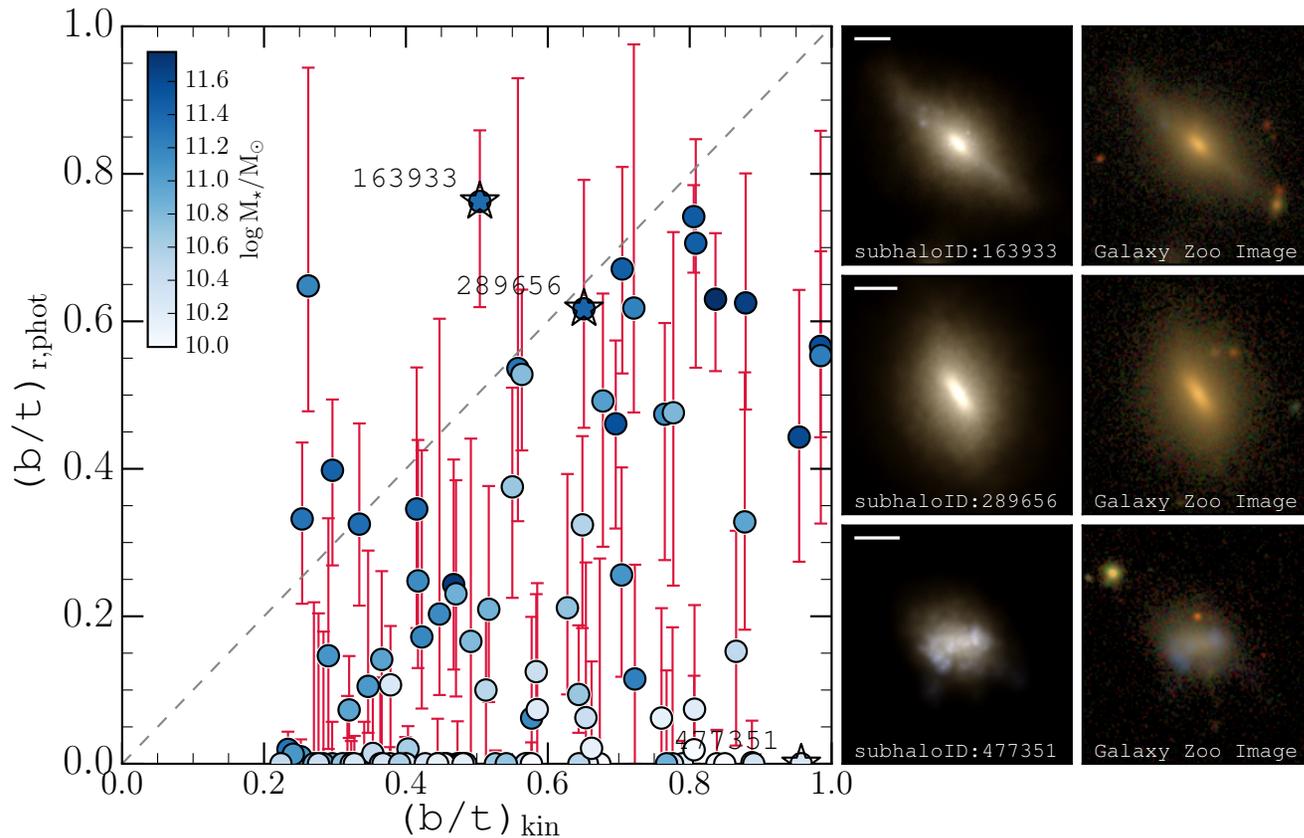


Figure 3.8: Photometric and kinematic bulge-to-total fractions taken from the decompositions results in the ASKA catalog for galaxies in the RIG sample. Points in the left panel show the median photometric $(b/t)_r$ over all placements and camera angles for each galaxy as a function of the kinematic (b/t) derived using Equation 3.1 from the orbits of the stellar particles in the simulation. Each point is colour-coded by the total stellar mass of the galaxy, indicated in the colourbar at the top left of the panel. Error bars show the 95% confidence interval about the medians. The grey dashed line shows the one-to-one relation on this plane. *gri*-colour composites of our synthetic images and Galaxy Zoo (Snyder et al., 2015) images in the right panels correspond to the labeled and points enclosed in star symbols in the left panel. White horizontal lines at the top left of the FoF images measure 10 kpc.

composites of our synthetic images and mock SDSS Galaxy Zoo visual classification images with realism from Snyder et al. (2015) for each highlighted RIG. Visual inspection of the morphology of the RIG with $(b/t)_{\text{phot}} > (b/t)_{\text{kin}}$ shows that it has a strong bulge component but that is embedded within a disc (edge-on in this camera angle). The uncertainties from the distribution of decomposition results are not consistent with the kinematic estimate. However, the photometric (b/t) fraction for this galaxy is reconcilable with its visual appearance. The photometric (b/t) fraction for the galaxy with $(b/t)_{\text{phot}} \approx (b/t)_{\text{kin}}$ in the middle row of images in Figure 3.8 is also visually reconcilable with the photometry. While the galaxy appears to contain a bar that may affect the photometric (b/t) , it is consistent with the kinematically derived quantity. The galaxy shown in the bottom row of images in Figure 3.8 is most intriguing. The galaxy shown in the bottom row has $(b/t)_{\text{phot}} = 0$ for all placements but $(b/t)_{\text{kin}} \approx 1$. However, there is no visual presence of a bulge or a disc – yet the kinematic information indicates that it is almost a pure bulge. Figure 3.8 shows that photometrically derived morphologies can achieve similar results to the kinematics. However, photometric (b/t) estimates for the majority of galaxies in the RIG sample are systematically lower than their kinematic counterparts.

The markers corresponding to each RIG in Figure 3.8 are colour-coded according to their total stellar masses. Colour-coding of the masses enables inspection of the dependence on stellar mass for the photometric and kinematic bulge-to-total estimates. As expected from Figure 3.7, only galaxies with stellar masses $\log M_{\star}/M_{\odot} \gtrsim 11$ contain appreciable photometric bulge fractions. Furthermore, galaxies with low stellar masses $\log M_{\star}/M_{\odot} \lesssim 10.5$ have the largest systematic errors between the photometric and kinematic estimates for (b/t) . The small number of photometric bulges at low masses and the presence of kinematic bulges corroborates with the bias expected from particle resolution. However, the galaxy shown in bottom right row of images in Figure 3.7 either does not generate confidence in the kinematic estimates for low-mass galaxies. Alternatively, the galaxy in the bottom right row of images in Figure 3.7 tells us that kinematics has nothing to do with visual or photometric morphology! Either way, the high kinematic bulge fractions of galaxies with no visual bulge such as seen in the bottom right row of images in Figure 3.7 make pinning particle resolution as the driving bias for reducing photometric (b/t) estimates more challenging. Still, the stronger correlation between kinematic and photometric (b/t) for galaxies with $\log M_{\star}/M_{\odot} \gtrsim 11$ and the systematically low photometric (b/t) for galaxies with masses $\log M_{\star}/M_{\odot} \lesssim 10.5$ presents a strong case for the suppression of photometric

bulge-to-total fractions by particle mass resolution limitations.

Future directions in investigating the photometry of bulges in simulations and its implications in Illustris will require stringent tests that examine radiative transfer schemes, particle resolution, and comparisons with observations. The facts that (1) high-mass galaxies in Illustris contain photo-bulges and (2) mock observations from high- resolution simulations also produce bulges (Christensen et al., 2014) support a scenario in which resolution plays a role in interpretation of bulges in mock observations of simulated galaxies. One seemingly basic observational test is whether observed galaxies also suffer from the same systematics between photometric and kinematic (b/t) as the simulations. However, the limiting factor for observations is that the velocities of resolved stars in galaxies beyond a small local volume of the observed universe are not accessible. For very nearby galaxies, stars that are intrinsically bright such as Red-Giant Branch (RGB) stars can be used to measure velocities and ascertain the stellar components in which they reside (e.g. Wheeler et al. 2015). However, the limitation of velocity estimates to those of RGB stars may bias kinematic decompositions depending on differences between the stellar populations within the physical components. Furthermore, kinematic decompositions based on the velocities of individual stars are limited to very small number of galaxies for which the stellar velocities have been measured (see references in Wheeler et al. 2015). New approaches that combine photometric decompositions with kinematic information from spatially resolved spectroscopy of large samples of galaxies with IFUs (e.g. KMOS: Wisnioski et al. 2015; MUSE: Contini et al. 2015; MaNGA: Bundy et al. 2015; SAMI: Fogarty et al. 2014; CALIFA: Sánchez et al. 2012) derived from morpho-kinematic decomposition tools such as GALPAK^{3D} (Bouché et al., 2015) will enable even more detailed comparisons between mock-observations of simulated galaxies and observations.

Chapter 4

Summary

In this thesis, I have described a new procedure for deriving image-based quantitative morphologies of simulated galaxies that enables fair comparisons with observations. I employ the new procedure in an analysis of galaxies from the Illustris simulation in a way that facilitates comparison with observed galaxies from the SDSS. The central tenet of a meaningful comparison between the observational properties of real and simulated galaxies is that the observational biases and methodology for measuring galaxy properties are consistent. Therefore, I apply an unprecedentedly comprehensive suite of observational realism to galaxies from the simulations and use the same image-based methodology for deriving galaxy properties that is used for observations. My methods are unique from previous attempts at comparisons between simulations and observations by combining three factors: (1) Using mock-observations of simulated galaxies to enable image-based comparisons with observations; (2) Applying a comprehensive suite of observational realism to the simulated galaxy images that facilitates unbiased comparisons with observations (3) Consistent methodology for derivation of parametric structural and photometric properties of observed and simulated galaxies.

4.1 Observational Realism and Morphological Decompositions

As described in Section 2.2.3, I first ensure consistent observational biases for measuring galaxy properties by applying a comprehensive suite of observational realism to the synthetic images of galaxies from the Illustris simulation. Next, paramet-

ric quantitative morphologies of the simulated galaxies are derived from bulge+disc decompositions. In brief, the following procedure was employed in the analyses of simulated galaxies:

- Selection of the SDSS fields into which simulated galaxies are inserted is determined by the projected positions of real galaxies in the SDSS – which ensures that positional biases on decomposition results from crowding, sky brightness, and PSF resolution are statistically the same as for observed galaxies.
- The flux from the simulated galaxies is convolved with the SDSS PSF corresponding to the specific location at which the synthetic images are inserted and signal shot noise is added to the synthetic image flux. Convolution with the reconstructed SDSS PSF for each placement ensures the same statistics for PSF resolution given that placement is assigned quasi-randomly from the positional distribution of real galaxies.
- `SEXTRACTOR` is used to perform deblending of galaxy flux with other sources and the sky. The parameters used in the deblending thresholds are the same as were used by Simard et al. (2011) in their analysis of galaxies from the SDSS. Consistency in the deblending is crucial to consistency in the decompositions (e.g. see Section 2.3.1 and Figure 1.7).
- 2-D parametric surface-brightness decompositions are performed with `GIM2D` – making the procedure for deriving structural and photometric properties of simulated galaxies completely consistent with observed galaxies. Furthermore, the parametric decompositions enable measurement of the structural properties of the physical components from the surface brightness distributions. Both bulge+disc decompositions and pure sérsic decompositions are performed for each galaxy for consistency with the catalogs of Simard et al. (2011) – though, in this work, only the bulge+disc decompositions are exploited in the comparisons with real galaxies.

4.2 Characterization of Biases

Several experiments were designed to characterize the biases in the decomposition results. The decomposition catalogs used in each experiment are described in Section

2.3.3 and summarized in Table 2.1. The results enabled quantification of statistical and systematic errors from observational biases and our decomposition pipeline. Many of the experiments were aimed at characterizing a specific bias, whilst taking judicious precautions to control other biases in each decomposition.

- *Stellar Light Distribution*: Results from the `SMOOTHING` catalog (Section 2.3.3) were used to characterize the biases from stellar light distribution (SLD) schemes used to create the synthetic images (Section 2.4.1). I showed that there are no systematic biases from the choice of SLD scheme on integrated magnitude or half-light radius in a representative sample of galaxies (RIGs). However, the decomposition results for a few diffuse galaxies show that they are systematically fainter and smaller for *all* SLD schemes relative to the corresponding measurements derived directly from the synthetic images – with systematic offsets as large as +3.0 magnitudes and -0.4 dex in half-light radius. The statistical errors in comparisons with alternative SLD schemes are roughly ± 0.2 mag in integrated magnitude and ± 0.04 dex in half-light radius – though much of this is attributed to the large scatter from the diffuse systems. Statistical errors on the bulge-to-total fractions in comparisons with the fiducial scheme were roughly ± 0.1 . However, a systematic trend of reduced (b/t) measurements were identified in comparisons of fixed kernel radius SLD schemes with respect to adaptive. Decomposition results for galaxies using SLD schemes with fixed kernel radii had reduced (b/t) by up to 0.6 relative to adaptive schemes. Detailed structural estimates may therefore be affected SLD schemes and the spatial resolution that they afford.
- *Camera Angle*: Results from the `CAMERAS` catalog (Section 2.3.3) were used to show that decomposition results from the RIGs in different projections yield largely consistent results for the majority of the sample (Section 2.4.2). Magnitude differences from decompositions of galaxies at different camera angles were typically < 0.1 mag and half-light radius differences < 0.15 dex. Diffuse galaxies generally had larger variations up to 2.0 mag in integrated magnitude and 0.6 dex in half-light radius. Bulge-to-total fraction differences appeared to depend on the median (b/t) across all camera angles. Galaxies that appear to be bulge- or disc-dominated from their median estimates have small variation in (b/t). But, galaxies with median $0.2 < (b/t)_{50\%} < 0.8$ show larger sensitivity to projection that appears to peak at $(b/t)_{50\%} \approx 0.5$. It appears that when sig-

nificant bulge *and* disc components are present within a galaxy, then projection significantly affects photometric decompositions of these components.

- *Environment and Crowding*: Results from the ASKA catalog were used to show that the decomposition results are largely robust to biases from crowding, sky brightness, and PSF resolution (Section 2.4.3). Statistical errors on integrated magnitude and half-light radius were < 0.05 mag and < 0.025 dex, respectively. However, diffuse galaxies have larger errors, up to 2 mag and 0.4 dex in magnitude and half-light radius. Statistical errors on the bulge-to-total fraction were related to the median (b/t) for the distribution of decompositions of each galaxy and were as large as 0.3. Median systematic errors on integrated magnitude and half-light radius were small, with $\Delta m_{r,50\%} < 0.05$ mag and $\Delta \log(rhl_{r,50\%}) < 0.05$ dex, for the majority of galaxies. Diffuse galaxies had large systematic offsets in integrated magnitude and half-light radius, with $\Delta m_{r,50\%}$ up to +2 mag and $\Delta \log(rhl_{r,50\%})$ down to -0.4 dex relative to the corresponding measurements derived directly from the synthetic images.
- *Internal Segmentation*: Inspection of the images, models, residuals, and science masks enabled the identification of internal segmentation by locally discrete substructure in the surface brightness distributions of diffuse galaxies (Section 2.4.4). In each of the experiments, diffuse galaxies had the largest systematic and random errors. Using the decomposition results from previous experiments, I found that the deblending in diffuse galaxies with bright substructure was highly sensitive to all observational biases which were the root cause of their large systematic and random errors.

Having characterized the biases in my experiments and catalogs, I performed decompositions of each camera angle of every galaxy in the Illustris synthetic image catalog of Torrey et al. (2015) to enable comparison with observations. First, I showed that roughly 30% of galaxies in Illustris were affected by internal segmentation (Section 2.4.5) and I have quantified the effect of this segmentation on estimates of size and flux using the DISTINCT catalog. Internal segmentation systematically reduced estimates of size (up to $\Delta \log rhl_r \approx -0.4$ dex) and flux (up to $\Delta m_r \approx +2.0$ mag fainter) for galaxies in Illustris and was a consistent bias in each experiment (i.e. no choice of SLD scheme reduced its effects in the decomposition results). However, decomposition results for integrated magnitude and half-light radius in the DISTINCT

catalog for the majority of galaxies were consistent with the properties derived from the synthetic images before the addition of observational biases.

4.3 Comparison with Real Galaxies

In Chapter 3, I employed the bulge+disc decomposition results from the DISTINCT catalog in an unprecedentedly fair comparison with the observed size-luminosity and (b/t) -stellar mass relations. Comparisons between simulated and observed galaxy size-luminosity relations have been previously performed (e.g. Brooks et al. 2011; Christensen et al. 2014; Snyder et al. 2015; Furlong et al. 2015) have been compromised by at least one of the following factors: (1) the limitations on statistical relevance in comparisons to observations due to small simulated galaxy samples; (2) inconsistent derivations of simulated and observed galaxy properties; (3) incomplete observational realism that biases the distributions of derived properties of simulated galaxies in comparisons with observations. I address each of these caveats in my methodology by using mock observations of a full population of simulated galaxies, applying comprehensive observational realism to enable an unbiased image-based comparison, and employing identical methods for deriving galaxy properties in simulated and observed galaxies.

- *Size-Luminosity Relation - First Look*: I matched Illustris galaxies by stellar mass to the stellar mass distribution of the SDSS – taking the SDSS galaxy population over $0.04 < z < 0.06$ with $\log M_*/M_\odot > 10$ as our comparison sample. In Section 3.1.2, I compared the size-luminosity relations of the SDSS and Illustris for the matched sample. Figure 3.1 showed that Illustris galaxies are generally larger and brighter for the same stellar masses as galaxies from the SDSS. Furthermore, the correlation between size and luminosity is not as strong in Illustris (and appears flat) relative to the SDSS relation. I concluded that such a discrepancy, which is particularly puzzling considering the negative systematic errors from internal segmentation, can only be caused by a morphological difference between Illustris and the SDSS.
- *Bulge and Disc Morphologies*: I compared the distributions of (b/t) as a function of total stellar mass using the mass-matched samples from the size-luminosity comparison. Figure 3.5 showed that Illustris is dominated by disc morphologies

at all masses in the sample – where the SDSS demonstrates diverse morphologies. I then showed that Illustris contains photo-bulges, but the relationship between stellar mass and (b/t) is more dichotomous than in the observations (i.e. Illustris contains too many discs at low mass and only high-mass galaxies contain appreciable bulge fractions). The size-luminosity relations of bulges and discs differ significantly. I concluded that matching the observed and simulated samples by stellar mass and by morphology would yield insight on whether morphological differences are affecting the comparison of their size-luminosity relations.

- *Size-Luminosity Relation – Impact of Morphology*: I revisited the size-luminosity relations of SDSS and Illustris – this time matching by stellar mass *and* (b/t) morphology by re-sampling SDSS galaxies to match the (b/t) -stellar mass distribution of Illustris from Figure 3.5. The comparison demonstrated that, indeed, the discrepancy in our previous comparison of the size-luminosity relations was that Illustris contained predominantly disc-dominated galaxies in that sample. By additionally matching by morphology, the agreement between the size-luminosity relations (which is essentially the disc size-luminosity relation) is significantly improved – leaving a reduced magnitude and size offset between the relations. The remaining offset is difficult to characterize without a detailed quantification of the effects of dust in the creation of the synthetic images and on our decomposition results. Given that a comparison that excludes dust has now been performed, future comparison that include dust will enable quantification of its biases.

4.4 The Photo-Bulge Defecit

Illustris contains too few bulges/spheroids at low stellar masses, $\log M_{\star}/M_{\odot} \lesssim 11$, relative to the observations and only has appreciable populations of galaxies with bulges at $\log M_{\star}/M_{\odot} \gtrsim 11$. I concluded that the driving mechanism behind the suppression of photometric bulges in Illustris is the particle resolution of the simulation – which could be tested in an image-based comparison of galaxies taken from the full-volume simulation of Illustris and zoom-in simulations with higher particle resolution (e.g. Sparre & Springel 2016).

Lastly, we compared our decomposition results for the photometric bulge-to-total

ratios with the bulge fractions derived from the kinematics. Confirming previous work using a larger sample and similar resolution (Scannapieco et al., 2010; Obreja et al., 2016), I show in Section 3.2.2 that the photometric estimates for (b/t) are systematically lower than the kinematic estimates. Furthermore, no discernible correlation between the photometric and kinematic estimates of (b/t) is seen. However, I find that galaxies with higher stellar masses, $\log M_{\star}/M_{\odot} \gtrsim 11$, have photometric (b/t) that are generally closer to the kinematic measurements. Although, inspection of several low-mass galaxies $\log M_{\star}/M_{\odot} \lesssim 10.5$ with high kinematic (b/t) indicates no visual presence of a photo-bulge – implying that (a) the spatial resolution is insufficient in these galaxies to resolve the bulge; (b) the kinematic estimate for the bulge does not reflect the true presence of a bulge; (c) there is no underlying connection between kinematics and visual or photometric morphology. A combination of (a) and (b) is also possible. In such a scenario, galaxies that are poorly resolved (both spatially in the images and by particles in the kinematics) may have reduced photometric estimates of (b/t) *and* have intrinsically large uncertainties in the kinematic estimates.

4.5 Future Directions

I will highlight a few of the immediate future directions that are accessible using the methodology I have described and demonstrated in this thesis.

4.5.1 Disc and Bulge Structural Properties

In this thesis, I placed particular emphasis on the validation of my methodology. Therefore, I used measurements that could easily be compared with properties derived directly from the synthetic images (i.e. integrated magnitude and half-light radius). These properties need not assume the analytic profiles of the physical components but do not employ the full gambit of structural information that are accessible from the bulge+disc or pure sérsic decompositions. Therefore, one direction that can be taken with our current catalogs is a comparison of the intrinsic properties of galactic discs taken from the separated disc components of galaxies. Additionally, analysis of the detailed structural properties of low-mass galaxies with bulges in Illustris and comparisons with their kinematics may yield further insights as to why photo-bulges are suppressed at low masses.

4.5.2 Photometric Estimates of Stellar Masses

Our comparisons match by stellar mass to the SDSS within a narrow volume of the observable universe. The stellar masses are derived from SED-template fitting to the results of bulge+disc decompositions in the SDSS *ugriz*-bands (Mendel et al., 2014). An interesting test would be to compare the known stellar masses from the simulation to the inferred masses from photometric decompositions and SED-fitting. Comparison of the results would be informative as to the adequacy of our mass-matching – which I took to be correct in this work. Furthermore, such a study would have implications for all comparisons in which stellar masses derived from observations are compared with simulations.

4.5.3 Redshift Dependence on Structural Estimates

In this first application of my methodology, I placed each galaxy from the $z = 0$ Illustris synthetic image catalog at a fixed distance in order to facilitate characterization of the other observational biases with each galaxy on level-ground with every other. However, using the RIG sample, a new experiment could be designed to test the biases associated with distance. Similar to the **SMOOTHING** and **CAMERAS** catalogs, several distance incarnations of each RIG would be inserted in the same uncrowded location in the SDSS in which decompositions would be performed. Comparison of the decomposition results for each RIG at various distances would inform on accessibility of the detailed structural properties of galaxies as a function of projected distance.

4.5.4 Zooming-in on Illustris

The conflicts between galaxies in Illustris and the observations yield intriguing questions as to the source of the morphological differences between the simulated and observed galaxy populations. To test the hypothesis that spatial or particle resolution is preventing adequate sampling of the bulge component, a deeper investigation of SLD schemes and comparisons with zoom-in simulations of Illustris are required. In particular, a recent zoom-in simulation of Illustris using the same hydrodynamic code and model suite improve the particle resolution by $40\times$ the resolution of the full volume (Sparre & Springel, 2016). The zoom-in will place the particle resolution of the zoomed-in low mass galaxies, $\log M_{\star}/M_{\odot} \lesssim 10.5$, on level-ground with galaxies at high mass in our current comparison – which appear to contain more substantial

numbers of galaxies with bulges. If bulges can be resolved in new synthetic images of the zoomed-in low-mass galaxies, then new constraints can be placed on the necessary resolution to resolve bulges for realistic comparisons with observations. Furthermore, the nature of whether photo-bulges genuinely correlate with the kinematic bulges is another test that may be performed with a comparison of both kinematic and photometric structural estimates of individual galaxies in the zoom-in and the full volume. Given that our lowest mass galaxies had the largest discrepancy in the kinematic and photometric bulge fractions, one could tackle the question of whether an increase by 400% in particle resolution can alleviate the discrepancy.

4.5.5 Application to other Simulations

The sizes derived from the particle information for galaxies are fundamentally biased in comparisons with observations. Comparisons that use mock-observations and observational realism are the only way to place observations and simulations on level-ground such that galaxy properties are derived consistently. The EAGLE simulation (Schaye et al., 2015) recently demonstrated that they reproduce the observed median evolution and scatter in the size-mass relation for actively star forming and passive galaxies over $0 < z < 2$ in a comparison with galaxy properties derived from HST CANDELS imaging (Patel et al., 2013; van Dokkum et al., 2014; van der Wel et al., 2014) using the *physical* information from the simulations (Furlong et al., 2015). This is a highly notable accomplishment of the EAGLE simulation. A very intriguing investigation would be to apply our bulge+disc decomposition and image realism methodology to galaxies from the EAGLE simulation for which synthetic images have also recently been developed for $z = 0$ (Trayford et al., 2015). With added evolutionary context through creation of synthetic images for various evolutionary snapshots in the simulation, two fundamental questions could be addressed with the results: (1) How well do the derived properties from mock observations compare with the observations at each evolutionary step? (2) How do the decomposition results compare with the physical properties of galaxies in EAGLE as a function of redshift (i.e. signal-to-noise)? Furthermore, because the comparison of Trayford et al. (2015) reproduced the observed sizes of galaxies with splitting of their sample into passive and actively star-forming galaxies by specific star-formation rates, the capacity to perform morphological splitting from photometry using my methodology will provide an orthogonal basis for the comparison.

Ultimately, the methodology I have piloted in this thesis is implemented in a way that is fully adaptable to observations from alternative instruments and scalable to even larger imaging surveys from future ground- and space-based facilities. Galaxy samples from the forthcoming Large Synoptic Survey Telescope (LSST) (LSST Science Collaboration et al., 2009) and Wide-Field InfraRed Survey (WFIRST) (Spergel et al., 2015) stand to be roughly two orders of magnitude larger than the largest current samples. In general, unbiased image-based comparisons of consistently derived photometric quantities of large populations of real and simulated galaxies stand to provide a wealth of feedback between computational and observational communities.

Bibliography

- Abadi, M. G., Navarro, J. F., Steinmetz, M., & Eke, V. R. 2003a, *Astrophysical Journal*, 591, 499
- 2003b, *Astrophysical Journal*, 597, 21
- Abazajian, K. N. et al. 2009, *Astrophysical Journal*, Supplement Series, 182, 543
- Abell, G. O. 1965, *Annual Review of Astronomy and Astrophysics*, 3, 1
- Abraham, R. G., Valdes, F., Yee, H. K. C., & van den Bergh, S. 1994, *Astrophysical Journal*, 432, 75
- Abraham, R. G., van den Bergh, S., Glazebrook, K., Ellis, R. S., Santiago, B. X., Surma, P., & Griffiths, R. E. 1996, *Astrophysical Journal*, Supplement Series, 107, 1
- Abraham, R. G., van den Bergh, S., & Nair, P. 2003, *Astrophysical Journal*, 588, 218
- Adelman-McCarthy, J. K. et al. 2008, *Astrophysical Journal*, Supplement Series, 175, 297
- 2006, *Astrophysical Journal*, Supplement Series, 162, 38
- Agertz, O. & Kravtsov, A. V. 2015, ArXiv e-prints
- Agertz, O., Teyssier, R., & Moore, B. 2011, *Monthly Notices of the RAS*, 410, 1391
- Allen, P. D., Driver, S. P., Graham, A. W., Cameron, E., Liske, J., & de Propriis, R. 2006, *Monthly Notices of the RAS*, 371, 2
- Alonso, M. S., Tissera, P. B., Coldwell, G., & Lambas, D. G. 2004, *Monthly Notices of the RAS*, 352, 1081

- Angulo, R. E., Springel, V., White, S. D. M., Jenkins, A., Baugh, C. M., & Frenk, C. S. 2012, *Monthly Notices of the RAS*, 426, 2046
- Annis, J. et al. 2014, *Astrophysical Journal*, 794, 120
- Athanassoula, E. 2005, *Monthly Notices of the RAS*, 358, 1477
- Aubourg, E. et al. 1993, *Nature*, 365, 623
- Aumer, M., White, S. D. M., Naab, T., & Scannapieco, C. 2013, *Monthly Notices of the RAS*, 434, 3142
- Baes, M., Verstappen, J., De Looze, I., Fritz, J., Saftly, W., Vidal Pérez, E., Stalevski, M., & Valcke, S. 2011, *Astrophysical Journal, Supplement Series*, 196, 22
- Bahcall, N. A. 1988, *Annual Review of Astronomy and Astrophysics*, 26, 631
- Baldry, I. K. et al. 2012, *Monthly Notices of the RAS*, 421, 621
- 2010, *Monthly Notices of the RAS*, 404, 86
- Bamford, S. P. et al. 2009, *Monthly Notices of the RAS*, 393, 1324
- Baugh, C. M., Cole, S., & Frenk, C. S. 1996, *Monthly Notices of the RAS*, 283, 1361
- Bender, R. & Moellenhoff, C. 1987, *Astronomy and Astrophysics*, 177, 71
- Bennett, C. L. et al. 2013, *Astrophysical Journal, Supplement Series*, 208, 20
- Bernardi, M., Hyde, J. B., Sheth, R. K., Miller, C. J., & Nichol, R. C. 2007, *Astronomical Journal*, 133, 1741
- Bershady, M. A., Jangren, A., & Conselice, C. J. 2000, *Astronomical Journal*, 119, 2645
- Bertin, E. & Arnouts, S. 1996, *Astronomy and Astrophysics, Supplement Series*, 117, 393
- Bertone, G., Hooper, D., & Silk, J. 2005, *Phys. Rep.*, 405, 279
- Blanton, M. R. et al. 2001, *Astronomical Journal*, 121, 2358
- Blanton, M. R. & Moustakas, J. 2009, *Annual Review of Astronomy and Astrophysics*, 47, 159

- Blanton, M. R. et al. 2005, *Astronomical Journal*, 129, 2562
- Blumenthal, G. R., Faber, S. M., Flores, R., & Primack, J. R. 1986, *Astrophysical Journal*, 301, 27
- Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, *Nature*, 311, 517
- Bouché, N., Carfantan, H., Schroetter, I., Michel-Dansac, L., & Contini, T. 2015, *Astronomical Journal*, 150, 92
- Bower, R. G., Benson, A. J., Malbon, R., Helly, J. C., Frenk, C. S., Baugh, C. M., Cole, S., & Lacey, C. G. 2006, *Monthly Notices of the RAS*, 370, 645
- Brook, C. B. et al. 2011, *Monthly Notices of the RAS*, 415, 1051
- Brook, C. B., Stinson, G., Gibson, B. K., Roškar, R., Wadsley, J., & Quinn, T. 2012a, *Monthly Notices of the RAS*, 419, 771
- Brook, C. B., Stinson, G., Gibson, B. K., Wadsley, J., & Quinn, T. 2012b, *Monthly Notices of the RAS*, 424, 1275
- Brooks, A. M. et al. 2011, *Astrophysical Journal*, 728, 51
- Bundy, K. et al. 2015, *Astrophysical Journal*, 798, 7
- Buta, R. J. 2011, ArXiv e-prints
- Calvi, R., Poggianti, B. M., Fasano, G., & Vulcani, B. 2012, *Monthly Notices of the RAS*, 419, L14
- Casteels, K. R. V. et al. 2013, *Monthly Notices of the RAS*, 429, 1051
- Charlot, S. & Fall, S. M. 2000, *Astrophysical Journal*, 539, 718
- Chincarini, G. & Vettolani, G. 1987, in *IAU Symposium*, Vol. 124, *Observational Cosmology*, ed. A. Hewitt, G. Burbidge, & L. Z. Fang, 275–300
- Christensen, C. R., Brooks, A. M., Fisher, D. B., Governato, F., McCleary, J., Quinn, T. R., Shen, S., & Wadsley, J. 2014, *Monthly Notices of the RAS*, 440, L51
- Cole, S., Lacey, C. G., Baugh, C. M., & Frenk, C. S. 2000, *Monthly Notices of the RAS*, 319, 168

- Colless, M. 1999, in *Large-Scale Structure in the Universe*, ed. G. Efstathiou & et al., 105
- Colless, M. et al. 2001, *Monthly Notices of the RAS*, 328, 1039
- Conroy, C., White, M., & Gunn, J. E. 2010, *Astrophysical Journal*, 708, 58
- Conselice, C. J. 1997, *Publications of the ASP*, 109, 1251
- 2003, *Astrophysical Journal*, Supplement Series, 147, 1
- 2014, *Annual Review of Astronomy and Astrophysics*, 52, 291
- Conselice, C. J., Bershadsky, M. A., & Gallagher, III, J. S. 2000a, *Astronomy and Astrophysics*, 354, L21
- Conselice, C. J., Bershadsky, M. A., & Jangren, A. 2000b, *Astrophysical Journal*, 529, 886
- Contini, T. et al. 2015, ArXiv e-prints
- Corwin, Jr., H. G., Buta, R. J., & de Vaucouleurs, G. 1994, *Astronomical Journal*, 108, 2128
- Courteau, S., Dutton, A. A., van den Bosch, F. C., MacArthur, L. A., Dekel, A., McIntosh, D. H., & Dale, D. A. 2007, *Astrophysical Journal*, 671, 203
- Crain, R. A. et al. 2015, *Monthly Notices of the RAS*, 450, 1937
- 2009, *Monthly Notices of the RAS*, 399, 1773
- Croft, R. A. C., Di Matteo, T., Springel, V., & Hernquist, L. 2009, *Monthly Notices of the RAS*, 400, 43
- Darg, D. W. et al. 2010, *Monthly Notices of the RAS*, 401, 1043
- Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *Astrophysical Journal*, 292, 371
- Davis, M. et al. 2003, in *Proc. SPIE*, Vol. 4834, *Discoveries and Research Prospects from 6- to 10-Meter-Class Telescopes II*, ed. P. Guhathakurta, 161–172

- Davis, M., Huchra, J., Latham, D. W., & Tonry, J. 1982, *Astrophysical Journal*, 253, 423
- de Souza, R. E., Gadotti, D. A., & dos Anjos, S. 2004, *Astrophysical Journal, Supplement Series*, 153, 411
- de Vaucouleurs, G. 1948, *Annales d'Astrophysique*, 11, 247
- 1953, *Monthly Notices of the RAS*, 113, 134
- 1959a, *Handbuch der Physik*, 53, 275
- 1959b, *Handbuch der Physik*, 53, 311
- Dekel, A. & Silk, J. 1986, *Astrophysical Journal*, 303, 39
- Deng, X.-F., Luo, C.-H., Jiang, P., & Ding, Y.-P. 2013, *Astrophysics and Space Science*, 347, 183
- Dopita, M. A. et al. 2006a, *Astrophysical Journal, Supplement Series*, 167, 177
- 2006b, *Astrophysical Journal*, 647, 244
- 2005, *Astrophysical Journal*, 619, 755
- Dressler, A. 1980, *Astrophysical Journal*, 236, 351
- 1984, *Annual Review of Astronomy and Astrophysics*, 22, 185
- Driver, S. P. et al. 2009, in *IAU Symposium, Vol. 254, The Galaxy Disk in Cosmological Context*, ed. J. Andersen, Nordströara, B. m, & J. Bland-Hawthorn, 469–474
- Dutton, A. A. et al. 2011, *Monthly Notices of the RAS*, 410, 1660
- Eisenstein, D. J. et al. 2011, *Astronomical Journal*, 142, 72
- Ellison, S. L., Mendel, J. T., Patton, D. R., & Scudder, J. M. 2013, *Monthly Notices of the RAS*, 435, 3627
- Ellison, S. L., Patton, D. R., Mendel, J. T., & Scudder, J. M. 2011, *Monthly Notices of the RAS*, 418, 2043
- Ellison, S. L., Patton, D. R., Simard, L., & McConnell, A. W. 2008, *Astronomical Journal*, 135, 1877

- Ellison, S. L., Patton, D. R., Simard, L., McConnachie, A. W., Baldry, I. K., & Mendel, J. T. 2010, *Monthly Notices of the RAS*, 407, 1514
- Emsellem, E., Monnet, G., & Bacon, R. 1994, *Astronomy and Astrophysics*, 285
- Faber, S. M. & Jackson, R. E. 1976, *Astrophysical Journal*, 204, 668
- Falco, E. E. et al. 1999, *Publications of the ASP*, 111, 438
- Fall, S. M. & Efstathiou, G. 1980, *Monthly Notices of the RAS*, 193, 189
- Fisher, D. B., Bolatto, A., Drory, N., Combes, F., Blitz, L., & Wong, T. 2013, *Astrophysical Journal*, 764, 174
- Fisher, D. B. & Drory, N. 2008, *Astronomical Journal*, 136, 773
- 2010, *Astrophysical Journal*, 716, 942
- Fogarty, L. M. R. et al. 2014, *Monthly Notices of the RAS*, 443, 485
- Folkes, S. et al. 1999, *Monthly Notices of the RAS*, 308, 459
- Fortson, L. et al. 2012, *Galaxy Zoo: Morphological Classification and Citizen Science*, ed. M. J. Way, J. D. Scargle, K. M. Ali, & A. N. Srivastava, 213–236
- Freeman, K. C. 1970, *Astrophysical Journal*, 160, 811
- 2007, *Astrophysics and Space Science Proceedings*, 3, 3
- Freeman, P. E., Izbicki, R., Lee, A. B., Newman, J. A., Conselice, C. J., Koekemoer, A. M., Lotz, J. M., & Mozena, M. 2013, *Monthly Notices of the RAS*, 434, 282
- Frenk, C. S., White, S. D. M., Efstathiou, G., & Davis, M. 1985, *Nature*, 317, 595
- Furlong, M. et al. 2015, *ArXiv e-prints*
- Gadotti, D. A. 2009, *Monthly Notices of the RAS*, 393, 1531
- Gadotti, D. A., Baes, M., & Falony, S. 2010, *Monthly Notices of the RAS*, 403, 2053
- Genel, S., Fall, S. M., Hernquist, L., Vogelsberger, M., Snyder, G. F., Rodriguez-Gomez, V., Sijacki, D., & Springel, V. 2015, *Astrophysical Journal*, Letters to the Editor, 804, L40

- Genel, S. et al. 2014, *Monthly Notices of the RAS*, 445, 175
- Gott, III, J. R., Dickinson, M., & Melott, A. L. 1986, *Astrophysical Journal*, 306, 341
- Governato, F. et al. 2004, *Astrophysical Journal*, 607, 688
- Graham, A. W., Driver, S. P., Petrosian, V., Conselice, C. J., Bershadsky, M. A., Crawford, S. M., & Goto, T. 2005, *Astronomical Journal*, 130, 1535
- Graham, A. W. & Worley, C. C. 2008, *Monthly Notices of the RAS*, 388, 1708
- Grogin, N. A. et al. 2011, *Astrophysical Journal*, Supplement Series, 197, 35
- Groth, E. J. & Peebles, P. J. E. 1977, *Astrophysical Journal*, 217, 385
- Groves, B., Dopita, M. A., Sutherland, R. S., Kewley, L. J., Fischera, J., Leitherer, C., Brandl, B., & van Breugel, W. 2008, *Astrophysical Journal*, Supplement Series, 176, 438
- Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, *Astrophysical Journal*, 742, 76
- Gunn, J. E. et al. 2006, *Astronomical Journal*, 131, 2332
- Guo, Q. et al. 2011, *Monthly Notices of the RAS*, 413, 101
- Hatton, S., Devriendt, J. E. G., Ninin, S., Bouchet, F. R., Guiderdoni, B., & Vibert, D. 2003, *Monthly Notices of the RAS*, 343, 75
- Häußler, B. et al. 2013, *Monthly Notices of the RAS*, 430, 330
- Hayward, C. C., Behroozi, P. S., Somerville, R. S., Primack, J. R., Moreno, J., & Wechsler, R. H. 2013a, *Monthly Notices of the RAS*, 434, 2572
- Hayward, C. C., Jonsson, P., Kereš, D., Magnelli, B., Hernquist, L., & Cox, T. J. 2012, *Monthly Notices of the RAS*, 424, 951
- Hayward, C. C., Narayanan, D., Kereš, D., Jonsson, P., Hopkins, P. F., Cox, T. J., & Hernquist, L. 2013b, *Monthly Notices of the RAS*, 428, 2529
- He, Y. Q., Xia, X. Y., Hao, C. N., Jing, Y. P., Mao, S., & Li, C. 2013, *Astrophysical Journal*, 773, 37

- Hinshaw, G. et al. 2013, *Astrophysical Journal*, Supplement Series, 208, 19
- Hogg, D. W. 1999, *ArXiv Astrophysics e-prints*
- Hopkins, P. F., Kereš, D., Oñorbe, J., Faucher-Giguère, C.-A., Quataert, E., Murray, N., & Bullock, J. S. 2014, *Monthly Notices of the RAS*, 445, 581
- Hubble, E. 1929, *Proceedings of the National Academy of Science*, 15, 168
- Hubble, E. P. 1926, *Astrophysical Journal*, 64
- 1936, *Realm of the Nebulae*
- Huchra, J. P. et al. 2012, *Astrophysical Journal*, Supplement Series, 199, 26
- Ivezic, Z. et al. 2008, *ArXiv e-prints*
- Jonsson, P. 2006, *Monthly Notices of the RAS*, 372, 2
- Jonsson, P., Groves, B. A., & Cox, T. J. 2010, *Monthly Notices of the RAS*, 403, 17
- Katz, N. & Gunn, J. E. 1991, *Astrophysical Journal*, 377, 365
- Katz, N., Hernquist, L., & Weinberg, D. H. 1992, *Astrophysical Journal*, Letters to the Editor, 399, L109
- Katz, N., Weinberg, D. H., & Hernquist, L. 1996, *Astrophysical Journal*, Supplement Series, 105, 19
- Kauffmann, G., Huang, M.-L., Moran, S., & Heckman, T. M. 2015, *Monthly Notices of the RAS*, 451, 878
- Kauffmann, G. & White, S. D. M. 1993, *Monthly Notices of the RAS*, 261
- Kauffmann, G., White, S. D. M., & Guiderdoni, B. 1993, *Monthly Notices of the RAS*, 264, 201
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, *Monthly Notices of the RAS*, 363, 2
- Kereš, D., Vogelsberger, M., Sijacki, D., Springel, V., & Hernquist, L. 2012, *Monthly Notices of the RAS*, 425, 2027

- Kim, J., Park, C., Rossi, G., Lee, S. M., & Gott, III, J. R. 2011, *Journal of Korean Astronomical Society*, 44, 217
- Kitchin, C. R. 1987, *Stars, nebulae and the interstellar medium. Observational physics and astrophysics*
- Komatsu, E. et al. 2009, *Astrophysical Journal, Supplement Series*, 180, 330
- Kormendy, J. 1977, *Astrophysical Journal*, 217, 406
- 1985, *Astrophysical Journal, Letters to the Editor*, 292, L9
- Kormendy, J. & Djorgovski, S. 1989, *Annual Review of Astronomy and Astrophysics*, 27, 235
- Kormendy, J. & Fisher, D. B. 2008, in *Astronomical Society of the Pacific Conference Series, Vol. 396, Formation and Evolution of Galaxy Disks*, ed. J. G. Funes & E. M. Corsini, 297
- Kormendy, J., Fisher, D. B., Cornell, M. E., & Bender, R. 2009, *Astrophysical Journal, Supplement Series*, 182, 216
- Kormendy, J. & Ho, L. C. 2013, *Annual Review of Astronomy and Astrophysics*, 51, 511
- Kormendy, J. & Kennicutt, Jr., R. C. 2004, *Annual Review of Astronomy and Astrophysics*, 42, 603
- Kuzio de Naray, R. 2008, in *Astronomical Society of the Pacific Conference Series, Vol. 396, Formation and Evolution of Galaxy Disks*, ed. J. G. Funes & E. M. Corsini, 445
- Lauer, T. R. et al. 2007, *Astrophysical Journal*, 662, 808
- Law, D. R., Steidel, C. C., Erb, D. K., Pettini, M., Reddy, N. A., Shapley, A. E., Adelberger, K. L., & Simenc, D. J. 2007, *Astrophysical Journal*, 656, 1
- Le Fèvre, O. et al. 2005, *Astronomy and Astrophysics*, 439, 845
- Leitherer, C., Ortiz Otálvaro, P. A., Bresolin, F., Kudritzki, R.-P., Lo Faro, B., Pauldrach, A. W. A., Pettini, M., & Rix, S. A. 2010, *Astrophysical Journal, Supplement Series*, 189, 309

- Leitherer, C. et al. 1999, *Astrophysical Journal*, Supplement Series, 123, 3
- Li, Y. et al. 2008, *Astrophysical Journal*, 678, 41
- Lilje, P. B. & Efstathiou, G. 1988, *Monthly Notices of the RAS*, 231, 635
- Lilly, S. J. et al. 2007, *Astrophysical Journal*, Supplement Series, 172, 70
- Lintott, C. J. et al. 2008, *Monthly Notices of the RAS*, 389, 1179
- Lisker, T., Grebel, E. K., Binggeli, B., & Glatt, K. 2007, *Astrophysical Journal*, 660, 1186
- Lotz, J. M., Primack, J., & Madau, P. 2004, *Astronomical Journal*, 128, 163
- LSST Science Collaboration et al. 2009, ArXiv e-prints
- Lupton, R., Gunn, J. E., Ivezić, Z., Knapp, G. R., & Kent, S. 2001, in *Astronomical Society of the Pacific Conference Series*, Vol. 238, *Astronomical Data Analysis Software and Systems X*, ed. F. R. Harnden, Jr., F. A. Primini, & H. E. Payne, 269
- Lupton, R. H., Ivezić, Z., Gunn, J. E., Knapp, G., Strauss, M. A., & Yasuda, N. 2002, in *Proc. SPIE*, Vol. 4836, *Survey and Other Telescope Technologies and Discoveries*, ed. J. A. Tyson & S. Wolff, 350–356
- Ma, C.-P., Greene, J. E., McConnell, N., Janish, R., Blakeslee, J. P., Thomas, J., & Murphy, J. D. 2014, *Astrophysical Journal*, 795, 158
- MacArthur, L. A., Courteau, S., & Holtzman, J. A. 2003, *Astrophysical Journal*, 582, 689
- MacArthur, L. A., Ellis, R. S., Treu, T., U, V., Bundy, K., & Moran, S. 2008, *Astrophysical Journal*, 680, 70
- Mandelbaum, R. et al. 2005, *Monthly Notices of the RAS*, 361, 1287
- Marinacci, F., Fraternali, F., Binney, J., Nipoti, C., Ciotti, L., & Londrillo, P. 2012, in *European Physical Journal Web of Conferences*, Vol. 19, *European Physical Journal Web of Conferences*, 08008

- Marinacci, F., Pakmor, R., & Springel, V. 2014, *Monthly Notices of the RAS*, 437, 1750
- Masters, K. L. et al. 2011, *Monthly Notices of the RAS*, 411, 2026
- McCarthy, I. G., Font, A. S., Crain, R. A., Deason, A. J., Schaye, J., & Theuns, T. 2012a, *Monthly Notices of the RAS*, 420, 2245
- McCarthy, I. G., Schaye, J., Bird, S., & Le Brun, A. M. C. 2016, *ArXiv e-prints*
- McCarthy, I. G., Schaye, J., Font, A. S., Theuns, T., Frenk, C. S., Crain, R. A., & Dalla Vecchia, C. 2012b, *Monthly Notices of the RAS*, 427, 379
- McKinnon, R., Torrey, P., & Vogelsberger, M. 2016, *Monthly Notices of the RAS*, 457, 3775
- Menanteau, F., Ford, H. C., Motta, V., Benítez, N., Martel, A. R., Blakeslee, J. P., & Infante, L. 2006, *Astronomical Journal*, 131, 208
- Mendel, J. T., Simard, L., Palmer, M., Ellison, S. L., & Patton, D. R. 2014, *Astrophysical Journal, Supplement Series*, 210, 3
- Miller, S. H., Bundy, K., Sullivan, M., Ellis, R. S., & Treu, T. 2011, *Astrophysical Journal*, 741, 115
- Moore, B. 1994, *Nature*, 370, 629
- Morgan, W. W. 1962, *Astrophysical Journal*, 135, 1
- Mortlock, A. et al. 2013, *Monthly Notices of the RAS*, 433, 1185
- Munshi, F. et al. 2013, *Astrophysical Journal*, 766, 56
- Murali, C., Katz, N., Hernquist, L., Weinberg, D. H., & Davé, R. 2002, *Astrophysical Journal*, 571, 1
- Nair, P. B. & Abraham, R. G. 2010, *Astrophysical Journal, Supplement Series*, 186, 427
- Nakamura, O., Fukugita, M., Yasuda, N., Loveday, J., Brinkmann, J., Schneider, D. P., Shimasaku, K., & SubbaRao, M. 2003, *Astronomical Journal*, 125, 1682

- Navarro, J. F. & Steinmetz, M. 2000, *Astrophysical Journal*, 538, 477
- Navarro, J. F. & White, S. D. M. 1994, *Monthly Notices of the RAS*, 267, 401
- Nelson, D. et al. 2015, ArXiv e-prints
- Newman, J. A. et al. 2013, *Astrophysical Journal*, Supplement Series, 208, 5
- Obreja, A., Stinson, G. S., Dutton, A. A., Macciò, A. V., Wang, L., & Kang, X. 2016, *Monthly Notices of the RAS*, 459, 467
- Ocvirk, P., Pichon, C., & Teyssier, R. 2008, *Monthly Notices of the RAS*, 390, 1326
- Okamoto, T., Eke, V. R., Frenk, C. S., & Jenkins, A. 2005, *Monthly Notices of the RAS*, 363, 1299
- Oort, J. H. 1983, *Annual Review of Astronomy and Astrophysics*, 21, 373
- Oppenheimer, B. D., Davé, R., Kereš, D., Fardal, M., Katz, N., Kollmeier, J. A., & Weinberg, D. H. 2010, *Monthly Notices of the RAS*, 406, 2325
- Ostriker, J. P. & Peebles, P. J. E. 1973, *Astrophysical Journal*, 186, 467
- Overzier, R., Lemson, G., Angulo, R. E., Bertin, E., Blaizot, J., Henriques, B. M. B., Marleau, G.-D., & White, S. D. M. 2013, *Monthly Notices of the RAS*, 428, 778
- Papovich, C., Dickinson, M., Giavalisco, M., Conselice, C. J., & Ferguson, H. C. 2005, *Astrophysical Journal*, 631, 101
- Papovich, C., Giavalisco, M., Dickinson, M., Conselice, C. J., & Ferguson, H. C. 2003, *Astrophysical Journal*, 598, 827
- Patel, S. G. et al. 2013, *Astrophysical Journal*, 766, 15
- Patton, D. R., Ellison, S. L., Simard, L., McConnachie, A. W., & Mendel, J. T. 2011, *Monthly Notices of the RAS*, 412, 591
- Patton, D. R., Torrey, P., Ellison, S. L., Mendel, J. T., & Scudder, J. M. 2013, *Monthly Notices of the RAS*, 433, L59
- Peebles, P. J. E. 1980, *The large-scale structure of the universe*

- Peng, C. Y., Ho, L. C., Impey, C. D., & Rix, H.-W. 2002, *Astronomical Journal*, 124, 266
- 2010, *Astronomical Journal*, 139, 2097
- Percival, W. J., Cole, S., Eisenstein, D. J., Nichol, R. C., Peacock, J. A., Pope, A. C., & Szalay, A. S. 2007, *Monthly Notices of the RAS*, 381, 1053
- Perez, J., Tissera, P., Padilla, N., Alonso, M. S., & Lambas, D. G. 2009, *Monthly Notices of the RAS*, 399, 1157
- Perlmutter, S. 2003, *Physics Today*, 56, 53
- Petrosian, V. 1976, *Astrophysical Journal, Letters to the Editor*, 209, L1
- Planck Collaboration et al. 2015, *ArXiv e-prints*
- Refregier, A. 2003, *Annual Review of Astronomy and Astrophysics*, 41, 645
- Reyes, R., Mandelbaum, R., Gunn, J. E., Pizagno, J., & Lackner, C. N. 2011, *Monthly Notices of the RAS*, 417, 2347
- Reynolds, J. H. 1916, *The Observatory*, 39, 174
- Riess, A. G. et al. 1998, *Astronomical Journal*, 116, 1009
- Rix, H.-W. & Zaritsky, D. 1995, *Astrophysical Journal*, 447, 82
- Robitaille, T. P. 2011, *Astronomy and Astrophysics*, 536, A79
- Rood, H. J. 1988, *Annual Review of Astronomy and Astrophysics*, 26, 245
- Rubin, V. C. & Ford, Jr., W. K. 1970, *Astrophysical Journal*, 159, 379
- Rubin, V. C., Ford, W. K. J., & Thonnard, N. 1980, *Astrophysical Journal*, 238, 471
- Sales, L. V., Navarro, J. F., Schaye, J., Dalla Vecchia, C., Springel, V., & Booth, C. M. 2010, *Monthly Notices of the RAS*, 409, 1541
- Sales, L. V., Navarro, J. F., Theuns, T., Schaye, J., White, S. D. M., Frenk, C. S., Crain, R. A., & Dalla Vecchia, C. 2012, *Monthly Notices of the RAS*, 423, 1544

- Sánchez, S. F. et al. 2012, *Astronomy and Astrophysics*, 538, A8
- Sandage, A. 1961, *The Hubble atlas of galaxies*
- 1975, *Classification and Stellar Content of Galaxies Obtained from Direct Photography*, ed. A. Sandage, M. Sandage, & J. Kristian (the University of Chicago Press), 1
- Sawala, T. et al. 2016, *Monthly Notices of the RAS*, 457, 1931
- Scannapieco, C., Gadotti, D. A., Jonsson, P., & White, S. D. M. 2010, *Monthly Notices of the RAS*, 407, L41
- Scannapieco, C. et al. 2012, *Monthly Notices of the RAS*, 423, 1726
- Schade, D., Carlberg, R. G., Yee, H. K. C., Lopez-Cruz, O., & Ellingson, E. 1996, *Astrophysical Journal, Letters to the Editor*, 464, L63
- Schade, D., Lilly, S. J., Crampton, D., Hammer, F., Le Fevre, O., & Tresse, L. 1995, *Astrophysical Journal, Letters to the Editor*, 451, L1
- Schawinski, K. et al. 2014, *Monthly Notices of the RAS*, 440, 889
- Schaye, J. et al. 2015, *Monthly Notices of the RAS*, 446, 521
- 2010, *Monthly Notices of the RAS*, 402, 1536
- Scudder, J. M., Ellison, S. L., Torrey, P., Patton, D. R., & Mendel, J. T. 2012, *Monthly Notices of the RAS*, 426, 549
- Sérsic, J. L. 1963, *Boletin de la Asociacion Argentina de Astronomia La Plata Argentina*, 6, 41
- Shane, C. D. & Wirtanen, C. A. 1954, *Astronomical Journal*, 59, 285
- Shen, S., Mo, H. J., White, S. D. M., Blanton, M. R., Kauffmann, G., Voges, W., Brinkmann, J., & Csabai, I. 2003, *Monthly Notices of the RAS*, 343, 978
- Simard, L. 1998, in *Astronomical Society of the Pacific Conference Series*, Vol. 145, *Astronomical Data Analysis Software and Systems VII*, ed. R. Albrecht, R. N. Hook, & H. A. Bushouse, 108

- Simard, L., Mendel, J. T., Patton, D. R., Ellison, S. L., & McConnachie, A. W. 2011, *Astrophysical Journal, Supplement Series*, 196, 11
- Simard, L. et al. 2002, *Astrophysical Journal, Supplement Series*, 142, 1
- Snyder, G. F. et al. 2015, *Monthly Notices of the RAS*, 454, 1886
- Somerville, R. S. & Davé, R. 2015, *Annual Review of Astronomy and Astrophysics*, 53, 51
- Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., & Hernquist, L. 2008, *Monthly Notices of the RAS*, 391, 481
- Somerville, R. S. & Primack, J. R. 1999, *Monthly Notices of the RAS*, 310, 1087
- Sparre, M. & Springel, V. 2016, *ArXiv e-prints*
- Spergel, D. et al. 2015, *ArXiv e-prints*
- Springel, V. & Hernquist, L. 2003, *Monthly Notices of the RAS*, 339, 289
- Springel, V. et al. 2005, *Nature*, 435, 629
- Steinacker, J., Baes, M., & Gordon, K. D. 2013, *Annual Review of Astronomy and Astrophysics*, 51, 63
- Stoughton, C. et al. 2002, *Astronomical Journal*, 123, 485
- Strateva, I. et al. 2001, *Astronomical Journal*, 122, 1861
- Takamiya, M. 1999, *Astrophysical Journal, Supplement Series*, 122, 109
- Tasca, L. A. M. & White, S. D. M. 2011, *Astronomy and Astrophysics*, 530, A106
- Tempel, E. 2011, PhD thesis, PhD thesis, *Dissertationes astronomiae Universitatis Tartuensis*, 104 pages, 2011, ISBN: 978-9949-19-627-2
- Teyssier, R. et al. 2009, *Astronomy and Astrophysics*, 497, 335
- Thompson, R., Davé, R., Huang, S., & Katz, N. 2015, *ArXiv e-prints*
- Toomre, A. & Toomre, J. 1972, *Astrophysical Journal*, 178, 623
- Torrey, P. et al. 2015, *Monthly Notices of the RAS*, 447, 2753

- Torrey, P., Vogelsberger, M., Genel, S., Sijacki, D., Springel, V., & Hernquist, L. 2014, *Monthly Notices of the RAS*, 438, 1985
- Torrey, P., Vogelsberger, M., Sijacki, D., Springel, V., & Hernquist, L. 2012, *Monthly Notices of the RAS*, 427, 2224
- Trayford, J. W., Theuns, T., Bower, R. G., Crain, R. A., Lagos, C. d. P., Schaller, M., & Schaye, J. 2016, *Monthly Notices of the RAS*
- Trayford, J. W. et al. 2015, *Monthly Notices of the RAS*, 452, 2879
- Trujillo, I. et al. 2004, *Astrophysical Journal*, 604, 521
- Tully, R. B., Mould, J. R., & Aaronson, M. 1982, *Astrophysical Journal*, 257, 527
- van den Bergh, S. 1998, *Galaxy Morphology and Classification*
- van der Kruit, P. C. & Searle, L. 1982, *Astronomy and Astrophysics*, 110, 61
- van der Wel, A. et al. 2014, *Astrophysical Journal*, 788, 28
- van Dokkum, P. G. et al. 2014, *Astrophysical Journal*, 791, 45
- van Zee, L. 2000, *Astronomical Journal*, 119, 2757
- Vázquez, G. A. & Leitherer, C. 2005, *Astrophysical Journal*, 621, 695
- Vika, M., Bamford, S. P., Häußler, B., Rojas, A. L., Borch, A., & Nichol, R. C. 2013, *Monthly Notices of the RAS*, 435, 623
- Vogelsberger, M., Genel, S., Sijacki, D., Torrey, P., Springel, V., & Hernquist, L. 2013, *Monthly Notices of the RAS*, 436, 3031
- Vogelsberger, M. et al. 2014a, *Nature*, 509, 177
- 2014b, *Monthly Notices of the RAS*, 444, 1518
- Vogelsberger, M., Sijacki, D., Kereš, D., Springel, V., & Hernquist, L. 2012, *Monthly Notices of the RAS*, 425, 3024
- von Hoerner, S. 1960, *ZAp*, 50
- Wadsley, J. W., Stadel, J., & Quinn, T. 2004, *New Astronomy*, 9, 137

- Weinberg, D. H., Hernquist, L., & Katz, N. 1997, *Astrophysical Journal*, 477, 8
- Wheeler, C., Pace, A. B., Bullock, J. S., Boylan-Kolchin, M., Onorbe, J., Fitts, A., Hopkins, P. F., & Keres, D. 2015, ArXiv e-prints
- White, S. D. M. & Rees, M. J. 1978, *Monthly Notices of the RAS*, 183, 341
- Willett, K. W. et al. 2013, *Monthly Notices of the RAS*, 435, 2835
- 2015, *Monthly Notices of the RAS*, 449, 820
- Wisnioski, E. et al. 2015, *Astrophysical Journal*, 799, 209
- Wolf, C., Meisenheimer, K., Rix, H.-W., Borch, A., Dye, S., & Kleinheinrich, M. 2003, *Astronomy and Astrophysics*, 401, 73
- Wyse, R. F. G., Gilmore, G., & Franx, M. 1997, *Annual Review of Astronomy and Astrophysics*, 35, 637
- Yasuda, N. et al. 2001, *Astronomical Journal*, 122, 1104
- York, D. G. 1982, *Annual Review of Astronomy and Astrophysics*, 20, 221
- York, D. G. et al. 2000, *Astronomical Journal*, 120, 1579