

STATMORPH-LSST: QUANTIFYING AND CORRECTING MORPHOLOGICAL BIASES IN GALAXY SURVEYS

ELIZAVETA SAZONOVA ^{1,2}, CAMERON R. MORGAN ^{1,2}, MICHAEL BALOGH ^{1,2}, MATÍAS BLAÑA ³, CARLOS G. BORNANCINI ^{4,5}, DARKO DONEVSKI ^{6,7}, ALISTER GRAHAM ⁸, HECTOR M. HERNANDEZ TOLEDO ⁹, BENNE W. HOLWERDA ¹⁰, JEYHAN S. KARTALTEPE ¹¹, GARRETH MARTIN ¹², WILLIAM J. PEARSON ⁶, ROSSELLA RAGUSA ¹³, VICENTE RODRIGUEZ-GOMEZ ¹⁴, MICHAEL J. RUTKOWSKI ¹⁵, JOSE ANTONIO VÁZQUEZ-MATA ⁹, AND ROGIER A. WINDHORST, ¹⁶

¹Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, ON, N2L 3G1 Canada

²Department of Physics and Astronomy, University of Waterloo, Waterloo, ON N2L 3G1, Canada

³Vicerrectoría de Investigación y Postgrado, Universidad de La Serena, La Serena 1700000, Chile

⁴Instituto de Astronomía Teórica y Experimental, (IATE, CONICET-UNC), Córdoba, Argentina

⁵Universidad Nacional de Córdoba, Observatorio Astronómico de Córdoba, Laprida 854, X5000BGR, Córdoba, Argentina

⁶National Center for Nuclear Research, Pasteura 7, 02-093 Warsaw, Poland

⁷SISSA, Via Bonomea 265, 34136 Trieste, Italy

⁸Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

⁹Universidad Nacional Autónoma de México, Instituto de Astronomía, A.P. 70-264, 04510 Ciudad de México, México

¹⁰University of Louisville, Department of Physics and Astronomy, 102 Natural Science Building, 40292 KY Louisville, USA

¹¹Laboratory for Multiwavelength Astrophysics, School of Physics and Astronomy, Rochester Institute of Technology, 84 Lomb Memorial Drive, Rochester, NY 14623, USA

¹²School of Physics and Astronomy, University of Nottingham, University Park, Nottingham NG7 2RD, UK

¹³INAF – Osservatorio Astronomico di Capodimonte, Salita Moiariello 16, 80131 Napoli, Italy

¹⁴Instituto de Radioastronomía y Astrofísica, Universidad Nacional Autónoma de México, A.P. 72-3, 58089 Morelia, México

¹⁵Minnesota State University, Mankato, Department of Physics and Astronomy, 141 Trafton Science Center N, Mankato, MN 56001, USA and

¹⁶School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287-6004, USA

Version November 14, 2025

ABSTRACT

Quantitative morphology provides a key probe of galaxy evolution across cosmic time and environments. However, these metrics can be biased by changes in imaging quality – resolution and depth – either across the survey area or the sample. To prepare for the upcoming Rubin LSST data, we investigate this bias for *all* metrics measured by **statmorph** and single-component Sérsic fitting with GALFIT. We find that geometrical measurements (ellipticity, axis ratio, Petrosian radius, and effective radius) are fairly robust at most depths and resolutions. Light concentration measurements (C , Gini, M_{20}) systematically decrease with resolution, leading low-mass or high-redshift bulge-dominated sources to appear indistinguishable from disks. Sérsic index n , while unbiased, suffers from a 20-40% uncertainty due to degeneracies in the Sérsic fit. Disturbance measurements (A , A_S , D) depend on signal-to-noise and are thus affected by noise and surface brightness dimming. We quantify this dependence for each parameter, offer empirical correction functions, and show that the evolution in C observed in JWST galaxies can be explained purely by observational biases. We propose two new measurements – isophotal asymmetry A_X and substructure St – that aim to resolve some of these biases. Finally, we provide a Python package **statmorph-lsst**^a implementing these changes and a full dataset that enables tests of custom functions^b.

Subject headings: Galaxy structure (622) — Galaxy morphology (582) — astronomy image processing (2306)

1. INTRODUCTION

The evolution of galaxies is driven by two main channels. The first is the evolution of the galaxies’ stellar population and gas content, driven by star formation in clouds of molecular gas. The second is the evolution of galactic structure, shaped both by star formation altering the potential well and gravitational interactions that can

redistribute the stars after they formed. In the local Universe, the two are closely linked: star-forming galaxies have primarily disk morphologies while most quiescent galaxies are either elliptical or lenticular with a prominent central bulge (e.g., [Strateva et al. 2001](#); [Schawinski et al. 2014](#)). However, the reason the two are linked is still debated: what physical processes tie star formation and galactic structure together so closely, and at what point in the evolution of the Universe do they become correlated?

One way to answer these questions is to determine

liza.sazonova@uwaterloo.ca

^a **statmorph-lsst** GitHub repository

^b Dataset DOI: [10.5281/zenodo.17585608](https://doi.org/10.5281/zenodo.17585608)

when and where the color-morphology bimodality is first established: by looking at the changes in structure in different environments (e.g., Dressler 1980; Holden et al. 2007) and across redshifts (e.g., Mortlock et al. 2013; Kartaltepe et al. 2023). To carry out these studies, we must robustly quantify the galaxy light distribution.

This is traditionally done using either a parametric approach, such as fitting a Sérsic profile to the light distribution (Sérsic 1963; Graham & Driver 2005), or a suite of non-parametric measurements: concentration, asymmetry, smoothness (CAS; Conselice 2003), Gini- M_{20} (Lotz et al. 2004), and others. These metrics are capable of quantifying the bulge strength¹ (e.g., Bruce et al. 2014), the degree of disturbance (e.g., Lotz et al. 2004), and more broadly the impact of a diverse range of internal and external processes, such as ram pressure stripping (e.g., Roberts & Parker 2020; Bellhouse et al. 2022), star formation knots (Ferreira et al. 2022a), dust lanes (Rutkowski et al. 2013; Sazonova et al. 2021), gas outflows (Deg et al. 2023), and even the cumulative baryonic feedback (Watkins et al. 2025; Martin et al. 2025). Since each parameter can be used for a wide variety of studies, these quantitative measurements are an essential tool in studying the evolution of galactic structure.

In the last few decades, the data from the Hubble Space Telescope (HST) and, more recently, the James Webb Space Telescope (JWST), allowed us to finally probe these topics across cosmic time – leading to more heated debates. Several recent studies found that high-redshift galaxies, in general, are more disk-like than bulge-like (e.g., Ferreira et al. 2022b, 2023; Kartaltepe et al. 2023; Morishita et al. 2024). On the other hand, a new population of massive compact galaxies emerged first at $z > 1$ and then a different population at $z \sim 5$, both non-existent in the local Universe (e.g., Poggianti et al. 2013; Matthee et al. 2024). However, quantitative morphological measurements make it difficult to distinguish between disks and spheroids in the early Universe, considering some high-redshift sources are essentially unresolved (e.g., Whalen et al. 2025). Finally, some morphological analyses point at the relative dearth of mergers at the highest redshifts – first at $z \sim 3$ with HST (Mortlock et al. 2013), then at $z \sim 5$ and even up to $z \sim 10$ with JWST (Ferreira et al. 2023; Ono et al. 2025). By nature of research, we are pushing our available data to the most distant, most compact, lowest mass galaxies, probing the extremes of our ability to image and resolve

galaxies themselves or their faint outskirts. While these studies demonstrate the potential of morphology metrics in helping us to understand the structure of galaxies in the early Universe, they also have to be considered with caution.

Most quantitative morphology metrics have been shown to be dependent in some way on the properties of the image, such as resolution and depth (e.g., Graham & Driver 2005; Holwerda et al. 2011; Bottrell et al. 2019). For example, Thorp et al. (2021) showed that the asymmetry parameter is sensitive to depth, and so it is difficult to detect mergers in galaxies with low signal-to-noise. These are especially important considerations in the era of large surveys, such as Vera Rubin Legacy Survey of Space and Time (LSST), and deep observations, such as JWST. Today we are probing the most extreme objects – the faintest, smallest, and the highest-redshift ones – pushing our metrics to their limit as galaxies span smaller angular sizes and suffer from cosmological surface brightness dimming, effectively making us unable to detect and measure faint features.

Moreover, quantitative morphology metrics are sensitive not only to the image itself, but also the implementation of the metric – the definition of the centre, the segmentation map, and the algorithm itself. For example, galaxies observed as part of SDSS survey have different asymmetry values in Pawlik et al. (2016) and Sazonova et al. (2021) catalogs, which use PawlikMorph and **statmorph** (Rodríguez-Gómez et al. 2019), respectively, so it is important to test the performance of individual codes as well as metrics. Since **statmorph** is the most commonly used code to measure non-parametric morphology that is free and open source, we decided to test the performance of all morphological measurements available in **statmorph** against changes in resolution and depth.

It is therefore crucial to come up with a common set of morphological metrics for all users, which 1) have a consistent implementation, 2) have a well-understood relationship with image quality, and 3) are publicly available and easy to implement. With such a set of metrics, catalogs can be compared across different instruments, surveys, or epochs of a single survey; and we can robustly study the evolution of galaxy structure across cosmic time without simply observing how redshift biases the metrics we chose.

This work is the first step in this task. In this paper, we systematically quantify how morphology metrics measured with **statmorph** and GALFIT depend on image resolution and depth. We explain in detail how **statmorph** metrics are calculated and the important caveats and pitfalls in their algorithm, and provide empirical corrections that can mitigate the image biases.

To do this, we used a dataset of 190 nearby galax-

¹ We use the term “bulge strength” to refer to any measure of how centrally concentrated the light is, rather than specifically the prominence of a physical bulge component or a Sérsic index. Similarly we use the term “early-type” to mean a galaxy with a centrally concentrated light profile regardless of its colour or the presence of a true separate bulge component, and “late-type” a galaxy with relatively more flux away from the center. We do not attempt to classify galaxies kinematically using these projected structural quantities.

ies observed by HST in the optical *I*-band filter, which provides an exquisite resolution of under 25 parsec/px at their distances. We chose objects with sufficiently deep imaging, reaching the *I*-band surface brightness limit $\mu_0 = 24$ mag/arcsec². We then degraded these images to varying resolutions and surface brightness limits to produce 60,000 augmented images, and evaluated the performance of the morphology metrics provided in *statmorph*.

The paper is organized as follows. Our dataset and the data processing steps we took prior to evaluating morphology are described in Section 2. Section 3 makes up the bulk of the paper: we define each morphology metric provided by *statmorph*, discuss its implementation details, show its dependence on resolution and depth, and for some provide an empirical correction term. In Section 4, we discuss our results and potential implications for studies of galaxy morphology across cosmic time, showing how some results can be seen as simply caused by a measurement bias. We summarize our results in Section 5. Throughout this work, we use Planck 2018 Λ CDM cosmology (Planck Collaboration et al. 2020) and AB magnitudes.

For convenience, Table 1 provides a look-up table listing each parameter studied in this work and the bias of this parameter across our sample, as well as in low-resolution and low-depth regimes. We encourage the readers to use this table to quickly jump to the parameters they are most interested in. A visual summary of our results is in Fig. 15 at the end of this paper.

2. DATA PROCESSING

2.1. Sample selection

The goal of this work is to test how the measured structure of galaxies depends on the image properties. One common approach is to use simulated galaxies (e.g., Thorp et al. 2021; Sazonova et al. 2024), which allows creating noiseless images to robustly test the impact of the sky background. However, simulated galaxies with large enough samples are produced either in cosmological simulations (e.g., Illustris TNG; Nelson et al. 2019) or by modelling simple analytic profiles (e.g., GalSim; Rowe et al. 2015). The first approach has a limited resolution defined by the softening length, while the second approach models only the large-scale light distribution, and so these simulations do not adequately resolve small-scale features present in a galaxy (e.g., star clusters, dust, foreground stars, and many more), which may affect the measurements. Instead, we opted to use real observations, similarly to Yu et al. (2023). Space-based imaging provides exquisite resolution of local galaxies, so we turned to archival HST observations of nearby systems. The depth of HST imaging is highly heterogeneous, so we built a sample of local galaxies with deep HST imag-

ing, thereby testing both a wide range of resolutions and depths.

Our analysis applies to changes in resolution and depth both across different surveys or instruments, or equivalently across different redshifts due to cosmological dimming and angular scale change. Traditionally, morphological studies focus on optical morphology, so we chose to restrict this study to the *HST* F814W (*I*) band. We do not apply, nor discuss, morphological K-corrections (e.g., Windhorst et al. 2002) here – but we stress the importance of choosing filters that map to the same rest-frame band when comparing morphologies of galaxies at different redshifts. While there is great potential in extending morphological studies to other wavelengths, it is beyond the scope of the current work.

The Revised New General Catalog catalog (RNGC/IC; Dreyer 1888, 1910; Sinnott 1988; Steinicke 2010) provides a list of large nearby galaxies and nebulae found within 200 Mpc. We cross-matched the RNGC/IC catalog with the Hubble Source Catalog V3 (HSCv3; Whitmore et al. 2016) to see which of these objects have been observed by the HST². We then queried the Mikulski Archive for Space Telescopes³ (MAST) for these observations, and further restricted our sample to galaxies with imaging deeper than the surface brightness limit $\mu_0 = 24$ mag/arcsec²⁴, which resulted in 169 sources. We downloaded the HST observations from the Hubble Legacy Archive⁵ (HLA).

We then visually inspected the images. We removed any sources that were clearly on the edge of the footprint and where the majority of the galaxy was not imaged, galaxies where the chip gap of the ACS detector passed through the nucleus of the source, as it makes any morphological measurements less reliable, and galaxies with severe imaging artifacts: invalid header information (25 galaxies). This left us with 146 galaxies from RNGC/IC. We then looked at each HST footprint to see if there are any other local galaxies in the image that are not included in the RNGC/IC, and cross-matched the positions of the candidate sources with the NED database⁶. We added all sources with known spectroscopic redshifts that are within 200 Mpc to our catalog (41 source, primarily from the Coma and Perseus clusters). These additional objects are almost all dwarf galaxies and satellites that

² We note that HSCv3 is not perfect and some of the galaxies from RNGC/IC did not have an HSCv3 cross-match, even though they have been observed by HST. We have included a couple of objects we knew about that had HST observations, but there may be more

³ MAST archive: <https://archive.stsci.edu/>

⁴ We define the depth or surface brightness limit, μ_0 , as a surface brightness of the 1σ level of the background. While it is common to define the depth as a $n\sigma$ detection of a point source for surveys, since morphological analyses aim to characterize extended emission, the surface brightness depth is more appropriate.

⁵ Hubble Legacy Archive: <https://hla.stsci.edu/>

⁶ NED Database: <https://ned.ipac.caltech.edu/>

TABLE 1: Table of Section 3 contents

p	Parameter Name	Sec.	x	y	Error [all data]	Error [low-res]	Error [low-SNR]
Geometric measurements (Sec. 3.1)							
^a $\mathbf{x}_0^A$	Centroid	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$4.40^{+12.8}_{-3.53}$	$16.6^{+11.1}_{-4.47}$	$4.34^{+12.3}_{-3.32}$
e	Ellipticity	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.02^{+0.09}_{-0.12}$	$-0.13^{+0.18}_{-0.15}$	$-0.02^{+0.14}_{-0.15}$
^a θ	Orientation	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.30^{+16.2}_{-14.6}$	$1.20^{+29.6}_{-22.3}$	$0.65^{+30.1}_{-25.2}$
^a $\theta_{e>0.3}$	Orientation ($e > 0.3$)	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.08^{+8.44}_{-6.17}$	$0.63^{+27.7}_{-20.5}$	$0.53^{+21.3}_{-13.8}$
^a θ^{Sersic}	Sérsic orientation	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.02^{+4.70}_{-4.82}$	$-0.21^{+25.0}_{-18}$	$-0.02^{+8.71}_{-7.09}$
^a $\theta^{\text{Sersic}}_{e>0.3}$	Sérsic orientation ($e > 0.3$)	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.04^{+1.98}_{-3.07}$	$-0.50^{+10.8}_{-11.5}$	$-0.07^{+4.00}_{-5.74}$
e^{Sersic}	Sérsic ellipticity	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.00^{+0.06}_{-0.04}$	$0.05^{+0.15}_{-0.10}$	$0.00^{+0.07}_{-0.05}$
Radius measurements (Sec. 3.2)							
$R_{p,o}^b$	Petrosian radius (circular)	3.2.1	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$0.08^{+1.07}_{-0.86}$	$0.70^{+1.40}_{-1.29}$	$-0.12^{+1.08}_{-1.68}$
$R_{p,e}$	Petrosian radius (elliptical)	3.2.1	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$0.01^{+1.23}_{-1.36}$	$0.47^{+1.81}_{-1.77}$	$0.01^{+1.23}_{-1.36}$
R_{20}	20% light radius	3.2.2	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$0.10^{+0.49}_{-0.14}$	$0.51^{+0.43}_{-0.28}$	$0.05^{+0.48}_{-0.19}$
R_{50}	Half-light radius	3.2.2	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$0.09^{+0.70}_{-0.24}$	$0.59^{+0.75}_{-0.48}$	$0.03^{+0.68}_{-0.48}$
R_{80}	80% light radius	3.2.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.06^{+0.82}_{-0.48}$	$1.07^{+1.42}_{-0.60}$	$-0.02^{+0.82}_{-1.00}$
R_{max}	Maximum radius	3.2.2	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$-0.75^{+3.81}_{-5.12}$	$-0.47^{+5.19}_{-5.71}$	$-4.09^{+2.60}_{-8.44}$
$R_{0.5}^{\text{Sersic}}$	Sérsic half-light radius	3.2.3	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$-0.04^{+0.71}_{-1.13}$	$-0.18^{+0.95}_{-1.32}$	$0.04^{+1.48}_{-1.06}$
Bulge strength measurements (Sec. 3.3)							
n	Sérsic index	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.07^{+0.54}_{-1.09}$	$-0.81^{+1.30}_{-2.13}$	$0.01^{+0.87}_{-0.86}$
C	Concentration	3.3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.22^{+0.29}_{-0.79}$	$-1.14^{+0.62}_{-0.56}$	$-0.22^{+0.32}_{-0.72}$
G	Gini coefficient	3.3.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.03^{+0.04}_{-0.07}$	$-0.08^{+0.08}_{-0.05}$	$-0.05^{+0.06}_{-0.07}$
M_{20}	Moment of the brightest 20%	3.3.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.15^{+0.49}_{-0.22}$	$0.63^{+0.37}_{-0.42}$	$0.20^{+0.52}_{-0.27}$
$B(G, M_{20})$	$G - M_{20}$ bulge strength	3.3.4	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.25^{+0.30}_{-0.55}$	$-0.82^{+0.50}_{-0.37}$	$-0.39^{+0.42}_{-0.60}$
Disturbance measurements (Sec. 3.4)							
^c $S(G, M_{20})$	$G - M_{20}$ disturbance	3.4.1	$\mathcal{R}_{\text{eff}}$	μ_0	$0.01^{+0.47}_{-0.52}$	$-0.01^{+1.02}_{-0.59}$	$-0.01^{+0.46}_{-0.64}$
A_{RMS}	RMS asymmetry	3.4.5	$\mathcal{R}_{\text{eff}}^a$	$\langle \text{SNR} \rangle$	$-0.01^{+0.83}_{-1.46}$	$0.15^{+1.05}_{-1.67}$	$-0.11^{+2.24}_{-2.19}$
A_{CAS}	CAS asymmetry	3.4.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.06^{+0.10}_{-0.21}$	$-0.03^{+0.08}_{-0.20}$	$-0.21^{+0.12}_{-0.18}$
^c A_S	Shape asymmetry	3.4.4	$\mathcal{R}_{\text{eff}}$	μ_0	$0.03^{+0.20}_{-0.14}$	$0.05^{+0.23}_{-0.17}$	$0.12^{+0.20}_{-0.18}$
A_X	Isophotal asymmetry	3.4.5	$\mathcal{R}_{\text{eff}}$	μ_0	—	—	—
A_o	Outer asymmetry	3.4.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.05^{+0.13}_{-0.20}$	$-0.02^{+0.12}_{-0.18}$	$-0.18^{+0.11}_{-0.23}$
Other measurements (Sec. 3.5)							
^c M	Multimode index	3.5.3	$\mathcal{R}_{\text{eff}}$	μ_0	$-0.01^{+0.44}_{-0.30}$	$-0.05^{+0.26}_{-0.16}$	$0.12^{+1.27}_{-0.38}$
I	Intensity index	3.5.4	$\mathcal{R}_{\text{eff}}$	μ_0	$-0.00^{+0.08}_{-0.14}$	$0.06^{+0.23}_{-0.19}$	$-0.00^{+0.10}_{-0.17}$
^c D	Deviation index	3.5.5	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.17^{+0.81}_{-0.41}$	$0.70^{+0.92}_{-0.58}$	$0.30^{+1.06}_{-0.77}$
S	Smoothness index	3.5.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$-0.00^{+0.03}_{-0.13}$	$-0.02^{+0.17}_{-0.35}$	$-0.06^{+0.05}_{-0.45}$
St	Substructure index	3.5.2	$\mathcal{R}_{\text{eff}}$	μ_0	$-0.01^{+0.04}_{-0.11}$	$-0.01^{+0.03}_{-0.07}$	$-0.05^{+0.05}_{-0.13}$

Notes – for each parameter, x and y are the most predictive resolution and signal-to-noise metrics found using mutual information regression (see Sec. 3). The error is the average bias from the baseline measurement obtained approximately at 100 pc/px depth and 23.5 mag/arcsec² surface brightness limit (see Sec. 3 for details). The final two columns give the bias in low-resolution (fewer than 5 elements) and low-SNR (fewer than 2 per pixel) subsamples.

^a Here we quantify the offset from the baseline in a random direction, so a correction is not possible.

^b We define the Petrosian radius in Sec. 3.2.1 setting $\eta = 0.2$.

^c The error quantiles are multiplied by 10.

TABLE 2: An excerpt of the local galaxy sample

Name	R.A. deg	Dec. deg	Dist. Mpc	$\mathcal{R}^a$ pc/px	R_p^b "	Detector
NGC 17	2.7775	-12.1078	83	20	12	ACS
NGC 59	3.8554	-21.4447	5	1	49	ACS
NGC 201	9.8954	0.8603	62	15	50	ACS
NGC 406	16.8517	-69.8761	21	5	73	ACS
NGC 384	16.8542	32.2928	60	14	7	WFC3
IC 1623 A	16.9450	-17.5067	85	21	15	ACS
NGC 404	17.3621	35.7183	4	1	76	ACS
NGC 449	19.0312	33.0889	67	16	15	WFC3
NGC 520	21.1446	3.7942	32	8	75	ACS
NGC 695	27.8092	22.5828	137	32	13	ACS
IC 1776	31.3129	6.1056	48	12	34	WFC3
NGC 1084	41.4992	-7.5778	20	5	76	ACS
NGC 1222	47.2371	-2.9556	34	7	22	WFC3
NGC 1260	49.3633	41.4053	81	15	23	WFC3
NGC 1289	49.7075	-1.9731	40	10	16	WFC3
NGC 1271	49.7967	41.3536	81	15	9	WFC3
NGC 1272	49.8387	41.4906	54	13	54	ACS
NGC 1281	50.0262	41.6297	61	11	20	WFC3
NGC 1311	50.0279	-52.1872	8	2	68	ACS
NGC 1344	52.0812	-31.0681	16	4	68	ACS
NGC 1345	52.3813	-17.7792	22	5	29	ACS
IC 1959	53.3004	-50.4122	9	2	64	ACS
IC 391	74.3421	78.1897	22	6	34	ACS
LEDA 16571	75.3973	-4.2571	446	90	4	ACS
NGC 1741 B	75.3975	-4.2633	58	15	21	ACS
NGC 1741	75.4091	-4.2596	61	14	21	ACS
IC 399	75.4333	-4.2883	56	14	10	ACS
NGC 2082	85.4621	-64.3014	17	4	49	ACS
...						

Notes – the full table of 187 galaxies included in this analysis is available in the online supplementary materials.

^a Angular resolution of the image in units of parsecs per pixel. The instrumental resolution is $0.05''/\text{px}$ for ACS and $0.04''/\text{px}$ for WFC3.

^b Petrosian radius (Sec. 3.2.1; Petrosian 1976)

were too faint to be included in the original RNGC/IC catalog, and allow us to study a wider range of possible morphologies, including satellite galaxies potentially undergoing gas stripping. Finally, this resulted in a catalog of 187 galaxies at distances smaller than 200 Mpc, imaged with HST with a depth of 24 mag/arcsec² or better. The sample spans a wide range of stellar masses, morphological types, and environments. We queried NED to get any missing *r*-band Petrosian radii R_p (see Sec. 3.2.1 for a definition), and spectroscopic distances. Our sample is summarized in Table 2. Appendix A shows the distributions of magnitudes and morphologies of the galaxies in our sample, confirming that we have a diverse set of galaxies spanning the entire Hubble sequence.

2.2. Image processing

For each galaxy, we created cutouts and ran the analysis semi-interactively. By default, we made cutouts spanning $3R_p$, increasing this size for some objects where the diffuse light extended close to the edges of the image. Since this is a local sample imaged by HST, in many cases the field of view was smaller than $3R_p$, so we used the entire footprint instead. We used the context information provided by the HLA pipeline to create masks for bad pixels and detector gaps, and the exposure time map alongside gain to get electron counts and hence the 1σ Poisson uncertainty map.

We estimated the background level of the images with PHOTUTILS (Bradley et al. 2020). However, many of our sources have a large angular extent comparable to the footprint of the detector, and so the image is dominated in some cases by the galaxy’s light, making a simple sigma-clipping procedure unreliable. Instead, we estimate the background iteratively. First, we estimate the background σ with sigma-clipping. Then, we run a 1σ source detection on the image and mask out any detected sources and re-run the sigma-clipping to get a new estimate for σ . We repeat this until the standard deviation estimates converge to within 5% (typically 4-5 times). Appendix B shows an example of this procedure on one galaxy. Finally, we use the **Background2D** with a box size spanning 1% of the image and a 5 pixel filter size to see if there are any large-scale gradients in the background. Typically, subtracting a single median background value is desirable in low-surface-brightness science to retain as much of faint diffuse light as possible. However, some HST images have variations in gain across the mosaic and hence effective depth across the detector. We visually inspect the **Background2D** fit and if we find background discontinuities across mosaicked chips we subtract the fit; otherwise we simply subtract a flat background value. We define the surface brightness of the image μ using 1σ of the background obtained with this procedure.

HLA does not provide variance maps for most images, so we constructed our own. The Poisson error of a raw image is a square root of the electron counts. To convert an image to counts we use the exposure time map, t_{exp} , gain g , and our background estimate via $f_{\text{sky}} \sim (gt_{\text{exp}} \times \sigma)^2$. We convert the image into electrons to calculate the Poisson error, and re-normalize the error array by the exposure time. We also save the exposure time map and the sky map to use in later steps.

We then created segmentation maps to mask out other foreground/background sources. Since our sources are large, there are more likely to be other sources overlapping with their light, and so an accurate segmentation map is crucial. We used a hot-cold approach described in Sazonova et al. (2021), similar to Galametz et al. (2013).




FIG. 1.—: A subset of augmentations performed for an example galaxy, NGC 17, observed in F814W (I band). The 1σ surface brightness limit μ_0 decreasing left to right, and the resolution $\mathcal{R}$ is degrading top to bottom. Foreground sources detected on each image are shown in white contours. As depth degrades, the large tidal tail to the southwest is lost in noise, while the internal disturbances are invisible in low-resolution imaging.

Sazonova et al. (2021) and Morgan et al. (2024) described the algorithm in more detail, but we summarize it in Appendix C.

Creating good-quality segmentation maps is particularly challenging for high-resolution observations of local galaxies, since a bright clump may be a foreground star or a star cluster within the galaxy itself. We aimed to mask any obvious stars and smaller background galaxies. To ensure a reliable source mask, we ran the segmentation procedure interactively, visually inspecting the resulting masks and changing the segmentation parameters when necessary until a good map was produced, decided by our visual examination and looking that 1) obvious stars are masked, 2) the source is not over-deblended, and 3) foreground objects away from the source are masked.

The goal of this work is to study how morphological parameters change based on image quality, and the segmentation procedure can dramatically change the quantitative morphology, e.g., if the main source is over-deblended, or if some contaminant is not masked. As the image quality changes, so does our ability to clearly distinguish different sources, and therefore the estimated segmentation maps. If only low-resolution imaging is available, it would be unrealistic to use a segmentation map obtained with higher-resolution data. Therefore, to mimic a real use case most closely, we create new segmentation maps each time the image quality is degraded. In most cases, we keep the same segmentation parameters from the original run, although in a few cases we change them to help deblending at the lowest resolutions. In cases of merging galaxies, the two sources are deblended in shallow imaging but not in deeper observations where

a tidal feature connects them. This is an intrinsic complication of trying to separate associated sources, and we did not adjust the algorithm to produce the same deblending results (see Sec. 3.5.3 for an example). We note that this adds an additional source of uncertainty to morphological measurements, since at different image depths what is defined as the “source” may change.

2.3. Image quality augmentations

Finally, for each galaxy, we create a series of images with a degraded image quality, changing both the resolution and depth – an example is shown in Fig. 1 for a subset of augmentations for NGC 17.

Resolution is defined both by two effects: the effective size of the point-spread function (PSF) and the physical pixel scale, and both of these can affect the measurements independently in different ways. However, most instruments set the pixel scale to sample the PSF by approximately two pixels for the PSF full width at half maximum (FWHM) to ensure Nyquist sampling. The exact ratio between the pixel scale and the PSF size varies from 1 (e.g., JWST) to 3~4 in surveys such as Sloan Digital Sky Survey (SDSS), LSST, or the Ultraviolet Near Infrared Optical Northern Survey (UNIONS). We chose to fix the PSF size to be twice the pixel scale ensuring Nyquist sampling (which is similar to the HST resolution). Therefore, when we say *resolution*, we interchangeably mean both the number of pixels, and the number of resolution elements set by the PSF. Investigating the relationship between the pixel size and the PSF size *independently* would be valuable, but is beyond the scope of this paper. We are planning to release an additional set of augmentations using an empirical LSST

PSF in the GitHub repository.

To change the image resolution, we first choose the new pixel size in parsecs per pixel ranging from 25 pc/px to 2,000 pc/px. Then, we convolve the original image with a Moffat PSF with a $\beta = 2$ and a FWHM that would correspond to two pixels at the new resolution. The convolution step smooths over the noise, which is unrealistic, as in real instruments the seeing is applied by the atmosphere or the telescope before the photons hit the detector with some Poissonian probability. Therefore, after the convolution, we resample each pixel from a Poisson distribution about that pixel’s value to mimic this effect. This degrades the raw signal-to-noise of the image, but the images are still deep enough for the rest of our analysis. Finally, we use SCIKIT-IMAGE (van der Walt et al. 2014) to resample the image from its original resolution to the new one. We also resample the mask, and calculate the new error array from the counts in the resampled image. Finally, we convolve the instrumental *HST* PSF simulated with TINYTIM (Krist et al. 2011) with the Moffat PSF to get the effective PSF of the new image, and store it for later analysis.

After the image is resampled to the new resolution, we calculate the new 1σ surface brightness limit (μ_0), since downscaling the image effectively increases the image depth, again using the iterative masking and sigma-clipping to calculate σ . Then we choose a new surface brightness limit level, in an evenly spaced grid ranging from 20 to 25.5 mag/arcsec². To decrease signal-to-noise, we add a Gaussian white noise field with 0 mean and a standard deviation σ' to match the desired depth. This effectively emulates a brighter sky with a higher Poisson error, rather than a lower exposure time. In the optical regime we are interested in, the impact of these two approaches is essentially the same, since the Poisson noise from the source is not a significant fraction of the total error. However, in UV or X-ray observations, there would be an important difference between larger exposure time or a brighter background, and so our results should be interpreted with caution in those regimes. Given the original background standard deviation σ^{og} , the original depth μ_{σ}^{og} , and the desired depth $\mu_{\sigma}^{\text{new}}$, the amount of additional noise needed is given by:

$$(\sigma')^2 = \sigma^2 \left(10^{2(\mu_{\sigma}^{\text{new}} - \mu_{\sigma}^{\text{og}})/2.5} - 1 \right) \quad (1)$$

we add this Gaussian noise to the rescaled image and the error array, re-run the segmentation routine.

For each galaxy, we attempt to create images for 18 physical resolutions from 25 to 2500 pc/px, and for 25 surface brightness limits from 20 to 26 mag/arcsec². If the image is too shallow, we limit this to the deepest surface brightness limit possible. If the image becomes smaller than 40 pixels, we stop degrading the resolution.

As a result, we produced a dataset of 64,000 images for 189 galaxies spanning a wide range of resolution and image depths. This dataset is available on Zenodo⁷, and can be well-suited for other analyses or for training deep learning models.

We then measure the quantitative morphologies described in Sec. 3 using *statmorph* (Rodríguez-Gomez et al. 2019) by passing the image, the error array as a weight map, the new PSF, and the segmentation map. In addition to the *statmorph* analysis, we calculate a new isophotal asymmetry (A_X ; Sec. 3.4.5), substructure (St ; Sec. 3.5.2), and run GALFIT (Sec. 3.2.3; Peng et al. 2002) to obtain single-Sérsic fits. Finally, we compute an effective signal-to-noise per pixel in the augmented image.

3. ANALYSIS OF MORPHOLOGY METRICS

In this section, we discuss the morphological parameters measured in our test suite by *statmorph*, GALFIT, and our custom algorithms. Table 1 lists all the parameters we studied in this work, and Fig. 15 provides a visual summary of our results – we encourage the readers to use these to find the parameters they are most interested in.

We measure each metric p on all augmented images, and compare its value to a tentative “baseline” value measured on images with resolution $\mathcal{R} = 100$ pc/px and surface brightness limit $\mu_0 = 23.5$ mag/arcsec²⁸. We use this “baseline” as a comparison, since it is impossible to measure a “true” value of a morphological parameter. Moreover, for parameters that do not converge and depend intrinsically on resolution and/or depth, the true value may not even be defined. Therefore, we opted to define a “baseline” which is the measurement you get when the source is reasonably well-resolved and faint features are visible.

For each metric, we consider that its measured value can depend on three variables: the baseline value p_{base} , resolution, and signal-to-noise. We do not consider that some metric may depend on a different metric (such as a centroid) for simplicity. The resolution dependence might be two-fold: a metric might depend on our ability to resolve some specific features with a given size (hence on $\mathcal{R}$ in pc/px) or on our ability to resolve some feature relative to a galaxy’s size. To distinguish these two cases, we also measure the effective resolution, $\mathcal{R}_{\text{eff}}$, defined as Petrosian radius (Sec. 3.2.1) divided by the pixel scale, or the number of resolution elements sampling the

⁷ Dataset DOI: [10.5281/zenodo.17585608](https://doi.org/10.5281/zenodo.17585608)

²⁸ We define a surface brightness limit μ_0 as the effective limit after accounting for cosmological dimming, $\mu_0 \equiv \mu - 7.5 \log(1+z)$, where μ is the 1σ surface brightness limit measured directly from the image. The factor of 7.5 arises from cosmological dimming and an AB magnitude system K-correction as described in Hogg et al. (2002). Since our sources are within 200 Mpc, μ and μ_0 are almost interchangeable.

galaxy⁹. Similarly, while some metrics depend on detecting low surface brightness features and hence μ_0 , others may depend more on the total flux and hence a signal-to-noise ratio, so we also measure a signal-to-noise per pixel (SNR_{px}) within $1.5 \times \text{Petrosian radius}$ ¹⁰. So in our modelling approach, p may depend on all of these factors, i.e. $p = p(p_{\text{base}}, \mu_0, \text{SNR}_{\text{px}}, \mathcal{R}, \mathcal{R}_{\text{eff}})$, but not on any other morphological parameters.

In practice, since $\mathcal{R}$ and $\mathcal{R}_{\text{eff}}$ as well as μ_0 and SNR_{px} are degenerate with one another, it is more useful to find which factors drive the measurement most. To do this, we use a mutual information score from SCIKIT-LEARN (Pedregosa et al. 2011). A mutual information score compares the joint distribution of two variables to their marginal distributions, and quantifies how much knowing one variable informs the other. The score is 0 if the two variables are independent, and 1 if one completely depends on the other. First, we compute the mutual information scores between p and each of p_{base} , μ_0 , SNR_{px} , $\mathcal{R}$, $\mathcal{R}_{\text{eff}}$. Then, we choose one variable between $(\mathcal{R}, \mathcal{R}_{\text{eff}})$, and one between $(\mu_0, \text{SNR}_{\text{px}})$ with the higher score as the more important one. In most cases, $\mathcal{R}_{\text{eff}}$ and SNR_{px} are more important than $\mathcal{R}$ and μ_0 .

We then choose the final “baseline” value for each galaxy. In cases where $\mathcal{R}_{\text{eff}}$ is more important to the measurements, it would be incorrect to choose the baseline as some arbitrary resolution in parsecs, since the *effective* resolution will change between larger and smaller galaxies at the same pixel scale. This is also true for SNR_{px} . Therefore, if these parameters are preferred, instead we selected the baseline based on $\mathcal{R}_{\text{eff}}$ and SNR_{px} . We chose the baseline of $\mathcal{R}_{\text{eff}}$ to be the median of the entire sample when $\mathcal{R} = 100$ pc/px, which was $\mathcal{R}_{\text{eff}} = 42$. Similarly, our baseline $\text{SNR}_{\text{px}} = 5.3$. We then reassigned the baseline augments for each galaxy.

Finally, for convenience, we provide fits to the function $p(p_{\text{base}}, \mathcal{R}, \mu, \mathcal{R}_{\text{eff}}, \text{SNR}_{\text{px}})$ obtained using the symbolic regression package `SymbolicRegression.jl` (Cranmer 2023). The details of the symbolic regression implementation are described in Appendix D. These convenience functions enable observers to convert the measurement they obtained on a given image to the expected values for an image of a different quality, and they are listed in Tab. D2. We caution that these fits are purely empirical rather than universal relationships. For most measurements, as the image quality decreases the scatter in p increases, so using these equations to find the expected p for a deeper or a better resolved image will introduce significant uncertainties in the estimate, and so should be done with care. Moreover, we assumed a Moffat PSF

with a two-pixel FWHM in our augmentations – for a different PSF, the fits are likely different, and so it is best to recalibrate the correction based on a known PSF. We would also not advise users to extrapolate our models beyond the fitted regime.

For each parameter p , we provide several analysis products. First, a mutual information score for each parameter and each variable, shown in the online materials¹¹. Second, a grid showing the variation in the parameter compared to the baseline as a function of resolution and depth. Third, a `SymbolicRegression` fit to the relationship between p and p_{base} (Tab. D2), and a distribution of corrected parameters. Finally, we also show how each parameter varies with depth and resolution for two example galaxies. Since there are a lot of figures, we only provide the most informative plots in this paper, and the remaining plots are available online. We explain the plots in detail in the following section describing the centroid (Sec. 3.1.1); the other sections adopt the same format.

We group the parameters roughly by their intended purpose: first, we discuss the simplest geometric measurements (image moments and ellipticity; Sec. 3.1), then radii and Sérsic modelling (3.2), then concentration or “bulge strength” measurements (C , G , M_{20} , Sec. 3.3), then asymmetry parameters (Sec. 3.4), and finally other metrics that do not fit well in these categories (smoothness, substructure, multimode-intensity-deviation; Sec. 3.5). We additionally define two new parameters: isophotal asymmetry (Sec. 3.4.5) and substructure (Sec. 3.5.2).

3.1. Moments of light

The simplest approach to quantifying a light distribution is to measure the moments of the image. For a 2D image, the moment is calculated independently in x and y directions, and so the moment of order i, j is

$$M_{ij} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^i y^j I(x, y) dx dy = \sum_x \sum_y x^i y^j I(x, y).$$

The 0th moment is then simply the total flux in the image. The centroid, i.e. a weighted average location of light, is the ratio of the first order moments and the total flux:

$$\mathbf{x}_0 = (x_0, y_0) = \left(\frac{M_{10}}{M_{00}}, \frac{M_{01}}{M_{00}} \right).$$

Finally, the second-order moments can be used to construct a covariance matrix for the light distribution,

$$\text{cov}(I) = \frac{1}{M_{00}} \begin{pmatrix} M_{20} & M_{11} \\ M_{11} & M_{02} \end{pmatrix},$$

which defines a Gaussian distribution with the same moments as the image. The eigenvalues of this covariance

⁹ Note that since the pixel size is always half the PSF FWHM, this is also the effective resolution relative to the PSF times two.

¹⁰ The Petrosian radius used here is measured on the baseline image.

¹¹ <https://github.com/astro-nova/statmorph-lsst>




FIG. 2.—: The measurement error compared to the baseline for the asymmetry centre ($\mathbf{x}_0^A$, **left**), ellipticity (e , **middle**), and orientation (θ , **right**), as a function of the average signal-to-noise per pixel $\langle \text{SNR} \rangle$ and the effective resolution $\mathcal{R}_{\text{eff}}$. For each galaxy and each image, the error is calculated as the difference in measurement compared to the baseline. $\mathbf{x}_0^A$ is normalized by the Petrosian radius (Sec. 3.2.1). The colored distribution shows the median error in each bin, while the gray contours show 1σ scatter. Both $\mathbf{x}_0^A$ and e are fairly robust to noise, with an increased uncertainty at $\langle \text{SNR} \rangle < 1$ but no systematic offsets. They are however biased at low resolution, with typical error in $\mathbf{x}_0^A$ being $0.2 \pm 0.07R_p$, and typical ellipticity error of 0.15 ± 0.12 . θ is less reliable in the low signal-to-noise regime, with an average $20 \pm 35^\circ$ error when $\langle \text{SNR} \rangle \sim 1$, and an over 10° scatter at all resolutions and depths. This error is mostly driven by circular sources, and if only inclined ($e > 0.3$) sources are considered, the error in θ is below 5° . Compared to Fig. 2, the inclusion of the PSF effect significantly improves e and θ measurements.

matrix give the semi-major and semi-minor lengths of the equivalent Gaussian, and so can be used to estimate the axis lengths, ellipticity, and orientation of the source. The size estimate is essentially the standard deviation of a Gaussian and so isn't particularly physically meaningful; however, we can use the axis ratio and the orientation measurements.

3.1.1. Centre

To calculate the covariance matrix, we first need to center the image, i.e. subtract the centroid position from the pixel coordinates. The easiest choice is to use the centroid calculated from the first moments. However, as noted by Rodriguez-Gomez et al. (2019), this often does not correspond to the brightest region of the galaxy, as asymmetries in the light distribution can shift the centroid away from the physical centre. We can get a more robust measurement of the centre using the pixel about which asymmetry is minimized (Sec. 3.4.2). STATMORPH calculates both of these centers, and therefore outputs elongation, ellipticity, and orientation measured about both coordinates, with suffixes `_centroid` and `_asymmetry`. For all other measurements the asymmetry centre is used by default. This centre typically differs from the centre determined during catalog creation, and while most surveys quote their astrometric precision, there are no studies of how well can we measure the centre in morphological pipelines.

The left panel of Fig. 2 shows how the asymmetry center, $\mathbf{x}_0^A$, depends on $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$. We calculated the centroid for each image and then found the distance between the measured and the baseline centers. We expressed the distance as a percentage of the baseline Petrosian radius (Sec. 3.2.1). We binned the images in groups of matched $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$ and calculated the median offset as well as its 16th and 84th percentiles. The colored plot shows a smoothed histogram of the median percentage offset ($\Delta\%$). The gray contours show the scatter in the offset, calculated as half the range between the 84th and 16th percentiles. This scatter can be interpreted as the expected error in an individual measurement. For example, at $\log \langle \text{SNR} \rangle = 0.75$ and $\log \mathcal{R}_{\text{eff}} = 0.5$, $\mathbf{x}_0^A$ is offset on average by $0.12R_p$ with a 6% spread about this value.

$\mathbf{x}_0^A$ is robust to changes in signal-to-noise: we don't find any systematic variations in the centre even in the lowest $\langle \text{SNR} \rangle$ bins, although the expected error increases to $>0.04R_p$ when $\langle \text{SNR} \rangle < 1$. However, the estimate of the centre becomes less reliable as the resolution degrades, with up to $0.20 \pm 0.07R_p$ offsets when there are fewer than 5 pixels spanning R_p . This result is unsurprising: at the lowest resolutions, one pixel spans larger distances, and so subpixel precision is required to estimate the centre well. Small errors in the subpixel centre estimate lead to significant offsets. This is especially true for the asym-

metry center, since minimizing asymmetry with subpixel precision leads to interpolation errors when the image is rotated about this centre (Sazonova et al. 2024).

For brevity, we do not show the results for the moment-based centroid, as they are almost identical, but these results are available in the online materials. The errors as the resolution degrades are comparable to those we found with the asymmetry center, but the centroid also depends on signal-to-noise, with up to $0.05 \pm 0.07 R_p$ offsets at fixed resolution and $\langle \text{SNR} \rangle < 1$. In most cases, the two centre estimates agree well, so it is reasonable to use either one, but the asymmetry-based centre is more robust to extended tidal features (Rodriguez-Gomez et al. 2019).

3.1.2. Ellipticity

Ellipticity quantifies how elongated or circular the source is by finding the ratio of the semi-major and the semi-minor axes (a and b). There are several different definitions of this ratio in literature: the ellipticity, $e = 1 - b/a$, elongation, a/b , or the axis ratio, $q = b/a$. They are perfectly degenerate with one another, so here we only discuss the ellipticity e . Ellipticity of 0 means a perfectly circular source, while ellipticity close to 1 means an elongated, or a cigar-shaped source. **statmorph** calculates ellipticity about the asymmetry centre and the centroid; here we only discuss the measurement relative to the asymmetry centre as the two are consistent in most cases.

The middle panel of Fig. 2 shows how e varies with $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$. e is largely insensitive to signal-to-noise, although the uncertainty in measurement increases to ± 0.12 when $\langle \text{SNR} \rangle < 1$. The offset in e is driven by resolution, with an up to 0.15 ± 0.12 underestimate when there are fewer than 5 pixels spanning the galaxy. This is likely because **statmorph** calculates e from the raw image moments without accounting for the PSF, which circularizes the incoming flux and reduces ellipticity. The issue is even more pronounced in Source Extractor, which first smooths the image with a small kernel, and then computes ellipticity. The discretization can also be an effect in the most extreme cases, where the entire source is spanned by a few pixels. In our set-up, the effect of the PSF is small, since the PSF is only 2 pixels wide; but in ground-based observations where the PSF is large relatively to the pixel scale, the effect may be more pronounced. This dependence on the PSF or pixel size also means that in a given observation, there is a *lower bound* on measured ellipticity, so extremely edge-on galaxies will appear more circularized (e.g., van der Wel et al. 2014).

We provide our best fit to correct for the resolution offset Appendix D, however we must caution that since the error is driven primarily by the PSF, the correction is

only applicable in cases where the PSF FWHM is roughly 2 pixels, and it is better to derive a correction based on your PSF. Alternatively, the best approach is to derive ellipticity from a parametric fit that includes a PSF model, e.g. through a Sérsic fit (Sec. 3.2.3). This is often done for studies of high-redshift galaxies (e.g., Vega-Ferrero et al. 2024; Pandya et al. 2024).

3.1.3. Orientation

The final parameter that uses the image moments is the orientation θ , i.e. the angle of inclination relative to the chosen centre. Orientation is measured starting from the vertical, going counter-clockwise, spanning from $-\pi$ to $+\pi$. In this work, we converted orientation from radians to degrees. Fig. 2 shows the orientation behavior in the right panel. The uncertainty in the orientation angle is over 10° for most resolutions and depths when compared to the baseline, increasing to 20° in low-resolution images and to 35° in noisy images. The error in orientation reaches 20° when $\langle \text{SNR} \rangle \sim 1$ and 5° when $\mathcal{R}_{\text{eff}} \sim 5$. However, most of this error comes from the intrinsic uncertainties in calculating a position error of round sources. When restricting our sample to only galaxies with $e > 0.3$, the error in orientation is $2 \pm 5^\circ$ in all but the lowest $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$ bins.

3.2. Radius measurements

Here we describe three galactic radii computed by **STATMORPH** and **GALFIT**: Petrosian, isophotal, and Sérsic. There are, however, many other ways to measure a galaxy’s size, some more physically motivated than others – such as looking at the break in the exponential disk profile (e.g., Hunter & Elmegreen 2006), a stellar mass cut-off (Chamba et al. 2022), or the H I mass threshold (e.g., Wang et al. 2016). While it is beyond the scope of our work to discuss these other radii here, we intend to analyze them in the future and upload the results to the online repository.

3.2.1. Petrosian radius

The Petrosian (1976) radius R_p is where the mean flux in an annulus centered on R_p is a fraction η of the mean flux up to R_p :

$$\frac{1}{4\pi R_p \delta r} \int_{R_p - \delta r}^{R_p + \delta r} I(r) dr = \eta \times \frac{1}{\pi R_p^2} \int_0^{R_p} I(r) dr, \quad (2)$$

where $I(r)$ is galaxy flux and δr is some small annulus width (1 pixel in **statmorph**). Most often, η is chosen to be 0.2¹². **STATMORPH** calculates R_p in circular and elliptical isophotes ($R_{p,o}$ and $R_{p,e}$), using the asymmetry

¹² This definition of η is actually the inverse of the original definition (Petrosian 1976; Graham & Driver 2005). We adopt the current definition since it is used by **statmorph** and SDSS.




FIG. 3.—: Same as Fig. 2, for three radius metrics: $R_{p,o}$ (left), R_{20} (middle), and Sérsic $R_{0.5}$ (right, discussed in Sec. 3.2.3). The bias and scatter are measured as a fraction of the baseline, rather than an absolute offset. $R_{p,o}$ and R_{20} are more correlated to the pixel scale $\mathcal{R}$ while $R_{0.5}^{\text{Sersic}}$ depends on the effective resolution $\mathcal{R}_{\text{eff}}$. The non-parametric radii are overestimated at low resolutions, up to 20% and 80% for $R_{p,o}$ and R_{20} respectively, but $R_{0.5}^{\text{Sersic}}$ is slightly underestimated (up to 10%) compared to the baseline. However, the uncertainty in $R_{0.5}^{\text{Sersic}}$ is larger than that for non-parametric radii.

centre (Sec. 3.1.1) and the ellipticity found using image moments (Sec. 3.1). Petrosian radius has been widely used as it is largely insensitive to the signal-to-noise ratio or the surface brightness limit (e.g., Petrosian 1976; Lotz et al. 2004) and so has been historically used in high-redshift science (e.g., Shen et al. 2003; Pérez-González et al. 2023), where the cosmological surface brightness dimming is important.

Fig. 3 shows the dependence of $R_{p,o}$ in kpc on $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$. R_p is a robust measurement in a low signal-to-noise regime, resulting in consistent radius estimates as long as $\langle \text{SNR} \rangle \geq 1$, and an up to $30 \pm 20\%$ underestimate in noisier images. There are only two other studies that investigate the dependence of R_p on signal-to-noise: Lotz et al. (2004) give a conservative estimate that R_p is reliable up to $\langle \text{SNR} \rangle \sim 5$, while Ren et al. (2024) find no dependence of R_p on image depth. Our results agree with Ren et al. (2024) in that on average, R_p is only underestimated in extremely low signal-to-noise images. However, the scatter in low- $\langle \text{SNR} \rangle$ bins reaches 20%, which means that measured R_p on noisy images carries about a 20% uncertainty, and should be treated with caution on an individual galaxy basis.

Similarly to Lotz et al. (2004), Ren et al. (2024) and Yu et al. (2023), we found that R_p is sensitive to resolution. It is stable at resolutions better than 500 pc/px (or, equivalently, 1 kpc-wide PSF). At lower resolutions, $R_{p,o}$ is overestimated by 15% with a 30% scatter. Our results are in complete agreement with Lotz et al. (2004), who also found that R_p is robust at resolutions below 500 pc/px and then begins to increase, albeit using a


much smaller sample of galaxies. Ren et al. (2024) show that the dependence of R_p on resolution arises from the PSF blurring the central flux to larger radii. Yu et al. (2023) also find that the error in R_p is always $0.77 \times \text{PSF FWHM}$. We cannot directly test the effect of the PSF separately from the pixel scale with our set-up. However, in Sec. 3.2.3 we show that the Sérsic effective radius, which is computed by accounting for the PSF, does not suffer from the same resolution effects, so we can indirectly conclude that this error is in fact driven by the PSF smearing.

The relationship between R_p and R_p^{base} we obtained with SymbolicRegression is:

$$R_p = R_p^{\text{base}} \tanh(3.2 \langle \text{SNR} \rangle) + \left(\frac{\mathcal{R}}{1000} \right)^{1.86}. \quad (3)$$

It is then possible to measure R_p and invert this relationship to estimate R_p , which we call “corrected R_p ”. This correction is different to the correction found in Yu et al. (2023): we have a signal-to-noise term, which depends on the baseline R_p and thus affects larger galaxies more, and a resolution term, which is flat with respect to the galaxy’s intrinsic size. However, the resolution term is not linear like in Yu et al. (2023). Fig. 4 shows the distribution of the measured R_p compared to the baseline on the left, and the corrected R_p on the right. The kernel density estimate of the entire dataset is shown in color, and we overplotted the distribution of $\langle \text{SNR} \rangle < 1$ measurements in white and $\mathcal{R} > 500$ pc ones in black. On the left panel, we see the same effect as in Fig. 3: R_p is overestimated at low resolutions. The correction allows

FIG. 4: The distribution of $R_{p,o}$ measurements compared to the baseline before (left) and after (right) a signal-to-noise and resolution correction from Eq. 3. Before the correction, $R_{p,o}$ is overestimated when resolution is worse than 500 pc/px (black contours), and slightly underestimated when $\langle \text{SNR} \rangle$ is below 2 (white contours). After applying the correction, we can recover the baseline $R_{p,o}$ for a wide range of resolutions and depths.



us to match R_p to the baseline value well in both the low- $\langle \text{SNR} \rangle$ and low- $\mathcal{R}_{\text{eff}}$ regimes. Our results contradict Ren et al. (2024) and Yu et al. (2023), who find that the bias in R_p is flat with respect to the intrinsic radius. Ren et al. (2024) find that the error is flat overall, while Yu et al. (2023) derive a correction that depends linearly on resolution. We find that the best correction includes all three of the input parameters, and can accurately recover the baseline R_p even in the lowest signal-to-noise or resolution regimes. However, since the resolution dependence likely arises from the PSF rather than discretization, it is best to define $\mathcal{R}$ as the width of the PSF rather than the pixel scale.

The behaviour we see when measuring R_p in elliptical rather than circular apertures is virtually identical, and we show the plots made for $R_{p,e}$ online. $R_{p,e}$ has the same noise dependence as $R_{p,o}$, unsurprisingly, since both the ellipticity and orientation are robust to noise (Sec. 3.1). In terms of resolution, $R_{p,e}$ is slightly better-behaved, with a $10 \pm 25\%$ error. The `SymbolicRegression` fit is also similar in spirit to the one we obtained for $R_{p,o}$, although the functional form is not exactly the same, highlighting the arbitrary nature of the fit: while we tried to set up the run in a physically motivated manner, there are infinitely many functional combinations that may provide similarly good fits to our data. All of the fit results for each parameter as shown in the Appendix D.

The final important consideration to note is that R_p is used by `statmorph` to define the outer edge of the galaxy and the total flux. This edge is then used downstream, e.g. in isophotal radii or asymmetry calculations. The outer edge is given by the `petro_extent_cas` parameter times R_p , which by default is 1.5. However, Graham et al. (2005) showed that for early-type galaxies, as much as 20% of the total flux is contained outside R_p , and so it is advisable to increase `petro_extent_cas` if the foreground contaminants are well-masked in the image.

3.2.2. R_{50} , R_{20} , and R_{80}

The half-light radius R_{50} , as well as R_{20} , and R_{80} , are isophotal radii measuring regions containing 50%, 20%, and 80% of the galaxy’s light respectively. Isophotes are calculated about the asymmetry centre ($\mathbf{x}_0^A$, Sec. 3.1.1). Calculating a percentile of a galaxy’s light requires knowing the extent of the *total* flux as well, so `STATMORPH` defines the total flux as the light contained within `petro_extent_cas` of $R_{p,o}$. R_{50} is commonly used as a half-light radius, while R_{20} and R_{80} are uncommon on their own and are mainly used to calculate the Concentration parameter (Sec. 3.3.1).

For readability of this paper, we do not show all fit results of these radii or all of the correction functions, which are available online. Here we show the error in R_{20} in the middle panel of Fig. 3 – qualitatively, the other parameters behave the same way but with a weaker bias. Perhaps unsurprisingly, since it depends on R_p , R_{20} is extremely sensitive to resolution. While R_p is overestimated by at most 30%, R_{20} is overestimated by more than 100% when there are fewer than 5 resolution elements. The measurement appears to slowly converge in well-resolved imaging although not entirely; there is an approximately 5% bias between our baseline images and images at higher resolutions. The behavior of R_{50} and R_{80} is qualitatively the same, with a bias reaching 50% and 25% at $\mathcal{R}_{\text{eff}} < 5$ respectively.

The only other study discussing R_{20} is Bershadsky et al. (2000), who also note that smaller percentile isophotes are more sensitive to resolution, but find that even R_{20} is stable at $\mathcal{R}_{\text{eff}} \sim 10$. However, Bershadsky et al. (2000) simply downsampled their images to investigate the resolution dependence, rather than also including the effect of the PSF. The stark difference between our results and those of Bershadsky et al. (2000) strongly indicate that the bias is caused by the PSF, and could be potentially remedied if the effect of the PSF was accounted for (as it is in the Sérsic fit, see below).

The `SymbolicRegression` fit to our data allowed us to correct for the resolution well (the fit parameters are in

Appendix D, and the corrected distribution is available online). However, since the bias appears to stem from the PSF, our fit would only work well in the identical set-ups, where the PSF is a 2-pixel Moffat. Therefore, we discourage using these fits in images with a different PSF without testing.

3.2.3. Sérsic fit & Effective radius

The Sérsic fit is one of the most widely used parametric ways to quantify galaxy structure. Sérsic (1963) generalized the de Vaucouleurs (1948) empirical profile to obtain a good fit to many different galaxies along a Hubble sequence:

$$I(r) = I_{0.5} \exp \left\{ -b_n \left[\left(\frac{r}{R_{0.5}^{\text{Sérsic}}} \right)^{1/n} - 1 \right] \right\} \quad (4)$$

where $I_{0.5}$ is the light intensity at the half-light radius $R_{0.5}^{\text{Sérsic}}$, n is the Sérsic index, and $b_n \approx 2n - 1/3$ and can be calculated exactly via Gamma functions (e.g., Graham & Driver 2005). The Sérsic profile describes an exponential light profile when $n = 1$ and a de Vaucouleurs (1948) profile of bright early-type galaxies when $n = 4$. In practice, in many applications the Sérsic profile is fit to the 2D light distribution, and so the fit has four additional free parameters: the centre (x_0, y_0) of the galaxy, the (constant) axis ratio q or ellipticity e , defined as $q \equiv b/a \equiv 1 - e$, and the orientation or position angle θ . These are the same parameters as those obtained from the moments (Sec. 3.1), except in this case a galaxy is modelled by a more appropriate profile than a Gaussian. One advantage of the Sérsic fit is that the effect of the PSF can be incorporated into the model by defining a model that is convolved with an input PSF, allowing to remedy some of the resolution effects.

There are several popular codes that optimize a 2D Sérsic profile: GALFIT (Peng et al. 2002), GALFITM (Häußler et al. 2022), GALIGHT (Ding et al. 2020), which is a wrapper for LENSTRONOMY (Birrer et al. 2021), Bayesian fitter `pysersic` (Pasha & Miller 2023), SOURCEEXTRACTOR++ (Bertin et al. 2020), ASTROPHOT (Stone et al. 2023), GALFIT-CORSAIR for fitting core-Sérsic models (Bonfini 2014), and more. An exhaustive comparison of these implementations is beyond our scope, but we caution that the fits obtained with different software may not be consistent with one another, although Ren et al. (2025) found a good agreement between four different fitters. One caveat we would like to highlight is that all of the software mentioned above fit the data to a *pixelized* model, which accounts for the binning of the profile flux onto square pixels, while the default Astropy (and thus `statmorph`) fitter does not do this, leading to systematically lower Sérsic indices for peaked profiles (a difference of $\Delta n \sim 1$ at $n \sim 4$). For

this reason, we run our tests using GALFIT rather than `statmorph`.

The Sérsic model is non-linear and so the fit may converge to a local minimum. Therefore, a good initial guess for the parameters is important. Moreover, it is useful to keep the number of free parameters to a minimum, especially when the data are not sufficient to constrain more complex fits. On one hand, Peng et al. (2010b) showed that galaxies are not well-described by a Sérsic profile when they are sufficiently resolved: nuclear structures, bars, spiral arms, and other features mean that a single Sérsic fit is often insufficient for nearby objects. Codes such as GALFIT (Peng et al. 2010b) allow the inclusion of any number of arbitrary models as well as Fourier perturbations. However, when the resolution is low and these features are smoothed over, as is the case for high-redshift objects, a single Sérsic fit is often better as additional modes are poorly constrained, and the derived parameters are highly degenerate. In our tests, we fit a single Sérsic profile convolved with a PSF for all galaxies to minimize the degeneracies between parameters. We use the `statmorph` estimates of the Sérsic parameters as initial guesses (as was done in Sazonova et al. 2021).

The dependence of $R_{0.5}^{\text{Sérsic}}$ on image properties is plotted in Fig. 3, and the advantages of accounting for the PSF become clear. While both R_p and R_{20} overestimate the true radius at low resolutions, $R_{0.5}^{\text{Sérsic}}$ is, on average, much closer to the baseline, with a $\sim 15\%$ difference between highest and lowest $\langle \text{SNR} \rangle$, slightly overestimating the radius in shallow imaging. The resolution dependence is weak. This acts as an indirect evidence that our results agree with Ren et al. (2024), who found that the error in radius measurements stems from the PSF smearing out the flux, not the pixelization itself. When the PSF is accounted for, we get much closer radius estimates to the baseline. However, $R_{0.5}^{\text{Sérsic}}$ is a much more uncertain measurement: at $\langle \text{SNR} \rangle \sim 1$ the scatter depends on the resolution and ranges from 50% to 15%, while it stays at 10% for R_{20} and 15% for R_p . This stems from the fact that $R_{0.5}^{\text{Sérsic}}$ requires fitting for other Sérsic parameters simultaneously, and there may not enough information to constrain the fit.

Fig. 5 shows the behaviour of other Sérsic parameters: n (left), e (middle), and θ (right). While $R_{0.5}^{\text{Sérsic}}$ is weakly affected by resolution, n has a higher resolution dependence, with a systematic bias of $\Delta n = -0.4$ on average in the least resolved galaxies. This bias is quite small given the typical range of n is between $n = 0$ for late-type and $n \gtrsim 4$ for early-type galaxies, but can be important in cases where $n \sim 2$. Ellipticity and orientation are robust at all resolution and noise levels, with a negligible (-0.04 ± 0.15) bias in e with fewer than 5 resolution elements. However, the scatter in n is large (± 1.5), which means that the measurement is uncertain




FIG. 5.—: Same as Fig. 2, for three Sérsic parameters: n (**left**), ellipticity (**middle**), and orientation (**right**). n is underestimated by up to 0.5 when there are fewer than 10 resolution elements per object, and it is uncertain, with ± 1.3 scatter at lowest resolution. However, n is robust to noise on average, with a ± 0.5 scatter at lowest $\langle \text{SNR} \rangle$. Both $e^{\text{Sérsic}}$ and $\theta^{\text{Sérsic}}$ are, on average, much more reliable than their non-parametric counterparts (Fig. 2), with negligible offsets from the baseline and increasing scatter in the lowest resolution and signal-to-noise regimes.




FIG. 6.—: same as Fig. 2, for bulge strength metrics concentration (C , first panel), Gini index (G , second panel), M_{20} (third panel) and the Gini- M_{20} bulge strength (last panel). All four metrics are biased with decreasing resolution, leading to lower bulge strength measurements with less $\mathcal{R}_{\text{eff}}$. Since $B(G, M_{20})$ is constructed from two resolution-dependent metrics, it is even more sensitive resolution, leading to a bias of ~ 1 at $\mathcal{R}_{\text{eff}} < 10$. While M_{20} and C are fairly insensitive to $\langle \text{SNR} \rangle$, G and hence $B(G, M_{20})$ are also underestimated at $\langle \text{SNR} \rangle < 3$.

for individual galaxies when the resolution is poor.

Part of the reason for the scatter in n and $R_{0.5}^{\text{Sersic}}$ is that these two parameters are degenerate with one another (e.g., [Graham et al. 1996](#); [Trujillo et al. 2001](#)), so varying them together can result in similarly well fitting profiles. GALFIT estimates the uncertainty in fitted parameters from the error image, but does not give uncertainties based on this degeneracy, so the errors are always underestimated. Bayesian codes, such as `pysersic` provide more robust error estimates, but are slower due to the need to run Monte Carlo chains. An alternative approach to quantify this degeneracy was proposed in [Trujillo et al. \(2001\)](#). We can define a “true” Sérsic profile, a “fitted” one, and a χ^2 between the two. Then, by fixing n , we can find the values of $R_{0.5}^{\text{Sersic}}$ and $I_{0.5}$ that minimize the χ^2 given a fixed n . While [Trujillo et al. \(2001\)](#) showed this relationship for $n_{\text{fit}} = 4$, we used a similar approach to analytically solve for the degeneracy between $R_{0.5}^{\text{Sersic}}$, n , and $I_{0.5}$ for arbitrary values of n in Appendix E.

We find that the degeneracy has the form

$$\frac{R_2}{R_1} = \exp 1.223 \left(\Delta n / n_1^{\sqrt{\mathcal{R}_{\min}}} \right) \quad (5)$$

where n_1 and R_1 are the “true” parameters, and n_2 , R_2 are the “fitted” ones. Here $\Delta n \equiv n_2 - n_1$. $\mathcal{R}_{\min}$ roughly corresponds to the image resolution, where we defined it as half the pixel scale.

This equation can be used to find the uncertainties on the fitted parameters, accounting for the degeneracies between them. If the best-fitting values are n_1, R_1 , then the *next*-best fitting $\Delta n, \Delta R$ can be computed with Eq. 5. These next-best-fits, of course, will result in a higher χ^2 . We can then set a limit on an acceptable χ^2 to find the range of parameters that all produce acceptable fits – this is an uncertainty intrinsic to the degeneracies of the Sérsic profile, given the image variance. In our analysis, we accept the fits with $p(\chi^2) > 0.05$, i.e. we calculate a 2σ intrinsic uncertainty on n and $R_{0.5}^{\text{Sersic}}$.

In Appendix E, we computed the χ^2 for these “next best fits”. We find that a typical uncertainty on $R_{0.5}^{\text{Sersic}}$ that is *intrinsic* to the Sérsic profile is 20% at $\langle \text{SNR} \rangle = 3$ and $\mathcal{R}_{\text{eff}} = 10$, comparable to the empirically measured scatter in Fig. 3. Similarly, we found that the intrinsic uncertainty in n ranges from 0.2 in late-type (true $n=1$) galaxies to 0.6 in early-type (true $n=4$) galaxies. This is a significantly higher error than what is usually reported by GALFIT, but still is an underestimate of the observed scatter in Fig. 5. Running a Monte Carlo-based fit will provide more robust uncertainty estimates specific to a given image, but this approach can still be used to estimate the joint uncertainties on n , $R_{0.5}^{\text{Sersic}}$, and $I_{0.5}$ much more realistically than a typical least-squares fit does.

Another consequence of this is that two or more com-


ponent Sérsic fits, such as bulge-disk decomposition, lead to degenerate parameters when n is allowed to vary. During decomposition, the central flux is dominated by the bulge component, while the outer flux is typically modelled as a disk. However, for both models, most of the constraining power lies in the inner regions where the signal-to-noise is highest. This means that the inner component is well-constrained, but for the disk component, $\mathcal{R}_{\min}$ is effectively set by the radius where the inner component is no longer dominant, rather than the resolution. As $\mathcal{R}_{\min}$ increases, the degeneracy between n and $R_{0.5}^{\text{Sersic}}$ becomes stronger, and n is essentially impossible to constrain. Therefore, we recommend that bulge-disk decomposition restricts the model to a conservative $n_{\text{bulge}} = 4$ and $n_{\text{disk}} = 1$ choice (e.g., [Mendel et al. 2014](#); [Bottrell et al. 2019](#); [Nedkova et al. 2024](#)). Varying the bulge’s Sérsic index while keeping an exponential disk is also possible (e.g., [Andredakis et al. 1995](#); [Graham 2001](#); [Allen et al. 2006](#); [Gadotti 2009](#); [Simard et al. 2011](#); [Dimauro et al. 2018](#); [Bottrell et al. 2019](#); [Lima-Dias et al. 2024](#)), but testing the intrinsic degeneracy of this choice is beyond the scope of this paper. Allowing both components to have a free value of n (e.g., [Lange et al. 2016](#); [Häußler et al. 2022](#)) will suffer from the greatest joint uncertainties of the components, since the slope of the disk-like component cannot be easily constrained in the center, which carries the most weight during the fit. [MacArthur et al. \(2003\)](#) investigated the reliability of the two-parameter fit using mock profiles and found 20% systemic errors in n for the free bulge component and only 5% for the disk, but did not discuss the degeneracy between the two components. A solution could be to model the components with a Bayesian approach, which accounts for degeneracies (e.g., [Robotham et al. 2017](#); [Argyle et al. 2018](#); [Casura et al. 2022](#); [Cook et al. 2025](#)).

3.3. Bulge strength measurements

3.3.1. Concentration

The concentration parameter C quantifies the relative brightness of the centre of a galaxy relative to its outskirts. C is measured by comparing two circular isophotal radii, R_{inner} and R_{outer} ([de Vaucouleurs 1977](#); [Okamura et al. 1984](#); [Kent 1985](#); [Bershady et al. 2000](#); [Conselice 2003](#)). Most commonly, the inner and outer radii are taken as the 20th and the 80th percentile isophotes, although different choices are possible: e.g., using a surface brightness threshold for the outer isophote as in [Kent \(1985\)](#) or a ratio of effective Sérsic radii as in [Trujillo et al. \(2001\)](#). [Bershady et al. \(2000\)](#) showed that the 20th and 80th isophotes are stable to changes in resolution and signal-to-noise, although we find the opposite in Sec. 3.2.2, and earlier [Graham et al. \(2001\)](#) also discussed the issues of 20th/80th isophotes. Still, the most

FIG. 7: Same as Fig. 4, showing the distribution of the measured (left) and corrected (right) C compared to the baseline. There is a bias where C is underestimated more for early-type galaxies at low $\mathcal{R}_{\text{eff}}$, and the scatter is larger for low $\langle \text{SNR} \rangle$. The $\langle \text{SNR} \rangle$ effect is mitigated by the correction. The $\mathcal{R}_{\text{eff}}$ bias is also alleviated, but the measurement uncertainties are amplified by the correction, leading to a larger scatter.



common definition of concentration, and one adopted by STATMORPH, is

$$C = 5 \log \frac{R_{80}}{R_{20}} \quad (6)$$

where the isophotes are found as described in 3.2.2.

The first panel in Fig. 6 shows the dependence of C on resolution and signal-to-noise. As seen in the figure, C is robust to changes in $\langle \text{SNR} \rangle$: there is no systematic offset in concentration until $\log \langle \text{SNR} \rangle$ reaches 0.2. This is expected since both R_{20} and R_{80} are not strongly affected by noise. Below this $\langle \text{SNR} \rangle$, C is slightly underestimated by ~ 0.2 . The scatter in C slightly increases as $\langle \text{SNR} \rangle$ decreases, reaching an uncertainty of ± 0.2 at low $\langle \text{SNR} \rangle$ at a fixed resolution. This is consistent with what Conselice et al. (2000), Lotz et al. (2004), and Ren et al. (2024) found, and we do not find the strong signal-to-noise dependence reported in Graham et al. (2001) and Trujillo et al. (2001).

Fig. 7 (left) shows the relationship between C and base C for the entire sample, the low- $\langle \text{SNR} \rangle$ low- $\mathcal{R}_{\text{eff}}$ regimes in red, white, and black respectively. The bias in C for poorly resolved galaxies depends strongly on the intrinsic concentration, with a higher bias for more bulge-dominated galaxies, reaching $\Delta C \sim 1.5$ when $\mathcal{R}_{\text{eff}} < 5$. Similarly, when $\langle \text{SNR} \rangle$ is low, early-type galaxies suffer a larger bias and more scatter in their C measurements, consistent with Wang et al. (2024), who found a bias of $\Delta C = 1$ for bulge-dominated sources. Therefore, at low resolution or poor $\langle \text{SNR} \rangle$, all galaxies will appear less bulge-dominated, and the distribution of bulge strengths will appear biased lower in agreement with previous studies (Yeom et al. 2017; Thorp et al. 2021; Yu et al. 2023; Ren et al. 2024; Wang et al. 2024)¹³ The dependence of

ΔC on intrinsic concentration has implications for studies of the galaxy morphology bimodality in the early Universe, where the resolution is poorer, as the fraction of centrally concentrated galaxies may appear to be lower due to this effect and can explain the decreasing C trends seen in recent JWST analyses (see Sec. 4).

The relationship between C and C_{base} we obtained with SymbolicRegression is

$$C = C_{\text{base}} \tanh(\log(0.4\mathcal{R}_{\text{eff}})) + \frac{3.3}{\sqrt{\mathcal{R}_{\text{eff}} - 0.76}}, \quad (7)$$

which reflects the dependence on intrinsic C_{base} . As in the case of R_p , we can then correct C measurements by calculating C_{base} from this equation. We are able to correct well for the bias in the low signal-to-noise galaxies (Fig. 7, right-hand panel). Our correction function also brings the measurements closer to C_{base} , although with a significant scatter introduced by the correction. Since at $\mathcal{R}_{\text{eff}} < 5$ the distribution of the raw C measurements (left) is so flat, any uncertainty in the measurements gets amplified, leading to a larger scatter in the corrected C (right). Moreover, since the correction mostly attempts to mitigate the effect of the PSF, our correction will be effective in images where the PSF FWHM is two pixels, and ideally a new correction should be derived for a range of different PSF sizes.

3.3.2. Gini index

The Gini index, G , originally proposed as a measure of wealth inequality, has been widely used to characterize the concentration of a galaxy's light since its introduction to astronomy in Abraham et al. (2003) and Lotz et al. (2004). The Gini index is defined as

$$G \equiv \frac{1}{2\bar{X}n(n-1)} \sum_i^n \sum_j^n |X_i - X_j|, \quad (8)$$

where n is the number of elements X_i and $\bar{X}$ is their mean. $G = 1$ means that all flux is concentrated in a

¹³ Conselice et al. (2000), Lanyon-Foster et al. (2012), and Yao et al. (2023) found no difference in C when artificially redshifting their sources. Given a wide range of studies that agree with our results, as well as the intrinsic dependence of C on isophotal radii which are resolution-dependent, we are confident that C does in fact depend on resolution.

single pixel, while $G = 0$ means that it is evenly distributed throughout the entire galaxy. In practice, this can be simplified to a single sum by first sorting X_i in ascending order, and then calculating G as

$$G = \frac{\sum_i^n (2i - n - 1)X_i}{n(n-1)\bar{X}} \quad (9)$$

By construction, the Gini index is meant to be used with strictly positive quantities – e.g., income, or a measure of impurity in machine learning applications (e.g., Pedregosa et al. 2011). If X_i is allowed to be negative, G is no longer bounded by the $(0, 1)$ range and interpreting the results becomes difficult. In astronomy, after the sky subtraction, the flux values are often negative, so to remedy this, G is most often calculated on the *absolute* values of flux. Therefore, the real definition becomes

$$G = \frac{\sum_i^n (2i - n - 1)|X_i|}{n(n-1)|\bar{X}|}. \quad (10)$$

Because of this absolute-valued definition, an interesting edge case is a Gini index of a white noise distribution, where $X_i \sim \mathcal{N}(0, \sigma^2)$. This is equivalent to the Gini index of a blank sky without any source flux. After the absolute value is taken, such a patch follows a half-normal distribution with mean $\mu = \sigma\sqrt{2/\pi}$. For a strictly positive distribution with known cumulative density function $F(x)$, the Gini index can be computed analytically, and for a half-normal distribution it is

$$G \equiv \frac{1}{\mu} \int_0^\infty F(x)(1 - F(x)) dx = \sqrt{2} - 1 \approx 0.41. \quad (11)$$

Lisker (2008) empirically found that as the signal-to-noise decreases, G approaches a constant value of ~ 0.4 . This is precisely due to approaching this limit of the Gini index of absolute-valued white noise.

In general, as more and more sky pixels are included in the aperture where G is calculated, the sum will become dominated by the sky, and so G will approach this value of 0.41. Therefore, G is not only sensitive to signal-to-noise, but also to the aperture size (Lotz et al. 2004; Lisker 2008; Rodriguez-Gomez et al. 2019). In practice, G is therefore calculated within an aperture defined by R_p , which is robust to changes in signal-to-noise (Sec. 3.2.1) and thus cosmological surface brightness dimming. In STATMORPH, an image is first smoothed by a filter with a size $R_p/5$, and then a new segmentation map is computed using the average SB at R_p as a threshold. G is then calculated only using the (un-smoothed) pixels included in this map. In their tests, Lisker (2008) confirm that using an aperture with a Petrosian radius has a sufficiently small contamination from sky pixels that G does not yet approach the noise-limited value of 0.41.

In Fig. 6 (second panel), we show that our results

are in agreement with previous discussion in Lotz et al. (2004), Holwerda et al. (2011) and Lisker (2008), where G is robust up to $\langle \text{SNR} \rangle \sim 2$. At lower $\langle \text{SNR} \rangle$, G begins to decrease towards the noise-limited value of 0.41. As we showed in Sec. 3.2.1, R_p is robust to changes in $\langle \text{SNR} \rangle$, only decreasing by $\sim 10\%$ as $\langle \text{SNR} \rangle \rightarrow 1$, and so as signal-to-noise decreases and R_p stays roughly constant, the pixels within R_p become more noise-dominated, leading to decreasing values of G . This explains the results in Petty et al. (2014), who use *HST* images to calculate G and find a flat distribution of $G \approx 0.45$ at high redshifts – their images were likely already in this noise-dominated regime, and caution should be taken when measuring G on noisy images. However, at medium signal-to-noise, the measurement is very reliable.

Similarly to concentration, G is much more sensitive to resolution than signal-to-noise: at more than 10 resolution elements per galaxy, G is stable, but with fewer than 5 it can be underestimated by up to 0.1 with a ± 0.05 scatter, which is significant given a small dynamic range of G values. This is consistent with the tests of G and resolution in Yeom et al. (2017) and Wang et al. (2024). In particular, Wang et al. (2024) – with their slightly different definition of $\mathcal{R}_{\text{eff}}$ – also find an increasing bias with resolution that flattens at $\mathcal{R}_{\text{eff}} > 6$, at slightly lower but similar resolutions than in our tests. We computed the Gini index for some simple analytic Sérsic profiles with a changing pixel scale and PSF size independently, and found that while both affect G , the PSF has a much stronger impact. As the PSF redistributes flux from the bright point sources outwards, it naturally decreases the Gini index. The correction derived with SymbolicRegression is:

$$G = G_{\text{base}} \left(0.83 - \frac{0.59}{\mathcal{R}_{\text{eff}}} \right) + 0.1 \tanh(\langle \text{SNR} \rangle) \quad (12)$$

The offset depends on both $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$, but the fit applies a flat correction to the noise-dependence. Therefore, comparing the *relative* Gini indices of galaxies at the same redshift observed with the same instrument is still possible provided the resolution is stable and the signal-to-noise is not so poor that G becomes noise-dominated. However, care must be taken when comparing galaxies with different effective resolutions.

3.3.3. M_{20}

M_{20} was first defined in Lotz et al. (2004) as an alternative measure of a galaxy’s concentration, and is often used in conjunction with the Gini index to find bulge-dominated or merging galaxies. It is based on the second moment of a galaxy’s light distribution. In general, we can find the second moment of the X brightest galaxy pixels by first sorting all the pixels by flux, and then finding the sum

$$m_X = \sum_i^X f_i ((x_i - x_0)^2 + (y_i - y_0)^2). \quad (13)$$

M_{20} is then defined as the moment of the brightest 20% of pixels divided by the total moment:

$$M_{20} \equiv \frac{m_{20\%}}{m_{100\%}}. \quad (14)$$

In practice, the total light cannot be known, and so $m_{100\%}$ is defined within some aperture. STATMORPH uses the Gini segmentation map, constructed with the average SB at R_p as a threshold, and so roughly corresponds to flux within R_p .

M_{20} is a purely photometric quantity similarly to C , and so is robust to changes in $\langle \text{SNR} \rangle$, as shown in the third panel of Fig. 6. As $\langle \text{SNR} \rangle$ decreases, the scatter in M_{20} increases as photometry becomes more uncertain, but there are no systematic biases in shallow imaging. This was also concluded in Ren et al. (2024), while Lotz et al. (2004) and Wang et al. (2024) found slightly higher values of M_{20} at very low signal-to-noise. The difference in our results is likely due to implementation; if the Petrosian-based segmentation map is used to compute M_{20} like in STATMORPH, the values are stable even at $\langle \text{SNR} \rangle < 1$.

However, M_{20} is biased at low resolutions: as $\mathcal{R}_{\text{eff}}$ decreases, M_{20} can increase by up to 0.6, which could lead a galaxy to be misclassified from bulge-dominated to disk-dominated. This is a natural consequence of the PSF: it distributes the flux from bright point sources, such as the central bulge and any star clusters, to large distances; therefore naturally increasing M_{20} by its definition. This was also the conclusion in Ren et al. (2024). The effect of the PSF is therefore important to keep in mind and correct for. The correction function we derived is

$$M_{20} = \left[0.74 - \frac{1}{\mathcal{R}_{\text{eff}}} \right] M_{20, \text{base}} - 0.5. \quad (15)$$

This correction depends only on $\mathcal{R}_{\text{eff}}$ as expected, and reduces the bias introduced by the PSF. However, since the scatter and hence the uncertainty in M_{20} at low resolutions is high, the correction will also amplify these uncertainties; moreover, it depends on the exact shape of the PSF and so ideally it should be re-derived with a PSF significantly different from a 2-pixel Moffat used in this work.

Finally, several studies that artificially redshifted their galaxies from $z \approx 0.5$ to $z > 2.5$ found no change in their M_{20} measurement (Petty et al. 2014; Yao et al. 2023; Salvador et al. 2024). We expect this stems from the fact that at $z > 1$, the decrease in angular size with proper distance is small, and so galaxies with similar intrinsic sizes will have similar apparent sizes. On the other

hand, cosmological surface brightness dimming does not strongly affect M_{20} as it is insensitive to signal-to-noise. However, observers should still be careful when comparing M_{20} measurements of galaxies with different *intrinsic* sizes, as it is the *effective* resolution that introduces the bias in M_{20} . As galaxies at higher redshifts are intrinsically smaller, their effective resolution is still lower and thus measured M_{20} is likely overestimated – see Sec. 4 for further discussion.

3.3.4. Gini- M_{20} bulge strength

Using subtle differences in otherwise similar metrics can be a powerful way to probe more detailed galaxy structure. Lotz et al. (2004) use G and M_{20} in tandem in this way. Both metrics probe how concentrated a galaxy light is, but G is agnostic to the centre of the galaxy, and measures bright pixels anywhere in the image. On the other hand, M_{20} is heavily weighted to the distance of these bright pixels from the centre of flux, making these two similar measurements actually quite different.

Regular early-type and late-type galaxies therefore form a “main sequence” on the $G - M_{20}$ plane. Bulge-dominated galaxies have bright pixels in the center, leading to high G and low M_{20} values. Star-forming disks have many bright knots of star formation, leading to low G and high M_{20} values instead, thus forming this “main sequence”. But there are rarer galaxies off the axis of this main sequence: for example, during a merger, when two galaxies have already a shared envelope but the nuclei have not yet coalesced, there are two bright nuclei – leading to high G – that are both removed from the common centre of light, leading to high M_{20} . Therefore, the $G - M_{20}$ plane can be used to both classify galaxies along the Hubble sequence, and identify a particular stage of mergers. Snyder et al. (2015b) and Rodriguez-Gomez et al. (2019) define a bulge strength and a merger statistic based on the distance along the main sequence and the offset away from it. Here we discuss the bulge strength statistic, computed in `statmorph` as

$$B(G, M_{20}) \equiv -0.693M_{20} + 4.95G - 3.96. \quad (16)$$

We note that this is simply an empirical relation derived using local galaxies in Snyder et al. (2015b). Sazonova et al. (2020) and Espejo Salcedo et al. (2025) show that at $z > 1$, galaxies still follow a “main sequence” in the $G - M_{20}$ plane, but the slope and the equation of this main sequence are different. Therefore it may be meaningful to redefine this statistic for each individual study, using a representative sample of early- and late-type galaxies at the same redshift with comparable imaging quality.

Perhaps unsurprisingly, we find that $B(G, M_{20})$ depends on $\langle \text{SNR} \rangle$ the same way as G does, with decreasing bulge strength at low signal-to-noise (Fig. 6). The de-

pendence of $B(G, M_{20})$ on resolution is however more striking. G decreases in low-resolution imaging, leading to galaxies being classified as more disk-like. M_{20} increases at low resolutions, also producing more disk-like morphologies. Therefore, $B(G, M_{20})$ absorbs both of these effects – as resolution decreases, $B(G, M_{20})$ decreases by up to 0.8. The dynamic range of this statistic on the Pan-STARRS sample used in [Rodríguez-Gómez et al. \(2019\)](#) is approximately $(-1, 1)$, and so a bias of -0.8 is a significant decrease. While the bias in G does not depend on G_{base} , the bias in M_{20} does. Therefore, even re-defining a new “ $G - M_{20}$ main sequence” in imaging with a different resolution will not fully mitigate this resolution dependence of $B(G, M_{20})$.

Similarly to C , a correction for $B(G, M_{20})$ can be derived directly from G and M_{20} . Alternatively, the best fit provided by `SymbolicRegression` is shown in Appendix D. For our PSF, this provides a good correction on average, but the scatter at $\mathcal{R}_{\text{eff}} < 5$ after the correction is high, reaching ± 1 – see the figure in the online materials.

3.4. Disturbance measurements

3.4.1. Gini- M_{20} disturbance

Following the Gini- M_{20} bulge strength, the disturbance in G - M_{20} plane is calculated as a vertical offset from the “main sequence” and corresponds to galaxies with concentrated but off-centered peaks, such as double nuclei. `statmorph` defines it as

$$S(G, M_{20}) = 0.139M_{20} + 0.99G - 0.327. \quad (17)$$

It is important to note that while $S(G, M_{20}) > 0$ indicates a deviation from the “main sequence” towards more unusual morphologies, high values being associated with mergers, a negative value of S does *not* mean that the galaxy is “less disturbed”, so we caution against over-interpreting negative values of this parameter.

The behaviour of $S(G, M_{20})$ at different image qualities is shown in the online materials for brevity. Interestingly, since $S(G, M_{20})$ measures the *perpendicular* distance in the G - M_{20} plane, the resolution dependencies of the two parameters mostly cancel out, so S is a stable measurement for the majority of galaxies in our sample, with only a ± 0.02 average deviation and a 0.1 average scatter across the entire parameter plane. However, there are only 20 galaxies in the sample with $S(G, M_{20}) > 1$. For these objects, there is a systematic decrease of roughly -0.1 at low resolutions, although this trend suffers from low-number statistics. Finally, if $B(G, M_{20})$ is re-defined at higher redshifts, then $S(G, M_{20})$ must also be re-defined accordingly ([Sazonova et al. 2020](#); [Espejo Salcedo et al. 2025](#)).

3.4.2. CAS Asymmetry

Asymmetry (A_{CAS}) is a part of the concentration-asymmetry-smoothness (CAS) system from [Conselice \(2003\)](#) and is one of the most widely used measurements of the irregularity of a galaxy. It is conceptually simple: an image is rotated by 180° about its centre, and then the rotated image is subtracted from the original. The asymmetry measurement is the sum of the residual normalized by the total flux of the galaxy.

In practice, this calculation involves many hidden nuances. The centre of rotation has a strong impact on the final value: while the original definition in [Schade et al. \(1995\)](#) used the centroid of the galaxy, [Conselice \(2003\)](#) find a centre which *minimizes* asymmetry. Moreover, the sky background itself has a non-zero asymmetry, both due to random fluctuations in flux and potential sky gradients, and so must be subtracted. A modern definition of A_{CAS} , and that adopted by `statmorph`, is:

$$A_{\text{CAS}} = \min_{\mathbf{x}_0} \frac{\sum |I_{ij} - I_{ij}^{180}|}{|I_{ij}|} - \langle A_{\text{CAS}, \text{bg}} \rangle \quad (18)$$


where I_{ij}^{180} is the flux rotated about $\mathbf{x}_0$, and $\langle A_{\text{CAS}, \text{bg}} \rangle$ is the average sky background found in an empty patch of the sky. However, while this definition is used in most codes, the exact implementation (e.g., between `statmorph` and MORFOMETRYKA) differs and so produces different measurements on the same set of images.

We have discussed this measurement in great detail in [Sazonova et al. \(2024\)](#), and so we keep the discussion in this work brief, and suggest the reader refers to our previous work, where we the impact of the sky subtraction, centering, and interpolation needed when the image is rotated on A_{CAS} . The primary conclusion in [Sazonova et al. \(2024\)](#) is that the absolute-valued definition makes a perfect sky subtraction impossible, and $\langle A_{\text{CAS}, \text{bg}} \rangle$ always overestimates the true contribution of the background to asymmetry. Therefore, this definition will *over*-subtract the background asymmetry, leading to systematically lower A_{CAS} values.

The dependence of A_{CAS} on resolution and depth in our current series of tests is shown in the left panel of Fig. 8, agreeing with our previous conclusions: asymmetry is systematically underestimated at low $\langle \text{SNR} \rangle$. Our results are qualitatively consistent with previous works (e.g., [Bottrell et al. 2019](#); [Thorp et al. 2021](#)). Similarly to [Thorp et al. \(2021\)](#), we also find a weak dependence on resolution, although it is not as significant as the dependence on $\langle \text{SNR} \rangle$.

There are several approaches to mitigate the severe dependence of A_{CAS} on signal-to-noise. While they are discussed at length in [Sazonova et al. \(2024\)](#), for consistency with other parameters, we also provide a `SymbolicRegression` correction in Appendix D. We also

FIG. 8: same as Fig. 2, for disturbance metrics: A_{CAS} (first panel; Conselice 2003) (first), A_{RMS} (second panel; Sazonova et al. 2024), shape asymmetry A_S (third panel; Pawlik et al. 2016), and A_{22} , asymmetry of the 22 mag/arcsec² isophote (last panel; this work). All asymmetry metrics except A_{RMS} are biased at decreasing $\langle \text{SNR} \rangle$, although A_{22} is robust as long as $\mu_0 > 22$. A_{RMS} , on the other hand, is much more sensitive to resolution than the other metrics.



test an *outer* asymmetry, A_o , defined as ab asymmetry where the inner pixels up to R_{50} are masked – these results are in the online materials and Appendix D. Again, we caution that while this correction achieves consistent measurements with the baseline *on average*, it introduces a large scatter for low- $\langle \text{SNR} \rangle$ objects and so it is better to either approach it with caution, or employ other metrics when $\langle \text{SNR} \rangle$ is variable or low (e.g., A_{RMS} , see below).

3.4.3. RMS Asymmetry

Root-mean-squared asymmetry (A_{RMS}) is an alternative definition of the same metric, where the residual flux is squared rather than absolute-valued. This metric was also proposed in Conselice (2003) but has historically gained less traction. It provides an important advantage: the sky contribution to asymmetry is separable from the object asymmetry since the square is better-behaved than an absolute value (Sazonova et al. 2024). Therefore, it is, by definition, much more robust to noise. A_{RMS} is defined as:

$$A_{\text{RMS}}^2 = \min_{\mathbf{x}_0} \frac{\sum (I_{ij} - I_{ij}^{180})^2}{I_{ij}^2} - \langle A_{\text{RMS,bg}}^2 \rangle \quad (19)$$

We have also discussed this metric in detail in Sazonova et al. (2024), but we still show the resolution and signal-to-noise plot in the second panel of Fig. 8. There is almost no dependence on signal to noise, aside from a small -0.1 average underestimate at $\langle \text{SNR} \rangle \sim 1$, and a small amount

of scatter as the signal-to-noise decreases. However, the metric is sensitive to resolution, with a ± 0.2 spread at different resolution levels. Since the PSF spreads out the flux, the sum of the squared fluxes becomes smaller, and so the measurement decreases when the PSF is large. Galaxies with higher intrinsic asymmetries are more affected by this. Our derived correction is:

$$A_{\text{RMS}} = A_{\text{rms,base}} \times 0.5 \log \mathcal{R}_{\text{eff}}. \quad (20)$$

While this correction provides a good fit to our data on average, it significantly overestimates A_{RMS} at low resolutions, $\mathcal{R}_{\text{eff}} < 5$ (see the online materials). Moreover, since the impact of the PSF depends on the shape of the PSF itself, we advice to derive a new correction if the user's PSF is significantly different from a 2-pixel Moffat used in this work, or explore deconvolution methods as described in Sazonova et al. (2024).

3.4.4. Shape asymmetry

CAS asymmetry and, to a larger extent, RMS asymmetry are more sensitive to bright disturbances such as star-forming clumps or multiple nuclei, since they are flux-weighted. However, in many cases, such as in tidal tails induced during late- or early-stage mergers, the asymmetric structures have a low surface brightness. Traditional metrics such as asymmetry and Gini- M_{20} disturbance are not necessarily designed to best detect these large but faint features. An alternative approach was




FIG. 9.—: The isophotal asymmetry contours for three galaxies: NGC 17 (**left**), NGC 201 (**centre**), and NGC 4600 (**right**) for a range of surface brightness thresholds from 20 to 25.5. NGC 17 is a post-merger and has strong tidal features, reflected in high $A_{>23}$. NGC 201 is a spiral galaxy with mild asymmetries induced at the medium SB level $A_{21.5}$ by the bar. Finally, NGC 4600 is a lenticular galaxy with low A_X except a small bump at $A_{20.5}$ caused by a star cluster. A_X as a function of X therefore is a powerful tool to trace which features cause which disturbances in a galaxy.

proposed in Pawlik et al. (2016): to measure an asymmetry of a binary detection map of a galaxy, without a flux weighting. This is a measure of *shape* asymmetry, A_S .

While there are different ways of constructing this binary detection maps, they are generally similar. `statmorph` implements A_S by first smoothing an image with a 3×3 boxcar filter, and then drawing a mask of all connected pixels above a 1σ threshold of the background noise. Then Eq. 18 is used to calculate the asymmetry using the same centre as that in A_{CAS} . Since the background pixels have a value of 0, a background term is not necessary.

Fig. 8 (third panel) shows the behaviour of A_S . While this metric is fairly robust to changes in resolution, it is sensitive to the surface brightness limit, μ_0 , leading to a systematic underestimate of $\Delta A_S \approx 0.15$ and a large scatter at $\mu_0 < 22$ compared to the baseline $\mu_0 = 23.5$. While this metric is often used to detect and time mergers or other disturbances (e.g., Pawlik et al. 2016; Snyder et al. 2015a; Foster et al. 2025; Yanagawa et al. 2025), there are no studies showing this behaviour which is crucial in detecting mergers in different surveys or at different redshifts. The intrinsic uncertainty in A_S compared to the baseline is too large compared to the bias to get an effective correction, but we provide the error statistics in Appendix D.

This dependence on noise is a natural consequence of the definition of A_S . Since the segmentation map used to

calculate it is defined using a 1σ background noise threshold, then the features used to calculate the asymmetry are the features that correspond to the surface brightness limit of the image. While a deep image, such as those expected from the LSST, will probe tidal features of nearby galaxies at $\mu_0 \approx 30$, an SDSS-like imaging of the same objects will probe the asymmetry of internal structures, such as the outskirts of the spiral arms. Therefore, A_S measured in these two cases will quantify the asymmetry of *different physical features*. Instead, we propose a similar, but a more physically motivated approach to quantifying a “shape asymmetry” discussed below.

3.4.5. Isophotal asymmetry

The goal of the new isophotal asymmetry (A_X) metrics is to alleviate the dependence of A_S on the background, instead providing features at a specific surface brightness or luminosity (and by proxy, mass) surface density.

It is defined in the same way as A_S , except we use a different threshold to compute the binary map. We define A_X as the asymmetry of a galaxy’s surface brightness $\text{SB} = X \text{ mag/arcsec}^2$ isophotes rather than the background noise level. This can be an even more powerful approach to trace disturbances in both the internal structure and the low surface brightness features of a galaxy. To measure an isophotal asymmetry, we simply create a mask of all pixels brighter than a particular SB threshold X . We then rotate this mask about the A_{CAS} centre and calculate the residual, divided by the size of the de-

tection mask. Unlike for shape asymmetry, we do not require the regions to be connected, since at a high surface brightness threshold this measurement probes disconnected clumps. This approach requires a good masking routine, ideally masking out any foreground stars. To make sure this measurement is physically meaningful, it should also be corrected for any cosmological dimming effects – in this work, we define rest-frame surface brightness μ_0 as $\mu_0 \equiv \mu - 2.5 \log(1+z)^3$, which includes the $(1+z)^4$ dimming and a K-correction arising bandpass shift in the AB magnitude system (Hogg et al. 2002).

Fig. 9 shows example isophotal contours for three galaxies: NGC 17, a recent merger, NGC 201, a spiral galaxy, and NGC 4600, a lenticular galaxy in the Virgo cluster. Each image is scaled to 25 parsec/pixel and has a surface brightness limit of 24 mag/arcsec². We constructed isophotes with a range of SB = 20 ~ 25.5 (light yellow to dark purple). Isophotal asymmetry measured using these different contours is shown below. The post-merger galaxy has a strong asymmetry at a low surface brightness levels – A_{22} and deeper – tracing the large tidal feature to the southwest. On the other hand, the spiral galaxy has symmetric faint isophotes and a large bright A_{21} , which traces the barred structure in the centre. Finally, the lenticular galaxy has low isophotal asymmetry at all surface brightnesses, slightly bumped up with a star cluster at $X \approx 20.5$. This shows the power of our proposed metric: a surface brightness, after correcting for cosmological dimming, maps directly to the luminosity density of the galaxy. Therefore measuring A_{21} or A_{24} probes the asymmetry of different features with different light, and hence mass surface density, and can help differentiate between tidal tails, double nuclei, knots of star formation or star clusters, and other structures. The profile of the galaxy where A_X is plotted against X has a potential to distinguish between different types of disturbances. It is also possible to define an integrated measure, such as a flux-weighted asymmetry of different isophotes, similar to the “twistiness” in Ryden et al. (1999). We plan to continue this exploration in an upcoming work.

The last panel in Fig. 8 shows the performance of A_{22} , the asymmetry measured using a 22 mag/arcsec² isophote. The measurement is not affected by the limiting depth by definition, as long as the depth exceeds the measured surface brightness. In practice, A_X is reliable when $X \lesssim \mu_0 - 0.5$, i.e. even if the sky surface brightness limit is shallower by up to 0.5 mag/arcsec² than the desired contour, A_X is still a good measurement. At very low resolutions, A_{22} is slightly underestimated. This is due to the loss of small, bright clumps as well as the PSF smearing effect as resolution decreases. High surface brightness measurements, such as A_{19} , will be affected more since they are more sensitive to individ-

ual clumps; and low surface brightness measurements are more robust to changes in resolution. In particular, in the low surface brightness regime probing the faintest tidal features, low resolution is *beneficial* since it increases the effective image depth while retaining these large features.

3.5. Other metrics

3.5.1. Smoothness

S , often referred to as either “smoothness” or “clumpiness” was originally defined in Conselice (2003) as part of the *CAS* system to quantify how uneven the light distribution is.

It is measured by convolving the image with a smoothing kernel, and then subtracting the smoothed profile from the original:

$$S = \frac{\sum \max(0, I_{ij} - I_{ij}^S)}{\sum I_{ij}} - \langle S_{bg} \rangle, \quad (21)$$

where I_{ij}^S is the smoothed image, and $\langle S_{bg} \rangle$ is the average clumpiness of the background. Since the nucleus of the galaxy is bright, any smoothing kernel would lead to a significant difference between smoothed and unsmoothed images in the centre, so the centre is excised from the calculation. `statmorph` therefore calculates S in an aperture between $0.25R_p$ to $1.5R_p$, using a uniform kernel of size $0.25R_p$. Pixels where the smoothed flux is *larger* than the original, so the residual is negative, are set to zero and thus ignored.

A high value of S typically corresponds to a galaxy with many bright star-forming regions, which disappear when smoothed over, giving rise to a large residual. However other small-scale structures such as bars and spiral arms will also contribute to the residual, so S measures the amount of substructure rather than simply clumps. Sometimes A and S are used in tandem to detect true mergers: Conselice (2003) suggests selecting $A > 0.3$ and $A > S$ to eliminate galaxies that are asymmetric due to their distribution of star-forming clumps. However, late-stage mergers such as those with embedded multiple nuclei will also have large values of S , since only one of the nuclei will be excised, so this threshold should be used with caution.

The first caveat of this measurement is that for small galaxies, with $R_p \approx 4$ pixels, the filter radius is 1 pixel or less, and so the image is not smoothed, leading to $S = 0$. This becomes important in large catalogs that include either low-mass or high-redshift sources with a small angular size. We recommend only using S measurements when $R_p > 12$ pixels, where the filtering step is reliable¹⁴.

¹⁴ One could also set `petro_fraction_cas` to a lower value, but this would affect the size of the excised region, which should ideally be well beyond the nucleus of the galaxy.




FIG. 10: same as Fig. 2, for smoothness (S ; Conselice 2003, left) and substructure (St ; Sec. 3.5.2, right). Both decrease at low $\langle \text{SNR} \rangle$ as clumps become more difficult to distinguish from the noise. The scatter in S however dominates the typical measurements of S at $\langle \text{SNR} \rangle < 3$, so this metric is noise-dominated in most cases. The scatter in St is smaller than its typical dynamic range, making it more reliable, but it has a resolution dependence absent in S .

The second caveat is that the background contribution to S can be large since the negative residual pixels are set to zero. Consider a Gaussian noise distribution with $I_{bg} \sim \mathcal{N}(0, \sigma)$. This noise is then passed through a uniform filter with area A to produce I_{bg}^S . Accounting for the covariance between the two, the residual then follows a distribution

$$I_{bg} - I_{bg}^S \sim \mathcal{N}\left(0, \sigma \sqrt{1 - \frac{1}{A}}\right). \quad (22)$$

Half of the residual pixels are set to zero and do not contribute to the expectation of the residual. The remaining ones follow a half-normal distribution, so the expectation of the residual becomes

$$\langle S_{\text{sky}} \rangle = \frac{\sigma}{2} \sqrt{\frac{2}{\pi}} \left(1 - \frac{1}{A}\right). \quad (23)$$

While this background term is separable and can be subtracted, unlike in the CAS asymmetry case (Sec. 3.4.2), the magnitude of the background term in practice is comparable to the magnitude of the residuals in the signal. Therefore, the random variance about the background median often dominates the real flux in the residual, and the measurement is noisy and background-dominated. In large surveys without exquisite depth, the background contribution is so significant that the smoothness values tend to vary randomly about 0 without a strong correlation to the underlying morphology of the source (e.g., Rodriguez-Gomez et al. 2019). We show the behavior of S with respect to the resolution and signal-to-noise in the left panel of Fig. 10 – there is a clear trend where S decreases, on average, in lower signal-to-noise regimes. But more importantly, the scatter increases as $\langle \text{SNR} \rangle$ drops and becomes more dominant than the systematic bias. This is due to the fact that the background subtraction noise dominates the measurement.

3.5.2. Smoothness alternative: Substructure

To address this issue, we propose a new metric, which we call *Substructure* (St). The motivation behind this

metric is the same as for the original smoothness. By applying a smoothing kernel to an image and looking at the residual, we are essentially performing an unsharp mask highlighting small-scale features – substructures – in the residual. Therefore, we repeat this process exactly like we did for smoothness, producing the same residual image.

The main difference is the treatment of the residual. Galaxy structures, such as clumps, spiral arms, bars, and others, are typically extended regions of several connected pixels. On the other hand, while the background also has spikes in the residual, they are disconnected as they arise from random noise. Therefore, our goal is to detect coherent residual structures that correspond to true galaxy features. We run `detect_sources` to identify any connected regions above 1σ of the residual noise. We use connectivity of 4 to exclude spurious connections in the noise, and choose a minimum area of 10 pixels as it is statistically unlikely to get 10 nearby pixels all above the noise threshold. We then compute the total flux contained in these detected clumps, normalized by the total flux in an aperture. Note that we use the image rather than the residual to calculate fluxes, unlike in the smoothness calculation. The background correction is similar: we compute the total flux in detected clumps in the background (if any), normalize by the total flux in the sky aperture, and then rescale the mean background substructure to the source aperture and subtract it. In most cases, due to our stringent detection criteria, the background substructure is 0. This is the same approach as what is used to detect clumps at high redshifts (e.g., Shibuya et al. 2016) and so our metric will be useful to evaluate the fraction of clumpy galaxies.

Fig. 11 shows an example of 8 galaxies in our sample, all imaged at 100 pc/px with $\mu_0 \sim 22.5$ mag/arcsec². The galaxies are sorted by increasing smoothness in the top row, and by increasing substructure in the bottom two rows. The bottom row shows the residual used in calculating both quantities, and orange regions show detected “clumps” that contribute to the substructure mea-




FIG. 11.—: Smoothness (S , **top**) and substructure (St , **middle**) measurements for 10 galaxies along with the unsharp residuals (**bottom**). All pixels in the residual are used to calculate S , and only contiguous regions (orange) are used in St . Galaxies are sorted by increasing S in the top row and St in the bottom two rows. St correlates better with a visual Hubble sequence while S varies a lot due to noise.

surement. The inner aperture is excised. First, it is clear in the images that the coherent regions detected have a range of morphologies: some are star-forming clumps, but some are spiral arms or the central bar which are also picked up by the unsharp mask – hence we opted to call the metric “substructure” rather than “clumpiness” to reflect the range of structures that contribute to it. Second, the substructure metric maps better to the diversity of galaxy types. When sorting galaxies by substructure (the middle row of Fig. 11), we see a clear Hubble sequence, starting from two early-type galaxies, then a barred early-type, and then a series of disks with clumps and spiral arms. On the other hand, smoothness is dominated by noise variability for all but the most clumpy objects and thus the sorting is less meaningful.

The right panel of Fig. 10 shows the dependence of substructure on depth and resolution. Similarly to smoothness, the parameter is very depth-dependent. This is to be expected: the calculation involves source detection above a noise level, and so the more noise there is, the fewer dim sources are detected, decreasing the overall flux contained in substructures. Similarly, as resolution decreases, substructures are more challenging to resolve. This is a natural consequence of what a metric aims to measure – however, for consistency across inhomogeneous surveys, one could adapt it to look for clumps that are brighter than a given limit, or train deep learning methods to ensure a homogeneous detection of clumps.

The substructure metric naturally traces bright star-forming clumps (e.g., Shibuya et al. 2016), as well as spiral arms, bars, and other sharp, high-contrast fea-

tures. This approach has many potential applications from detecting clumps and globular clusters to characterizing their flux ratios, colours, and relative distribution. This metric can be especially powerful when combined with integral field spectroscopy, which would trace the differences in ionization mechanisms, stellar populations, and metallicity within these clumps. In a future work, we plan to use the substructure metric in tandem with the visual classifications of galaxies observed with the MaNGA survey (Vázquez-Mata et al. 2025) to go beyond the Hubble sequence presented in Fig. 11.

3.5.3. Multimode

The multimode (M), intensity (I), and deviation (D) – together, MID – statistics were developed by Freeman et al. (2013) as an alternative way to probe disturbed and clumpy flux distributions. *statmorph* uses a slightly different definition of the statistics from Peth et al. (2016), which we adopt here as well.

First, a new segmentation map is generated to calculate the MID parameters. This map is defined similarly to the Petrosian radius: a surface brightness threshold q_{SB} for source detection is chosen such that the average surface brightness in the detected region is $1/\eta \times q_{SB}$. As for R_p , the segmentation process is repeated iteratively until the optimal q_{SB} is found such that $\eta = 0.2$.

Multimode finds two brightest peaks inside the segmentation map and compares their relative areas. First, we define some intensity quantile q_x . Then, construct a binary map – similar to the shape asymmetry map – where all pixels brighter than q_x are set to 1 and the rest




FIG. 12.—: Measurements of multimode (M ; Freeman et al. 2013; Peth et al. 2016) for a merger (**top**) and a spiral galaxy (**bottom**) in three regimes: deep and high-resolution (**left**), shallow low-resolution (**middle**), and shallow low-resolution (**right**). In the best imaging, M distinguishes the merger from the normal galaxy. In shallow high-resolution imaging, the second nucleus of the merger is deblended, leading to $M = 0$. In low- $\langle\text{SNR}\rangle$ imaging, a noise spike in the spiral galaxy leads to a spurious $M = 1$.

to 0. If q_x is low (e.g., $x = 0\%$), this selects all pixels inside the segmentation map and creates a single continuous region. If q_x is high (e.g., $x = 99\%$), a small region in the nucleus is created. For a simple analytic Sérsic profile with a single peak, regardless of q_x , there will only be one region. However, for multimodal sources such as mergers or even spiral galaxies, there exists a value of q_x where the map contains several disconnected regions. We then sort these regions by their areas and calculate the ratio between the areas of the largest and second largest regions:

$$R(x) = \frac{A_2(x)}{A_1(x)} \times A_2(x).$$

If there is only one region on the binary map, $R(x) = 0$. This ratio is calculated iteratively until an optimal $\hat{x}$ is found such that $R(\hat{x})$ is maximized. Multimode is then defined as¹⁵

$$M \equiv A_2(\hat{x})/A_1(\hat{x}).$$

This is a difficult optimization problem since it involves two steps: finding the optimal segmentation map and then $\hat{x}$. $R(x)$ can have many local maxima, so **statmorph** implements a brute-force method, increasing x by 2%, followed by a basin-hopping optimization step.

For a galaxy with two equally bright nuclei, $M = 1$;

and an early-type galaxy will have $M \approx 0$ – not necessarily 0 due to possible noise spikes.

In practice, for the majority of images in our test series, M is low ($M < 0.1$ for 77% of the sample and $M = 0$ for 20%). While in theory M could pick up star-forming clumps, they are much smaller and fainter than the central region of the galaxy, leading to low M . Since low- M images dominate our dataset, we do not provide a resolution-depth grid for M here, but it is available in the online materials. Instead, we discuss the general behaviour of the metric.

Fig. 12 shows an example multimode calculation for two galaxies: a merger with two nuclei in the top row and a spiral galaxy in the bottom, for a range of resolution and surface brightness limits. The two regions A_2 and A_1 for optimized $\hat{x}$ are in purple and yellow contours, and the masked foreground sources are hatched. We show the calculation in three regimes: deep and high-resolution, shallow high-resolution, and shallow poor resolution. In the first regime, the metric behaves as expected, with the merger having high M and the spiral galaxy essentially 0. The middle and third columns show edge cases: where M drops to 0 for a clearly disturbed galaxy, and where M increases to 1 for an undetected source.

In general, in sources with double nuclei where M is high, the measurement is robust to changes in signal-to-noise and resolution up to a certain limit. This limit depends on the intrinsic properties of a galaxy. For example, in the top left panel in Fig. 12, the two nuclei are enveloped in a single halo, leading to a high M . As signal-to-noise decreases, eventually the halo flux becomes lower than the detection threshold used in the main segmentation algorithm, and the two nuclei are deblended; so M drops from 0.47 to 0. In our example, this happens at $\mu_0 = 21$ mag/arcsec², but the exact threshold depends on the flux ratio between the nuclei and the halo, and so will change from one galaxy to another. Similarly, as resolution decreases, the two peaks become more difficult to resolve, and eventually they blend into a single peak with $M \approx 0$. For the galaxy in Fig. 12 this happens at $\mathcal{R} \sim 900$ pc/px. However, the exact resolution where this happens depends on the surface brightness of the galaxy, the flux ratios of the two peaks, and their physical separation, and so will also vary for different mergers. Until these thresholds in $\langle\text{SNR}\rangle$ and $\mathcal{R}_{\text{eff}}$ are reached, the measurement is quite robust to changes in image properties – but the problem is that these thresholds are impossible to define for all galaxies simultaneously. One way to reframe the multimode metric is that it measures a particular combination of merger mass ratios and times where the two nuclei are bright, separable, but not deblendable.

Finally, for extremely noisy images, where $\langle\text{SNR}\rangle < 1$, the optimizer tends to find a high value of $\hat{x}$ which leaves

¹⁵ Note that this is the definition from Peth et al. (2016) where M is a dimensionless quantity. In Freeman et al. (2013), $M \equiv R(\hat{x})$.




FIG. 13.—: The intensity and deviation (I and D ; Freeman et al. 2013) watershed map for an example spiral galaxy at different resolutions: 50, 75, and 100 pc/px on the left, centre, and right, respectively. While the changes in resolution are relatively small and visually the same features are present in the galaxy, the watershed maps are markedly different. The peak in the 50 pc/px case does not seem to correspond to a physical feature but leads to a bifurcation in the watershed and hence a high I measurement. D is stable across all three images despite the differences in the watershed maps.

several single-pixel noise spikes in the binary map, in addition to a (possible) real peak in the nucleus. Then the optimizer returns $M = 1/A_1$. In the most noise-dominated cases, *all* detected regions, including A_1 , are simply due to noise, which typically gives $M = 1, 0.5$, or $1/3$. These fractional M values should be treated carefully, especially if $\langle \text{SNR} \rangle$ is low. An example of this is shown in the bottom right panel of Fig. 12, where only two pixels are identified in the binary map, resulting in $M = 1$. One improvement to the algorithm could be setting a minimum area for the contiguous regions in the binary map, making it more robust to spurious noise detections; however, this is only an issue in the extremely low signal-to-noise regime.

3.5.4. Intensity

Intensity (I) is similar to multimode in that it measures the relative strengths of two peaks, but it is weighted by flux rather than area of the two regions. First, an image is smoothed by a Gaussian kernel (with $\sigma = 1$ px in `statmorph`). Then we find local maxima in the image (within the MID segmentation map) and create a watershed map for each local peak. In practice, there are many ways to identify local maxima, and the original algorithms described in Freeman et al. (2013) and Peth et al. (2016) are not explicit in how these maxima are defined, so I measurements produced by different codes likely differ. `statmorph` uses the SCIKIT-IMAGE library to find local maxima with `peak_local_max` and create the watershed map with `watershed`. Fluxes in each segment are then summed. Then intensity statistic is defined as

$$I \equiv I_2/I_1,$$

where I_1 and I_2 are the fluxes of the brightest and the




FIG. 14.—: same as Fig. 2, for deviation (D ; Freeman et al. 2013). There is a small bias in D at $\langle \text{SNR} \rangle < 2$ and $\mathcal{R}_{\text{eff}} < 10$, but the bias is lower than the intrinsic scatter or the typical values expected of disturbed sources ($D > 0.1$), making D a reliable metric overall.

second brightest segments, respectively.

In practice, this metric is extremely sensitive to resolution and its output does not necessarily have the intended meaning, for example for late-type galaxies. A large fraction of images in our sample ($> 50\%$) have intensity of 0, so we do not provide a resolution-depth grid, although it is available online. Fig. 13 shows an example galaxy imaged at 50, 75, and 100 pc/px, with the intensity peaks and watershed contours outlined. At these resolution levels, visually, the same features in the galaxy are present. However in the 50 pc/px case, the watershed algorithm identifies secondary peaks in the central region, while in the other cases it does not, leading to different intensity values despite a relatively small change in resolution. The peaks identified by the algorithm also do not correspond to any visually obvious peaks. The number of tertiary peaks also depends strongly on the resolution of each image, although these do not contribute to the overall intensity calculation.

The crux of the issue here is that SCIKIT-IMAGE simply applies a 3×3 maximum filter to the image and labels any pixel where the input and the max-filtered arrays are equal as a peak. This means, for example, that in a noise-dominated array every 3×3 box will contain a peak. While this issue is partly mitigated by the MID segmentation map, the algorithm still identifies many spurious peaks in the low signal-to-noise regions, and the robustness of peaks found in the centre is hard to determine. SCIKIT-LEARN provides additional parameters that can enforce a minimum flux ratio and/or distance between

adjacent peaks. Setting these parameters to physically meaningful quantities could help alleviate some of the stochasticity of the peak finder, but this investigation is beyond the scope of our paper, and they are not currently used in `statmorph`.

Another, more subtle issue is that watershed segmentation is not the best choice for calculating the total fluxes associated with each component. As seen in Fig. 13, the watershed algorithm essentially draws a boundary between two peaks that is equidistant from both of them. Even a small clump – not visible by eye in our example – sufficiently near the centre of the galaxy forces the watershed to split the source into two halves, leading to a high $I \approx 0.4$. However, for a more physical meaning, most of the flux should be attributed to the main galaxy halo, not the clump. This requires using more sophisticated decomposition methods, such as doing it in frequency space (e.g., Sattari et al. 2023; Kalita et al. 2025b), using the smoothness/substructure (Sec. 3.5.2), or modelling the clumps alongside the main Sérsic profile (e.g., Kalita et al. 2025a). With the current approach, I cannot distinguish between a real bimodality in the light distribution – e.g., two nuclei – and a small clump causing a bifurcation in the watershed map. Adding in the information from other metrics such as multimode can help break this degeneracy.

3.5.5. Deviation

Deviation (D) is the final statistic of the trio, derived directly from the intensity measurement. It is simply the distance between the image centroid and the brightest peak I_1 found earlier, normalized by the “radius” of the MID segmentaion map of area N defined as $\sqrt{\pi/N}$. The peak I_1 typically corresponds to the brightest pixel (nucleus) of a galaxy, while the centroid may be offset due to asymmetries in the light distribution, so D is essentially a measure of asymmetry, similar to $S(G, M_{20})$.

While the intensity watershed map may be unstable, the brightest peak is always identified well in it, so the position of I_1 is reliable. The centroid is also a robust metric (see Sec. 3.1.1), so D is robust to changes in signal-to-noise and resolution as shown in Fig. 14. There is a small positive bias in D measured on low-resolution or shallow images, however, this bias is comparable to intrinsic scatter, and is lower than typical D values expected for disturbed objects ($D > 0.1$). The slight increase in D is due to poorer deblending of foreground sources at low resolutions. For true double nuclei or mergers, D actually decreases at low resolutions where the two brightness peaks blend in and are no longer identified as separate intensity peaks.

4. DISCUSSION & IMPLICATIONS

In this work, we analyzed all morphological metrics measured by the `statmorph` package in terms of their dependence on resolution and depth. In Fig. 15, we offer a concise review of our findings: some metrics depend on resolution, some on signal-to-noise, some on both, and some metrics are robust to changes in either. It is important to consider which metrics you are using based on your science needs, which we roughly split into three broad areas.

Large surveys: large surveys, such as the SDSS, LSST, Roman, UNIONS, or Euclid typically have a homogeneous resolution and depth, so morphological catalogs should be consistent across the survey area, up to variations in seeing and effective exposure time. However, surveys that are carried out over several epochs with different data releases, each release increasing the image depth. Therefore, some parameters – such as A , A_S – will not be consistent between different data releases.

Individual programs / archival data: archival observations, gathered from different PIs with different science goals, by construction have inhomogeneous exposure times and therefore image depth. So parameters such as A and S would not be consistent across observations. On the other hand, in archives of space-based imaging where the resolution is consistent, metrics such as C , G , M_{20} , and A_{RMS} will be robust from one image to another.

4.1. Evolution of bias from $0.5 < z < 3.5$

In both of these modes, the main difficulty arises when studying objects at different cosmic times – whether found in the same survey footprint or in archival observations. Even with a consistent angular scale, the *effective* resolution in terms of galactic radii changes as objects further away appear smaller. Similarly, the same surface brightness limit probes different physical surface brightnesses in $z = 0$ and $z = 2$ objects due to cosmological dimming. For lower-redshift objects ($z \lesssim 0.2$), the surface brightness dimming is negligible while the resolution change is significant, so parameters such as C , G , M_{20} will be affected more. For high-redshift objects ($z > 1$) the inverse is true as the angular diameter dependence on distance flattens. However, if we account for changes in intrinsic luminosity and radii of galaxies at a fixed mass as was done in Whitney et al. (2021) and Yu et al. (2023), due to the rapid size evolution, the effective resolution $\mathcal{R}_{\text{eff}}$ does change as a function of redshift even at $z > 1$ despite the almost constant angular scale.

In Fig. 16, we show the observed changes in six commonly used morphological parameters: C , A , $B(G, M_{20})$, R_p , $R_{0.5}^{\text{Sérsic}}$, and n due to biases caused by the losses

Geometric		
$\mathbf{x}_0$	Centroid	3.1.1
e	Ellipticity	3.1
e_{Sersic}	Sérsic e	3.2.3
θ	Orientation	3.1
θ_{Sersic}	Sérsic θ	3.2.3
R_p	Petrosian R	3.2.1
$R_{20,50,80}$	Isophotal R	3.2.2
$R_{0.5}^{\text{Sersic}}$	Sérsic R	3.2.3
Bulge strength		
n	Sérsic index	3.2.3
C	Concentration	3.3.1
G	Gini index	3.3.2
M_{20}	M_{20}	3.3.3
$B(G, M_{20})$	Bulge strength	3.3.4
Disturbance		
$S(G, M_{20})$	Disturbance	3.4.1
A	Asymmetry	3.4.2
A_O	Outer A	3.4.2
A_S	Shape A	3.4.4
A_{RMS}	RMS A	3.4.3
S	Smoothness	3.5.1
M	Multimode	3.5.3
I	Intensity	3.5.4
D	Deviation	3.5.5
New in this work		
A_X	Isophotal A	3.4.5
St	Substructure	3.5.2




FIG. 15.—: A summary of the parameters studies in this work, roughly broken down into four regions: resolution-dependent, depth-dependent, robust, and unreliable. The table on the left provides references to the sections where each parameter is discussed.

in depth and resolution if the same source is redshifted from $z = 0.5$ to $z = 3$. We do not consider changes in wavelength for this toy example as we assume the filter is changed appropriately to probe the same rest-frame wavelength. Therefore, our analysis in Fig. 16 aims to show the differences in morphological measurements arising *purely from biases*. The evolution of intrinsic properties – stellar structure, dust content, star forming regions, and others – would be seen as trends on top of the ones caused by biases.

For each parameter, we show the changes due to cosmological dimming and angular size change in the top panel. In the bottom panel, we also include an intrinsic evolution in size and luminosity: $R \sim R_0(1+z)^{-0.71}$ (Ormerod et al. 2024) and $L \sim L_0(1+z)$ (Yu et al. 2023). As galaxies of a given mass are intrinsically smaller at higher redshifts, the physical resolution $\mathcal{R}$ in parsecs per pixel does not change much, but the effective resolution in $\mathcal{R}_{\text{eff}}$ does. For each galaxy in our sample, we assume that the baseline measurement is taken with JWST NIRCcam at $z = 0.5$ with resolution 200 pc/px and depth $\mu_0 = 23.5$ mag/arcsec². Then, for a given redshift, we calculated the change in resolution ($\mathcal{R}$, top or $\mathcal{R}_{\text{eff}}$, bottom) and depth (μ_0 , top or $\langle \text{SNR} \rangle$, bottom) arising from cosmological effects (top and bottom) and intrinsic size/luminosity evolution (bottom only). For each galaxy and redshift, we then found

a run from our suite that best matches that resolution and depth. Fig. 16 shows how the median and $\pm 1\sigma$ quantiles of the entire sample change with redshift in black. In red/blue we show the median and $\pm 1\sigma$ quantiles for galaxies with high or low baseline values of a given parameter, respectively, using the following thresholds: $C = 3.5$, $A = 0.2$, $B = 0$, $R_p = 4$, $R_{0.5} = 4$, $n = 2.5$.

As seen in Fig. 16, when size evolution is accounted for, there is significant evolution in C , A , and $B(G, M_{20})$ that arises due to observational biases rather than a real evolution in these parameters. C and B will decrease by 1 each from $z = 0.5$ to $z = 3.5$, while for late-type C remains constant on average and B decreases by 0.5. A decreases by 0.5 on average for disturbed galaxies and 0.1 for undisturbed ones. Galaxies at higher redshifts appear less bulge-dominated and less disturbed than they would at lower redshift due to insufficient resolution and depth. On the other hand, both radii measurements and n are fairly robust across this redshift range, although there is large scatter at high z , which reflects the increasing uncertainty in these measurements.

These biases might explain some of the puzzling results from morphological analyses in the past. On one hand, we expect that major mergers are more frequent at earlier times from the hierarchical structure formation (e.g., Rodríguez-Gomez et al. 2015). On another, repeated analyses of HST and JWST data found that




FIG. 16.—: The change in six common morphological metrics as a result of moving the same galaxies to a higher redshift (**top**), and also accounting for a size and luminosity evolution of galaxies with redshift (**bottom**). The black, pink, and blue lines show the median of the entire sample, and high- p and low- p subsamples for each parameter p , respectively. C , A and $B(G, M_{20})$ parameters suffer a systematic bias that needs to be corrected for when looking at structural evolution across redshifts. R_p , $R_{0.5}^{\text{Sersic}}$, and n are largely unbiased but the uncertainty increases as the effective resolution and depth degrade. We assume a size evolution $R \sim R_0(1+z)^{-0.71}$ (Ormerod et al. 2024) and luminosity evolution $L \sim L_0(1+z)$ (Yu et al. 2023).

asymmetry is, on average, low and constant with redshift (Ferreira et al. 2023; Yao et al. 2023; Ren et al. 2024; Costantin et al. 2025), while Mortlock et al. (2013) and Ferreira et al. (2022b, 2023) find that up to 40% of galaxies are visually peculiar. Considering the strong bias in asymmetry we see in Fig. 16, we expect that the *true* asymmetry increases significantly with redshift, but this signal is washed away in observational effects. Accounting for the loss in image quality for distant galaxies, Sazonova et al. (in prep.) find a strong increase in non-parametric disturbance between $z=0.5$ and $z=3$.

Similarly, several studies find that concentration and bulge strength are lower at high z (Yao et al. 2023; Ren et al. 2024; Costantin et al. 2025), while the Sérsic index does not evolve significantly (Ferreira et al. 2023; Sun et al. 2024; Martorano et al. 2025). Concentration and Sérsic indices are correlated by construction (Trujillo et al. 2001), so this result appears puzzling at first – but Fig. 16 shows that while n is not affected by imaging

biases, C is. The decrease in C by ~ 1 over the $0.5 < z < 3$ redshift range for early-type galaxies is completely consistent with the findings of Ren et al. (2024), who showed that C decreases by 1 among the quiescent massive galaxies, but not star-forming, and Law et al. (2012), who found that C is constant with z in star-forming galaxies. Finally, Lotz et al. (2008) used $B(G, M_{20})$ measured with HST data at $0.2 < 1.2$ to split galaxies into early- and late-type, and found that the fraction of bulge-dominated (E, S0, Sa) galaxies increases from 21% at $z=1.2$ to 44% at $z=0.3$. However, this redshift range covers a significant change in the effective resolution, which in our tests leads to $\Delta B \sim 1$. Accounting for this, the evolution of the E/S0/Sa fraction could be much weaker, which agrees with the recent results from Kartaltepe et al. (2023) that 30% of visually classified galaxies at $z < 5$ have a spheroid. Our results indicate that the apparent decline in bulge strength at $z > 3$ may largely arise from observational biases, rather than true structural evolution.

One way to study parameters such as CAS, G, and M_{20} across cosmic time is to measure the observed evolution and subtract off the expected bias induced from changing image quality – either obtained from redshifting sources (e.g., Whitney et al. 2021) or by fitting analytical solutions in terms of $\langle \text{SNR} \rangle$ and $\mathcal{R}$ (e.g., Pawlik et al. 2018). We fitted for this dependence using symbolic regression, and provided correction functions for each parameter (Appendix D). However, our correction assumes that the PSF is a 2 pixel-wide Moffat – for different PSF sizes, it is best to recalibrate the correction terms.

Another approach is to artificially degrade the imaging across the sample such that the physical scale and surface brightness limits are matched for each galaxy, regardless of the instrument taken or the redshift of the source. Sazonova et al. (in prep.) took this approach for galaxies in JWST JADES survey, and found a strong evolution between $0.5 < z < 3$, where high-redshift galaxies have distinctly more disturbed structures.

Finally, machine learning approaches can be used to remedy the imaging effects – for example, applying meaningful augmentations that represent the loss of resolution and signal-to-noise in the training step can produce a network that is invariant to observational biases (e.g., Vega-Ferrero et al. 2024). We plan to use our augmentation dataset to train a deep learning network that is invariant to the changes in signal-to-noise and resolution considered in this work, and then train this model on data from different telescopes and observations to be invariant to realistic imaging artifacts.

4.2. Multi-wavelength morphology

Throughout this work we focused only on galaxy structure in the rest-frame I -band, which traces intermediate stellar populations and dust. In our redshift analysis, we assumed that the observer is able to switch bands to match the rest-frame wavelength. If instead a survey is designed in a single filter, then higher-redshift sources will move further to the rest-frame UV, and the observations will trace younger stellar populations. When observed in the UV, even early-type galaxies can have late-type or clumpy morphology (Windhorst et al. 2002; Mager et al. 2018). Rutkowski et al. (2025) in particular found that 15% of $0.5 < z < 1.5$ quiescent galaxies still have residual star formation and so would appear to have structures of late-type galaxies when viewed in the UV, where the older stellar populations are not visible. It may be possible to correct for these effects with “morphological k -corrections” (Windhorst et al. 2002), but this would require *assuming* some relationship between UV and optical structures, and so is not advisable for populations of high-redshift or irregular galaxies where this relationship is not known a priori.

A more common approach is to design high-redshift

surveys using imaging from multiple bands, such that each band traces the same rest-frame filter in a given redshift range. This should also be done carefully, since one would need to account both for the change in the central wavelength and the bandwidth, so that redder (tracing higher-redshift) bands are wider than the bluer (tracing lower-redshift) ones. For example, the filters on the *HST* and *JWST* telescopes are designed with this in mind. Finally, when analyzing populations that include quiescent galaxies, it may be more important to not match the central wavelength precisely, but to ensure that the filter does not straddle the 4000Å break since the contribution blueward of the break will be much weaker.

However, while we must be careful to compare like rest-frame images, in general it is valuable to consider the structure of galaxies in different wavelengths, from UV all the way to radio. Different wavelengths of light are produced different emission mechanisms and hence correspond to different galaxy components: UV, optical, and near-IR trace young, intermediate, and old stellar populations respectively. Dust attenuates light to varying extents, so a comparison between these wavelengths can also constrain dust properties. Mid- and far-IR highlight dust emission itself as well as emission from an AGN, if present. Finally, radio observations can trace atomic and molecular hydrogen gas content of the galaxy at different densities, and IFU observations can do so for ionized gas. Comparing how galaxy structure changes across all of these components by measuring morphology can be an extremely powerful tool to tie together various physical processes and their role in redistributing the gas, dust, and stellar content of a galaxy. With this approach we can start to constrain quenching processes more precisely: for example, a galaxy with asymmetries in the UV and IR is likely a post-merger, while disturbances that are only visible in the UV are likely caused by star formation. Mergers also induce changing signatures on the timescales of $\sim 1\text{Gyr}$, which can be seen as changes on the CAS/G/ M_{20} plane (Snyder et al. 2015b; Taylor-Mager et al. 2007).

Given the results of our work, care must be especially taken when comparing structural parameters in different wavelength regimes, since multi-wavelength observations typically have varying resolution and signal-to-noise levels. While we did not explicitly run the tests for varying wavelengths, our results should be applicable to this use case as well, since the main difference between UV and IR imaging is still simply the depth and resolution. When biases arising from imaging quality are accounted for, any changes in the measured structural parameters can be used to detect changes in actual substructures. For example, in Sazonova et al. (2021) we used this approach to detect disturbances in dust substructure in post-starburst galaxies by comparing *HST* B , I , and H -

band imaging. In a similar vein, GALFITM incorporates different wavelength observations to study how the Sérsic profile changes from one band to another, which can be interpreted as changes in structure of different stellar populations (Häußler et al. 2022).

The final consideration is that extending this analysis to multiple wavelengths will greatly increase the dimensionality of the problem – for 30 or more morphological parameters and 10 bands, each galaxy would have 300 structural measurements. This goes far beyond diagnostic diagrams traditionally used in morphological studies (e.g., Conselice 2003; Lotz et al. 2004). These large datasets would be perfectly suited for machine learning and dimensionality reduction algorithm, where the entire parameter space is compressed to a few important features which are then used to classify galaxies or probe their physical processes. For example, Mukundan et al. (2024) and Holwerda et al. (2025) used k-nearest neighbours classification on the space of structural measurements to classify galaxies into different Hubble types and irregulars. Aguilar-Argüello et al. (2025) used XGBoost for the same task, and provided feature importances for each morphological metric they considered. In Sazonova et al. (in prep.), we use the Uniform Manifold Approximation & Projection (UMAP) technique to reduce the dimensionality of the space of **statmorph** outputs to two important new variables, and study the evolution of galaxies in the UMAP phase space. These studies still focused on single-band imaging, but they offer a promising way towards tackling a larger-dimensional dataset of multiwavelength observations.

4.3. Final remarks

This work is intended to prepare for the morphological analysis of galaxies in the Rubin LSST. The changes to parameters such as smoothness/substructure, Sérsic index, and shape asymmetry discussed here will be incorporated in a **statmorph-lsst** public package, which will be used to construct catalogs in the upcoming Rubin data releases. This package is publicly available on a GitHub repository¹⁶ and contributions from other researchers interested in this effort are welcome. Finally, we also make available the entire suite of 64,000 augmented images we used for these tests on Zenodo, and we developing a tool for researchers to test the performance of custom metrics and construct their own corrections. The tool will be linked to the **statmorph-lsst** repository.

5. SUMMARY

Using a sample of 189 local galaxies, we created a dataset of 64,000 images at varying physical resolutions and surface brightness limits. This dataset can enable

a wide range of further analyses and machine learning applications, and it is available online for public use on Zenodo¹⁷.

We then investigated how every morphological parameter measured by **statmorph** depends on the signal-to-noise and resolution of the galaxy. Our results are broadly summarized in Fig. 15:

1. Geometrical parameters – ellipticity, orientation, Petrosian radius R_p , and $R_{0.5}^{\text{Sérsic}}$ – are robust to changes in image properties and can be used consistently across different surveys and redshifts (Sec. 3.1 and 3.2).
2. Bulge strength parameters – Gini, M_{20} , C – depend on the effective resolution $\mathcal{R}_{\text{eff}}$, defined as the ratio of R_p to pixel scale. These will be biased when resolution changes, either from the changing image quality, or from cosmological effects (Sec. 3.3). This means that high-resolution imaging, such as that from *Euclid* or JWST, is required to accurately measure these parameters at $z > 1$.
3. Disturbance parameters – A , A_S , substructure St – depend on signal-to-noise and thus will be inconsistent across different surveys, or for objects at different redshifts due to the cosmological surface brightness dimming. A_{RMS} , introduced in Sazonova et al. (2024), does not depend on signal-to-noise but is sensitive to resolution. (Sec. 3.4 and 3.5).
4. Smoothness and intensity are unreliable parameters: first due to the strong noise contribution, and second due to its watershed definition that bifurcates the galaxy flux. Multimode is robust to changes in $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$ up to a point and detects double nuclei well, but the threshold where it becomes unreliable varies from one galaxy to another and so is difficult to generalize. Deviation, although defined using Intensity, is a robust metric for a wide range of $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$ (Sec. 3.5).
5. GALFIT Sérsic fitting results generally do not suffer from a strong bias. However, we quantify an intrinsic uncertainty in Sérsic parameters caused by the degeneracies in the fit to be on the order of 20% for $R_{0.5}^{\text{Sérsic}}$ and up to 40% for n (Appendix E). The default **statmorph** implementation of the Sérsic fit will also lead to an underestimate in n compared to other libraries.

We applied our analysis to see how would the morphology of galaxies appear to evolve from $z = 0.5$ to $z = 3.5$ due to observational biases alone, and found that the observed decrease in the numbers of centrally concentrated

¹⁶ **statmorph-lsst** GitHub repository

¹⁷ Dataset DOI: 10.5281/zenodo.17585608

galaxies can be explained by the bias in the concentration index. We also found that the true number of disturbed or asymmetric galaxies is likely even higher than what is currently reported.

We implemented several changes in the new `statmorph-lsst` package, which will be used in the analysis of Rubin LSST data:

1. Introduced a new isophotal asymmetry, A_X , similar to A_S from Pawlik et al. (2016). This is an asymmetry of the X mag/arcsec² isophote. Since it is defined as a function of surface brightness, it is not dependent noise as long as the isophote is above the surface brightness limit, and has a mild resolution dependence when X is low. One advantage of A_X is that it probes physical features of a particular luminosity (and by proxy, mass) surface density.
2. Introduced a new substructure metric St , defined similarly to S , but removing the noise contribution by requiring the substructure clumps to be connected. This metric allows detecting high-contrast features such as clumps, spiral regions, and bars, and together with multi-wavelength data can be used to further dissect properties of star-forming regions – their flux ratios, colours, metallicities, and more.
3. Added the functionality to first degrade the image to the desired physical resolution and depth, and then compute the parameters to ensure consistency across different masses and redshifts.

This `statmorph-lsst` package is available on GitHub and contributions are welcome. In addition to this package, we published our suite of galaxy augmentations on Zenodo, and will release a web-app that allows testing

new, custom functions before they are implemented in `statmorph-lsst`. We encourage other morphology experts to try our suite and contribute to our package and the webapp!

ACKNOWLEDGEMENTS

We thank the University of Waterloo for their support of the Canada Rubin Fellowship program. ES thanks Roan Haggard, Ana Ennis, and the members of the unofficial “galaxies office” for many fruitful scientific discussions throughout this project. CRM acknowledges support from an Ontario Graduate Scholarship. D.D acknowledges support from the NAWA (Bekker grant BPN /BEK/2024/1/00029/DEC/1). HMHT acknowledges financial support from DGAPA-PAPIIT project AG101725 and CONAHCYT project CF-2023-G-1052. RAW acknowledges support from NASA JWST Interdisciplinary Scientist grants NAG5-12460, NNX14AN10G and 80NSSC18K0200 from GSFC.

Data Access: [Zenodo DOI: 10.5281/zenodo.17585608](https://zenodo.org/doi/10.5281/zenodo.17585608), `statmorph-lsst`

Facilities: HST (WFC3, ACS, WFPC2), MAST, HLA, NED

Software: `statmorph` (Rodriguez-Gomez et al. 2019), `photutils` (Bradley et al. 2020), `astropy` (Astropy Collaboration et al. 2013, 2018; Collaboration et al. 2022), `GALSIM` (Rowe et al. 2015), `GALFIT` (Peng et al. 2002, 2010a), `Scikit-Image` (van der Walt et al. 2014), `Scikit-Learn` (Pedregosa et al. 2011), `Matplotlib` (Hunter 2007), `pandas` (Reback et al. 2020), `NumPy` (Harris et al. 2020), `SciPy` (Virtanen et al. 2020), `Julia`, `SymbolicRegression.jl` (Cranmer 2023)

REFERENCES

- Abraham R. G., et al., 2003, *ApJ*, 588, 218
Aguilar-Argüello G., et al., 2025, *MNRAS*, 537, 876
Allen P. D., et al., 2006, *MNRAS*, 371, 2
Andrae R., et al., 2011, *MNRAS*, 411, 385
Andredakis Y. C., et al., 1995, *MNRAS*, 275, 874
Argyle J. J., et al., 2018, *MNRAS*, 479, 3076
Arnouts S., et al., 2007, *A&A*, 476, 137
Astropy Collaboration et al., 2013, *A&A*, 558, A33
Astropy Collaboration et al., 2018, *AJ*, 156, 123
Bellhouse C., et al., 2022, *ApJ*, 937, 18
Bershady M. A., et al., 2000, *AJ*, 119, 2645
Bertin E., et al., 2020, in Pizzo R., et al., eds, *Astronomical Society of the Pacific Conference Series Vol. 527*, *Astronomical Data Analysis Software and Systems XXIX*. p. 461, <https://ui.adsabs.harvard.edu/abs/2020ASPC..527..461B>
Birrer S., et al., 2021, *The Journal of Open Source Software*, 6, 3283
Bonfini P., 2014, *PASP*, 126, 935
Bottrell C., et al., 2019, *MNRAS*, 490, 5390
Bradley L., et al., 2020, `astropy/photutils`: 1.0.1, [doi:10.5281/zenodo.596036](https://doi.org/10.5281/zenodo.596036), <https://ui.adsabs.harvard.edu/abs/2020zndo....596036B>
Bruce V. A., et al., 2014, *MNRAS*, 444, 1001
Casura S., et al., 2022, *MNRAS*, 516, 942
Chamba N., et al., 2022, *A&A*, 667, A87
Collaboration A., et al., 2022, *ApJ*, 935, 167
Conselice C. J., 2003, *ApJS*, 147, 1
Conselice C. J., et al., 2000, *ApJ*, 529, 886
Cook R. H. W., et al., 2025, *MNRAS*, 539, 2829
Costantin L., et al., 2025, *A&A*, 699, A360
Cranmer M., 2023, *arXiv e-prints*, p. arXiv:2305.01582
Deg N., et al., 2023, *MNRAS*, 523, 4340
Dimauro P., et al., 2018, *MNRAS*, 478, 5410
Ding X., et al., 2020, *The Astrophysical Journal*, 888, 37
Djorgovski S., Davis M., 1987, *ApJ*, 313, 59
Dressler A., 1980, *ApJ*, 236, 351
Dreyer J. L. E., 1888, *MmRAS*, 49, 1
Dreyer J. L. E., 1910, *MmRAS*, 59, 105
Espejo Salcedo J. M., et al., 2025, *A&A*, 700, A42
Ferreira L., et al., 2022a, *ApJ*, 931, 34
Ferreira L., et al., 2022b, *ApJ*, 938, L2
Ferreira L., et al., 2023, *ApJ*, 955, 94
Foster L. M., et al., 2025, *ApJ*, 982, 120
Freeman P. E., et al., 2013, *MNRAS*, 434, 282
Gadotti D. A., 2009, *MNRAS*, 393, 1531
Galametz A., et al., 2013, *ApJS*, 206, 10
Graham A. W., 1998, *MNRAS*, 295, 933
Graham A. W., 2001, *AJ*, 121, 820

- Graham A. W., Driver S. P., 2005, *PASA*, 22, 118
- Graham A., et al., 1996, *ApJ*, 465, 534
- Graham A. W., et al., 2001, *AJ*, 122, 1707
- Graham A. W., et al., 2005, *AJ*, 130, 1535
- Harris C. R., et al., 2020, *Nature*, 585, 357
- Häukler B., et al., 2022, *A&A*, 664, A92
- Hogg D. W., et al., 2002, *arXiv e-prints*, pp astro-ph/0210394
- Holden B. P., et al., 2007, *ApJ*, 670, 190
- Holwerda B. W., et al., 2011, *MNRAS*, 416, 2401
- Holwerda B. W., et al., 2025, *PASA*, 42, e028
- Hunter J. D., 2007, *Computing in Science & Engineering*, 9, 90
- Hunter D. A., Elmegreen B. G., 2006, *ApJS*, 162, 49
- Kalita B. S., et al., 2025a, *MNRAS*, 536, 3090
- Kalita B. S., et al., 2025b, *MNRAS*, 537, 402
- Kartaltepe J. S., et al., 2023, *ApJ*, 946, L15
- Kent S. M., 1985, *ApJS*, 59, 115
- Krist J. E., et al., 2011, 20 years of Hubble Space Telescope optical modeling using Tiny Tim. p. 81270J, doi:10.1117/12.892762, <https://ui.adsabs.harvard.edu/abs/2011SPIE.8127E..0JK>
- Lange R., et al., 2016, *MNRAS*, 462, 1470
- Lanyon-Foster M. M., et al., 2012, *MNRAS*, 424, 1852
- Law D. R., et al., 2012, *ApJ*, 745, 85
- Lima-Dias C., et al., 2024, *MNRAS*, 527, 5792
- Lisker T., 2008, *ApJS*, 179, 319
- Lotz J. M., et al., 2004, *AJ*, 128, 163
- Lotz J. M., et al., 2008, *MNRAS*, 391, 1137
- MacArthur L. A., et al., 2003, *ApJ*, 582, 689
- Mager V. A., et al., 2018, *ApJ*, 864, 123
- Martin G., et al., 2025, *MNRAS*, 541, 1831
- Martorano M., et al., 2025, *A&A*, 694, A76
- Matthee J., et al., 2024, *ApJ*, 963, 129
- Mendel J. T., et al., 2014, *ApJS*, 210, 3
- Morgan C. R., et al., 2024, *A&A*, 691, A20
- Morishita T., et al., 2024, *ApJ*, 963, 9
- Mortlock A., et al., 2013, *MNRAS*, 433, 1185
- Mukundan K., et al., 2024, *MNRAS*, 533, 292
- Nedkova K. V., et al., 2024, *MNRAS*, 532, 3747
- Nelson D., et al., 2019, *Computational Astrophysics and Cosmology*, 6, 2
- Okamura S., et al., 1984, *ApJ*, 280, 7
- Ono Y., et al., 2025, *arXiv e-prints*, p. arXiv:2502.08885
- Ormerod K., et al., 2024, *MNRAS*, 527, 6110
- Pandya V., et al., 2024, *ApJ*, 963, 54
- Pasha I., Miller T. B., 2023, *The Journal of Open Source Software*, 8, 5703
- Pawlik M. M., et al., 2016, *MNRAS*, 456, 3032
- Pawlik M. M., et al., 2018, *MNRAS*, 477, 1708
- Pedregosa F., et al., 2011, *Journal of Machine Learning Research*, 12, 2825
- Peng C. Y., et al., 2002, *AJ*, 124, 266
- Peng C. Y., et al., 2010a, *AJ*, 139, 2097
- Peng Y.-j., et al., 2010b, *ApJ*, 721, 193
- Pérez-González P. G., et al., 2023, *ApJ*, 946, L16
- Peth M. A., et al., 2016, *MNRAS*, 458, 963
- Petrosian V., 1976, *ApJ*, 210, L53
- Petty S. M., et al., 2014, *AJ*, 148, 111
- Planck Collaboration et al., 2020, *A&A*, 641, A6
- Poggianti B. M., et al., 2013, *ApJ*, 777, 125
- Reback J., et al., 2020, pandas-dev/pandas: Pandas 1.1.3, doi:10.5281/ZENODO.4067057
- Ren J., et al., 2024, *ApJ*, 969, 4
- Ren J., et al., 2025, *ApJ*, 982, 200
- Roberts I. D., Parker L. C., 2020, *MNRAS*, 495, 554
- Robotham A. S. G., et al., 2017, *MNRAS*, 466, 1513
- Rodriguez-Gomez V., et al., 2015, *MNRAS*, 449, 49
- Rodriguez-Gomez V., et al., 2019, *MNRAS*, 483, 4140
- Rowe B. T. P., et al., 2015, *Astronomy and Computing*, 10, 121
- Rutkowski M. J., et al., 2013, *AJ*, 146, 11
- Rutkowski M. J., et al., 2025, *ApJ*, 983, L32
- Ryden B. S., et al., 1999, *ApJ*, 517, 650
- Salvador D., et al., 2024, *A&A*, 684, A166
- Sattari Z., et al., 2023, *ApJ*, 951, 147
- Sazonova E., et al., 2020, *ApJ*, 899, 85
- Sazonova E., et al., 2021, *ApJ*, 919, 134
- Sazonova E., et al., 2024, *The Open Journal of Astrophysics*, 7, 77
- Schade D., et al., 1995, *ApJ*, 451, L1
- Schawinski K., et al., 2014, *MNRAS*, 440, 889
- Sérsic J. L., 1963, Boletín de la Asociacion Argentina de Astronomia La Plata Argentina, 6, 41
- Shen S., et al., 2003, *MNRAS*, 343, 978
- Shibuya T., et al., 2016, *ApJ*, 821, 72
- Simard L., et al., 2011, *ApJS*, 196, 11
- Sinnott R. W., 1988, NGC 2000.0: The Complete New General Catalogue and Index Catalogues of Nebulae and Star Clusters by J. L. E. Dreyer. <https://ui.adsabs.harvard.edu/abs/1988cngc.book.....S>
- Snyder G. F., et al., 2015a, *MNRAS*, 451, 4290
- Snyder G. F., et al., 2015b, *MNRAS*, 454, 1886
- Steinicke W., 2010, Observing and Cataloguing Nebulae and Star Clusters. <https://ui.adsabs.harvard.edu/abs/2010ocns.book.....S>
- Stone C. J., et al., 2023, *MNRAS*, 525, 6377
- Strateva I., et al., 2001, *AJ*, 122, 1861
- Sun W., et al., 2024, *ApJ*, 960, 104
- Taylor-Mager V. A., et al., 2007, *ApJ*, 659, 162
- Thorp M. D., et al., 2021, *MNRAS*, 507, 886
- Trujillo I., et al., 2001, *MNRAS*, 326, 869
- Vázquez-Mata J. A., et al., 2025, *MNRAS*, 544, 1056
- Vega-Ferrero J., et al., 2024, *ApJ*, 961, 51
- Virtanen P., et al., 2020, *Nature Methods*, 17, 261
- Wang J., et al., 2016, *MNRAS*, 460, 2143
- Wang J.-H., et al., 2024, *A&A*, 686, A100
- Watkins A. E., et al., 2025, *MNRAS*, 537, 3499
- Whalen K. E., et al., 2025, *arXiv e-prints*, p. arXiv:2509.21236
- Whitmore B. C., et al., 2016, *AJ*, 151, 134
- Whitney A., et al., 2021, *ApJ*, 919, 139
- Windhorst R. A., et al., 2002, *ApJS*, 143, 113
- Yanagawa A., et al., 2025, *PASJ*
- Yao Y., et al., 2023, *ApJ*, 954, 113
- Yeom B.-S., et al., 2017, *Journal of Astronomy and Space Sciences*, 34, 183
- Yu S.-Y., et al., 2023, *A&A*, 676, A74
- de Vaucouleurs G., 1948, Annales d'Astrophysique, 11, 247
- de Vaucouleurs G., 1977, in Tinsley B. M., Larson D. Campbell R. B. G., eds, Evolution of Galaxies and Stellar Populations. p. 43, <https://ui.adsabs.harvard.edu/abs/1977egsp.conf....43D>
- van der Walt S., et al., 2014, *PeerJ*, 2, e453
- van der Wel A., et al., 2014, *ApJ*, 788, 28

200 Mpc. This adds some lower-mass galaxies to our sample that were not originally in the RNGC/IC.

To demonstrate that we are able to probe a range of morphologies, we show the distributions of some of the parameters studies in this paper, obtained from our “baseline” images (100 pc/px and $\mu_0 = 23.5$ mag/arcsec²) in Fig. A1.

In particular, we show the Sérsic half-light radius as a function of F814W (*I*-band) absolute magnitude (upper left), the Gini- M_{20} plane (upper right), Sérsic n and concentration (bottom left), and two asymmetry measures (bottom right).

We use the size and absolute magnitude as a proxy for the size-mass relation since we do not have mass measurements for some galaxy in our sample. However, it is clear that we probe a wide range of galaxy sizes and mag-

APPENDICES

A. THE DIVERSITY OF MORPHOLOGIES IN THE RNGC/IC SAMPLE

Since the NGC and the IC catalogues are some of the earliest large compilations of extragalactic sources, they are biased towards galaxies that are nearby and bright. Because of this, the RNGC/IC catalog does not contain, for example, nearby dwarf galaxies that were not detected at the time. Due to this limitation we cannot analyze the behaviour of morphological metrics on the faintest dwarf galaxies, but we attempt to remedy this bias by inspecting the *HST* tiles for any nearby sources that also have a spectroscopic redshift placing them within




FIG. A1.—: The distribution of several structural measurements in our sample of 189 RNGC/IC and associated galaxies: Sérsic $R_{0.5}$ versus absolute F814W magnitude (top left), Gini- M_{20} (top right), concentration-Sérsic n (bottom left), and for two asymmetry metrics (bottom right). The galaxies analyzed in this work sample the distribution of different morphologies along the Hubble sequence well, including both early-type, late-type galaxies and mergers with a range of absolute magnitudes and sizes.

nitudes, spanning several orders of absolute magnitudes and physical sizes. In the Gini- M_{20} diagram, the galaxies sample the entire parameter space between early-type and Sa, late-type, and irregular galaxies. We have several merging objects with high asymmetry values, and a good sample of concentration indices. Therefore we are confident that we can appropriately sample the dynamic range of galaxy morphologies, at least outside the dwarf regime.

B. BACKGROUND ESTIMATION

To estimate the background of each image, we use an improved iterative sigma-clipping technique, which we demonstrate in Fig. B1 below. An example galaxy NGC 1084, which spans most of the *HST* chip, is shown in a log scaling in the first panel. In the second panel, we show the same galaxy with a linear scale centered on the background median, which highlights the extended low surface brightness envelope of this source.

Sigma-clipping on its own is already iterative, where all pixels above 3σ are masked, the standard deviation is re-estimated, and so on until convergence. The third panel in Fig. B1 shows the pixels that are masked using the default sigma-clipping implementation. However, the

source light dominates most of the pixels and so it is not fully clipped. The halo of the galaxy is not fully masked in this way, which leads us to underestimate the surface brightness limit of the image.

Instead, we run source detection to mask out any sources and estimate the background this way. We smooth the image by a 5-pixel Gaussian kernel, and detect any sources above a 1σ threshold with a minimum area of 1% of the image size. For each detection, we expand the source mask by 10% of its size. Finally, we mask out all the sources and re-run sigma-clipping to get new background statistics. With the new σ estimate, we repeat the source detection step, and continue doing so until σ converges to within 5% or 30 iterations have elapsed. The last panel of Fig. B1 shows the source mask obtained from our procedure: it captures most of the diffuse light of the galaxy, and provides us with a more accurate σ and hence surface brightness limit estimate.

C. SEGMENTATION

Since the sources in our sample are extremely local and have large angular extents, we needed to implement a careful segmentation and masking routine to remove as many foreground stars and background galaxies as possible. To do this, we follow a hot-and-cold approach, similar to Sazonova et al. (2021) and Morgan et al. (2024). Here we describe our routine, which is available in the `statmorph-lsst/test_suite` package. Fig. C1 shows the steps of the algorithm for one galaxy, where segmentation was particularly challenging.

To ensure high-quality masks, we perform the segmentation routine iteratively by visually inspecting the image and changing parameters until a satisfactory map is produced. Parameters that can be varied in our routine, with their typical values, are listed in Table C1.

First, we load the image and interpolate over any masked or missing pixels other within the detector footprint. This is necessary because the ACS detector has a chip gap in the middle, often passing through the galaxy of interest, and simply masking the chip gap causes the segmentation map to split into two distinct regions, effectively masking out one half of the galaxy. Then, for large images (over 300 pixels), we apply a Gaussian smoothing filter with $\sigma = 2$ to average over the noise spikes. For every task involving source detection, we choose 4-connected regions rather than 8 commonly used elsewhere (e.g., in Source Extractor Arnouts et al. 2007). We find that a stronger adjacency requirement means fewer spurious noise regions are marked as sources. With this in mind, we proceed with the hot-cold segmentation.

1. **Hot mode:** we detect all bright sources above a high `hot_threshold` of the background RMS with




FIG. B1.—: A comparison of background estimation via sigma-clipping (**third** panel) and our iterative procedure (**fourth** panel) for an example galaxy shown in log (**first**) and linear (**second**) scales. Simple sigma-clipping fails to mask the diffuse light of the source since it dominates the chip. Using iterations of sigma-clipping and source detection better masks the faint halo, allowing us to get a more accurate measurement of the background noise level and hence the surface brightness limit.




FIG. C1.—: An example of our “hot-cold” segmentation approach to the image of NGC 4874, the second brightest galaxy of the Coma cluster (**first** panel). NGC 4874 is surrounded by other galaxies, which need to be carefully deblended. In the “hot” step, all sources with SNR higher than 30 are detected and masked (**second** panel). Then, we detect all sources above 1.5 SNR in the “cold” step and deblend the central galaxy into subcomponents (**third** panel). Finally we mask any subcomponent with a greater than 20% overlap with the “hot” mask (**fourth** panel and pink contours in the first panel).

TABLE C1: Segmentation parameters

Param.	Description	Typical value	Range
<code>hot_thresh</code>	Signal-to-noise threshold in the “hot” mode	30	5~1000
<code>cold_thresh</code>	Signal-to-noise threshold in the “cold” mode	1.5	1.5~5
<code>min_area</code>	Minimum area of a detected segment, in pixels	5	1~5
<code>contrast</code>	Contrast needed to deblend adjacent regions in the central source	0.001	0.001~1
<code>overlap_thresh</code>	Minimum overlap between a deblended “cold” segment and the “hot” map to mask out that segment	0.2	0.05~0.4

some `min_area`. This step usually selects the central source and other bright sources: foreground stars and other galaxies. We add all sources except the central to a mask¹⁸. We then grow the mask based on the area of each detected “hot” source, so that larger detections are grown by a larger

amount, since they typically have diffuse light extending further.

2. **Cold mode:** next, we detect any sources with a lower `cold_threshold` on signal-to-noise and the same `min_area`. This selects the main source as well as any other sources significantly present in the image.

3. **Deblending:** the central “cold” source typically

¹⁸ This could, in principle, also mask out bright star clusters, but we opted to be cautious and mark any unresolved point source as we have no way of distinguishing them from foreground stars.

consists of our target galaxy and several embedded contaminants. The brightest peaks of contaminants are already masked in the “hot” step, but their diffuse light can still bleed into the main source. We therefore deblend the central source with a high **contrast**. We iterate through every deblended segment and calculate the fraction of that segment’s area that overlaps with a “hot” mask. If there is sufficient overlap, set by **overlap_thresh**, then the segment is contaminated and we mask it.

4. **Mask contaminants:** we now grow the constructed mask to ensure the faint light from the contaminants is masked as well. We grow each masked segment proportionally to its size using a single convolution step. By default, **detect_sources** returns a map where the pixels corresponding to a source have a value equal to its label. Instead, we assign the pixel values to be the area of each segment. We then convolve this map with a Gaussian filter, using a size that is 5% of the image. Finally, we assign all pixels with a flat value over 0.05 to a mask. Since different segments have a different area, then after a convolution with a Gaussian, larger segments extent further until they reach this threshold, and therefore the mask grows more. Finally, we return this mask, and the **cold mode** segmentation map where only the central source is left.

D. SYMBOLIC REGRESSION.JL

We use **SymbolicRegression.jl** package (Cranmer 2023) to estimate the correction functions between our observed and baseline parameters. The advantage of symbolic regression is that it optimizes both the form of the function itself, and any free parameters simultaneously. It works via a “genetic modification” scheme: it builds a tree that consists of the input variables, free parameters, and operators. At each iteration, the algorithm can either change the tree (by adding, removing, or swapping nodes) or numerically optimize all free parameters. As the iteration progresses, the tree grows, representing more complex functions. At the end of each iteration, a “hall of fame” of the functions with the lowest loss is saved and carried over to the next iteration.

The objective of the **SymbolicRegression** optimizer is to minimize the loss of the training set, but it also computes the complexity of each proposed function, where the complexity is a number of variables, parameters, and operators used. For example, $f(x) = 3 \times \exp(2 \times x)$ would have a complexity 6. The hall of fame stores best-performing functions with different complexities, from lowest to highest, and calculates a *score*: the relative improvement in loss compares to the function with the next

lowest complexity. The user then uses the loss and the score in tandem to select the best-fitting function. The user can set up the **SymbolicRegression** run to fit their science purpose, and many steps of the algorithm can be changed.

One of the most important steps is the selection of the function space to search. Most morphological parameters are bound between zero and some value by definition, and many converge to the baseline value. Looking at the difference between baseline and observed measurements as a function of resolution or depth, we often found that a sigmoid described the relationship well. To capture these behaviours, we fit a function of the form

$$p = p_{\text{base}} \times f(x, y) + g(x, y) \quad (\text{D1})$$

where x and y are the measures of resolution and depth selected with the mutual information score. This function captures possible variations in the measurement that depend on the baseline value and any linear offsets. We allowed f and g to consist of basic algebraic operations ($+$, $-$, $\times$, $\div$) as well as $\exp$, $\tanh$, $\sqrt{\cdot}$, and $\log$, to capture sigmoid behaviour. We set up constraints so that $\exp$, $\tanh$, and $\sqrt{\cdot}$ cannot be nested within each other or inside the $\log$. Finally, we set the complexity of $+$ and $\times$ to 0 to allow the code to normalize the resolution or depth metrics with a lower penalty. With these rules, $f(x) = 3 \times \exp(2 \times x)$ would have a complexity of 4.

We chose the loss function to be an L1 loss rather than an L2 (root-mean-square) loss, since it is more robust to outliers and so provided us with more stable fits. To train the model, we selected 70 random galaxies from the sample with the most obvious outliers removed. We trained the model for 1000 iterations. Table D1 shows all of the other parameters of **SymbolicRegression** needed to reproduce our results. At the end of the training cycle, we tested the highest-complexity model scoring over 10^{-2} . Note that complexity is highly penalized in the score, so this approach still always preferred low-complexity models. We analyzed the distribution of corrected measurements for the overall sample as well as low-resolution and low-SNR subsets (as in Fig. 4). If the chosen model did not debias these subsamples, we tested the next most complex model, until a satisfactory correction was found for poor image quality.

Tab. D2 shows the derived corrections and the 16/50/84th percentiles of the error distribution $p - p_{\text{base}}$ for each parameter. Note that for some parameters, the intrinsic scatter in the distribution dominated any systematic biases, so we did not fit a correction as this would just amplify the scatter.

TABLE D1: `SymbolicRegression.jl` parameters

Parameter	Description	Value
N	# of galaxies in the training set	70
f	Functional form of the fit	$p_{\text{pred}} = p_{\text{base}} \times f(x, y) + g(x, y)$
<code>elementwise_loss</code>	Loss function during optimization, $\mathcal{L}(p_{\text{pred}}, p_{\text{true}})$	<code>L1DistLoss</code> $\equiv p_{\text{pred}} - p_{\text{true}} $
<code>binary_operators</code>	Operators acting on two arguments, $a \otimes b$	$+, -, \times, \div$
<code>unary_operators</code>	Operators acting on a single argument, $F(a)$	$\exp, \tanh, \log_{10}, 1/x, \sqrt{\cdot}$
<code>maxsize</code>	Maximum model complexity (see text)	25
<code>complexity_of_operators</code>	See text	0 for $\times$ and $+$, 1 for the rest
<code>niterations</code>	Number of competition iterations	1000
<code>ncycles_per_iteration</code>	Number of mutations per iteration	1000
<code>population_size</code>	Number of functions in each iteration	50

E. DEGENERACIES IN THE SÉRSIC FIT

The fact that the parameters in the Sérsic model are degenerate with one another has been known for a long time (e.g., [Graham et al. 1996](#); [Graham 1998](#); [Peng et al. 2010b](#); [Andrae et al. 2011](#)), however little has been done to *quantify* this degeneracy. An insightful work by [Trujillo et al. \(2001\)](#) does this for one case: an idealized galaxy described by some arbitrary Sérsic profile with some parameters n_1, R_1, I_{01} being fit with a fixed $n_2 = 4$ profile and some free R_2 and I_{02} . They solve for the degeneracy between the fitted parameters I_{02}, R_2 as a function of n_1, n_2 , and R_1 , and show that the degeneracy explains some of the tightness of the correlation between R and $\langle I \rangle$, the two axis of the fundamental plane of elliptical galaxies ([Djorgovski & Davis 1987](#)).

We extend this analysis for a general Sérsic fit, where n_2 may also vary. The question we ask is then: given a galaxy well-described by some (n_1, R_1, I_{01}) , if we try to fit it with some n_2 , what values for R_2 and I_{02} give the best fit? In other words, what is the degeneracy between Δn , ΔR , and ΔI_0 ?

To answer this, we want to minimize the total error between the profiles measured between some inner radius $r_{\min}$ and outer radius $r_{\max}$. We define the error as

$$E = \int_{r_{\min}}^{r_{\max}} (I_1(r) - I_2(r))^2 dr \quad (\text{E1})$$

where $I_1(r)$ is the true profile and $I_2(r)$ is the profile we are trying to optimize. This is slightly different from the [Trujillo et al. \(2001\)](#) approach, which defines an error somewhat akin to a χ^2 and normalize the square difference by $I_2(r)$, however this leads to numerical instability when $r \rightarrow \infty$. Finally, [Trujillo et al. \(2001\)](#) chose $r_{\min} = 0$ and solve for different $r_{\max}$. While it is true that we cannot know the profile out to infinity from an image, the intensity in the outskirts is negligible compared to the central one, and so $r_{\max}$ does not significantly affect the solution. On the other hand, due to pixelization, we do not know the true central flux, so we

choose to keep $r_{\min}$ non-zero in our analysis. Therefore, we want to minimize

$$E = \int_{r_{\min}}^{\infty} I_1^2(r) + I_2^2(r) - 2I_1(r)I_2(r) dr \quad (\text{E2})$$

For simplicity of the derivation, we use an alternative (but an equivalent) definition of the Sérsic profile:

$$I(r) = I_0 \exp \left[-b_n \left(\frac{r}{R_e} \right)^{1/n} \right] \quad (\text{E3})$$

where I_0 is the intensity at the centre of the galaxy rather than at effective radius R_e . The two are related to one another as $I(0) = I(R_e)e^{b_n}$.

To find the optimal values of R and I_0 , we want to solve for $\partial E / \partial R_2 = 0$ and $\partial E / \partial I_{02} = 0$. First, solving for $\partial E / \partial I_{02}$, we get

$$\begin{aligned} \frac{\partial E}{\partial I_{02}} &= 2I_{02} \int_{r_{\min}}^{\infty} \exp \left[-b_2 \left(\frac{r}{R_2} \right)^{\frac{1}{n_2}} \right] dr \\ &\quad - 2I_{01} \int_{r_{\min}}^{\infty} \exp \left[-b_2 \left(\frac{r}{R_2} \right)^{\frac{1}{n_2}} - b_1 \left(\frac{r}{R_1} \right)^{\frac{1}{n_1}} \right] dr \\ &= 0. \end{aligned} \quad (\text{E4})$$

The first term has a closed-form solution. The second term needs to be integrated numerically. For convenience, following [Trujillo et al. \(2001\)](#), we define this term as

$$S(r) \equiv \exp \left[-b_2 \left(\frac{r}{R_2} \right)^{\frac{1}{n_2}} - b_1 \left(\frac{r}{R_1} \right)^{\frac{1}{n_1}} \right], \quad (\text{E5})$$

note, however, that our $S(r)$ differs from the one in [Trujillo et al. \(2001\)](#) since both terms in the exponential are always negative, and therefore this exponential always vanishes at infinity; while the one in [Trujillo et al.](#)

TABLE D2: SymbolicRegression.jl fit results

p	Sec.	x	y	$f(x, y)$	$g(x, y)$	Error [all data]		Error [low-res]		Error [low-SNR]	
						Orig.	Corr.	Orig.	Corr.	Orig.	Corr.
Geometric measurements (Sec. 3.1)											
$\mathbf{a} \mathbf{x}_0^A$	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$4.40^{+12.8}_{-3.53}$	–	$16.6^{+11.1}_{-4.47}$	–	$4.34^{+12.3}_{-3.32}$	–
e	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh(\log x)$	0.031	$-0.02^{+0.09}_{-0.12}$	$-0.00^{+0.12}_{-0.10}$	$-0.13^{+0.18}_{-0.15}$	$-0.00^{+0.31}_{-0.15}$	$-0.02^{+0.14}_{-0.15}$	$-0.00^{+0.16}_{-0.15}$
$\mathbf{a} \theta$	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$0.30^{+16.2}_{-14.6}$	–	$1.20^{+29.6}_{-22.3}$	–	$0.65^{+30.1}_{-25.2}$	–
$\mathbf{a} \theta_{e>0.3}$	3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$0.08^{+8.44}_{-6.17}$	–	$0.63^{+27.7}_{-20.5}$	–	$0.53^{+21.3}_{-13.8}$	–
$\mathbf{a} \theta^{\text{Sersic}}$	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$-0.02^{+4.70}_{-4.82}$	–	$-0.21^{+25.0}_{-18}$	–	$-0.02^{+8.71}_{-7.09}$	–
$\mathbf{a} \theta^{\text{Sersic}}_{e>0.3}$	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$-0.04^{+1.98}_{-3.07}$	–	$-0.50^{+10.8}_{-11.5}$	–	$-0.07^{+4.00}_{-5.74}$	–
$\mathbf{b} e^{\text{Sersic}}$	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$0.00^{+0.06}_{-0.04}$	–	$0.05^{+0.15}_{-0.10}$	–	$0.00^{+0.07}_{-0.05}$	–
Radius measurements (Sec. 3.2)											
$R_{p,o}$	3.2.1	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$\tanh(3.21y)$	$(x/1000)^{1.86}$	$0.08^{+1.07}_{-0.86}$	$0.00^{+0.77}_{-0.81}$	$0.70^{+1.40}_{-1.29}$	$0.08^{+1.07}_{-1.26}$	$-0.12^{+1.08}_{-1.68}$	$-0.04^{+1.18}_{-1.26}$
$R_{p,e}$	3.2.1	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$\tanh(y + 0.51)$	$0.19 \log x$	$0.01^{+1.23}_{-1.36}$	$-0.02^{+1.83}_{-0.86}$	$0.47^{+1.81}_{-1.77}$	$0.43^{+2.41}_{-1.40}$	$0.01^{+1.23}_{-1.36}$	$-0.02^{+1.83}_{-0.86}$
R_{20}	3.2.2	$\mathcal{R}$	$\langle \text{SNR} \rangle$	0.91	$(x/1300)^{1.12}$	$0.10^{+0.49}_{-0.14}$	$0.01^{+0.16}_{-0.14}$	$0.51^{+0.43}_{-0.28}$	$0.03^{+0.25}_{-0.26}$	$0.05^{+0.48}_{-0.19}$	$-0.00^{+0.19}_{-0.21}$
R_{50}	3.2.2	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$\tanh(3.21y)$	$(x/1000)^{1.86}$	$0.09^{+0.70}_{-0.24}$	$-0.01^{+0.30}_{-0.38}$	$0.59^{+0.75}_{-0.48}$	$-0.05^{+0.41}_{-0.66}$	$0.03^{+0.68}_{-0.48}$	$0.01^{+0.62}_{-0.44}$
R_{80}	3.2.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh(3.35y)$	$3.04/x$	$0.06^{+0.82}_{-0.48}$	$-0.06^{+0.67}_{-0.38}$	$1.07^{+1.42}_{-0.60}$	$0.24^{+1.35}_{-0.52}$	$-0.02^{+0.82}_{-1.00}$	$-0.04^{+1.04}_{-0.63}$
R_{max}	3.2.2	$\mathcal{R}$	$\langle \text{SNR} \rangle$	$\log(y + 1.62)$	1.08	$-0.75^{+3.81}_{-5.12}$	$0.05^{+4.27}_{-2.58}$	$-0.47^{+5.19}_{-5.71}$	$0.66^{+5.06}_{-3.03}$	$-4.09^{+2.60}_{-8.44}$	$0.03^{+6.01}_{-4.80}$
$\mathbf{b} R_{0.5}^{\text{Sersic}}$	3.2.3	$\mathcal{R}$	$\langle \text{SNR} \rangle$	–	–	$-0.04^{+0.71}_{-1.13}$	–	$-0.18^{+0.95}_{-1.32}$	–	$0.04^{+1.48}_{-1.06}$	–
Bulge strength measurements (Sec. 3.3)											
n	3.2.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh(0.2x)$	0	$-0.07^{+0.54}_{-1.09}$	$-0.02^{+0.71}_{-0.91}$	$-0.81^{+1.30}_{-2.13}$	$0.25^{+2.15}_{-2.10}$	$0.01^{+0.87}_{-0.86}$	$0.03^{+0.99}_{-0.78}$
C	3.3.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh(\log 0.4x))$	$\frac{3.3}{\sqrt{x-0.76}}$	$-0.22^{+0.29}_{-0.79}$	$0.00^{+0.36}_{-0.38}$	$-1.14^{+0.62}_{-0.56}$	$0.12^{+1.86}_{-1.07}$	$-0.22^{+0.32}_{-0.72}$	$-0.06^{+0.46}_{-0.49}$
G	3.3.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.83 - 0.59/x$	$0.1 \tanh y$	$-0.03^{+0.04}_{-0.07}$	$-0.00^{+0.05}_{-0.04}$	$-0.08^{+0.08}_{-0.05}$	$0.01^{+0.14}_{-0.06}$	$-0.05^{+0.06}_{-0.07}$	$0.00^{+0.05}_{-0.06}$
M_{20}	3.3.3	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$0.74 - 1/x$	–0.5	$0.15^{+0.49}_{-0.22}$	$0.01^{+0.45}_{-0.29}$	$0.63^{+0.37}_{-0.42}$	$-0.10^{+0.65}_{-0.52}$	$0.20^{+0.52}_{-0.27}$	$0.12^{+0.63}_{-0.35}$
$B(G, M_{20})$	3.3.4	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh x/15$	$-\frac{2.1}{x \tanh y}$	$-0.25^{+0.30}_{-0.55}$	$-0.02^{+0.32}_{-0.44}$	$-0.82^{+0.50}_{-0.37}$	$0.25^{+1.67}_{-0.84}$	$-0.39^{+0.42}_{-0.60}$	$-0.16^{+0.40}_{-0.52}$
Disturbance measurements (Sec. 3.4)											
$\mathbf{b}^c S(G, M_{20})$	3.4.1 ^{bc}	$\mathcal{R}_{\text{eff}}$	μ_0	–	–	$0.01^{+0.47}_{-0.52}$	–	$-0.01^{+1.02}_{-0.59}$	–	$-0.01^{+0.46}_{-0.64}$	–
A_{RMS}	3.4.5	$\mathcal{R}_{\text{eff}}^{\mathbf{a}}$	$\langle \text{SNR} \rangle$	$0.5 \log(x)$	0	$-0.01^{+0.83}_{-1.46}$	$0.39^{+1.52}_{-0.89}$	$0.15^{+1.05}_{-1.67}$	$1.65^{+4.80}_{-1.49}$	$-0.11^{+2.24}_{-2.19}$	$0.23^{+2.50}_{-1.97}$
$\mathbf{d} A_{\text{RMS}}$	3.4.5	$\mathcal{R}_{\text{eff}}^{\mathbf{a}}$	$\langle \text{SNR} \rangle$	$0.55 \log(x - 1.5)$	$0.2/x$	$-0.01^{+0.83}_{-1.46}$	$0.07^{+1.17}_{-1.06}$	$-0.25^{+0.78}_{-2.16}$	$0.00^{+2.54}_{-1.53}$	$-0.11^{+2.24}_{-2.19}$	$-0.08^{+2.28}_{-1.97}$
A_{CAS}	3.4.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh\left(\frac{xy}{138}\right)$	$-\frac{0.295}{\exp(0.65y)}$	$-0.06^{+0.10}_{-0.21}$	$0.03^{+0.22}_{-0.26}$	$-0.03^{+0.08}_{-0.20}$	$0.00^{+0.59}_{-1.46}$	$-0.21^{+0.12}_{-0.18}$	$0.02^{+0.87}_{-1.08}$
$\mathbf{b}^c A_S$	3.4.4	$\mathcal{R}_{\text{eff}}$	μ_0	–	–	$0.03^{+0.20}_{-0.14}$	–	$0.05^{+0.23}_{-0.17}$	–	$0.12^{+0.20}_{-0.18}$	–
A_{22}	3.4.5	$\mathcal{R}_{\text{eff}}$	μ_0	0.94	$1/y$	$0.04^{+0.47}_{-0.09}$	$0.01^{+0.50}_{-0.10}$	$0.13^{+0.73}_{-0.13}$	$0.10^{+0.77}_{-0.14}$	$0.17^{+0.61}_{-0.20}$	$0.14^{+0.65}_{-0.20}$
A_o	3.4.2	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\log y$	$0.14(\log y - 0.8)$	$-0.05^{+0.13}_{-0.20}$	$0.00^{+0.21}_{-0.25}$	$-0.02^{+0.12}_{-0.18}$	$-0.02^{+0.21}_{-0.29}$	$-0.18^{+0.11}_{-0.23}$	$-0.06^{+0.57}_{-0.56}$
Other measurements (Sec. 3.5)											
$\mathbf{b}^c M$	3.5.3	$\mathcal{R}_{\text{eff}}$	μ_0	–	–	$-0.01^{+0.44}_{-0.30}$	–	$-0.05^{+0.26}_{-0.16}$	–	$0.12^{+1.27}_{-0.38}$	–
$\mathbf{b} I$	3.5.4	$\mathcal{R}_{\text{eff}}$	μ_0	–	–	$-0.00^{+0.08}_{-0.14}$	–	$0.06^{+0.23}_{-0.19}$	–	$-0.00^{+0.10}_{-0.17}$	–
$\mathbf{c} D$	3.5.5	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	$\tanh(x/14) - 0.12$	$\frac{0.53}{x+1.55}$	$0.17^{+0.81}_{-0.41}$	$0.00^{+1.59}_{-0.76}$	$0.70^{+0.92}_{-0.58}$	$-0.40^{+6.51}_{-3.94}$	$0.30^{+1.06}_{-0.77}$	$0.19^{+2.00}_{-1.00}$
$\mathbf{b} S$	3.5.1	$\mathcal{R}_{\text{eff}}$	$\langle \text{SNR} \rangle$	–	–	$-0.00^{+0.03}_{-0.13}$	–	$-0.02^{+0.17}_{-0.35}$	–	$-0.06^{+0.05}_{-0.45}$	–
St	3.5.2	$\mathcal{R}_{\text{eff}}$	μ_0	–	–	$-0.01^{+0.04}_{-0.11}$	–	$-0.01^{+0.03}_{-0.07}$	–	$-0.05^{+0.05}_{-0.13}$	–

Notes – for each parameter, x and y are found using mutual information regression. We then compute the 16th, 50th, and 84th quantiles of the error ($p - p_{\text{base}}$) distribution for the entire sample, and low-resolution / low-SNR subsamples. We repeat this for ($p_{\text{corr}} - p_{\text{base}}$) where a SymbolicRegression correction is fitted.

^a Here we quantify the offset from the baseline in a random direction, so a correction is not possible.

^b There resolution/depth bias is significantly smaller than the average scatter, so a bias correction is not needed.

^c The error quantiles are multiplied by 10.

^d Better correction for low-resolution images but a similar performance overall.

(2001) becomes positive at large r and forces the integral to diverge if $n_2 < n_1$. Substituting in the solution of the first term and rearranging the equation, we get an expression for I_{02}/I_{01} :

$$\frac{I_{02}}{I_{01}} = \frac{2^{n_2} b_2^{n_2}}{n_2 R_2} \Gamma^{-1} \left[n_2, 2b_2 \left(\frac{r_{\min}}{R_2} \right)^{\frac{1}{n_2}} \right] \int_{r_{\min}}^{\infty} S(r) dr \quad (\text{E6})$$

where Γ is an upper incomplete Gamma function defined as $\Gamma(t, x) \equiv \int_x^{\infty} t^{s-1} e^{-t} dt$ and $\Gamma^{-1} \equiv 1/\Gamma$. For simplicity, we also define $X \equiv 2b_2 \left(\frac{r_{\min}}{R_2} \right)^{\frac{1}{n_2}}$.

We then follow the same steps using $\partial E / \partial R_2 = 0$. This produces another expression for I_{02}/I_{01} :

$$\frac{I_{02}}{I_{01}} = \frac{2^{n_2+1} b_2^{n_2+1}}{n_2 R_2} \Gamma^{-1} [n_2 + 1, X] \times \int_{r_{\min}}^{\infty} \left(\frac{r}{R_2} \right)^{\frac{1}{n_2}} S(r) dr. \quad (\text{E7})$$

With these two derivatives, we now have an implicit expression for R_2 that is independent from I_{02} , and so by solving it numerically, we can evaluate the dependence of the best-fit R_2 on n_2 . Similarly, we can then substitute this R_2 to get the dependence of I_{02} on n_2 . The expression for R_2 is:

$$\int_{r_{\min}}^{\infty} S(r) \left[1 - 2b_2 \frac{\Gamma[n_2, X]}{\Gamma[n_2 + 1, X]} \left(\frac{r}{R_2} \right)^{\frac{1}{n_2}} \right] dr = 0 \quad (\text{E8})$$

This expression is almost equivalent to the one in Trujillo et al. (2001), except the functional form of $S(r)$ is different, n_2 is free, and the lower bound of the integral is allowed to vary. This integral does not have a closed form solution and has to be solved numerically. We use a SciPy root finder by setting a search bracket of $(r_{\min}, R_1)$ if $n_2 < n_1$ and (R_1, ∞) if $n_2 > n_1$.

To test this solution in a realistic use-case, we did the following: we first generated a perfect GALSIM Sérsic profiles with a range of chosen n_1 and set $R_1 = 10$, $I_{01} = 100$. We chose some pixel scale to discretize the image to, which we now define as $\mathcal{R}_{\text{eff}} \equiv R_0/\text{scale}$, similar to the primary tests of this paper. We defined $r_{\min} \equiv \text{scale}/2$. Then, we added a varying amount of sky and applied Poisson noise to achieve different levels of $\langle \text{SNR} \rangle$, calculated in a $2R_1$ aperture. Finally, for a wide range of n_2 , we found the best-fitting R_2 and I_2 and calculated the χ^2 of the fit.

In Fig. E1, we show how the ratios of best-fitting R_2/R_1 and I_{02}/I_{01} change as a function of $\Delta n \equiv n_2 - n_1$ for two test cases: $n_1 = 2$ (dashed) and $n_1 = 4$ (solid). We vary the effective resolution, showed in color from the lowest-resolution (dark) to a relatively high resolution of

20 pixels per R_1 (bright). In all cases, the relationship between R_2/R_1 or I_{02}/I_{01} and Δn is log-linear. In other words, the degeneracy is perfectly described by the functional form

$$\frac{R_2}{R_1} = e^{\alpha \Delta n} \quad \text{and} \quad \frac{I_{02}}{I_{01}} = e^{-\beta \Delta n},$$

where α and β depend on n_1 and $r_{\min}$. The relationship between α and β is independent of n_2 and only weakly dependent on the resolution, as seen in the right-most panel of Fig. E1, and is approximately

$$\frac{I_{02}}{I_{01}} \approx \left(\frac{R_2}{R_1} \right)^{-1.65},$$

plotted as a thin grey line in the third panel of Fig. E1. As Trujillo et al. (2001) note, the degeneracy between these two quantities partially explains the tightness in the empirical relationship between R_e and $\langle I \rangle$ which makes up one of the axes on the fundamental plane of elliptical galaxies.

To find the relationship between α , $r_{\min}$, and n_0 , we used `SymbolicRegression.jl`. We computed α for all of our test cases and used only n_0 and $r_{\min}$ to find an equation for α – the fitter did not have access to R_1 , R_2 , or Δn . The best-scoring functional form is

$$\alpha = 1.223/n_0^{\sqrt{r_{\min}}}.$$

Using this fit, we compute predicted α and then the corresponding R_2/R_1 . The mean squared error on the ratio of effective radii is only 0.0005, indicating an excellent fit to the data, especially considering that the fitter was not given radius values themselves. Therefore, we conclude that the degeneracy between the best-fitting radius R_2 given n_2 is:

$$\frac{R_2}{R_1} \approx \exp \left[1.223 \Delta n / n_0^{\sqrt{r_{\min}}} \right]$$

This is a purely analytic degeneracy between R and n that arises from the definition of the Sérsic model itself. Values of R_2 computed in this way will minimize the squared error between an observed Sérsic galaxy (n_1, R_1, I_{01}) and a fitted model with a fixed n_2 . However, obviously, *the best fit* will be achieved when $n_2 = n_1$. Therefore, it makes sense to then ask a question: how far can n_2 change and still produce a good fit?

To investigate this, we opted to compute a more empirically motivated goodness-of-fit using the χ^2 of the simulated images. We then found the *maximum* Δn range which resulted in a $p(\chi^2) > 0.05$, which would be an acceptable fit. This corresponds to a reduced χ^2 between 1.05 and 1.3, depending on the number of degrees of freedom, i.e. the size of the image.




FIG. E1.—: The relationship between a fixed $\Delta n \equiv n_2 - n_1$ and best-fitting R_2 (left), I_{02} (center), as well as between R_2 and I_{02} (right). The degeneracy in the fit is log-linear and depends on the resolution (i.e. $r_{\min}$) and the original Sérsic index n_1 . The degeneracy is $R_2/R_1 = \exp[1.223\Delta n/n_1^{r_{\min}}]$ and $I_{02}/I_{01} = (R_2/R_1)^{-1.65}$. However, while these values of R_2 and I_2 give the best fit given a fixed n_2 , the minimized χ^2 of the fit may still be too large for a given Δn ; the largest allowed Δn are shown in Fig. E2




FIG. E2.—: The uncertainty in the Sérsic index and effective radius arising from the degeneracies in the analytic profile. Shown are the maximum Δn (left) and hence $\Delta R/R_1$ (right) allowed by $p(\chi^2) > 0.05$ for an $n_1 = 4$ profile, as a function of resolution and depth. At low $\langle \text{SNR} \rangle$, the uncertainty in ΔR reaches 30%, which explains the scatter in Sérsic radius measurements on the observed dataset in Fig. 3. The uncertainty in Δn at medium resolutions is $\pm 0.25 \sim \pm 1$, depending on $\langle \text{SNR} \rangle$, slightly lower than the scatter in Fig. 5

Fig. E2 shows the maximum values of $\Delta R/R$ (left) and Δn (right) that pass these criteria as a function of $\langle \text{SNR} \rangle$ and $\mathcal{R}_{\text{eff}}$ for a true profile with $n = 4$. This is essentially a measure of uncertainty in R and n that is inherent to the Sérsic profile, and so it should be treated as a *lower bound* on the uncertainty in these quantities. In realistic use cases, since galaxies are complex structures with features besides a simple one-component Sérsic, real χ^2 will be higher than in our idealized case, and poorer fits

might have to be accepted, which increases the intrinsic uncertainty on R and n .

Nevertheless, even considering these as lower bounds, the uncertainty in R for $\langle \text{SNR} \rangle \sim 5$ is 15% and can reach 50% at lower $\langle \text{SNR} \rangle$, which is in excellent agreement with the scatter in Sérsic radius at low signal-to-noise in our observational tests (Fig. 3). The analytic uncertainty is a somewhat weaker function of resolution than what we saw in Fig. 3, likely due to the complex structure of




FIG. E3.—: The uncertainty Δn and $\Delta R/R_1$, defined the same way as in Fig. E2, for a mock galaxy with $\langle \text{SNR} \rangle = 3$ and $\mathcal{R}_{\text{eff}} = 10$, as a function of n_1 . While the uncertainty in n grows from 0.1 to 0.6 at high Sérsic indices, the uncertainty in the effective Radius stays roughly constant at 20% for an image of this quality, independent of n_1 .

real galaxies. Nevertheless, our analytic approach perfectly predicts the observed scatter in the Sérsic radius measurements as a function of $\langle \text{SNR} \rangle$, which is quite remarkable.

The intrinsic uncertainty in n predicted from our model is somewhat lower than the observational scatter seen in Fig. 5, where the scatter was $\Delta n \sim 1$ at $\mathcal{R}_{\text{eff}} = 10$ for all $\langle \text{SNR} \rangle$. In contrast, the model allows stricter constraints on Δn when $\langle \text{SNR} \rangle$ is high, reaching a ± 0.15 uncertainty. This discrepancy could also be due to the underlying structure of realistic galaxies: e.g., a concentrated nuclear starburst could significantly increase the n estimate when sufficiently resolved. Even our simple model, however, still provides a good indication of the *lower bound* on the uncertainty in n one should expect.

Finally, the range of $\Delta R/R$ and Δn allowed by our χ^2 estimates varies as a function of the intrinsic profiles, i.e. the original n_1 . More early-type profiles are more degenerate with one another, while lower- n profiles are better constrained. To demonstrate this, we plotted the allowed range of $\Delta R/R_1$ (orange shading) and Δn (blue shading) for images with $\langle \text{SNR} \rangle = 3$ and $\mathcal{R}_{\text{eff}} = 10$ as a function of n_1 . Interestingly, while Δn increases with n_1 as expected, the $\Delta R/R_1$ estimate is fairly constant with n_1 , staying at approximately 20%.

This paper was built using the Open Journal of Astrophysics L^AT_EX template. The OJA is a journal which provides fast and easy peer review for new papers in the **astro-ph** section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.