University of Louisville

ThinkIR: The University of Louisville's Institutional Repository

8-31-2022

Galaxy And Mass Assembly (GAMA): bulge-disc decomposition of KiDS data in the nearby Universe

Sarah Casura

Jochen Liske

Aaron S G Robotham

Sarah Brough

Simon P. Driver

See next page for additional authors

Follow this and additional works at: https://ir.library.louisville.edu/faculty

Part of the Astrophysics and Astronomy Commons

Original Publication Information

ThinkIR Citation

Authors

Sarah Casura, Jochen Liske, Aaron S G Robotham, Sarah Brough, Simon P. Driver, Alister W. Graham, Boris Häußler, Benne Holwerda, Andrew M. Hopkins, Lee S. Kelvin, Amanda J. Moffett, Dan S. Taranu, and Edward N. Taylor

Galaxy And Mass Assembly (GAMA): bulge-disc decomposition of KiDS data in the nearby Universe

Sarah Casura [1]★ Jochen Liske [1], Aaron S. G. Robotham [2], Sarah Brough [3], Simon P. Driver [2], Alister W. Graham,[4] Boris Häußler [5], Benne W. Holwerda,[6] Andrew M. Hopkins,[7] Lee S. Kelvin [8], Amanda J. Moffett,[9] Dan S. Taranu[8,2] and Edward N. Taylor[4]

[1]*Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany*
[2]*ICRAR, M468, University of Western Australia, Crawley, WA 6009, Australia*
[3]*School of Physics, University of New South Wales, NSW 2052, Australia*
[4]*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia*
[5]*European Southern Observatory, Alonso de Cordova 3107, Casilla 19001, Santiago, Chile*
[6]*Department of Physics and Astronomy, 102 Natural Science Building, University of Louisville, Louisville KY 40292, USA*
[7]*Australian Astronomical Optics, Macquarie University, 105 Delhi Rd, North Ryde, NSW 2113, Australia*
[8]*Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA*
[9]*Department of Physics and Astronomy, University of North Georgia, 3820 Mundy Mill Rd., Oakwood GA 30566, USA*

ABSTRACT

We derive single Sérsic fits and bulge-disc decompositions for 13 096 galaxies at redshifts $z < 0.08$ in the GAMA II equatorial survey regions in the Kilo-Degree Survey (KiDS) g, r, and i bands. The surface brightness fitting is performed using the Bayesian two-dimensional profile fitting code PROFIT. We fit three models to each galaxy in each band independently with a fully automated Markov chain Monte Carlo analysis: a single Sérsic model, a Sérsic plus exponential and a point source plus exponential. After fitting the galaxies, we perform model selection and flag galaxies for which none of our models are appropriate (mainly mergers/Irregular galaxies). The fit quality is assessed by visual inspections, comparison to previous works, comparison of independent fits of galaxies in the overlap regions between KiDS tiles and bespoke simulations. The latter two are also used for a detailed investigation of systematic error sources. We find that our fit results are robust across various galaxy types and image qualities with minimal biases. Errors given by the MCMC underestimate the true errors typically by factors 2–3. Automated model selection criteria are accurate to > 90 per cent as calibrated by visual inspection of a subsample of galaxies. We also present $g-r$ component colours and the corresponding colour–magnitude diagram, consistent with previous works despite our increased fit flexibility. Such reliable structural parameters for the components of a diverse sample of galaxies across multiple bands will be integral to various studies of galaxy properties and evolution. All results are integrated into the GAMA database.

Key words: methods: statistical – catalogues – galaxies: fundamental parameters – galaxies: structure; techniques: photometric.

1 INTRODUCTION

The quantitative modelling of galaxy surface brightness distributions has a long history dating back to de Vaucouleurs (1948), Sérsic (1963) and even earlier works; see Graham (2013) for a review of the development of light profile models. While the early works focused on azimuthally averaged galaxy profiling with just a single functional form (e.g. Kormendy 1977), modern codes allow users to decompose galaxies into several distinct components and to take into account the full two-dimensional information. To this end, there are many different techniques, methods, and code packages, all of which have become increasingly sophisticated as the quality and quantity of available astronomical data have grown. Broadly, they can be divided into parametric and non-parametric modelling as well as one-dimensional and two-dimensional methods. Which

of these is most appropriate to use depends on the science case and the available data. This work falls into the regime of large-scale automated analyses of galaxies with often barely resolved components, for which we want to obtain structural parameters that are easily comparable between galaxies. Hence, two-dimensional parametric analysis is most appropriate (see also the discussion in Robotham et al. 2017 and references therein). Examples of such two-dimensional, parametric fitting tools used for large-scale automated analyses include GIM2D (Simard et al. 2002), BUDDA (de Souza, Gadotti & dos Anjos 2004), GALFIT3 (Peng et al. 2010), GALFITM (Vika et al. 2013), IMFIT (Erwin 2015), PROFIT (Robotham et al. 2017), and PHI (Argyle et al. 2018). Each of these tools comes with its own advantages and disadvantages, which goes to show how difficult the problem of galaxy modelling is, especially when automated for large samples of a very diverse galaxy population. Usually, some form of post-processing is needed to assess the influence of systematic errors, judge the convergence,

★ E-mail: sarah.casura@uni-hamburg.de

exclude bad fits and identify the most appropriate model to use for each galaxy. This can be achieved via visual inspection (for small enough samples), logical filters, frequentist statistics such as the *F*-test, Bayesian inference, or similar methods (see e.g. Allen et al. 2006; Gadotti 2009; Simard et al. 2011; Vika et al. 2014; Meert, Vikram & Bernardi 2015; Lange et al. 2016; Méndez-Abreu et al. 2017).

Despite the associated difficulties (e.g. convergence and quality of fit metrics), many authors have performed two-dimensional surface brightness profile fitting for large numbers of galaxies, modelling the radial light profile as a simple functional form, most often a Sérsic function (Blanton et al. 2003, 2005; Barden et al. 2005; Trujillo et al. 2006; Hyde & Bernardi 2009; La Barbera et al. 2010; Kelvin et al. 2012; van der Wel et al. 2012; Häußler et al. 2013; Shibuya, Ouchi & Harikane 2015; Sánchez-Janssen et al. 2016, to name just a few). The results of such analyses have been used to derive a number of key relations between different galaxy properties, their formation and evolutionary history, and interactions with the environment. For example, many works have studied the distribution of, and relation between, size and mass or luminosity for different galaxy types (split by e.g. Sérsic index or colour), sometimes including morphology, surface brightness, internal velocity, environment, wavelength, colour, or redshift effects (e.g. Shen et al. 2003; Barden et al. 2005; Blanton et al. 2005; Trujillo et al. 2006; Hyde & Bernardi 2009; La Barbera et al. 2010; Kelvin et al. 2014; van der Wel et al. 2014; Lange et al. 2015; Shibuya et al. 2015; Kawinwanichakij et al. 2021; Nedkova et al. 2021).

With improving data quality of surveys, the galaxy fitting community has increasingly shifted towards fitting more than one component, i.e. to perform bulge-disc decomposition. While some authors, such as Gadotti (2009), Salo et al. (2015), or Gao & Ho (2017) also account for bars, central point sources, spiral arms or other additional morphological features, most works focus on the bulge and disc. The focus on only two components is especially true when running automated analyses of large samples, since in many cases the data quality is not sufficient to meaningfully constrain more than one or two components, or it would require extensive manual tuning based on visual inspection. From a more physical point of view, the majority of the stellar mass in the local Universe resides in ellipticals, discs, and classical bulges, with pseudo-bulges and bars only contributing a few per cent (Gadotti 2009). Hence, for automated analyses it is common practice to fit only two components, where the term 'bulge' is used to describe the central component, irrespective of whether it is a classical bulge, pseudo-bulge, bar, lens, active galactic nucleus (AGN), or a mixture thereof, while 'disc' refers to a more extended component with typically lower surface brightness and potential additional structure such as spiral arms, breaks, flares, or rings.

Examples of large bulge-disc decomposition studies include Simard et al. (2002, 2011), Allen et al. (2006), Benson et al. (2007), Gadotti (2009), Lackner & Gunn (2012), Fernández Lorenzo et al. (2014), Head et al. (2014), Mendel et al. (2014), Vika et al. (2014), Meert et al. (2015), Meert, Vikram & Bernardi (2016), Kennedy et al. (2016), Kim et al. (2016), Lange et al. (2016), Dimauro et al. (2018), Bottrell et al. (2019), Cook et al. (2019), Barsanti et al. (2021), and Häußler et al. (2022). Such catalogues can then be used to determine the relative numbers of different galaxy components as well as their luminosity or stellar mass functions, size–mass or size–luminosity relations, including their redshift evolution and dependence on other properties of the galaxy and its environment (similar to the studies of entire galaxies mentioned earlier). This has been done by Driver et al. (2007b), Dutton et al. (2011), Tasca et al. (2014), Kennedy

et al. (2016), Lange et al. (2016), Moffett et al. (2016), Dimauro et al. (2019), and many others.

In addition, quantitative measures for the components of galaxies aid the comparison of observational data to theory and simulations. Bulges and discs are often decisively different not only in their visual appearance but also in their structure, dynamics, stellar populations, gas and dust content, and are thought to have different formation pathways (Cole et al. 2000; Cook, Lapi & Granato 2009; Driver et al. 2013; Lange et al. 2016; Dimauro et al. 2018; Lagos et al. 2018; Oh et al. 2020). Consequently, bulge-disc decomposition studies provide stringent constraints on the formation and evolutionary histories of galaxies and their physical properties that are not easily measured directly such as the dark matter halo, the build-up of stellar mass (in different components) over time, or merger histories (examples include Driver et al. 2013; Bottrell et al. 2017; Bluck et al. 2019; Rodriguez-Gomez et al. 2019; de Graaff et al. 2022). Hence, consistently measuring the structure of the stellar components is essential to make full use of current and future large-scale observational surveys such as the Kilo-Degree Survey (KiDS; de Jong et al. 2013) or the Legacy Survey of Space and Time (LSST; Ivezić et al. 2019), and of cosmological hydrodynamical simulations such as Illustris (Vogelsberger et al. 2014) and IllustrisTNG (The Next Generation; Pillepich et al. 2018) or Evolution and Assembly of GaLaxies and their Environments (EAGLE; Schaye et al. 2015).

In the present study, we obtain single Sérsic fits and bulge-disc decompositions for 13 096 GAMA galaxies in the KiDS *g*, *r*, and *i* bands. We choose PROFIT as our modelling software due to its Bayesian nature (allowing full MCMC treatment including more realistic error estimates), its suitability to large-scale automated analyses and its ability, in combination with PROFOUND, to serve as a fully self-contained package covering all steps of the analysis from image segmentation through to model fitting. We supplement this functionality with our own routines for the rejection of unsuitable fits, model selection, and a characterization of systematic uncertainties. The resulting catalogue has already been used to aid the kinematic bulge-disc decomposition of a sample of galaxies in the Sydney-AAO Multi-object Integral-field spectroscopy (SAMI) Galaxy Survey (Oh et al. 2020) and to examine the properties of galaxy groups (Cluver et al. 2020), with many more studies in progress.

Our own plans include deriving the stellar mass functions of bulges and discs, studying component colours and trends of other Sérsic parameters with wavelength, and constraining the nature and distribution of dust in galaxy discs. The latter can be achieved by comparing the distribution of bulges and discs in the luminosity–size plane to dust radiative transfer models such as those presented in Popescu et al. (2011) and preceding papers of this series (similar to the analysis performed by Driver et al. 2007a albeit with more and better data and at several wavelengths). For these science aims, we are most interested in obtaining structural parameters that are directly comparable amongst each other, i.e. consistent within the data set; and correctly represent the statistical properties of the entire sample, with less emphasis placed on capturing all aspects of the detailed structure of individual galaxies. Correspondingly, we choose to model a maximum of two components for each galaxy and use the terms 'bulge' and 'disc' in their widest senses, in line with previous automated decompositions of large samples. In particular the 'bulges' we obtain are often mixtures of classical or pseudo-bulges, bars, lenses, and AGN. Similarly, we place more emphasis on the central, high surface brightness regions of galaxies by modelling only a relatively tight region around each galaxy of interest. While most of the fits we obtain are not perfect (because galaxies are more

complex than two simple components), they do achieve the above named aims and are comparable to similar studies.

In Section 2 we describe our data (GAMA and KiDS), sample selection, and code (PROFIT and PROFOUND), and discuss in detail the distinguishing features of this study compared to previous work. Section 3 then presents the pipeline we developed for the bulge-disc decomposition, including preparatory work and post-processing, before we show our main results in Section 4. Sections 5 and 6 focus on the quality control of the fits by comparison to previous work and a detailed investigation into systematic uncertainties and biases from simulations and the overlap sample. We conclude with a summary and information on catalogue access in Section 7. We assume a standard cosmology of $H_0 = 70 \, \mathrm{km \, s^{-1} \, Mpc^{-1}}$, $\Omega_m = 0.3$, and $\Omega_\lambda = 0.7$ throughout.

2 DATA, SAMPLE, AND CODE

2.1 GAMA

The Galaxy and Mass Assembly (GAMA)[1] survey is a large low-redshift spectroscopic survey covering $\sim 238\,000$ galaxies in $286 \, \mathrm{deg}^2$ of sky (split into five survey regions) out to a redshift of approximately 0.6 and a depth of $r < 19.8$ mag. The observations were taken using the AAOmega spectrograph on the *Anglo-Australian Telescope* and were completed in 2014. The survey strategy and spectroscopic data reduction are described in detail in Driver et al. (2009, 2011), Baldry et al. (2010, 2014), Robotham et al. (2010), Hopkins et al. (2013), and Liske et al. (2015).

In addition to the spectroscopic data, the GAMA team collected imaging data on the same galaxies from a number of independent surveys in more than 20 bands with wavelengths between 1 nm and 1 m. Details of the imaging surveys and the photometric data reduction are given in Liske et al. (2015), Driver et al. (2016, 2022), and references therein. The combined spectroscopic and multiwavelength photometric data at this depth, resolution, and completeness provide a unique opportunity to study a variety of properties of the low-redshift galaxy population.

In this work, we focus on the KiDS g, r, and i-band imaging data (see Section 2.2) in the GAMA II equatorial survey regions, which are three regions of size $12° \times 5°$ located along the equator at 9, 12, and approximately 14.5 h in right ascencion (the G09, G12, and G15 regions). For our sample selection, we make use of the equatorial input catalogue[2] (EqInputCat:TilingCatv46; Baldry et al. 2010) and the most recent version of the redshifts originally described by Baldry et al. (2012; LocalFlowCorrection:DistancesFramesv14), see details in Section 2.3. For the stellar mass–size relation (Section 5.3), we also use the Data Release (DR) 3 version of the stellar mass catalogue first presented in Taylor et al. (2011; StellarMasses:StellarMassesv19); for the comparison to previous work (Section 5.4) we use the single Sérsic fits of Kelvin et al. (2012; SersicPhotometry:SersicCatSDSSv09); and in order to correct galaxy colours for Galactic extinction, we use the corresponding table provided along with the equatorial input catalogue (EqInput-

Cat:GalacticExtinctionv03). All of these catalogues can be obtained from the GAMA database[1].

2.2 KiDS

The Kilo-Degree Survey (KiDS, de Jong et al. 2013) is a wide-field imaging survey in the Southern sky using the *VLT Survey Telescope* (*VST*) at the ESO Paranal Observatory. 1350 deg^2 are mapped in the optical broad-band filters u, g, r, i; while the VISTA Kilo-degree INfrared Galaxy (VIKING) Survey (Edge et al. 2013) provides the corresponding near-infrared data in the Z, Y, J, H, K_s bands. The GAMA II equatorial survey regions have been covered as of DR3.0.

KiDS provides $\sim 1° \times 1°$ science tiles calibrated to absolute values of flux with associated weight maps (inverse variance) and binary masks. The science tiles are composed of five dithers (four in u) totalling 1000, 900, 1800, and 1200 s exposure time in u, g, r, i, with all dithers aligned in the right ascencion and declination axes (i.e. no rotational dithers). The r-band observations are performed during the best seeing conditions in dark time; while g, u, and i have progressively worse seeing and i is additionally taken during grey time or bright moon. During co-addition, the dithers across all four bands are re-gridded on to a common pixel scale of 0.2 arcsec. The magnitude zeropoint of the science tiles is close to zero with small corrections given in the image headers. The r-band point spread function (PSF) size is typically 0.7 arcsec and the limiting magnitudes in u, g, r, i are ~ 24.2, 25.1, 25.0, 23.7 mag, respectively (5σ in a 2 arcsec aperture). This high image quality, depth, survey size, and wide wavelength coverage in combination with VIKING make KiDS data unique. For details, see Kuijken et al. (2019).

For this work, we use the g, r, and i-band science tiles, weight maps, and masks from KiDS DR4.0 (Kuijken et al. 2019), which are publicly available[3] for our selected sample of galaxies (Section 2.3). We plan to extend the analysis to include the KiDS u and the VIKING Z, Y, J, H, and K_s bands in a future work.

2.3 Sample selection

Our main sample consists of all GAMA II equatorial region main survey targets with a reliable redshift in the range $0.005 < z < 0.08$, which are a total of 12 958 objects.[4] In addition, we include all 2404 targets of the 'GAMA sample' of the SAMI Galaxy Survey[5] (Bryant et al. 2015), the majority of which are already in our main sample. The combination of both results in the full sample of 13 096 unique physical objects, which were imaged a total of 14 966 times in each of the KiDS g, r, and i bands due to small overlap regions between the tiles. 11301, 1742, 31, and 22 objects were imaged once, twice, three, and four times, respectively. We keep these multiple data matches to the same physical object separate during all processing steps to serve as an internal consistency check.

2.4 PROFIT

We fit the surface brightness distributions of our sample of galaxies using PROFIT[6] (v1.3.2) which is a free and open-source, fully

---

Bayesian two-dimensional profile fitting code (Robotham et al. 2017). PROFIT offers great flexibility: there are several built-in profiles to choose from, it is easy to add several components of the same or different profiles, there is a choice of likelihood calculations and optimization algorithms that can be used (various downhill gradient options, genetic algorithms, over 60 variants of Markov chain Monte Carlo methods; MCMC), parameters can be fitted in linear or logarithmic space, it is possible to add complex priors for each, as well as constraints relating several parameters; and much more. The pixel integrations are performed using a standalone C++library (LIBPROFIT), making it both faster and more accurate than other commonly used algorithms such as GALFIT (Peng et al. 2010; see detailed comparison in Robotham et al. 2017). This allows us to fit galaxies with the computationally more expensive MCMC algorithms, overcoming the main problems of downhill gradient based optimizers: their susceptibility to initial guesses and their inability to easily derive realistic error estimates (e.g. Lange et al. 2016). This makes PROFIT highly suitable for the decomposition of large sets of galaxies with little user intervention.

2.5 PROFOUND

PROFIT (Section 2.4) requires a number of inputs apart from the (sky-subtracted) science image and the chosen model to fit, most importantly initial parameter guesses, a segmentation map specifying which pixels to fit, a sigma (error) map, and a PSF image. To provide these inputs in a robust and consistent manner, the sister package PROFOUND[7] (Robotham et al. 2018) was developed, which also serves as a stand-alone source finding and image analysis tool. The main novelties of PROFOUND compared to other commonly used free and open-source packages such as SOURCE EXTRACTOR (Bertin & Arnouts 1996) are that, rather than elliptical apertures, PROFOUND uses dilated 'segments' (collections of pixels of arbitrary shape) with watershed de-blending across saddle-points in flux. This means that the flux from each pixel is attributed to exactly one source (or the background) and apertures are never overlapping or nested. It also allows for extracting more complex object shapes than ellipses while still capturing the total flux due to the segment dilation (expansion) process. This makes it less prone to catastrophic segmentation failures (such as fragmentation of bright sources or blending of several sources into one aperture), reducing the need for manual intervention or multiple runs with 'hot' and 'cold' deblending settings, hence making PROFOUND particularly suitable for large-scale automated analysis of deep extragalactic surveys (Davies et al. 2018; Robotham et al. 2018; Bellstedt et al. 2020).

Apart from the segmentation map, the main function of the package, PROFOUNDPROFOUND, also returns estimated sky and sky-RMS maps (if not given as inputs) and a wealth of ancillary data including a list of segments and their properties such as their size, ellipticity, and the flux contained. The latter is particularly useful to obtain reasonable initial parameter guesses for galaxy fitting; or for identifying certain types of sources (e.g. stars for PSF estimation). The package also contains many additional functions for further image analysis and processing, all within the same framework. In addition, combining PROFOUND with PROFIT allows the user to estimate a PSF (see Section 3.1.5), hence entirely removing any dependence on external tools. Finally, both packages come with comprehensive documentation and many extended examples and vignettes which serve as great resources for newcomers to the fields of source extraction and galaxy fitting.

We use PROFOUND (v1.9.2) along with PROFIT (v1.3.2) for all preparatory steps (image segmentation/source identification, sky subtraction, initial parameter estimates, and PSF determination; see Section 3.1 for details) producing the inputs needed for the galaxy fitting with PROFIT.

3 BULGE-DISC DECOMPOSITION PIPELINE

We use the free and open-source R programming language (R Core Team 2020) for all scripting.

3.1 Preparatory steps

3.1.1 Cutouts and masking

KiDS imaging tiles are registered to the same pixel grid across all four bands (with matching weight maps and masks), such that a joint analysis of the bands is straightforward. They are also aligned such that the x-axis corresponds to right ascension (RA) and the y-axis to declination (Dec). Hence, we obtain a 400 arcsec × 400 arcsec cutout of the KiDS tile, associated weight map and mask for each object in our sample and for each of the three KiDS bands we used (g, r, i). The masks of all three bands are then combined and all pixels which have a value greater than 0 in any of the masks are excluded from analysis. This results in approximately 20 per cent of all pixels being masked out. This large fraction of masking is primarily due to the reflection haloes of bright stars that are also clearly visible in the data (see de Jong et al. 2015 for details). We combine the masks in this way to ensure that the pixels used for analysis are exactly the same in all bands and so the results are most directly comparable between bands. Objects for which the central pixel is masked ($\sim$ 20 per cent of all galaxies) are skipped in the galaxy fitting.

3.1.2 Image segmentation

We perform image segmentation in order to determine which pixels to fit for each of our objects, identify other nearby sources, improve the background subtraction, and obtain reasonable initial guesses for the galaxy parameters. This is performed on the joint g, r, i cutouts with PROFOUND in several steps.

First, we add the cutouts in the g, r, and i bands using inverse variance weighting and compute the joint weight map. We then estimate the (joint *gri*) sky by running the stacked image through PROFOUNDPROFOUND passing in the correct magnitude zeropoint, mask, and weight, but leaving SKYCUT on its default of 1. This means that all pixels with a flux at least 1σ above the median are progressively assigned to segments (collections of pixels belonging to an object) using an iterative process: starting with the brightest pixel in the image, segments are grown by adding neighbouring pixels with lower flux; new segments are started when a pixel shows more flux than its neighbours (within some tolerance) or when all neighbouring pixels above the SKYCUT value have been assigned. Once all pixels above SKYCUT have been assigned, the resulting segments are additionally expanded until flux convergence is reached. For more details, see Robotham et al. (2018).

Along the way, PROFOUND estimates the sky background several times since object detection relies on accurate background subtrac-
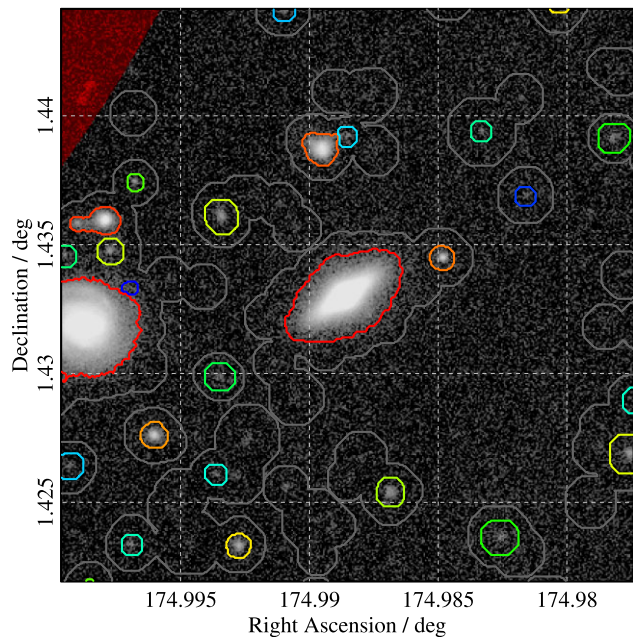
Figure 1. The PROFOUND segmentation map obtained for the galaxy 396 740 overlaid on the KiDS *r*-band image. Note this is only a cutout of the full segmentation map showing the central 100 arcsec × 100 arcsec. Identified objects (segments) are shown with contours, coloured from red to blue according to the flux contained. Grey contours indicate the more dilated segmentation map used for the background subtraction. Masked areas are shaded red.

tion and vice versa.[8] For the final sky estimate, the already-dilated segments are expanded even further to ensure that no object flux will bias the background determination. This very aggressive object mask is indicated with grey contours in Fig. 1. We use it for the joint-*gri* sky estimate here and also for the band-specific background determination detailed in Section 3.1.3 (performed in the same way).

For the galaxy fitting, however, we decided to use tighter segments that do not push that deeply into the sky. Besides speeding up the fit, this naturally results in the best possible fit to the inner, high signal-to-noise regions of the galaxy that we are most interested in, and reduces the sensitivity to background subtraction problems, flux from the wings of other objects, and features that cannot be captured by our models such as disc breaks and flares, and edge-on discs requiring the inclined disc model (van der Kruit & Searle 1981). Note, however, that this choice comes with some trade-offs, most notably that the fit frequently over-predicts the flux outside the segment boundary. We address this in more detail in Section 3.3.3.

To obtain these tighter segmentation maps, we run PROFOUNDPRO-FOUND again with the sky now fixed and a higher SKYCUT value of 2. This means that only pixels with a flux at least 2σ above the background level are considered in the segmentation, which ensures that fewer noise fluctuations are 'detected' and segment borders are smooth. In order to capture all flux of the galaxy wings, the segment for the object of interest (only) is then expanded further (using PROFOUNDMAKESEGIMDILATE) such that its area increases by typically around 30 per cent. This last step also ensures unbiased smooth borders of the segment since it is entirely independent of noise fluctuations. The resulting segmentation map is indicated with

coloured contours in Fig. 1 and is used for galaxy fitting in all bands, so that exactly the same pixels are fitted in each band (the segmentation statistics are of course re-calculated in each band).

3.1.3 Background subtraction

KiDS tiles are background-subtracted already; however, we opt to use the sky estimated by PROFOUND to even out inhomogeneities on smaller scales. For this, we split our 400 arcsec × 400 arcsec cutout into 16 square boxes and mask out all objects using the aggressively dilated segmentation map indicated with grey contours in Fig. 1 (cf. Section 3.1.2). The sky is then estimated as the median of the remaining (background) pixels in each box independently; and the solutions between the boxes interpolated with a bicubic spline.[9] This is done for each band independently; however, the segmentation maps used to mask out objects are the same in all bands.

This procedure for the background subtraction was chosen after extensive testing during pipeline development. In short, we found that the PROFOUND sky adopted here does not subtract object wings while still homogenizing the background well enough to avoid having to fit it along with the object of interest (introducing possible parameter degeneracies). It also decreases the sensitivity of the fit to the chosen segment size.

3.1.4 Sigma maps

Once the image segmentation and background subtraction is completed, we also calculate the sigma (error) map for each cutout (independently in each band). This is a combination of the KiDS weight map (where $\sigma = 1/\sqrt{\text{weight}}$) and the object shot noise. The latter is estimated as $\sqrt{N}$, where N is the number of photons per pixel (using positive-valued pixels only). This, in turn, is obtained by converting the image into counts using the gain provided in the meta-data associated with each KiDS tile.

3.1.5 PSF estimation

PSF fitting is performed on the background-subtracted 400 arcsec × 400 arcsec cutouts with corresponding masks and sigma maps in each band. The segmentation statistics returned by PROFOUND are used to identify isolated stars (round, bright, small, and highly concentrated objects with few nearby segments). More details on the star candidate selection are given in Appendix A. These objects are then fitted with a Moffat function using PROFIT; fitting all parameters except boxyness, i.e. the position, magnitude, full width at half-maximum (FWHM), concentration index, axial ratio, and position angle. Scale parameters are fitted in logarithmic space, a Normal likelihood function is used, initial guesses are taken from the segmentation statistics, and we use the BFGS algorithm from optim (R Core Team 2020), which is a fast downhill gradient optimization using a quasi-Newton method published simultaneously by Broyden (1970), Fletcher (1970), Goldfarb (1970), and Shanno (1970).

Some of the objects fitted above may not actually be suitable for PSF estimation as they can be too faint or bright (close to saturation), have irregular features, bad pixels, or additional small objects included in the fitting segment. Unsuitable objects are excluded by a combination of hard cuts in reduced chi-square (χ^2_ν),

---

[8]The sky variance can also be estimated, but in our case this is already provided as the KiDS weight maps and given to PROFOUND as an input.

[9]This is done by PROFOUNDPROFOUND internally; with the box size and the order of the interpolation spline being some of the variables we set.
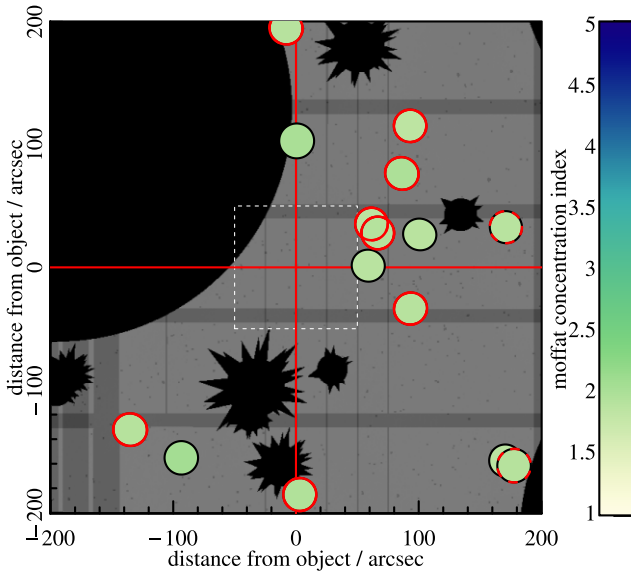
Figure 2. The result of the PSF fitting for the galaxy 396740 in the KiDS *r* band with the dashed white square indicating the cutout shown in Fig. 1. The greyscale image shows the *r*-band weight map with lighter colours meaning higher weight. Masked areas from the stacked *gri*-masks are shown in black (zero weight). The vertical and horizontal red lines indicate the position of the object of interest (galaxy 396740) and split the image into its four quadrants. All fitted PSFs are shown as coloured ellipses with the size (FWHM multiplied by 20), axial ratio, orientation angle, and concentration index (colour) taken from the fitted Moffat parameters. Stars selected for estimating the final model PSF have red borders; dashed red borders mean a fit was classified as suitable, but not selected because the maximum of eight stars was already reached.

position and magnitude relative to the PROFOUND estimates and an iterative 2σ-clipping in FWHM, concentration index, angle, and axial ratio. Again, more details can be found in Appendix A. Finally, we take the median of the Moffat parameters of a maximum of eight suitable stars (the closest two from each quadrant where possible to ensure an even distribution around the position of interest) and use these Moffat parameters to create a model PSF image. The size of the PSF image is adjusted to include at least 99 per cent of the total flux; or to a maximum of the median segment size within which the stars were fitted, with pixels in the corners of the image set to zero to avoid having a rectangular PSF.

Fig. 2 shows an example diagnostic plot of the PSF fitting result.

3.1.6 Outputs

For the fitting, we are only interested in the central galaxy and the closest neighbouring sources (for potential simultaneous fitting and to gain a better overview during visual inspection). Hence we do not save the entire 2000 × 2000 pixel cutouts used in the preparatory work as that would unnecessarily waste storage space and computational time used on reading and writing files. Instead, the image, corresponding mask, segmentation map, sigma map, and sky image are cut down to the smallest possible size that includes the object of interest (centred) and its neighbouring (touching) segments before saving. These five files, the model PSF image, and some ancillary information such as the segment statistics are the main outputs of the preparatory work pipeline and serve as inputs for the galaxy fitting, which we describe in the next section.

3.2 Galaxy fitting

3.2.1 Inputs and models

We use the Bayesian code PROFIT (Robotham et al. 2017) to perform two-dimensional multicomponent surface brightness modelling in each band independently, assuming elliptical geometry and a (combination of) Sérsic function(s) as the radial profile. The Sérsic (1963) function is described by three main parameters: the Sérsic index *n* giving the overall shape (with special cases $n = 0.5$: Gaussian; $n = 1$: exponential and $n = 4$: de Vaucouleurs (1948) profile), the effective radius R_e including half of the total flux, and the overall normalization which we specify as total magnitude *m*. In addition, in two dimensions the axial ratio *b/a* gives the ratio of the minor to the major axis of the elliptical model and the position angle PA its orientation, while *x* and *y* are used to define the position in RA and Dec. Throughout the paper, R_e refers to the effective radius along the major axis of the elliptical model. The Sérsic model is detailed in Graham & Driver (2005).

The data inputs for PROFIT are a background-subtracted image, corresponding mask, segmentation map, sigma map, and PSF. All of these are obtained during the preparatory steps (Section 3.1). On the modelling side, the main choices are the profile(s) to fit with initial parameter guesses and priors, the likelihood function to use, the fitting algorithm and convergence criteria; which are detailed in Sections 3.2.2–3.2.5. In short, we choose to fit each object with three different models in a four-step procedure:

(i) Single component Sérsic fits with initial guesses from segmentation statistics.

(ii) Double component Sérsic bulge plus exponential disc fits with initial guesses from single component fits.

(iii) Double component re-fits for a subset of galaxies which seemed to have the bulge and disc components swapped in step (ii), see Section 3.2.6.

(iv) '1.5-component' point source bulge plus exponential disc fits with initial guesses from double component fits.

Note that, for brevity, we will call the central component 'bulge' throughout this paper, even if it may not be a classical bulge. In particular, we do not distinguish classical bulges from pseudo-bulges, bars, AGNs, nuclear discs, combinations thereof, or anything else that may emit light near the centre of a galaxy. Hence, we also use the term 'bulge' for 1.5-component fits where the central component is unresolved and for double component central components with low Sérsic index and/or low axial ratios.

To implement our three models, we make use of two of the many models built into PROFIT, namely the Sérsic and point source models. We fit all parameters except boxyness (i.e. we do not allow deviations of components from an elliptical shape) and, for the double and 1.5-component models, tying the positions of the two components together. Exponential discs are implemented using a Sérsic profile with the Sérsic index fixed to 1. This leaves seven free parameters for our single Sérsic and 1.5-component models and 11 free parameters of the double component fits, which are summarized in Table 1. Scale parameters (Sérsic index, effective radius and axial ratio) are treated in logarithmic space throughout, i.e. the actual fitting parameters are $\log_{10}(X)$ for scale parameters *X*.

The 1.5-component model is needed for around 15–30 per cent of our double component systems where the bulge is too small relative to the image resolution to meaningfully constrain its Sérsic parameters (the exact number depends on the band due to the different PSF sizes). With the point source profile, at least we can determine the existence

Table 1. The fitting parameters for each of our three models.

Parameter	Single	Double		1.5-comp.	
		Bulge	Disc	Bulge	Disc
x-centre	Free	Free		Free	
y-centre	Free	Free		Free	
m	Free	Free	Free	Free	Free
$\log_{10}(R_e)$	Free	Free	Free	N/A	Free
$\log_{10}(n)$	Free	Free	1	N/A	1
$\log_{10}(b/a)$	Free	Free	Free	N/A	Free
PA	Free	Free	Free	N/A	Free
boxyness	0	0	0	N/A	0

of a second component and constrain its magnitude and hence the bulge-to-total (or AGN-to-total, bar-to-total, etc.) flux ratio.

If the centre of an object is masked or the PSF estimation failed (which happens if large fractions of the surrounding area are masked), then the object is skipped and no fits are obtained. This affects approximately 20 per cent of the galaxies. All other objects are fitted with all three models; and the best model is selected subsequently (see Section 3.3.2 on details of the model selection and Section 4.2 for the corresponding statistics).

3.2.2 Initial guesses

Since we use MCMC algorithms, our fits do not strongly depend on the initial guesses. However, reasonable starting parameters are still required for convergence within finite computing times.

The initial guesses for the single component Sérsic parameters are obtained directly from the segmentation statistics output by PRO-FOUNDPROFOUND (Section 2.5) where we use the position, magnitude, effective radius (R50), axial ratio, and angle as given; and the inverse of the concentration (1/con) for the Sérsic index.

For the double component fits, we convert the single component fits into initial guesses as follows: the position is taken unchanged, the magnitude of the single component fit is split equally between the two components, the bulge, and disc effective radii are taken as 1/2 and 1 times the single component effective radius, respectively, the Sérsic index of the bulge is set to 4 and its axial ratio to 1 (round), the disc axial ratio is set to the axial ratio of the single component fit and the position angles of both components are taken as that of the single component fit.

Initial guesses for the 1.5-component fits are taken from the double component fits (after making sure the components are not swapped, see Section 3.2.6), where the bulge magnitude is used as the point source magnitude and the disc parameters are taken unchanged.

3.2.3 Priors, intervals, constraints

All parameters are limited to fixed intervals. In addition, there can be constraints between parameters (such that, e.g. the bulge and disc positions can be tied together). If a (trial) parameter is outside the bounds of its interval or constraint during any step of the fitting process, PROFIT moves it back on to the limit before the likelihood is evaluated.

The limits for single-component fits are given in Table 2. In addition, the position angle is constrained such that if it leaves its interval, it is not just moved back on to the limit but jumps back 180° (which is the same angle, just more in the centre of the fitting interval).

Table 2. The fitting limits for single-component fits.

Parameter(s)	Lower limit	Upper limit
x- and y-centre	0	Cutout side length
Magnitude	10	25
Effective radius	0.5 pixels	$\sqrt{2}$ cutout side length
Sérsic index	0.1	20
Axial ratio	0.05	1
Position angle	$-90°$	$270°$

There are no additional priors or constraints for single component fits. This means that in effect, we use unnormalized uniform priors which are 1 everywhere in the respective interval and zero otherwise. For scale parameters (which are fitted in logarithmic space) the priors are uniform in logarithmic space, corresponding to Jeffreys (1946), i.e. uninformative, priors.

The limits and constraints for double and 1.5-component fits are the same as for the single component fits (for both bulges and discs), except for the magnitude where the individual component magnitudes have infinity as their upper limit and instead the *total* magnitude is constrained to be within the magnitude limits. This is most consistent and also allows the fitting procedure to discard one of the two components for systems which can equally well be fitted with a single Sérsic function (we then take this into account in the model selection).

Note that the above procedure results in unnormalized likelihoods. The lack of normalization does not impede our analysis because the only time when we compare likelihoods is during model selection, where we effectively fold the normalization into the calibration during visual inspection (Section 3.3.2).

3.2.4 Likelihood function

We use a Normal likelihood function for all fits. We have tested a t-distribution likelihood function which is less sensitive to outliers/unfittable regions; but found that the Normal likelihood function is better suited to our needs for several reasons.

First of all, the t-distribution fits often preferred to use the freedom of the bulge parameters to fit disc features instead (e.g. rings, bumps, flares, etc. that cannot be captured by the exponential model), treating the bulge as an outlier since the t-distribution prefers a few strong outliers (the bulge pixels) over many weak ones.

Secondly, the t-distribution fits fail for galaxies which are perfectly fitted by the model since then the errors truly are distributed Normally. This is a relatively common occurrence.

Hence some galaxies (~ 20 per cent) need to be fitted with a Normal distribution anyway, which, third, makes model selection much harder since the likelihood values obtained with different likelihood functions cannot easily be compared to each other.

3.2.5 Fit and convergence

All fits are performed on the sky-subtracted image within the galaxy segment only using the CONVERGEFIT function from the ALLSTARFIT package (Taranu 2022). This function uses a combination of different downhill gradient algorithms available in the NLOPTR package (Johnson 2017) followed by several MCMC fits with LAPLACESDEMON (Statisticat 2018) until convergence is reached.

The downhill gradient algorithms are used first to improve the initial guesses. The MCMC chain is not very sensitive to the initial guesses, but converges much faster if starting closer to the peak of the likelihood. Once the MCMC chains have converged, 2000 further

likelihood points are collected to ensure a stationary sample for the subsequent analysis of the galaxy.

We only fit the primary object of interest. While simultaneously fitting neighbouring sources is possible in PROFIT and might have improved the fit on a few objects, the effects are generally small since the galaxies we study are not in highly crowded fields and the segmentation process usually excludes the vast majority of the flux from other sources. This is especially true since we use tight fitting segments within which the galaxy flux is dominant (cf. Section 3.1.2); and considering that the watershed algorithm of PROFOUND cleanly separates even overlapping sources, so neighbours are automatically masked (Section 2.5). Hence we opted for the simpler and computationally cheaper option of just fitting the main objects. We confirm that this does not lead to major biases in Section 6.2.

An example fit for an object which is well-represented by our two-component model is shown in Fig. 3.

3.2.6 *Component swapping*

Approximately 20–30 per cent of the double component fits have their bulge and disc components swapped, i.e. the exponential component fitting the central region and the Sérsic component fitting the wings (this is a common problem in galaxy fitting, first pointed out by Allen et al. 2006). In particular, the freedom of the Sérsic component is often used to fit discs that do not follow pure exponential profiles while at the same time being the dominant component in terms of flux (which is the case for most galaxies). To solve this problem, we devised an empirical swapping procedure guided by the visual inspection of a subsample of our galaxies.

First, we select the galaxies that are most likely to have swapped components based on a cut in the plane of the ratio of Sérsic indices and the ratio of effective radii for the single component fits and the bulge of the double component fits. The reasoning for choosing this parameter space to calibrate the cut was that we would generally expect bulges alone to be more concentrated (i.e. smaller effective radius and higher Sérsic index) than when mixed with their respective discs in the single Sérsic fits. This results in approximately 30 per cent of our sample being flagged as possibly swapped, which we then re-fit in a second step.

The re-fit is performed in exactly the same way as the original fit, except that we now use the results of the previous double-component fit as initial guesses, swapping around the bulge and disc components (except for the bulge Sérsic index for which we use a value of 4). While the MCMC chain is less sensitive to initial guesses than a downhill gradient algorithm, it will still show some dependency for finite run-times. In particular, in our double component model the two components are *nearly* interchangeable with the only difference being the Sérsic index (fixed to 1 for the disc, free for the bulge). Hence there will always be two high maxima in likelihood space, which are far apart in the 11-dimensional parameter space. Moving from one to the other would require changing nine parameters (all except position) at once in the right direction and hence is statistically unlikely. Therefore, we assist the code in finding the other maximum by manually swapping the initial guesses.

In approximately 5 per cent of all re-fits, the code still converges on the same fit as before the swapping, but in most cases we find another likelihood maximum which corresponds to the bulge and disc components being reversed. As a third step we then select between the old and the new fit to obtain the physically more appropriate one. For this we first check whether either of the fits has a bulge Sérsic index

smaller than 2 *and* a bulge effective radius at least 10 per cent larger than the disc effective radius *and* a bulge-to-total ratio above 0.7 (i.e. the 'bulge' component is close to exponential, larger than the disc and contains the majority of the flux). If this is the case for only one of the fits, we choose the other one. If it is true for both or neither of the fits, then we apply our main criterion, which is that we choose the fit with the higher absolute value of bulge flux in the central pixel. These selection criteria are again based on visual inspection guided by the notion that we expect the bulge to be smaller and steeper than the disc and have proven to work very well. Note that the fit we select in this way is the one that is physically better motivated (i.e. with the bulge at the centre), and not necessarily the one which is statistically better.

After this procedure, the number of galaxies which still have the bulge and disc components swapped (and are classified as double component fits in model selection) is reduced to $\sim 1-2$ per cent.

3.3 Post-processing

3.3.1 *Flagging of bad fits*

After all three models have been fitted to all objects, we run them through our outlier flagging process (separately in each band). Each model is treated separately first; they are then combined during the model selection (Section 3.3.2).

The criteria for flagging bad fits (outliers) are: a very irregular fitting segment, an extreme bulge-to-total flux ratio, numerical integration problems, a parameter hitting its fit limits, poor χ^2 statistics, a large distance between the input and fitted positions, and a small fraction of model flux within the fitting segment. Additionally, there are some cautionary flags that identify fits which should be treated with extra care. All criteria are derived from and calibrated against visual inspection and described in more detail below. For orientation, we give the percentage of affected *r*-band fits in parentheses for each criterion. Note, however, that bad fits tend to fall into multiple of these categories, so the total number of bad fits is smaller than the sum of flagged objects in each category. Overall, approximately 9 per cent, 11 per cent, and 10 per cent of all non-skipped fits are flagged in the *g*, *r*, and *i* bands, respectively (after model selection, Section 3.3.2).

Very irregular segment (5.4 per cent): We calculate the difference between the magnitude of the model contained within the segment and the magnitude contained within the 'segment radius', which is defined as the maximum distance between the centre of the fit and the edge of segment. Objects where this magnitude difference is larger than 0.3 are flagged, as this is an indication for irregular segments (shredded, partly masked or cut off by another object for example). Note this criterion, as expected, often shows overlap with the criterion on the fraction of model flux contained within the segment (see description below).

Extreme bulge-to-total ratio (0.1 per cent): We flag double component and 1.5-component fits with a bulge-to-total ratio smaller than 0.001 or larger than 0.999 because in these cases the second component has negligible flux and a single component fit is better suited.

Numerical integration problems (0.2 per cent): PROFIT includes an oversampling scheme for accurate pixel flux integration where pixels containing steep flux gradients are recursively oversampled up to an oversampling factor of 4096; in the central pixel even up to $\sim 10^9$ (for more details see Robotham et al. 2017). However, for very extreme model parameters, even this procedure may not be accurate enough anymore, leading to significant errors in the pixel flux calculations. This could be improved by changing the default
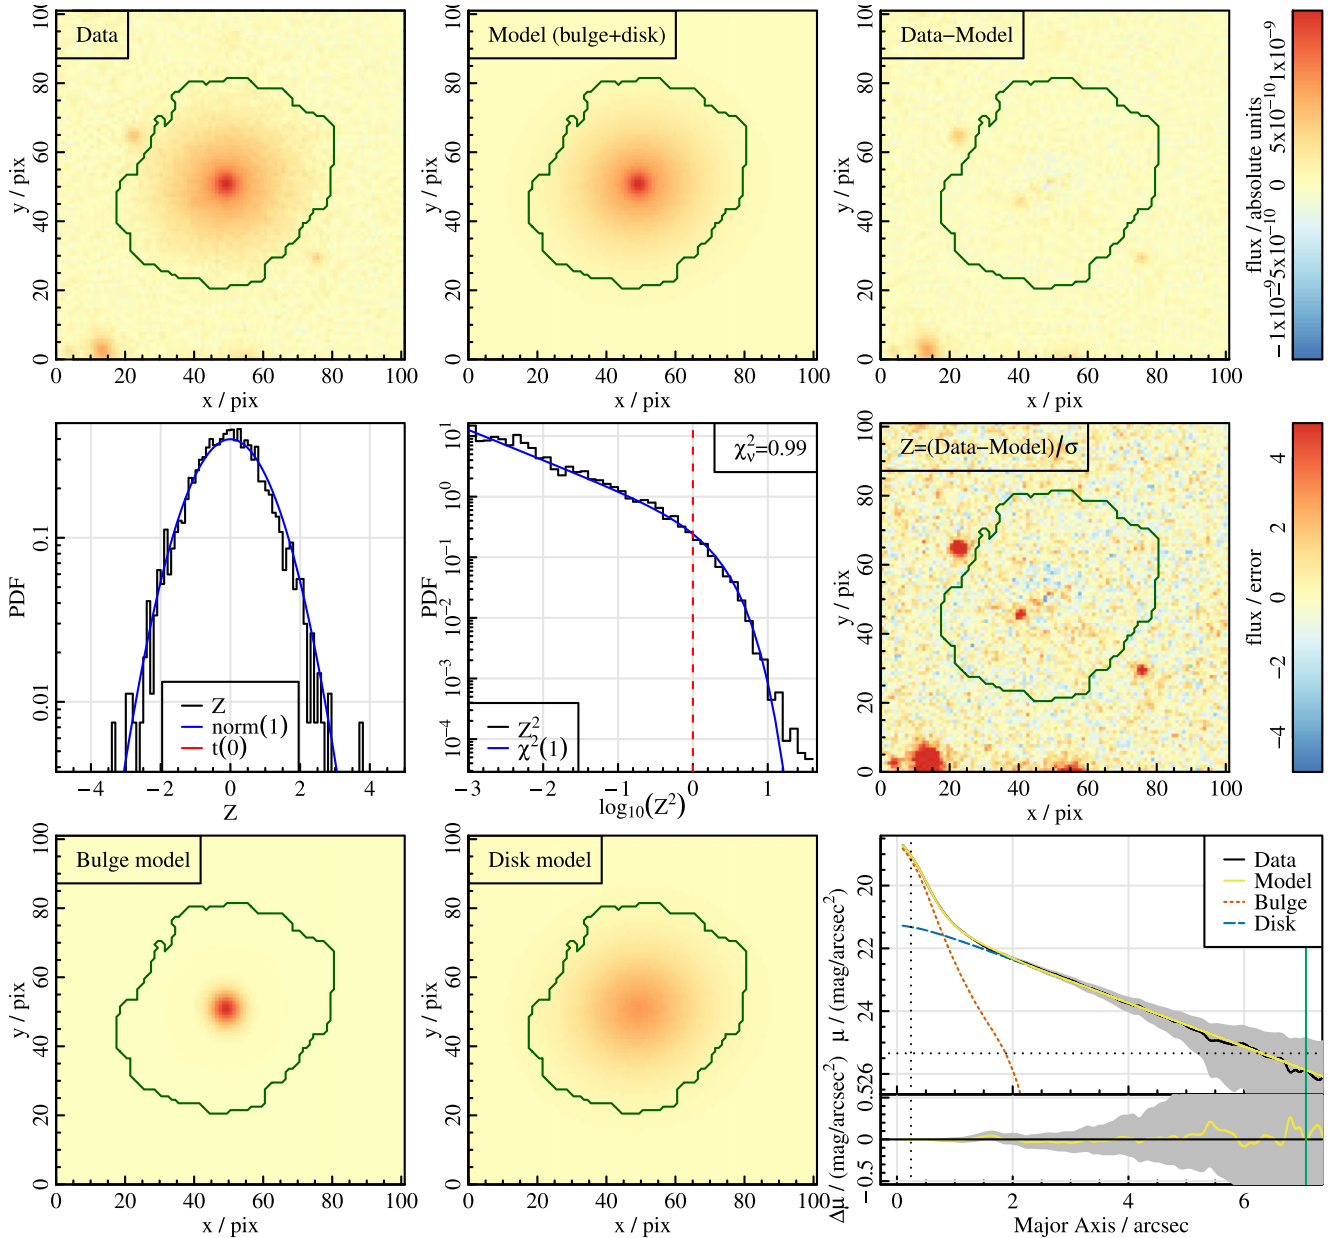
Figure 3. The result of the two-component (Sérsic bulge plus exponential disc) fit for the galaxy 611 298 in the KiDS r band. **Top row:** The data, two-component model and residual between them shown in absolute values of flux given by the colour bar on the right. The green contour indicates the segment used for fitting. Note that the flux scaling here is non-linear and optimized to increase visibility of galaxy features, but is the same in all three panels. **Middle row:** Goodness of fit statistics. The right-hand panel is the normalized residual Z (colour bar on the right) capped at $\pm 5\sigma$. The left-hand panel is the one-dimensional distribution (measured probability density function, PDF) of Z within the segment; with blue and red curves showing a Normal distribution with a standard deviation of 1 and a Students-t distribution with the relevant degrees of freedom for comparison. The middle panel shows the measured PDF of Z^2 compared to a χ^2-distribution of 1 degree of freedom (blue). The reduced chi-square, χ^2_ν (sum over Z^2 divided by the degrees of freedom of the fit) is given in the top right corner. **Bottom row:** The bulge and disc models in two dimensions on the same flux scale as the top row; and the bulge, disc, and total model compared against the data in one-dimensional form (azimuthally averaged over elliptical annuli). The FWHM of the PSF and the approximate 1σ surface brightness limits are indicated by vertical and horizontal dotted lines for orientation. The vertical solid green line indicates the segment radius beyond which our model is extrapolated. The pixel scale is 0.2 arcsec for KiDS data, i.e. 1 arcsec corresponds to 5 pix.

oversampling values to achieve higher accuracy (at the cost of increased computational time); however, we opted for simply excluding those cases since usually this only happens for unresolved bulges which are better represented by the 1.5-component fits anyway.

Parameter hitting limit (5.8 per cent): We flag objects where the magnitude, effective radius, or Sérsic index hit either of their limits (cf. Section 3.2.3); or the axial ratio hit its lower limit (for double component fits this applies to both components individually). The axial ratio upper limit is not flagged because fits are allowed to be exactly round, but there is a cautionary flag for all objects which hit any of its parameter limits (6.5 per cent). We also add a cautionary flag for suspiciously small or large errors on any parameter, where 'suspicious' is defined as being an outlier in the respective distribution of errors (2.1 per cent).

Poor χ^2 statistics (0.1 per cent): We flag fits with a χ_ν^2 larger than 80; or where the χ^2 in the central pixel is more than 1000 times larger than the average χ^2 per pixel since that is an indication that the bulge was not fitted.

Large distance between input and output position (0.3 per cent): We flag fits with a distance between the input and output position of more than 2 arcsec (10 pix), which are usually highly asymmetrical objects, mergers, objects with very nearby other objects (especially small objects embedded in the wings of much larger objects), or objects in regions of the image with unmasked instrumental effects. Often the fitted object then is not the one that we intended to fit. There is also a cautionary flag for offsets above 1 arcsec (1.3 per cent).

Small fraction of model flux within fitting segment (1.4 per cent): We flag fits where the amount of model flux (of any component) that falls within the fitting segment is less than 20 per cent. With so little flux to work on PROFIT cannot constrain the parameters well anymore and these are often objects which are cut off by a masked region (e.g. a bright star) or other nearby objects. There is a cautionary flag for objects where the fraction of model flux (of any component) that falls within the segment is less than 50 per cent (9.3 per cent).

3.3.2 Model selection

In Bayesian analysis, model selection is performed by computing the posterior odds ratio between two models, $O_{1:2}$, which is the ratio of the probabilities of the models given the same set of data and background information. With the help of Bayes' theorem and assuming a prior odds ratio of 1, this becomes the Bayes factor (see e.g. Sivia & Skilling 2006 for a detailed treatment):

$$O_{1:2} = \frac{\mathrm{p}(data|M_1, I)}{\mathrm{p}(data|M_2, I)}, \tag{1}$$

where p($data|M_1, I$) is the probability of the data given Model 1 and any background information I (the marginalized likelihood of Model 1). In practice, these probabilities are often difficult to compute because they require marginalizing (i.e. integrating) over all model parameters.

Hence, many information criteria tests have been developed which are based on the (non-marginalized) likelihood (or χ^2) combined with some penalty term depending on the number of model parameters. This penalty term serves to judge whether a more complicated model is justified and takes the role of Ockham's factor (which is automatically included in equation 1 due to the integration over all parameters). Commonly used tests include the Akaike information criterion (AIC; Akaike 1974), the Bayesian information criterion (BIC; Schwarz 1978), or the deviance information criterion (DIC; Spiegelhalter et al. 2002). We choose to use the deviance information criterion, which is usually recommended over the AIC or BIC in Bayesian analysis (Hilbe, de Souza & Ishida 2017) and straightforward to compute from an MCMC output. Brief tests using the BIC or the estimated log marginal likelihood output by LAPLACESDEMON showed similar results.

The DIC is a direct output of the LAPLACESDEMON function (see Section 3.2.5) and is defined as

$$\mathrm{DIC} = Dev + pD = Dev + \mathrm{var}(Dev)/2, \tag{2}$$

where pD is a measure of the number of free parameters in the model and $Dev = -2 \times$log-likelihood is the deviance. In theory, then, if the DIC difference ΔDIC between two models is negative, the first model is preferred and if it is positive, then the second model is preferred; with differences larger than approximately 4

being considered meaningful (Hilbe et al. 2017). However, for the case of galaxy fitting where many features are present that cannot be captured by the model (bars, spiral arms, disc breaks or flares, tidal tails, mergers, foreground objects, etc.), we want to choose the model that we consider physically more appropriate rather than better in a strictly statistical sense. This requires visual classification, logical filters, detailed simulations, or a manual calibration of the ΔDIC cut (or whichever other chosen diagnostic) by visual inspection of a representative sub-sample (e.g. Allen et al. 2006; Simard et al. 2011; Vika et al. 2014; Argyle et al. 2018; Kruk et al. 2018). We choose the latter approach, which has the added advantages that we do not need to worry about normalizing our likelihoods (cf. Section 3.2.3); hence, circumventing dependencies of the results on prior widths; nor the fact that our pixel values are correlated (due to the PSF) – these effects are simply folded into the visual calibration.

We use a random sample of $\sim$700 non-skipped objects per band (i.e. $\sim$2000 objects in total) for the calibration; and a further 1000 r-band objects that were previously inspected for cross-checking the results. In addition, our model selection procedure takes into account some of the outlier flagging (Section 3.3.1). For each of the $\sim$700 objects in each band, SC visually inspected the fits of all three models and classified the object into one of the categories: 'single component', '1.5-component', 'double component', 'not sure if 1.5 or double component', 'not sure at all', 'unfittable' (outlier). We then calculate the DIC differences between all three models (i.e. ΔDIC$_{1-1.5}$, ΔDIC$_{1-2}$, and ΔDIC$_{1.5-2}$) and calibrate them for model selection in two steps: first, we select between single component fit or not; of the ones that are not single component fits we then select between double component or 1.5-component fits.

For the first step of model selection calibration, the ΔDIC$_{1-1.5}$ and ΔDIC$_{1-2}$ cuts are optimized such that the minimum number of fits is classified wrongly. 'Wrong' in this case means a fit was manually classified as 'single' but is now a double/1.5; or a fit was manually classified as '1.5', 'double', or 'not sure if 1.5 or double' but is now a single. 'Unfittable' and 'not sure at all' cases are ignored. For the second step of model selection calibration, the ΔDIC$_{1.5-2}$ cut is optimized in the same way; where 'wrong' now means that the fit was manually classified as '1.5' but is now a double or vice versa, with all other categories being ignored. For the two steps of the calibration, we bootstrap the manual sample 1000 and 500 times respectively and repeat the optimization to get an estimate of the error on the chosen DIC cuts. These errors are chosen as the 1σ quantiles (i.e. they contain the central 68 per cent of DIC cut distributions). Our calibrated DIC cuts hence all have a median, a lower limit, and an upper limit. Any object within these limits is flagged as unsure in the model selection, i.e. the DIC differences are not conclusive for this object.

To perform the actual model selection, the calibrated DIC cuts in each band are then applied to the entire sample, again in a two-step procedure: the single component fit is selected if neither of the 1.5 or double component fits are significantly better (as indicated by the DIC differences). Double component fits need to be significantly better than 1.5-component fits, too. In all cases, if the DIC difference is very clear, we do the model selection first; then flag objects as outliers if needed.[10] In the unsure region of the DIC difference, we

---

[10]This means that it is possible (and not uncommon) that a galaxy which is classified as an outlier has a non-flagged fit in another model (but the fit that was chosen was significantly better than the other one, despite it being an outlier).

Table 3. The confusion matrix for our model selection based on a DIC difference cut compared against visual inspection for the *r* band. All values are in per cent of the total number of visually inspected *r*-band galaxies. Bold font highlights galaxies classified correctly, while grey shows those that were ignored during the calibration.

	Number of components		
Visual classification	1	1.5	2
'single'	**41.6**	0	2.7
'1.5'	2.2	**2.4**	0.9
'double'	3.1	0.1	**9.2**
'1.5 or double'	0.3	**0.6**	**3.0**
'unsure'	16.1	0.4	13.1
'unfittable'	0.9	0.6	2.7

choose the model that is not flagged as outlier; if neither is flagged, the DIC cut is applied.

Compared against visual inspection (keeping in mind that visual classification is not free of errors either), roughly 7 per cent, 9 per cent, and 6 per cent of the galaxies end up in the wrong category in total in the *g*, *r*, and *i* bands, respectively (in both steps of model selection combined, ignoring cases which were visually classified as 'unsure'). Table 3 gives the detailed confusion matrix for the *r* band. Note that we do not consider the success of the outlier flagging here, so for outliers we show what the galaxy would have been classified as if it were not flagged (absolute value of the NCOMP column in our catalogue). We highlight those galaxies that are correctly classified in bold and show those that were ignored during the model selection calibration process in grey font. The remaining (black) numbers add up to the 9 per cent quoted above. Corresponding confusion matrices for the *g* and *i* bands are given in Appendix B (Tables B1 and B2). Note that since we minimize the *total* number of fits classified wrongly, there is a slight bias against the rarer categories in the automated model selection. For example, the relative fraction of true 1.5-component objects (as per the visual inspection) that is classified wrongly by the automated selection is higher simply because 1.5-component objects are much rarer than single or double component objects.

The accuracy of the model selection is also confirmed using simulations, to the extent to which our simulations allow us to do so (see Section 6.3 for details).

3.3.3 Truncating to segment radii

As detailed in Section 3.1.2, we produce segmentation maps that define the fitting region, meaning that only pixels within the fitting segment are considered during the evaluation of the likelihood of the model (equivalent to giving all pixels outside the segment zero weight in the fit). We choose tight fitting segments (cf. Section 3.1.2) in order to obtain the best possible fit in the inner, high signal-to-noise ratio regions of the galaxies, and be less sensitive to disc breaks, flares, nearby other objects, sky subtraction problems, and similar. The disadvantage of this approach is that profiles are not necessarily forced to zero for large radii, i.e. our Sérsic fits often show unphysically large effective radii combined with high Sérsic indices.

To mitigate this effect, we define a 'segment radius' for each galaxy segment, which is simply the maximum distance between the fitted galaxy centre and the edge of the segment and can be understood as the upper limit to within which our model is valid. We then calculate the 'segment magnitude', m_{seg}, which is the magnitude

of the (intrinsic, not PSF-convolved) profile integrated to the segment radius (rather than infinity); and the 'segment effective radius', $R_{e, seg}$, which is the radius containing half of the flux defined by the segment magnitude. These values (and quantities derived from them, such as segment bulge-to-total flux ratios) are provided in the catalogue (labelled *_SEGRAD) and we strongly recommend using these instead of the Sérsic values integrated to infinity whenever they are available. For a direct parameter comparison to other works, the values in those catalogues should also be appropriately truncated.

In the following, we explain this recommendation in more detail; with further points to note in Sections 5.4 and 6.4.

Fig. 4 illustrates the effects produced by our tight fitting segments, how to mitigate those by truncating the magnitude, and effective radius appropriately; and the circumstances under which this correction is necessary. For two example galaxies – 21 4264 and 3896 188 – we show a detailed comparison of our single Sérsic fit to a fit using a larger segment and to the fit obtained in Kelvin et al. (2012). We present a more general (statistical) comparison of our fit results to those of Kelvin et al. (2012) in Section 5.4, where we also give more details on how their fits were derived. For the purposes of the analysis in this section it suffices to say that Kelvin et al. (2012) used much larger fitting regions than we do, while the remaining analysis is in many ways analogous to ours (although they use different data, code, and procedures in detail).

Focussing on the left half of Fig. 4 first, the first six panels (top two rows) show the KiDS *r*-band data, our single Sérsic model, the residual, and various goodness of fit statistics as described in the caption of Fig. 3. Panels 7–12 (rows 3 and 4) show the same for a larger fitting segment as indicated by the green contour. Panels 13–18 again show the same for the Kelvin et al. (2012) fit, where we note that this was originally performed on *r*-band Sloan Digital Sky Survey (SDSS; York et al. 2000) data but is now evaluated on the *r*-band KiDS data. The Kelvin et al. (2012) fits were performed on cutouts larger than the size shown here, i.e. they include all visible pixels (and more) in the fit. Note that the reduced chi-square value quoted in the bottom middle panel of each set of plots always is evaluated within the smallest segment so that they can be directly compared. Finally, the bottom two panels show a direct comparison of the one-dimensional profiles of all three fits, which we will now study in detail.

In the top panel of this one-dimensional plot, we show the surface brightness (azimuthally averaged over elliptical annuli) against the projected major axis for the data (solid black line with grey uncertainty region), our model fit for the fiducial segment (dashed blue line) and the larger segment (dash–dotted pink line) and the Kelvin et al. (2012) model fit (dotted orange line). The vertical green solid and dashed lines indicate the segment radii (for the two segment sizes respectively) beyond which our model is an extrapolation. The vertical dotted line shows the half width at half-maximum (HWHM) of the PSF and the horizontal dotted line is the 1σ surface brightness limit of the data. The inset in the bottom left of this plot shows the fitted magnitude *m*, effective radius R_e in arcseconds, and Sérsic index *n* values for our and the Kelvin et al. (2012) fits; and the corresponding segment-radius-truncated values for *m* and R_e. Finally, the bottom panel shows the difference between all three models and the data (with errors): our fiducial fit in blue with a dashed line, the fit in the larger segment in pink with a dash–dotted line and the Kelvin et al. (2012) fit in orange with a dotted line.

Clearly, our model is a better fit to the inner regions of the galaxy than the Kelvin et al. (2012) fit (out to about 2 arcsec, also evident
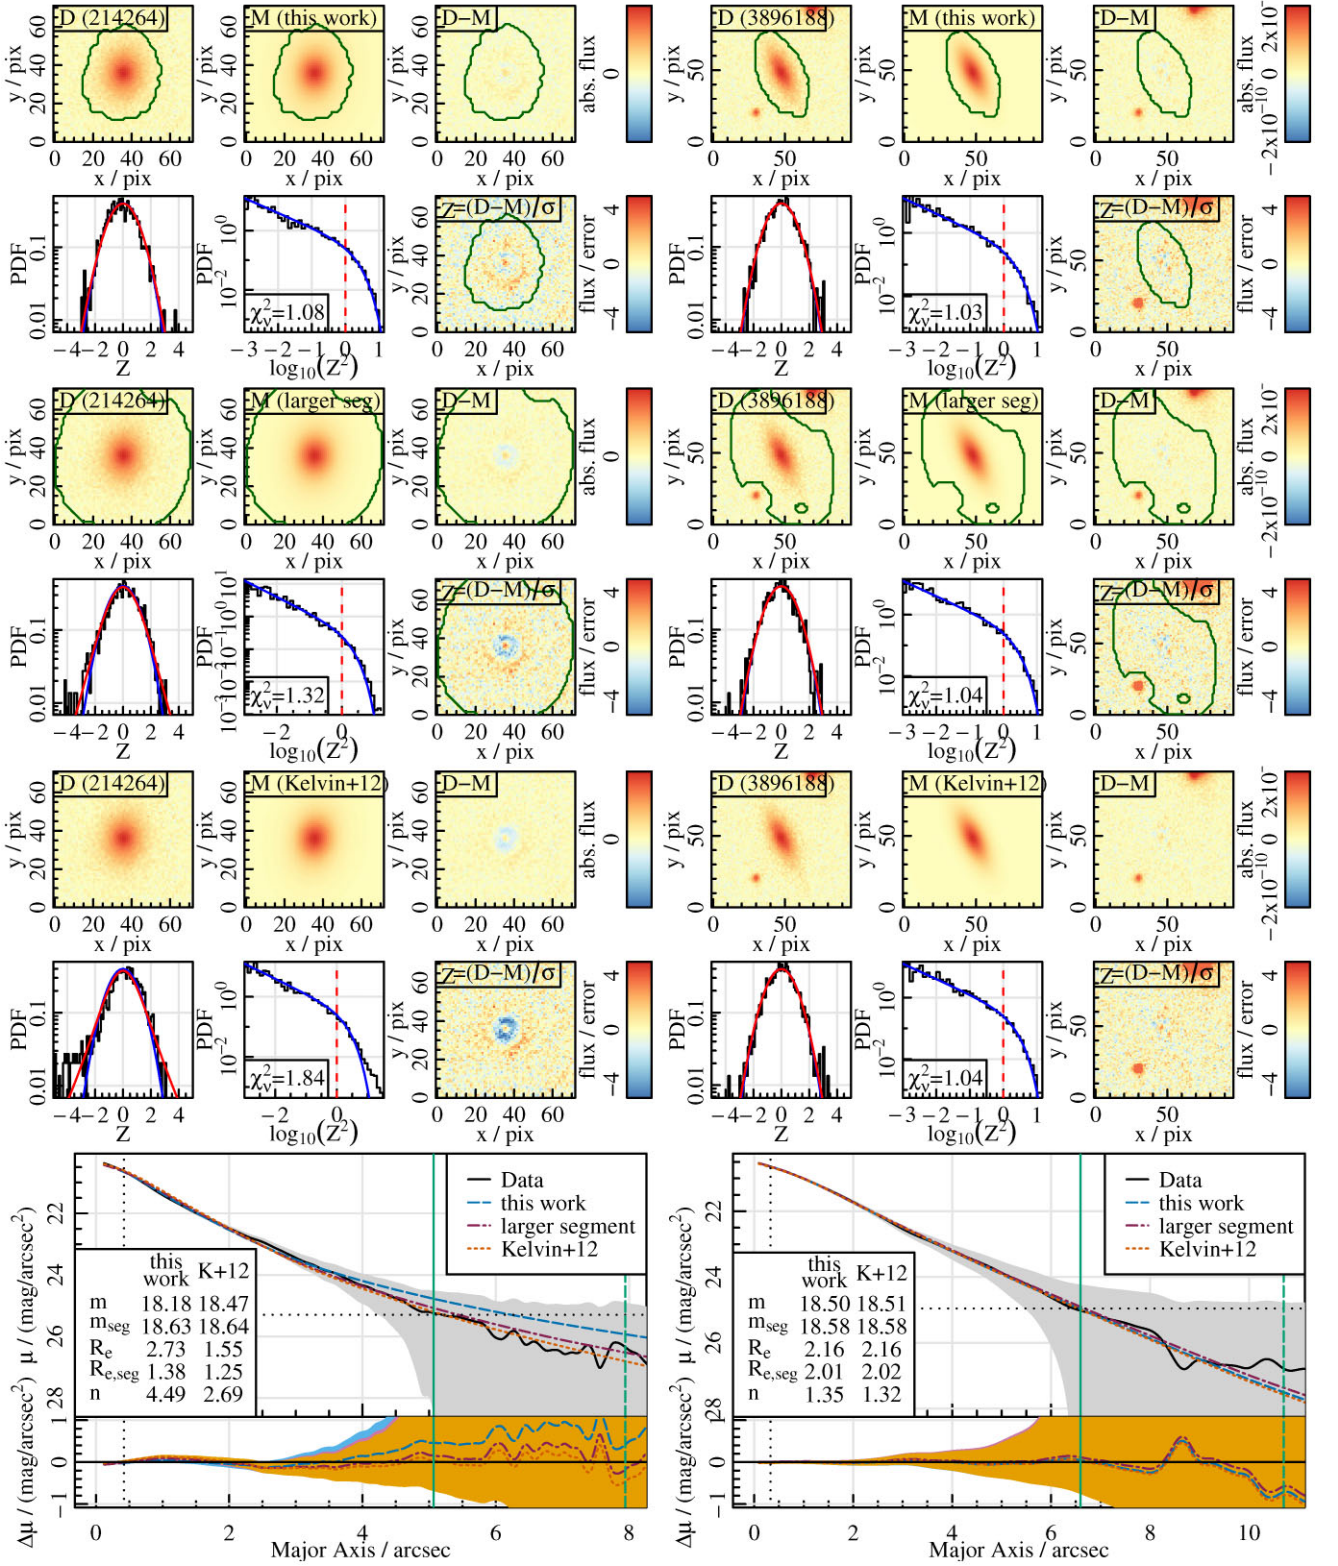
Figure 4. **Left:** Detailed comparison of our single-Sérsic fit, our fit using a larger segment, and the Kelvin et al. (2012) fit to the galaxy 214 264, which in reality is a 1.5-component system. **Right:** The same for galaxy 3896 188, which is well-described by a single Sérsic component. **Top two rows:** Our fit to the KiDS *r*-band data with panels the same as those in Fig. 3. **Rows 3 and 4:** The fit we obtained by using a larger fitting segment as indicated. **Rows 5 and 6:** The Kelvin et al. (2012) fits (originally performed on SDSS *r*-band data) evaluated on the KiDS *r*-band data, which use a fitting region larger than the cutout shown. **Bottom two panels:** Direct comparison of the one-dimensional profiles, see text for details. The vertical green solid and dotted lines indicate the segment radii for the two segment sizes. The vertical dotted line shows the HWHM of the PSF; the horizontal dotted line is the 1σ surface brightness limit of the data. The quoted values for the effective radius are in arcseconds.

from the two-dimensional plots and from the reduced χ^2-value within the segment decreasing from 1.84 to 1.08), owing to the higher Sérsic index which better represents the steep bulge at the centre. However, it has a large effective radius and considerable amounts of model flux at large radii which are not observed in the data. In particular in the region beyond the segment radius, where our model is merely extrapolated, it is clearly oversubtracting the data (also visible in the two-dimensional plots). Correspondingly, the truncated segment quantities differ substantially from the fitted Sérsic values. The Kelvin et al. (2012) fits, instead, use a larger fitting region and hence follow the data out to larger radii, which results in a worse fit of the central regions but does not contain such large amounts of excess flux beyond the surface brightness limit. Hence, truncating to segment radii has a smaller effect on the parameter values. The truncated values for both models are then in reasonable agreement with each other, except for the Sérsic index, for which no truncated version exists as it would be unclear how to define such a value. Our fit in the larger segment is in between the two others in all respects, since it has a fitting region intermediate to the other two.

Note that the differences only come about when the model is not (in a formal statistical sense) a good representation of the data, i.e. when there is a need to compromise between fitting different regions. In the case of the left side of Fig. 4, the galaxy shown is better described by a 1.5-component model ($m^B = 20.47$, $m^D = 18.79$, $R_e^D = 1.89$ arcsec), although in general there are many objects in our sample for which even a two-component model cannot capture all aspects of the data. For comparison, in the right half of Fig. 4, we show a galaxy that is well-described by a single Sérsic model: here, both our and the Kelvin et al. (2012) fits arrive at virtually the same solution despite the different fitting regions. In fact, all three models and the data are nearly indistinguishable all the way down to the 1σ surface brightness limit.

In short, there is no perfect way to fit a Sérsic function to an object which intrinsically does *not* have a pure Sérsic profile. For such objects, which unfortunately comprise the majority of our sample, the fitted parameters will always depend on the exact fitting region used as well as the quality of the data (its depth in particular). Most previous work, including Kelvin et al. (2012), opted to use large fitting regions in order to include enough sky pixels to ensure that the profiles are constrained to approach zero flux at large radii (although a Sérsic function technically never reaches zero exactly). Here, we choose a different approach by using smaller fitting segments. This means that the profiles are not constrained to approach zero flux at large radii. Instead more emphasis is placed on adequately representing the inner regions of the galaxies. We choose this approach since it is most appropriate for our science case, where we are primarily interested in comparing the high signal-to-noise regions of galaxies from the same data set amongst each other. In addition, it decreases the sensitivity of our fits to deviations from a Sérsic profile in the low surface brightness wings of objects (arguably no galaxy truly follows a Sérsic profile to infinity) as well as nearby other objects and inaccuracies in the sky subtraction. We stress that this means that our parameters are not directly comparable to other works using larger fitting segments. In particular, our Sérsic indices tend to be systematically higher (see Section 5.4) since high Sérsic indices result in high amounts of flux at large radii and are hence suppressed when constraining the models to zero flux at large radii. Magnitudes and effective radii can be compared to those of other studies by truncating to segment radii.

4 RESULTS

4.1 `BDDecomp` DMU

Our main result is the `BDDecomp` DMU in the GAMA database. It contains eight catalogues: `BDInputs` (three times, one for each band) with the most important outputs of the preparatory work pipeline (segmentation, PSF estimation, initial guesses), `BDModelsAll` (three times, one for each band) with the output from the actual galaxy fitting and post-processing (model selection, flagging of bad fits and truncating to segment radii) and `BDModels` which combines the most important columns of the three `BDModelsAll` tables and has a few additional joint columns (mainly joint model selection). Finally, the table `BDModelsAlt` presents the same information as `BDModels` just with the three bands arranged in rows instead of columns. Each table is accompanied by comprehensive documentation including descriptions of all columns, details on the processing steps and practical tips for using the catalogue. The DMU also provides all input data used for the fitting (i.e. image cutouts, masks, error maps, segmentation maps, sky estimates, PSFs) as well as various diagnostic plots of the fit results on the GAMA file server, where detailed descriptions of these files can be found.

In the following sections we present an overview over the contents of the main catalogue, `BDModels`.

4.2 Catalogue statistics

Table 4 gives an overview of the fit and post-processing results. Starting with our full sample (13 096 galaxies from the combination of our main and SAMI samples, see Section 2.3), we show how the number of galaxies evolves through all steps of the pipeline. The results are split per-band and per-model where necessary. At some steps, we also include percentages of galaxies lost or remaining (grey font). In short, we lose nearly 20 per cent of our sample to masking and a further almost 10 per cent to the flagging of bad fits; where the former is a random subset while the latter preferentially affects certain types of galaxies (e.g. mergers and irregulars).

Note that we used stacked *gri* images for segmentation and masking, but then treated the galaxies independently in all bands except for the model selection, where we performed both a per-band and a joint version. Therefore, the column 'joint *gri*' always gives the number of galaxies that were 'good' in all three bands (hence why numbers are generally lower), except for the model selection, where it shows the results of the joint model selection (cf. Section 3.3.2).

Fig. 5 visualizes the most important information given in Table 4, namely the final number of objects classified in each category: lighter bars in the background refer to individual fits (total 14 966) with the number of unique galaxies (total 13 096) overplotted. When several fits to the same galaxy were classified in different categories, we allocate it to the highest of those,[11] which is consistent with Table 4. NCOMP = −999 means the object was skipped (not fitted) because it is masked or the PSF estimation failed (usually because of large masked areas in the immediate vicinity of the object). NCOMP = 1,

---

[11]This means that a galaxy is classified as 'outlier' if *all* fits to it are outliers and it is 'skipped' only if *all* fits are skipped. Galaxies with good fits are allocated to the most complex model of the available fits (assuming that one of the images was deeper and allowed to constrain more components than the other(s)), while within the outlier categories we allocate it to the simplest model. Note that in Table 4 we only show the total number of flagged fits and do not split them into the different outlier categories.

Table 4. Fit results: numbers (black) and percentages (grey) of galaxies remaining or lost at each step in our pipeline, split per-band, and per-model where necessary.[1]

Band	g			r			i			Joint gri		
Model (components)	1	1.5	2	1	1.5	2	1	1.5	2	1	1.5	2
Number of:												
(1) Unique objects (galaxies)						13 096						
(2) Images (independent fits)						14 966						
(3) Images not masked						11 989						
lost due to masking (per cent)						20						
(4) Successful PSFs		11 838			11 872			11 946			11 683	
lost due to PSF fails (per cent)		1			0.8			0.3			2	
(5) Successful fits	11 837	11 837	11 831	11 872	11 870	11 861	11 946	11 943	11 945	11 682	11 678	11 665
lost due to fit fails (per cent)	<0.01	<0.01	0.05	0	0.01	0.07	0	0.02	<0.01	<0.01	0.03	0.12
(6) Fits not flagged	10 951	7122	8022	11 025	8164	8759	11 086	7620	7775	10 680	6446	5870
not flagged/successful (per cent)	93	60	68	93	69	74	93	64	65	91	55	50
(7) Selected fits	8294	740	1743	7061	585	2935	7411	662	2663	7308	621	2009
selected/successful (per cent)	70	6	15	59	5	25	62	6	22	63	5	17
total number (per band) of:												
(8) Good | flagged | skipped fits		10 777 | 1061 | 3128			10 581 | 1291 | 3094			10 736 | 1210 | 3020			9938 | 1745 | 3283	
good | f. | s./all images (per cent)		72 | 7 | 21			71 | 9 | 21			72 | 8 | 20			66 | 12 | 22	
(9) Good | flagged | skipped gal.		9722 | 935 | 2439			9545 | 1145 | 2406			9687 | 1059 | 2350			8998 | 1559 | 2539	
good | f. | s./unique objects (per cent)		74 | 7 | 19			73 | 9 | 18			74 | 8 | 18			69 | 12 | 19	

Notes/explanations of each step: (1) The full sample results from the combination of our main and SAMI samples (Section 2.3). (2) Some galaxies have been imaged more than once due to overlap regions between KiDS tiles. These duplicate observations of the same physical objects are treated independently throughout our pipeline (Section 2.3). (3) We use the associated KiDS masks, combining the three bands. Images for which the central galaxy pixel is masked ($\sim$ 20 per cent) are skipped during the fitting (Section 3.1.1). (4) For each image in each band, a PSF is then estimated by fitting nearby stars. If the PSF estimation fails, the galaxy is skipped during the fit (Section 3.1.5). Note that technically, we estimate PSFs also for galaxies that are masked in step (3), but we do not list those here. (5) For each non-masked image with a successful PSF estimate, we attempt three fits: a single Sérsic (1), a pointsource + exponential (1.5) and a Sérsic + exponential (2). Very rarely, the fit attempts fail with an error (Section 3.2.5). (6) Each fit (for each model independently) is passed through our outlier flagging process, identifying bad fits (Section 3.3.1). (7) Of the non-flagged (i.e. good) fits, we then select the most appropriate one during model selection (Section 3.3.2). (8) Summing up the selected fits for each model (step 7) gives the total number of good fits. The difference between the good and successful fits (step 5) stems from the outlier flagging. Skipped fits are due to masking, PSF or fit fails (steps 3, 4, 5). The sum of good, flagged and skipped fits gives the total number of independent fits (step 2). (9) Removing duplicate observations for the same physical objects gives the number of good, flagged, and skipped galaxies, which sum to the number of unique objects (step 1). Here we always use the best available result for each galaxy, i.e. it is counted as 'good' if at least one of the multiple observations was 'good'. [1] Based on information given in the *_BDQUAL_FLAG, *_OUTLIER_FLAG, and *_NCOMP columns of the BDModels catalogue.
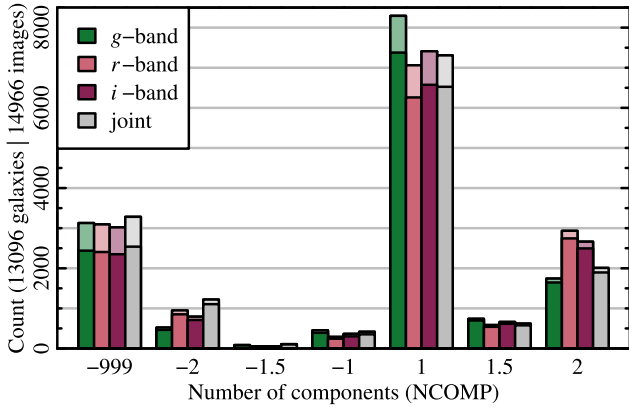
Figure 5. The number of components assigned in our model selection procedure for individual bands and the joint analysis. 1, 1.5, and 2 mean single Sérsic, point source bulge + exponential disc and Sérsic bulge + exponential disc models, respectively. Negative values indicate that the chosen (best) fit was flagged as unreliable (mostly irregular or partly masked galaxies). −999 is assigned to skipped fits, either because the galaxy centre is masked (most cases) or because the PSF estimation failed. The lighter (higher) bars show the number of images, whereas the shorter bars in the foreground indicate the number of unique objects. See text and Table 4 for details.

1.5, or 2 indicates that this is a good fit classified as single, 1.5 or double component fit. NCOMP = −1, −1.5, or −2 indicates that this is a bad fit (outlier) which would have been classified as single, 1.5 or double component fit if it were not an outlier (most often these are mergers/irregular galaxies for which our models are not appropriate; or galaxies that are partly masked). We keep these three classes separate since automated outlier identification can never be perfect; and what should be considered a bad fit will depend on the use case. The flagging of fits is hence only intended as a guide and all available information in the catalogue is retained for all fitted

objects. Fig. 3 shows an example two-component fit. Examples for a single, 1.5-component and outlier can be found in Appendix C.

4.3 Parameter distributions

Fig. 6 shows the distribution of the main parameters – magnitude, effective radius and Sérsic index – in all three bands (*g*, *r*, and *i*) for single Sérsic fits, bulges, and discs. The single Sérsic fit distributions are shown for all galaxies with NCOMP > 0 (i.e. all non-outliers) in black and for those galaxies which were actually classified as single component systems (NCOMP = 1) in yellow. Red dotted and blue dashed lines show bulges and discs, respectively. For discs we show the 1.5-component fits and double component fits combined (i.e. the 1.5-component parameters for objects with NCOMP = 1.5 and double component parameters for those with NCOMP = 2 added into one histogram); the Sérsic index is not shown since it was fixed to 1. Bulge magnitudes are also shown for 1.5 and double component fits combined; effective radii and Sérsic indices are only shown for the double component fits since they do not exist in the point source model. The legend indicates the numbers of objects in each histogram, which can also be inferred from Table 4. Magnitudes and effective radii are truncated at the segment radii which we found to give more robust results than using the Sérsic values extrapolated to infinity (see Sections 3.3.3, 5.4, and 6.4).

The first thing apparent from Fig. 6 is that the distributions in the three bands are generally very similar, which is reassuring given that the fits were performed independently. Looking at the first column, the single Sérsic number counts increase up to a sharp drop just before 20 mag in all bands, which is not surprising given the GAMA survey limit of 19.8 mag. The faintest of these objects are all classified as single component galaxies (yellow line is on top of black line), while some of the brighter objects are successfully decomposed into bulges and discs. Discs are generally slightly brighter than bulges. The bulges show a second, smaller peak at very faint magnitudes which we found to be the ones from the 1.5-component fits (unresolved,
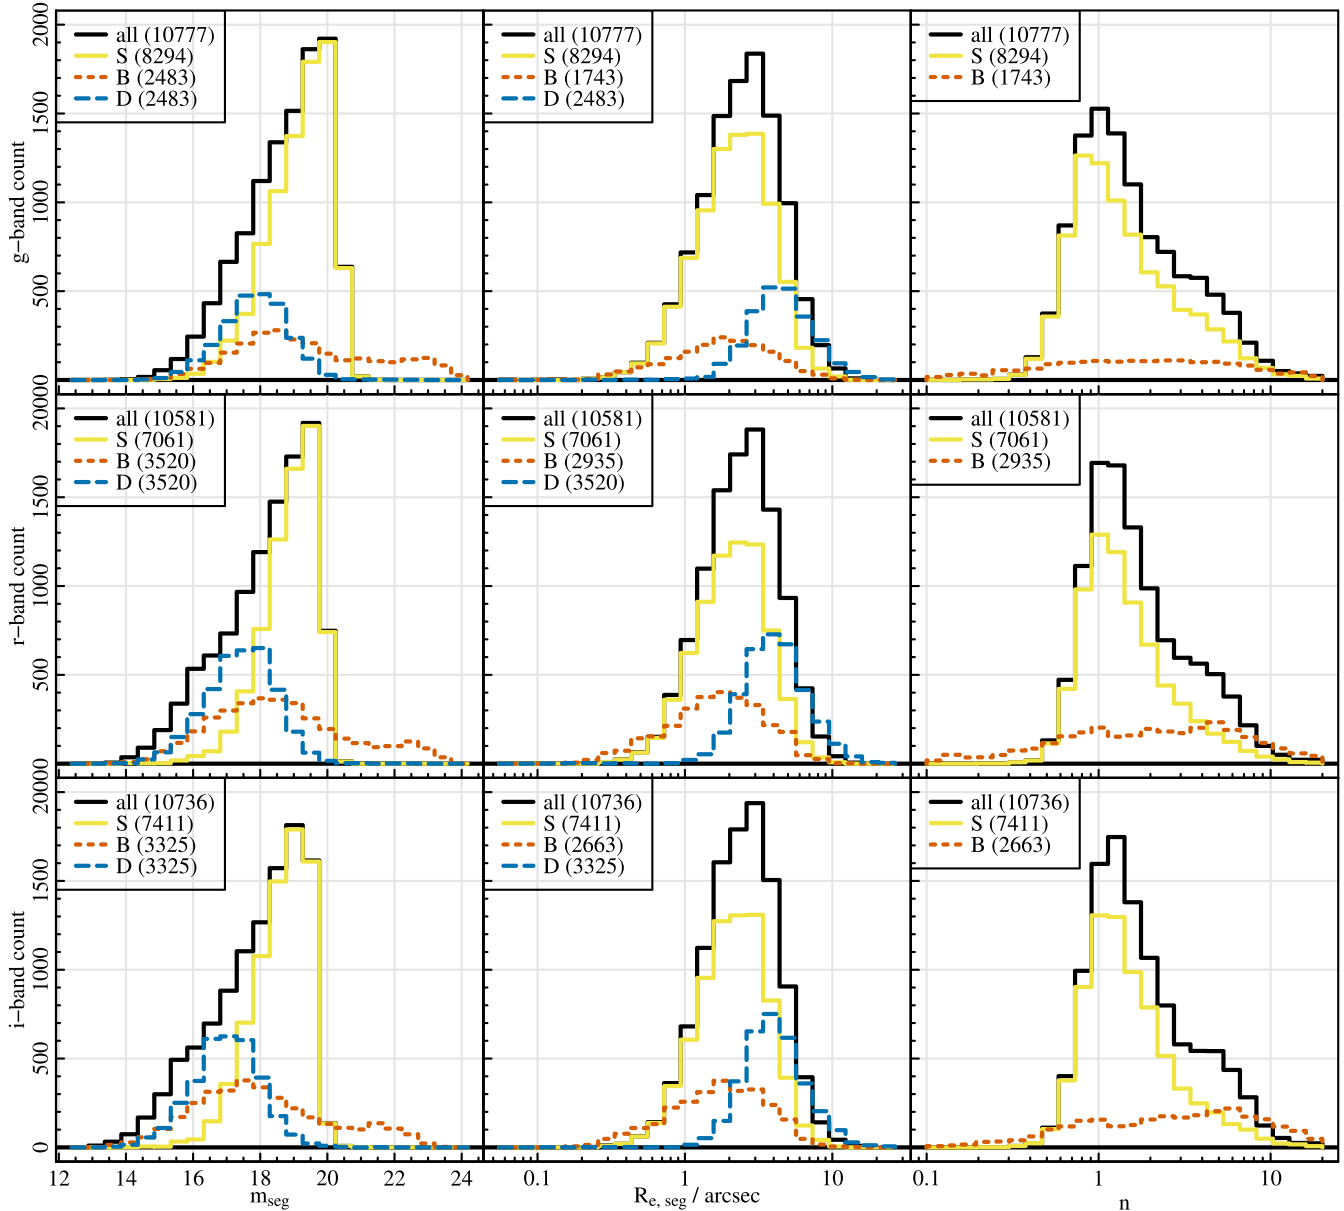
Figure 6. The distribution of the main parameters (limited to segment radii) for all bands and models. Left, middle, and right columns show magnitude, effective radius, and Sérsic index while top, middle, and bottom rows show the *g*, *r*, and *i* band, respectively. The solid yellow lines are the single Sérsic values for those galaxies which were classified as single component systems, dotted red and dashed blue lines show bulges and discs, respectively, for those objects classified as 1.5 or double component systems. For reference the solid black line shows the single Sérsic fits for all galaxies with NCOMP > 0 (i.e. including those classified as 1.5 or double component systems). The number of objects in each histogram is given in the legends, where the number of bulges and discs differs for effective radii and Sérsic indices because these parameters do not exist for 1.5-component fits (point source bulge). We do not show disc Sérsic indices since they were fixed to 1 (exponential).

faint bulges). There is a slight trend for magnitudes to become brighter moving from *g* to *r* to *i* for all components, as expected from typical galaxy colours. We investigate the colours further in Section 4.4.

From the middle panels it becomes obvious that bulges tend to be smaller than discs by a factor ∼2, while single Sérsic fits span a wide range of sizes. Similar to the trend observed in the magnitudes, the smallest objects are classified as single component systems, while some of the larger galaxies can be successfully decomposed.

The Sérsic indices of single component systems show a clear peak around a value of 1 (exponential), a sharp drop-off at lower

values, and a longer tail towards higher values. Interestingly, the single Sérsic distributions showing all systems (black lines) have a secondary 'bump' around a value of 4 or 5 (classical de Vaucouleurs bulge), which is not apparent in those galaxies classified as single component systems (yellow line). Hence most of those high Sérsic index objects were indeed found to contain bulges and were classified as double component systems. The bulges themselves show a wide range of Sérsic indices with (at least in *r* and *i* bands) a slightly double-peaked nature around values of 1 and 4–6. At this point, we would like to remind the reader that we use the term 'bulge' to refer to all kinds of central components of galaxies, including classical
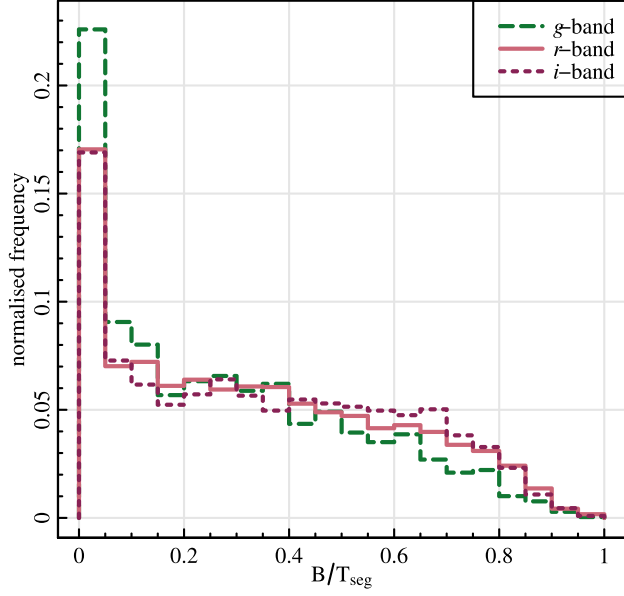
Figure 7. The distribution of the bulge to total flux ratio (limited to segment radii) for the 1.5 and double component fits in all bands. Dashed green, solid light red, and dotted dark pink lines refer to the *g*, *r*, and *i* bands, respectively. The histograms have been normalized by their respective total number of fits (cf. Fig. 6) to make the bands directly comparable.

Figure 8. The difference between the single Sérsic magnitude and the total magnitude derived from the double or 1.5-component fits for those galaxies that were classified as such, for all bands (all magnitudes limited to segment radii). The scatter plot shows the difference between the two magnitudes against the single Sérsic magnitude for all three bands with the running medians and 1σ-percentiles overplotted. The top and right-hand panels show the respective marginal distributions. Dashed green, solid light red, and dotted dark pink lines refer to the *g*, *r*, and *i* bands, respectively.

bulges, pseudo-bulges, bars, and AGN (cf. Section 3.2.1). Hence the 'bulge' distribution will include a variety of physical components and their combinations, leading to the wide spread of values. In addition, the Sérsic index tends to be the parameter with the largest uncertainty, with typical galaxies showing relative errors on their bulge Sérsic index of 1–10 per cent, adding further scatter to the distribution.

Since the bulge to total flux ratio is a derived parameter that is frequently of interest, we additionally show it in Fig. 7 for all three bands; for those galaxies that were classified as a 1.5 or double component fit in the respective band. The majority of systems have intermediate values of B/T with only a few per cent at the extreme end above 0.8. The secondary peak at very low B/T values around 0.02 stems from the 1.5-component fits. The B/T ratio generally increases from *g* to *r* to *i*, as expected (see Section 5.2).

Finally, as a first consistency check, we show the difference between the single Sérsic magnitude and the total magnitude derived from the double or 1.5-component fits, all limited to segment radii, in Fig. 8. The distributions for all three bands are highly peaked around zero, with the vast majority of objects having total magnitudes consistent with the single Sérsic magnitudes within 0.1 mag (over the entire magnitude range). We only show galaxies that were classified as 1.5 or double component fits here to ensure reliable bulge and disc magnitudes, but note that the distribution is very similar when including objects classified as single Sérsic fits. This indicates that the total magnitude is well-constrained even in the case when the individual component magnitudes are not (see also Section 4.4).

4.4 Galaxy and component colours

The left-hand panel of Fig. 9 shows the distribution of $g-r$ colours for galaxies and their components. The colours are corrected for Galactic extinction, but not for dust attenuation in the emitting galaxy. The Galactic extinction was obtained from v03 of the
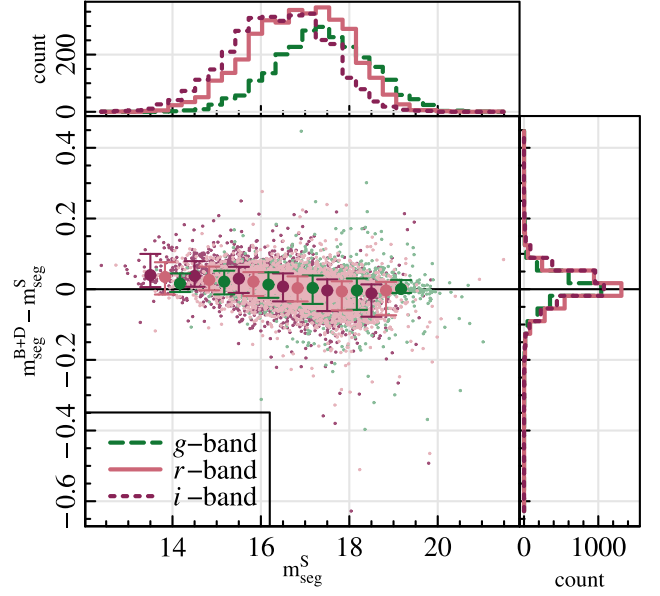
GalacticExtinction catalogue accompanying the equatorial input catalogue on the GAMA database. We focus on $g-r$ colours here since those bands are the deepest KiDS exposures. Results for $g-i$ are qualitatively similar, albeit a bit more noisy.

The solid black line shows the colour distribution for all single Sérsic fits that were not classified as outliers in the joint model selection (NCOMP > 0). It is clearly bimodal, with redder colours typically belonging to higher Sérsic index objects as indicated by the thinner dark red and light blue lines splitting the distribution at $n = 2.5$ (in the *r* band). Not entirely surprisingly (given the distribution of Sérsic indices in Fig. 6), the distribution of single Sérsic objects actually classified as such (NCOMP = 1, solid yellow line) mostly follows the distribution of low Sérsic index objects; while the high Sérsic index objects tend to be classified as double component systems. For the latter, we show total colours with a dash–dotted green line, bulge colours with a dotted red, and disc colours with a dashed blue line. As expected, bulges tend to be redder than discs, although the scatter is large.

The right-hand panel of Fig. 9 shows the corresponding colour–magnitude diagram. Colours and absolute magnitudes are both corrected for Galactic extinction but not for dust attenuation in the emitting galaxy. The absolute magnitude was calculated using the distance modulus provided in v14 of the DistancesFrames catalogue from the GAMA database which we also used to obtain redshifts for the sample selection. Again, we focus on $g-r$ colour and absolute *r*-band magnitude, M_r, but note that results using $g-i$ and/or M_i are qualitatively similar.

The grey density plot in the background shows the single Sérsic fits for all non-outlier (NCOMP > 0) galaxies, corresponding to the black line in the left-hand panel of Fig. 9. The bimodality of the distribution is even clearer here, with the red sequence and blue
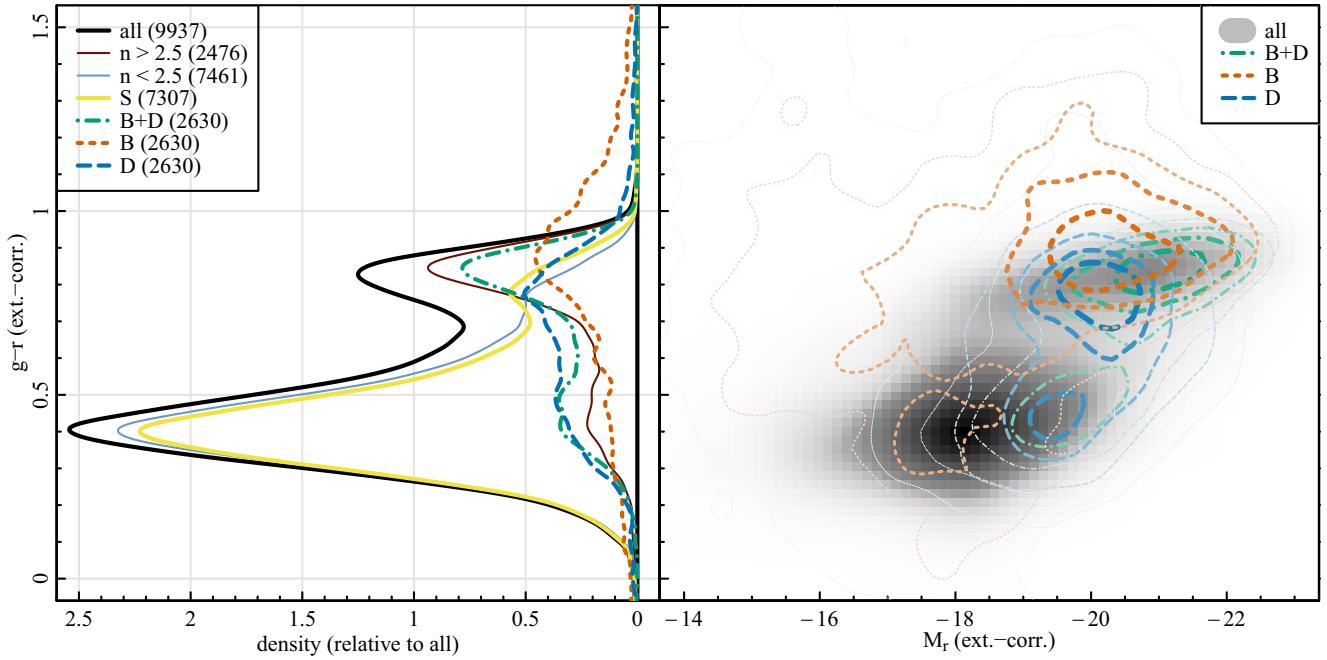
Figure 9. Left-hand panel: The Galactic extinction-corrected $g-r$ colour distributions (limited to segment radii) for galaxies and their components. The colour coding of the lines is the same as for Fig. 6, although with a few additions: The solid black line shows single Sérsic fits for all galaxies with NCOMP > 0 in the joint model selection; the thinner dark red and light blue solid lines split this sample into those with $n > 2.5$ and $n < 2.5$ in the r band. The solid yellow line gives the single Sérsic values for those galaxies which were classified as single component systems, dotted red and dashed blue lines show bulges and discs, respectively, for those objects classified as 1.5 or double component systems (always in the joint model selection). The dot–dashed green histogram gives the total galaxy colour (derived from the addition of bulge and disc flux) for 1.5 and double component systems. The number of objects in each histogram is given in the legend. **Right-hand panel:** The Galactic extinction-corrected $g-r$ versus M_r colour–magnitude diagram (limited to segment radii) for galaxies and their components. The colour coding of the lines is the same as for the left-hand panel. Contours include 10, 25, 50, 75, and 90 per cent of the sample.

cloud being well separated. The green contours indicate the part of the sample that was classified as 1.5 or double component object;[12] as expected, this is concentrated towards the bright end of the galaxy distribution and hence encompasses mostly galaxies located in the red sequence. Correspondingly, bulges and discs are both relatively red, with bulges on average slightly redder than the total galaxies and discs slightly bluer (while both components – obviously – are fainter than the total galaxy). However, both components show a large scatter and overlap with each other: both faint blue bulges exist as well as bright red discs.

A detailed study of component colours and the different populations in the right-hand panel of Fig. 9 is beyond the scope of this work. However, we note that the total galaxy colours show much less scatter; indicating that the scatter results from a different splitting of the light into bulge and disc components in the g and r bands, while the total amount of light is well-constrained (cf. also Fig. 8). A further brief investigation into extreme systems (blue bulges with red discs and also excessively red bulges with very blue discs) suggests that they are caused by a variety of remaining uncertainties in our analysis, e.g. swapped components in one of the two bands (Section 3.2.6), small faint bulges that are barely detected in the r band and missed in the g band, the 'bulge' component dominating both small and large radii in one of the two bands (cf. Section 4.5.2) or failures in the

flagging of bad fits, all combined with model selection uncertainties and the necessity of joint model selection to compromise between the bands. While each of these processes by itself only affects a small number of galaxies, in sum across both bands they do reach the 10–20 per cent level. Still, on average our colours do follow the expected trends, as we show in Section 5.2 with an overview of similar studies in the literature. We will study the colours of galaxies and their components in more detail in forthcoming work, also including further bands (KiDS u and VIKING Z, Y, J, H, K_s) and taking full account of inclination effects due to dust in the emitting galaxies (see e.g. Driver et al. 2008). We will then also assess trends in other parameters, such as the component effective radii, with wavelength.

4.5 Catalogue limitations

We finish this section by pointing out a few limitations of our results that users of the catalogue should be aware of.

4.5.1 Model limitations

All of our models are axially symmetric and monotonically decreasing in intensity from the centre. We are unable to capture asymmetries such as spiral arms, offset bulges, tidal tails, mergers, star-forming regions etc.; or disc features such as rings, bumps, truncations, or flares. If such features are present in the data, they may bias or skew the model parameters. We also remind the reader that when we talk about 'bulges', what we really mean are the central components. This could be a classical bulge, a pseudo-bulge, an AGN, a bar, or any

---

[12]To be precise, the green contours were derived by adding the respective bulge and disc fluxes of the 1.5 or double component objects (for consistency with the bulge and disc contours), while the grey density plot is based on the single Sérsic fits (for robustness at low magnitudes). These two versions of the total galaxy magnitude are generally very similar as evidenced by Fig. 8.

combination (sometimes resulting in the model trying to fit a mixture between e.g. a bar and a bulge). We make no attempt to distinguish between these cases.

4.5.2 Model selection caveats

Model selection is accurate to > 90 per cent compared to what could be achieved by visual classification (Section 3.3.2). However, it is important to note that our aim in the model selection is to determine which one of our three models is most appropriate to use for the given data; and not how many physically distinct components an object consists of. The reason for this is that for a given galaxy, the data quality will strongly influence how many fitting parameters can be meaningfully constrained and using more model parameters will inevitably overfit the data and lead to unphysical results. Hence, even in the joint model selection, we base our visual classification on the fit and residuals in individual bands (which is what we fit to), rather than e.g. colour images. Due to the different depths and resolutions of the KiDS bands, it is hence common for the same galaxy to be classified as double component in one band, but single component in another.

In an attempt to make fitting parameters more directly comparable across bands, we introduced the joint model selection (Section 3.3.2), yet this is necessarily a compromise between the different bands. For example, we lose bulges that are resolved in the *r* band but not in *g* and *i* bands due to the larger PSFs; or there may be some ill-constrained *i*-band fitting parameters for an extended low-surface brightness disc that is visible in *r* and *g* but not in the shallower *i*-band image. There are also more skipped fits and outliers in the joint model selection than in the band-specific ones because all objects that are skipped or flagged in at least one of the three bands is skipped or flagged in the joint model selection.

Irrespective of the result of the model selection, we provide all fitted parameters of all models in the catalogue (along with the postage stamps of all fits and a flag indicating the preferred model). This allows users to perform their own selection if desired; but also requires care as not all provided parameters will be meaningful. While single Sérsic fits to double component objects are mostly reasonable; double component fits to true single component galaxies will have unconstrained and hence potentially unphysical parameters for at least one of the components.

We are also aware of a population of objects that are classified as double component fits but have the bulge component dominating both the centre and the outskirts, with the disc only dominating at intermediate radii or even staying 'below' the bulge at all radii. We believe these are essentially single component systems that do not follow a Sérsic law (e.g. Sérsic index would be higher at centre than outskirts); and so the freedom of the disc is used to offset this. This population is easily identifiable by the high bulge-to-total ratio (B/T $\gtrsim$ 0.6 or 0.7). The single Sérsic fits may be more appropriate to use in these cases (see also the discussion of this issue in Allen et al. 2006). An example is shown in Fig. C4 in the appendix.

4.5.3 Drawbacks of tight fitting segments

As detailed in Section 3.3.3, we use relatively tight segments around the galaxies for fitting, which results in the best possible fit of the inner regions of the galaxy but can lead to large, unphysical wings. Hence we recommend using only integrated properties, i.e. the summed flux/magnitude within the region that was fitted and the corresponding effective radii and bulge-to-total ratios as

given by the corresponding ∗_SEGRAD properties in the catalogue. For comparisons to other catalogues using larger fitting segments, their profiles should also be appropriately truncated (see details in Sections 3.3.3 and 5.4).

4.5.4 Sources of systematic uncertainties

We provide errors for each fitted parameter in the catalogue including our best estimate of systematic uncertainties taken from Table 6. However, we do not apply the (small) bias corrections given in the same table, since they are only applicable to large random samples of our galaxies and not to individual objects. In addition, we would like to point out that the systematic errors were estimated from single Sérsic *r*-band fits. We expect that individual components as well as the *g* and *i* bands are affected by similar systematics, but we did not test for this. Also, there are some systematic uncertainties that we do not account for in our simulations, most obviously galaxy features that cannot be captured by our models. For these reasons, the given errors should still be considered as lower limits of the true errors.

4.5.5 GAMA-KiDS RA/Dec offset

We observed an average offset between the input and output (fitted) positions of galaxies in both RA and Dec of approx. 0.4 pix (0.08 arcsec). This is due to an offset between the GAMA (SDSS) and KiDS positions; the same offset can be seen when comparing the KiDS source catalogue with the Gaia catalogue; see also fig. 15 in Kuijken et al. (2019). We correct for this during the outlier rejection, but give the original (uncorrected) fitted values for position otherwise.

4.5.6 Completeness limits

Due to our sample selection (Section 2.3), our spectroscopic completeness is 100 per cent and even the faintest objects in our sample are well-resolved and bright enough in KiDS data to allow for robust single Sérsic fits. However, there is a systemic limit to the component magnitude in that the samples of bulges and discs with magnitudes fainter than the GAMA limit ($r < 19.8$ mag) are incomplete. For example, a bulge with a magnitude of 22 mag in the *r* band will only be contained in our sample if the corresponding disc is bright enough such that the total magnitude is below 19.8 mag; hence the sample of bulges with 22 mag is incomplete. This applies almost exclusively to the faint bulges from the 1.5-component fits as can be seen in the first column of Fig. 6.

5 COMPARISON TO PREVIOUS WORKS

After presenting some of the contents of our main catalogue, we now turn towards demonstrating its robustness. We start with a comparison to previous works in this section (including work on the same galaxy sample); then describe additional internal consistency checks and a detailed study of biases and systematic errors with bespoke simulations in Section 6.

5.1 Comparison of catalogue statistics

As a first check, we compare our model selection statistics to those of other bulge-disc decomposition works, although care must be taken in judging these results since they will depend on the sample selection, data quality and observational band.

Table 5. Comparison of our catalogue statistics to previous works (in the *r* band unless stated otherwise).

Reference	Single (per cent)	Double (per cent)	Unsuitable (per cent)	Notes
This work	47	23	30	Double including 1.5-comp.; unsuitable including masking (20 per cent)
Barsanti et al. (2021)	47	28	25	Cluster S0 galaxies
Meert et al. (2015)	44	39	17	Larger sample up to higher redshift but smaller magnitude range
Head et al. (2014)	19	35	46	*g* band; early-type sample; more stringent criteria for 'good' fits
Lackner & Gunn (2012)	35	29	36	Single corresponding to pure exponential or de Vaucouleurs; unsuitable corresponding to their 'Sérsic' category
Simard et al. (2011)	73	26	1	Unsuitable corresponding to failure rate of fitting routine; no selection of 'good' fits given
Allen et al. (2006)	43	34	23	*B* band; unsuitable galaxies excluded through cuts in redshift, galaxy size and surface brightness

Table 5 summarizes the corresponding percentages including a few notes on the most important differences of the quoted works to ours (more details on the majority of these studies are given in Section 5.2). In short, for the automated decomposition of large samples of galaxies in the *r* band, most authors – including ourselves – class roughly half of all galaxies as being well-represented by a single Sérsic model, with the other half split approximately evenly into double component fits and objects unsuitable for fitting with such simple models.

5.2 Comparison of component colours to literature

$g-r$ colours of galaxy components, such as those we present in Fig. 9 and Section 4.4, are not found frequently in the literature, although a number of authors have presented bulge-disc decompositions in several bands. For example, Simard et al. (2011) perform bulge-disc decompositions for a large sample of galaxies in the SDSS *g* and *r* bands but only present colour–magnitude diagrams for the total galaxies (their figs 9 and 10). These are visually comparable to our total galaxy colours as indicated by the dot–dashed green contours in the right-hand panel of fig. 9. Mendel et al. (2014) add the SDSS *u*, *i*, and *z* bands to the analysis of Simard et al. (2011) and present component masses in *ugriz* but also do not study component colours.

Similarly, Meert et al. (2015) present a large *r*-band catalogue which is extended to include the *g* and *i* bands in Meert et al. (2016). Colour–magnitude diagrams, however, are again only presented for total galaxies, with the authors noting that component colours can be calculated from their catalogue but should be used with care since they are subject to large uncertainties.

More recently, Dimauro et al. (2018) provide (UVJ) component colours in their catalogue but defer their study to future work; while Bottrell et al. (2019) present *ugriz* colour–magnitude diagrams for total galaxies (colour coded by B/T); but again not for individual components.

Among the first to show component colours for a large sample of galaxies were Lackner & Gunn (2012) in their study of ∼70 000 $z <$ 0.05 SDSS galaxies in the *g*, *r*, and *i* bands. However, in contrast to our fits, their *g* and *i*-band fits are not independent. Instead, in order to decrease the noisiness of the colours, the structural parameters are taken from the *r* band and only the magnitude is adjusted. Additionally, Lackner & Gunn (2012; along with e.g. Mendel et al. 2014) fix the Sérsic index of the bulge to either 1 or 4 for their double component fits to limit the number of free parameters since the data is insufficient to constrain the bulge light profile. Keeping these differences in mind, their fig. 32 showing the $g-r$ versus M_r colour–magnitude diagram for bulges and discs as contours superimposed on the greyscale background for all galaxies can be compared to the right-hand panel of our Fig. 9 (for a more detailed description of Fig. 9, see Section 4.4). In general, both plots are very similar: the grey background shows a large blue cloud and a well-separated red sequence. The double component fits populate the red sequence, green valley and the brighter part of the blue cloud. The bulges tend to be slightly redder than the red sequence but with a large scatter especially at the faint end. Discs spread from the red sequence towards the green valley with a smaller population also in the blue cloud.[13] Lackner & Gunn (2012) also note the large scatter in colour for bulges in particular, despite their fitting constraints and lower redshift limit. Hence it is not surprising that even with the higher quality KiDS data and our new fitting routines, we get a large scatter in component colours, especially since we leave the bulge Sérsic index free and perform independent fits in both bands. The latter in particular can lead to very extreme colours since it is not guaranteed that the 'bulge' and 'disc' models actually fit the same features in both images (in particular when there are additional features present that are not fully captured by the models; see also Section 4.4).

Kim et al. (2016) found similar difficulties when performing *g* and *r*-band decompositions on ∼10 000 large bright and approximately face-on SDSS galaxies. While they leave the Sérsic index free as we do, the *g*-band structural parameters are again taken from the *r*-band fits with only the magnitudes adjusted. Despite this, they find it necessary to remove almost 40 per cent of their sample after fitting because they show excessively red bulge colours (combined with low B/T values in the *r*-band). After this cut, the $g-r$ versus M_r colour–magnitude diagram for bulges shown in their Fig. 7 is slightly less noisy than ours (Fig. 9), although still comparable. Kim et al. (2016) did not study the properties of the discs in their sample.

One of the most direct comparisons to make is with Kennedy et al. (2016) who study GAMA galaxies in the G09 region (a subset of our sample) in the *ugrizYJHK* bands from the SDSS (York et al. 2000) and the *United Kingdom Infra-red Telescope* Infrared Deep Sky Survey (UKIDSS, Lawrence et al. 2007). They use the MEGAMORPH multiband fitting method with GALAPAGOS-2 and GALFITM (Häußler et al. 2013; Vika et al. 2013) to perform simultaneous Sérsic plus exponential fits across all nine bands. The structural parameters are constrained to be the same in all bands, with only the component magnitudes allowed to vary freely, therefore providing robust colours. While the paper focusses on studying $u-r$ colours, the corresponding catalogue on the GAMA database

---

[13]For reference, their cyan contours represent 6684 galaxies with a bulge Sérsic index of 1 and the magenta contours show 14 042 objects with a bulge Sérsic index of 4. Also note that their *x*-axis is reversed with respect to ours.
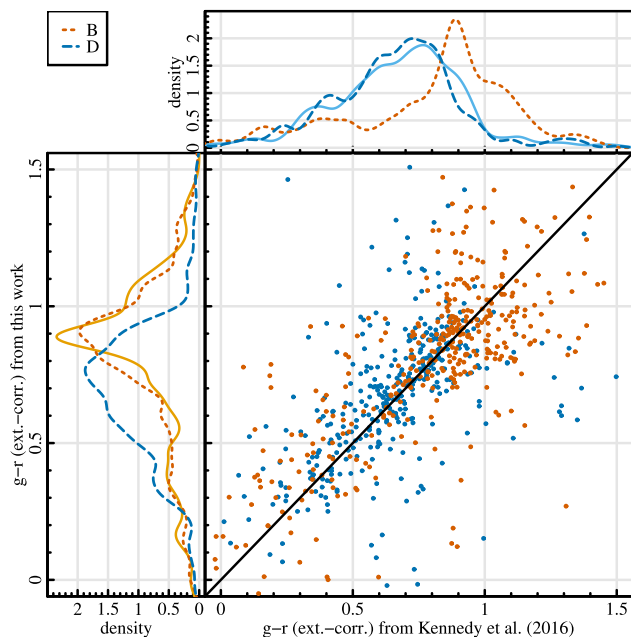
Figure 10. Our Galactic extinction-corrected $g-r$ component colours (limited to segment radii) compared against those from Kennedy et al. (2016) for a subsample of 390 objects that appear in both catalogues and were classified as double component fits in the joint model selection. The scatter plot shows the direct comparison, while the density plots show the respective distributions in both catalogues (ours on the left, Kennedy et al. 2016 on the top). Bulges are again shown in red with dotted lines and discs in blue with dashed lines. To aid the direct comparison of the distributions, the lighter solid lines also show the Kennedy et al. (2016) bulge distribution on the left and our disc distribution on the top.

(`MegaMorph:MegaMorphCatv01`) contains the information for the fits in all nine bands such that $g-r$ colours can easily be derived.

This comparison is shown in Fig. 10 for those galaxies that were present in both catalogues and classified as double component fits (NCOMP = 2) in the joint model selection of our fits (Kennedy et al. 2016 do not perform model selection nor outlier rejection). In addition to the scatter plot directly comparing the component colours, we show the corresponding distributions in the left (this work) and top (Kennedy et al. 2016) panels of Fig. 10. As always, bulges are shown in red (points and dotted lines) and discs in blue (points and dashed lines). To aid the direct comparison of the distributions, we additionally show the Kennedy et al. (2016) bulge colour distribution in the left-hand panel as a solid orange line and the disc distribution from this work in the top panel as a light blue solid line. Component colours are all corrected for Galactic extinction and limited to segment radii for our fits.

Despite the large scatter, it can be seen that our component colours are generally in agreement with those from Kennedy et al. (2016) with no systematic differences. The scatter in both catalogues is also comparable, although Kennedy et al. (2016) perform multiband simultaneous fits with fixed structural parameters that should lead to more robust component colours. This advantage of their work seems to be balanced by advantages of our work, such as the improved data quality of KiDS, the robustness of the fitting procedure with PROFOUND and PROFIT and our post-processing steps (in particular outlier rejection and model selection).

In addition to the large $g-r$ component colour studies discussed above, there are a number of publications focussing on the $g-i$

colours of bulges and discs for samples ranging between ∼100 and ∼1000 objects (i.e. roughly a factor of 10 smaller than ours), namely Gadotti (2009), Head et al. (2014), Vika et al. (2014), Fernández Lorenzo et al. (2014), Cook et al. (2019), and Barsanti et al. (2021). We briefly compare our work to their results here, noting that all above authors have more stringent constraints on their fits than we do and also report problems in deriving bulge colours. For example, Fernández Lorenzo et al. (2014), although they fit the galaxies in both bands, use fixed aperture photometry to derive more stable bulge colours. Vika et al. (2014), while performing *simultaneous* multiband fits, do not allow for any variation of structural parameters (except magnitudes) with wavelength. Head et al. (2014), in addition to varying magnitudes, allow for a trend in disc sizes with wavelength in approximately 30 per cent of their sample, noting that this leads to increased scatter. Cook et al. (2019), who use PROFIT like this work, allow discs to deviate slightly from the exponential profile but fix all bulges to be exactly round (axial ratio of 1), again only allow magnitudes and disc sizes to vary between bands and employ a sophisticated, visually-guided re-fitting procedure to obtain physically meaningful fits for 'difficult' objects. Barsanti et al. (2021), also employing PROFIT, additionally allow for differing bulge sizes and Sérsic indices in the different bands (but fixing bulge and disc axial ratios and position angles and performing model selection in the r band only), but class approximately half of their double component fits as 'unreliable'. Gadotti (2009), fitting bulges, bars, and discs to a sample of face-on, visually-selected 'well-behaved' galaxies refrain from automated fitting and instead treat each galaxy individually.

After these notes on the inherent difficulties associated with deriving component colours, we can now turn to the corresponding results: Head et al. (2014), in their study of early-type red-sequence galaxies in the Coma cluster, measure an average $g-i$ difference between bulges and discs of 0.09 ± 0.01 mag. Similarly, Barsanti et al. (2021) find a bulge-disc $g-i$ difference of 0.11 ± 0.02 mag for their sample of S0 cluster galaxies. Fernández Lorenzo et al. (2014), on the other hand, have a sample of mostly late-type spirals (with B/T<0.1 for ∼66 per cent of their objects) and find a difference of 0.29 mag in the median $g-i$ bulge and disc colours, i.e. a factor of ∼3 larger. In line with this, Vika et al. (2014) report that the bulge and disc colours are similar for early-type galaxies but differ significantly for late-types. The $g-i$ differences for the different morphological classes given in their table two range from 0.03 ± 0.04 mag for ellipticals to 0.28 ± 0.06 mag for late-type spirals; with the overall average (comprising approximately two thirds late-types) being 0.19 ± 0.04. Similarly, the average $g-i$ colour difference of the Gadotti (2009) sample of varying galaxy types amounts to 0.18 ± 0.04 (from the online-version of their table 2).

Our results are perfectly in line with this: the median bulge-disc $g-i$ colour difference for our 1.5 or double component fits is 0.17 ± 0.01 mag, consistent with the Vika et al. (2014) and Gadotti (2009) results. Limiting to objects with a total $g-i > 1$ (red-sequence galaxies) reduces the value to 0.14 ± 0.01 mag; while focussing on two-component fits only (excluding 1.5-component fits) yields 0.10 ± 0.02 mag, in agreement with Head et al. (2014) and Barsanti et al. (2021). This is because our double component galaxies lie predominantly on the red sequence, as can be seen in Fig. 9. 1.5-component fits on the other hand, have very small (namely unresolved) bulges by definition and hence belong to the class of late-type spirals. In fact, 87 per cent of the 1.5-component objects have a (g-band) B/T ratio less than 0.1, with the median value as low as 0.02. Computing the bulge-disc $g-i$ colour difference for this sample of objects yields a value of 0.46 ± 0.02, suggesting

that the trend described in Vika et al. (2014) continues at very low B/T.

From all of these comparisons we conclude that our component colours – although noisy – are in line with previous work. In order to increase the colour robustness while preserving the ability to capture physical trends with wavelength (i.e. not fixing the structural parameters to be the same in all bands), a simultaneous fit in all bands is needed. This has many advantages as shown by the MEGAMORPH project team using GALAPAGOS and GALFITM (Häußler et al. 2013, 2022; Vika et al. 2013), especially for automated analyses, since it naturally ensures smooth wavelength trends while preserving physical variation and additionally allows more robust fits to fainter magnitudes. With PROFIT v2.0.0, released in February 2021, now supporting a multiband fitting mode, this is certainly an interesting avenue to explore in future work and could provide a valuable alternative.

5.3 Comparison to size–stellar mass relations of Lange et al. (2015)

Fig. 11 shows the size–stellar mass relation obtained from our *r*-band single Sérsic fits in combination with the redshifts and distance moduli of v14 of the `DistancesFrames` catalogue (Baldry et al. 2012) and v19 of the `StellarMasses` catalogue (Taylor et al. 2011); both from the GAMA database. The aperture-derived stellar masses have been scaled to match the Sérsic total flux using the `fluxscale` keyword provided in the `StellarMasses` catalogue. The *g*-band and *i*-band results are very similar to those from the *r* band and hence we do not show them here.

The sample is limited to objects which were not flagged during our outlier rejection (Section 3.3.1) and split into early- and late-type galaxies according to our fitted Sérsic index ($n \lessgtr 2.5$; analogous to Lange et al. 2015). We also limit the redshift range to $0.0001 < z < 0.06$ and the stellar mass range to $M_* > 10^9\,\mathrm{M_\odot}$, thus avoiding the need for volume corrections. For comparison, we show the $M_* - R_e$ relations obtained by Lange et al. (2015) by fitting a single power law to the single component *r*-band fits of Kelvin et al. (2012; pre-release of `SersicPhotometry:SersicCatSDSSv09`) combined with an earlier version of the stellar masses catalogue of Taylor et al. (2011; `StellarMasses:StellarMassesv16`). We note that the stellar masses did not change much between v16 and v19: the mean and standard deviation of $\Delta\log_{10}(M_*/\mathrm{M_\odot})$ are 0.006 and 0.07, respectively for our sample. The two panels show the results obtained with effective radii taken directly from the Sérsic fits (top; extrapolated to infinity by definition) or limited to the segment radius within which they were fitted (bottom; see Section 3.3.3 for further explanation). To guide the eye, we also show the running median and its error for our data; where the error is taken as the 1σ-quantile divided by the square-root of data points within that bin (usually smaller than the size of the data points).

In both cases, the slope of the mass–size relation obtained from our data agrees well with the fit results of Lange et al. (2015). There is a clear offset in the absolute sizes, but those will inherently depend on the exact definition of the size measurement at hand as well as the (depth of the) data used. Already calculating effective radii within the segments within which we fitted for them (bottom panel) brings our results much closer to those of Lange et al. (2015); although the measurements are then not directly comparable to their fits anymore since they use Sérsic values extrapolated to infinity (which will, in turn, depend on the segment size used for fitting). We now discuss these issues further by directly comparing our fits to those of Kelvin et al. (2012), which the Lange et al. (2015) results were based on.
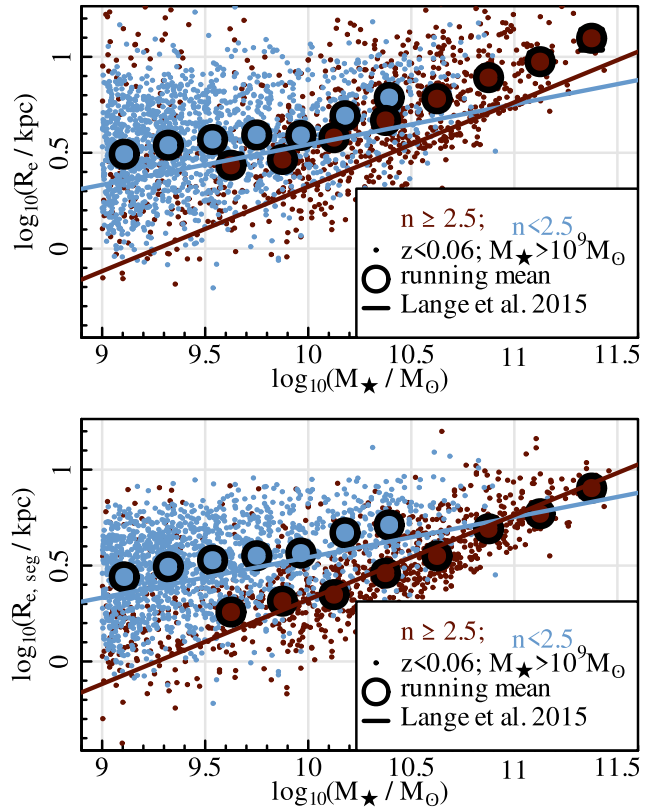

Figure 11. The size–stellar mass relation for our *r*-band fits (dots) compared to the Lange et al. (2015) fits (lines). The sizes are obtained from our single Sérsic effective radii (top panel: extrapolated to infinity; bottom panel: limited to segment radius) and the distance moduli provided in the `DistancesFrames` catalogue originally described by Baldry et al. (2012). The stellar masses are taken from the most recent version of the `StellarMasses` catalogue initially presented in Taylor et al. (2011). The sample is limited to the redshift range $0.0001 < z < 0.06$ (redshifts also from `DistancesFrames`) and the stellar mass range $M_* > 10^9\,\mathrm{M_\odot}$. Large circles with error bars indicate the running median with its error (usually smaller than the data point). Solid lines show the single exponential $M_* - R_e$ relation fits obtained by Lange et al. (2015) for their single component *r*-band sample, split by a Sérsic index cut at $n = 2.5$ (taken from their tables 2 and 3).

5.4 Comparison to single Sérsic fits of Kelvin et al. (2012)

To further investigate the size offset observed in Section 5.3, we directly compared our fits to those of Kelvin et al. (2012; `SersicPhotometry:SersicCatSDSSv09` on the GAMA database). Since Kelvin et al. (2012) do not provide double component fits, the analysis is limited to single Sérsic fits. We again use the *r* band as an example for the discussion, but note that results are very similar for the *g* and *i* bands.

The Kelvin et al. (2012) fits are based on the Structural Investigation of Galaxies via Model Analysis (SIGMA) code applied to data from SDSS DR7. SIGMA is a wrapper around SOURCE EXTRACTOR, PSF EXTRACTOR, and GALFIT3 performing similar steps to what we do in our pipeline (Section 3), i.e. source identification, background subtraction, PSF estimation, and 2D model fits to the surface brightness profile of the galaxies. The differences lie in the data and code used, where we upgrade SDSS to KiDS, SOURCE EXTRACTOR to PROFOUND, PSF EXTRACTOR to a combination of PROFOUND and PROFIT and GALFIT to PROFIT; with all the

advantages described in Sections 2.2, 2.4, and 2.5. In addition, we also perform multicomponent fits and model selection. For the comparison to the Kelvin et al. (2012) results, we focus on the three most important single Sérsic fit parameters: magnitude m, Sérsic index n, and effective radius R_e, which tend to be the least 'well-behaved' (position, axial ratio, and angle are generally more easily constrained and uncorrelated).

Fig. 12 shows the difference between our fits and the Kelvin et al. (2012) fits for these three parameters. For the magnitudes, we show the difference [Kelvin et al. (2012) fits – our fits], while for the effective radius and Sérsic index (scale parameters) we show the quotient [Kelvin et al. (2012) fits / our fits] on the y-axis; always plotted against our fitted values on the x-axis (in logarithmic space for scale parameters). Again, we show the results for Sérsic parameters extrapolated to infinity (top) and for the magnitude and effective radii calculated within the segment radius (see Section 3.3.3), where we limit both our fits and those of Kelvin et al. (2012) to our fitting segment (which are generally smaller than the fitting regions used in Kelvin et al. (2012)) to obtain directly comparable results.

For the Sérsic parameters extrapolated to infinity (top panels), large differences can be seen in all fitted parameters, including clear systematic trends across the parameter space. This shows once again that fitted Sérsic parameters are not directly comparable given the differences in the data, code, and processing steps with a wealth of potentially different systematic uncertainties (Section 3.3.3). However, when we limit the analysis to our segment sizes (bottom panels), the fits become much more comparable. On average, now, our fits are ~0.03 mag brighter and approximately 7 per cent larger than the Kelvin et al. (2012) fits to the same galaxies, which is not surprising given the increased depth and resolution of KiDS compared to SDSS and the numerous sources of different systematic uncertainties (e.g. differing sky subtraction and PSF estimation). Also, there are fewer trends across the parameter space, indicating that systematic differences arise mainly from the extrapolation to infinity.

The slight exception to this is the Sérsic index, which still shows some trends. The reason for this is that the Sérsic index, unlike the magnitude and effective radius, cannot easily be corrected for different fitting segment sizes. The Kelvin et al. (2012) fits, which were performed within larger fitting segments than our fits, will hence inevitably have to compromise more between the inner and outer regions of the galaxy to be fitted (unless the light profile truly follows a single Sérsic profile with no deviations out to very large radii, which is rarely the case). Our tight fitting segments, on the other hand, will result in better fits to the inner regions of the galaxy at the expense of producing unphysical wings when extrapolated beyond the fitting region (Section 3.3.3). In a sense, fitted Sérsic indices are hence always a weighted average (or compromise) across a range of radii and their absolute values will never be directly comparable between catalogues unless the fitting regions are exactly the same (or the galaxies studied happen to follow perfect Sérsic profiles).

5.5 Comparison to Lange et al. (2016) double component fits

Double component fits for samples of galaxies overlapping with our sample have been performed by Lange et al. (2016), Häußler et al. (2022) and Robotham, Bellstedt & Driver (2022). The latter two include a comparison to the fits presented in this work (figs 19 and 20 in Robotham et al. 2022 and fig. B2 in Häußler et al. 2022), including a discussion of observed differences. We refer the reader to these works and do not repeat the comparisons here. Instead we
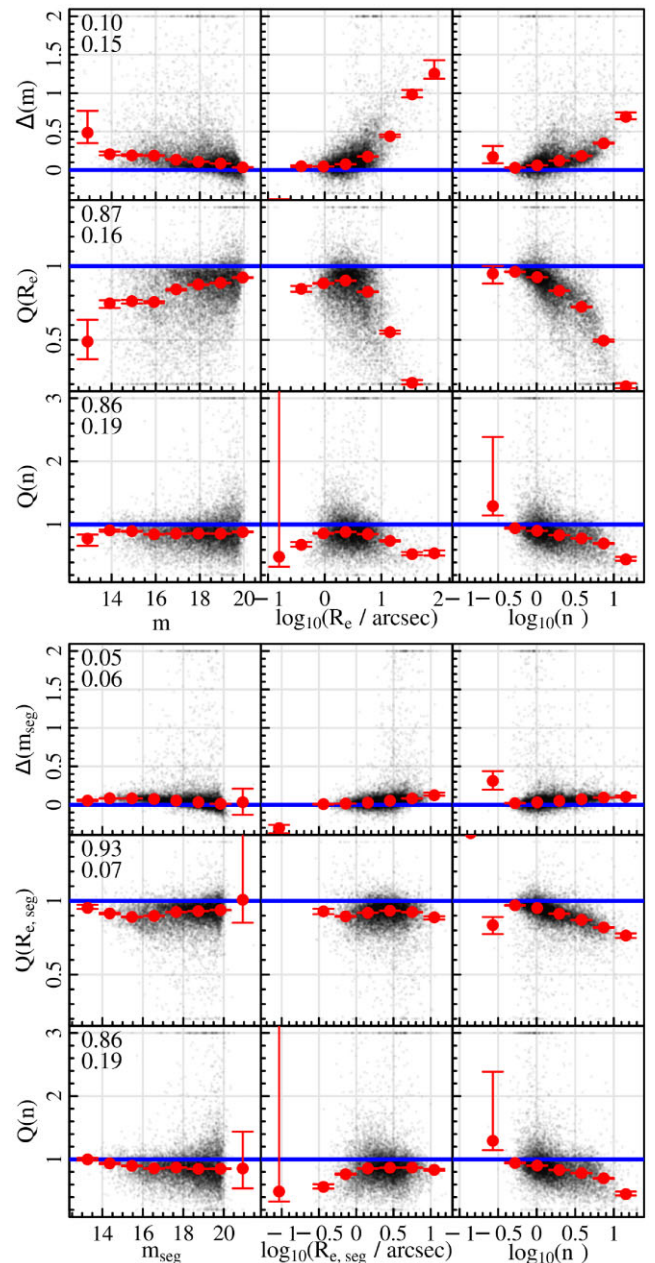
Figure 12. The difference Δ or quotient Q (for scale parameters) between the Kelvin et al. (2012) fits and our fits plotted against our fits for the three most important single-Sérsic parameters magnitude, effective radius, and Sérsic index in the r band. The top panels show the Sérsic parameters extrapolated to infinity, while for the bottom panels we calculated the magnitude and radius within the segment radius for both our and the Kelvin et al. (2012) fits. Outliers are clipped to the plotting interval; which is the same in both cases. Black dots show all fits, red dots with error bars show the running median and its error in evenly spaced bins, and horizontal blue lines indicate no difference between the fits. The numbers in the top left corners of the first column of panels show the median and 1σ-quantile of the respective distribution in the y-direction (which is identical for all panels of a row). The sample is limited to fits that are available in the Kelvin et al. (2012) catalogue and were classified as single component in our model selection.

focus on a comparison to Lange et al. (2016), for which the fits can be obtained from the GAMA wiki.

The Lange et al. (2016) fits are performed on SDSS *r*-band data for GAMA galaxies with redshifts $z < 0.06$, which is a subsample of our sample. They use SIGMA (Section 5.4) to fit each galaxy with a single Sérsic and a double Sérsic function, starting from a grid of starting values in both cases to overcome the susceptibility of GALFIT to initial guesses and to obtain more realistic values for the fit uncertainties. Note that the disc Sérsic index is left as a free fitting parameter in contrast to our work where we fix it to 1 (exponential).

Fig. 13 shows the direct comparison of our double component fits to those obtained by Lange et al. (2016) for the five most important double component parameters (limited to segment radii), similar to the single Sérsic comparison against Kelvin et al. (2012) in Fig. 12 and Section 5.4. We note that the comparison to the Lange et al. (2016) single Sérsic fits is very similar to that for Kelvin et al. (2012) except for the smaller sample size, so we do not show it here. The scatter is larger for all parameters in Fig. 13 compared to Fig. 12. This is expected since fitting two components will intrinsically show higher uncertainties and more degeneracies than fitting a single Sérsic function. There is an average offset and clear trends across the parameter space for the bulge and disc magnitudes (first two rows of Fig. 13). On average, our bulges are ~ 0.3 mag fainter than those of Lange et al. (2016), with discs being correspondingly brighter. Small and faint bulges are most severely affected. This is most likely caused by the lower resolution of the SDSS data compared to KiDS: the same trend is observed by Häußler et al. (2022) when comparing our fits to their fits performed on SDSS data (left two panels of the second row of their fig. B2), but not when comparing our fits to their fits performed on KiDS data (right two panels of the second row of their fig. B2).

The effective radii and bulge Sérsic index show reduced overall offsets in Fig. 13 as compared to Fig. 12. The median values are consistent with no difference for most bins with reasonable numbers of data points. The bulge effective radii show an average offset of less than 1 per cent (albeit with large scatter), while our discs tend to be slightly larger (by around 6 per cent) than those of Lange et al. (2016). The bulge Sérsic indices are also consistent within 1 per cent on average, despite their large scatter. This reduction of systematic differences compared to the single Sérsic fits in Fig. 12 is most likely due to the generally better model fit of the double component model, which reduces the effects of the different fitting segment sizes (since those are only relevant for galaxies which cannot be accurately represented by the model, cf. Sections 3.3.3 and 5.4).

6 SYSTEMATIC UNCERTAINTIES AND BIASES FROM SIMULATIONS AND THE OVERLAP SAMPLE

The MCMC chain errors returned by the fitting procedure do not include systematic uncertainties which arise due to galaxy features not accounted for in the models, nearby other objects, imperfect PSF estimation, background subtraction inaccuracies and similar effects. For an individual galaxy, the presence of such 'features' will systematically shift the fitted parameters away from the true values, thus introducing a bias. For a statistically large enough sample of galaxies, however, most of these effects are expected to cancel out on average since they are random from one galaxy to the next (e.g. nearby other sources shifting the fitted positions). These 'random systematics' can – for statistical samples – be accounted for by simply increasing the given parameter errors such that in most cases, the true

values are included in the credible intervals again. Such systematics can be studied using overlap sample galaxies, i.e. those that appeared in more than one KiDS tile (cf. Section 2.3). In addition, there can be 'one-sided effects' that lead to an overall bias across the sample, e.g. due to excess flux from nearby objects. These can only be detected using simulations. In the following, we study both of these effects using our bespoke simulations, the overlap sample of real galaxies, and the overlap sample of simulated galaxies; where we refer to the random systematics as 'error underestimates' and to the one-sided effects as 'biases'.

The final corrections for both of these effects are listed in Table 6. In short, biases are very small ($\lesssim 1$ per cent), while systematic errors are a factor of 2–3 larger than the random MCMC errors alone. The error underestimate corrections are also applied to the released catalogues, while the bias corrections are not since they are only valid for a large random subset of our galaxies and not for individual objects. Note that the systematic error studies were carried out on v03 of the BDDecomp DMU, while the remainder of this work describes v04 (both on the GAMA database). However, since v04 is statistically identical to v03, the results can directly be transferred. We would also like to point out that we focus on single Sérsic *r*-band fits in this section. We expect individual components in the 1.5 and double component fits as well as the *g* and *i* bands to be affected by similar systematics. Effects are likely to become worse for fainter and/or less well-resolved objects (i.e. bulges in particular). A more detailed investigation into the systematics of double component fits will be included in future work.

6.1 Overlap sample comparison

As an internal consistency check, we compared the fit results obtained from multiple observations of the same physical object (Section 2.3) in Fig. 14. The plots are very similar to the ones in the bottom panels of Fig. 12 (see description in Section 5.4; but note the different *y*-axis scale), except that we now show the differences between two of our own fits to different KiDS images of the same galaxy (in the overlap region between the KiDS tiles). Hence all fits shown in Fig. 14 are based on KiDS data and use the exact same pipeline for analysis, though the different observations are treated entirely independently. We always use the deeper image as the reference image (the image depth at the edge of KiDS tiles can vary greatly depending on the number of dithers – between 1 and 5 – that cover the area), but evaluate the magnitude and effective radius within whichever of the two segment radii is smaller to avoid extrapolation and obtain consistent results.

As expected, there are very little differences between the two fits to the same galaxy; and there are no systematic trends across the parameter space. The running median is consistent with 0 (or 1, for scale parameters) in almost all bins, which shows that there are no inherent systematic differences in our fits related to image depth. This holds true despite the segment radii being systematically larger for the deeper images (as expected), as long as both fits are evaluated within the same region before comparison.

6.2 Simulations: parameter recovery

As a final test of our pipeline, we ran simulations where we insert single Sérsic model galaxies convolved with an appropriate PSF at random locations in the KiDS data. To obtain a realistic distribution of parameters for the model galaxies (including correlations), we use the fitted single Sérsic parameters of a random sample of 1000 *r*-band galaxies that were not classified as outliers. The PSF to
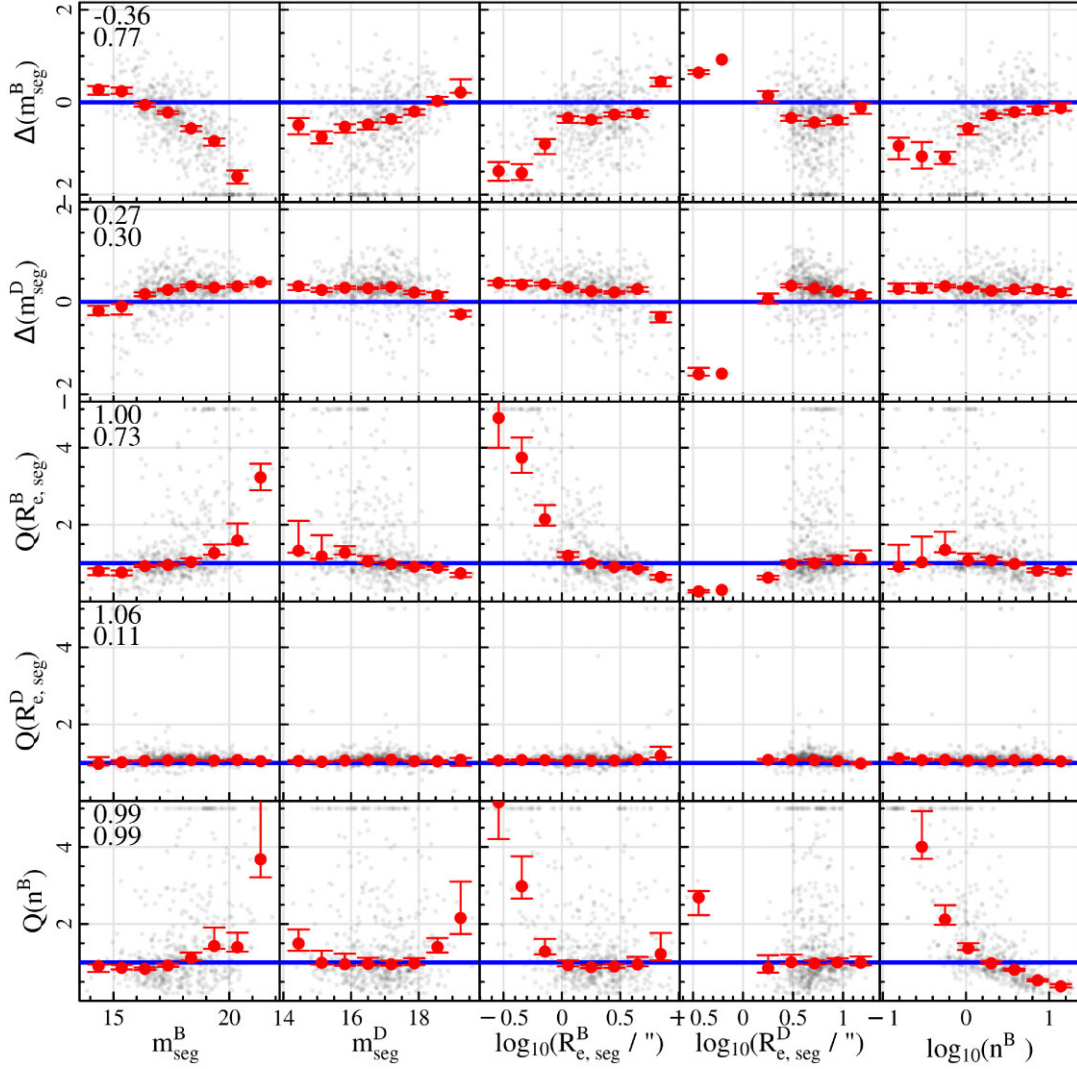
Figure 13. The difference Δ or quotient Q (for scale parameters) between the Lange et al. (2016) fits and our fits plotted against our fits for the five most important double component parameters bulge magnitude m^B, disc magnitude m^D, bulge effective radius R_e^B, disc effective radius R_e^D, and bulge Sérsic index n^B in the r band. Magnitudes and effective radii are limited to segment radii and all values are clipped to the plotting intervals. Black dots show all fits, red dots with error bars show the running median, and its error in evenly spaced bins and horizontal blue lines indicate no difference between the fits. The numbers in the top left corners of the first column of panels show the median and 1σ-quantile of the respective distribution in the y-direction (which is identical for all panels of a row). The sample is limited to fits for which both bulge and disc measurements were available in the Lange et al. (2016) catalogue, that were classified as double component object in our model selection and that were not flagged as unreliable in either catalogue.

convolve with is taken as the model PSF that was fitted to the nearest real galaxy (at the position where the model galaxy is inserted), which is close to the real PSF at that image location. We then simply add the PSF-convolved galaxy to the KiDS data and run the resulting image through our entire pipeline (segmentation, sky subtraction, PSF estimation, galaxy fitting, outlier flagging, model selection).

In this way we are able to check for intrinsic biases in our entire pipeline, with three exceptions: issues due to galaxy features not represented by our models (bars, spiral arms, disc breaks, mergers, etc.), problems in the data processing performed by the KiDS team (if any), and deviations of the true PSF from a Moffat function.

In Fig. 15 we show the corresponding plots to Figs 12 and 14; where on the x-axis we now have the true (input) parameters of our simulated galaxies and on the y-axis the difference between the fitted and the true values; both limited to segment radii.

Generally, all parameters are recovered well, although the magnitudes show a slight offset of ~ 0.01 mag (with corresponding trends in effective radius and Sérsic index since these parameters are correlated); worsening for faint objects. This offset is driven by a number of galaxies scattering to very low values, i.e. where the fit attributes significantly more flux to the galaxy than what we put into the simulation. Visual inspection of these simulated objects revealed that all of them have additional flux from other objects included in the segmentation maps. Fig. 16 shows an example, where the difference between the fitted and the true magnitude is -0.17.

Nearby objects affect approximately 5–10 per cent of our simulated fits in this way. Since it is a one-sided effect (there are no sources with negative flux), it results in a slight overall bias across the sample. This is expected to occur at a similar level also in the fits to real galaxies and could only be improved by simultaneously
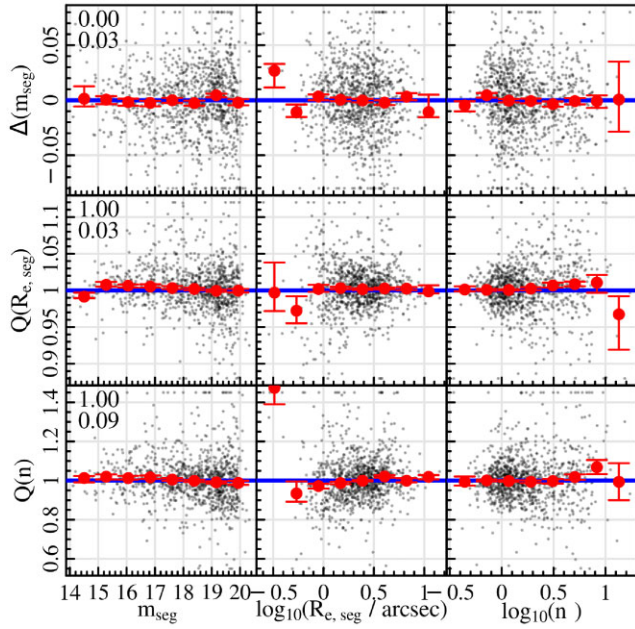
Figure 14. Similar to the bottom panels of Fig. 12 but now showing an internal consistency check of our catalogue using galaxies that were imaged (and successfully fitted) at least twice in the KiDS *r* band. We pass these duplicate observations of the same physical objects through our pipeline independently and then compare the fit to the shallower image with the fit to the deeper (higher signal-to-noise ratio) image. Note the different plotting ranges relative to Fig. 12 (on the *y*-axis in particular).
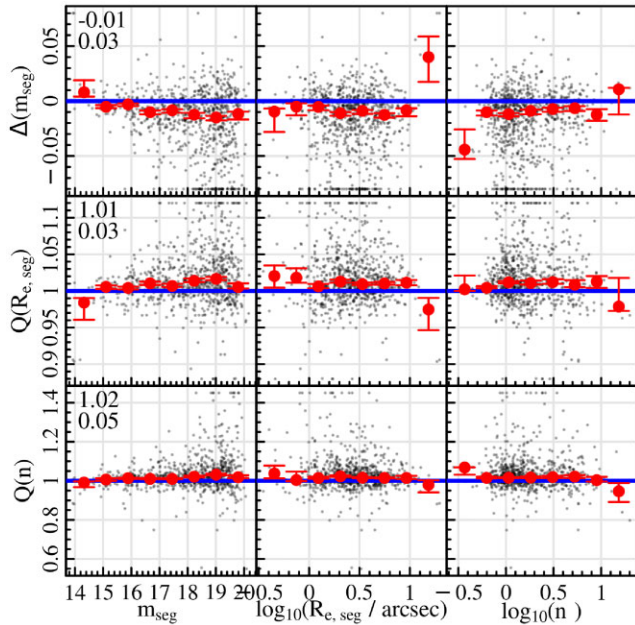

Figure 16. An example fit to a simulated galaxy where the difference between the true and the fitted magnitude is large due to the wings of a nearby bright object and a small faint object included in the segmentation map. Panels are the same as top and middle row in Fig. 3.

Figure 17. An example fit to a simulated galaxy which was flagged as a bad fit due to a nearby masked area from a bright star chopping up the segmentation map. Panels are the same as top and middle row in Fig. 3.

6.3 Simulations: model selection accuracy

Since we know that all of our input galaxies were perfect single Sérsic systems, the model selection and outlier rejection statistics can be used to judge the failure rate of these routines (cf. also Sections 3.3 and 4.2). We simulated 1000 galaxies at random locations; which resulted in 1126 objects to be fit (due to the overlap regions between KiDS tiles). Of these, 262 (23 per cent) were skipped; which is similar to the fraction of skipped fits for real galaxies, as expected since the main reason for this are the KiDS masks. Of the remaining objects, 94 per cent are classified as single component fits, three per cent are 1.5 or double component fits and 3 per cent are flagged as outliers.

The number of outliers is significantly less than the 11 per cent of real *r*-band galaxies flagged (Section 3.3.1) because in the simulations all galaxies are intrinsically 'well-behaved'. Fig. 17 shows an example of the most commonly occurring reason for being flagged as an outlier (in the simulations), namely the mask of a nearby bright star chopping up the segmentation map.

The fact that 97 per cent of non-outlier simulated galaxies are correctly classified as single component fits confirms that model selection is accurate provided the galaxy can be unambiguously assigned to the single Sérsic model (cf. Section 3.3.2). We visually inspected the 3 per cent 1.5 and double component fits and found nearby interfering objects in all of them. Fig. 18 shows an example,

Figure 15. Similar to Fig. 14 but comparing the fitted parameters of simulated images to the true (input) parameters; both limited to the segment radius.

fitting nearby sources (see also the discussion of this issue in Häussler et al. 2007). However, for this work we decided against this option as explained in Section 3.2.5. We may revisit this decision in future work.
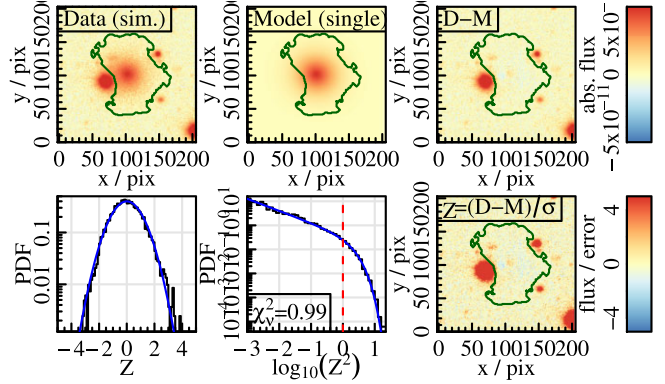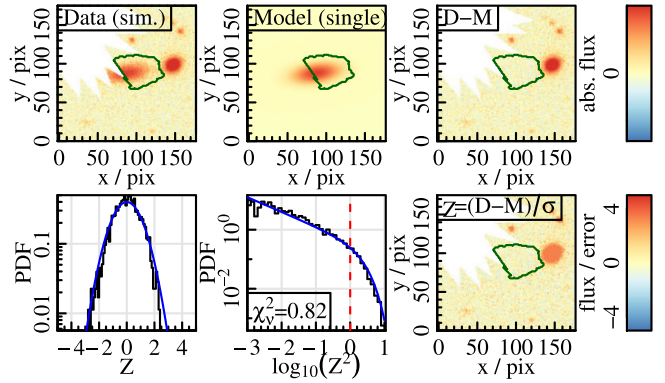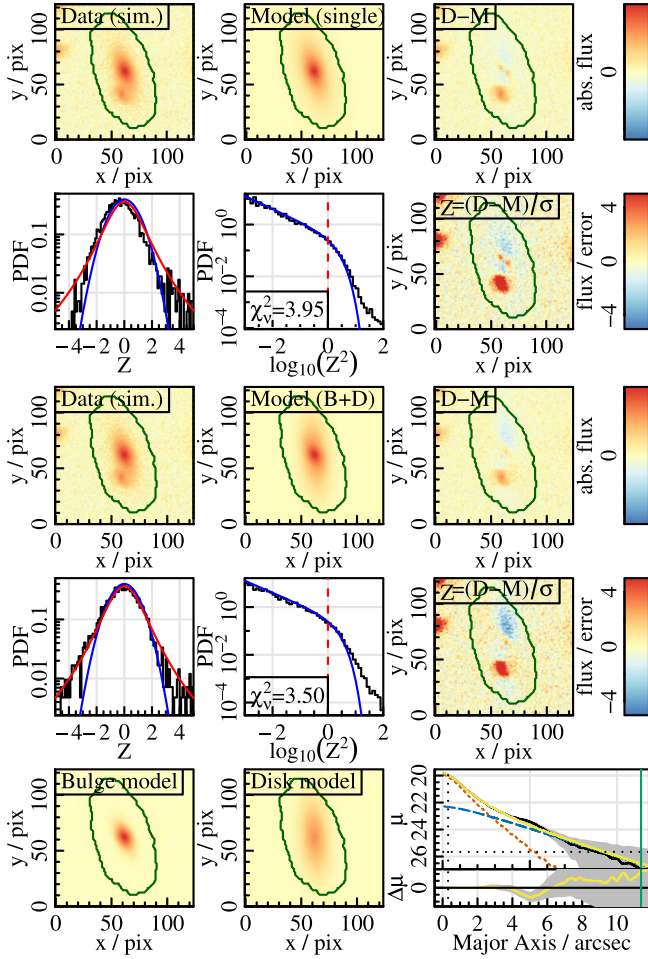
Figure 18. An example fit to a simulated galaxy which was classified as a double component fit. Single component fit at the top; double component fit at the bottom. Panels are the same as in Fig. 3.

where the fit attempts to capture the additional 'features' with the freedom of a second component. Note that since we only simulated single Sérsic objects, we cannot comment on the accuracy of the model selection procedure for double component objects here. However, our model selection procedure was optimized on all types of real galaxies (not just single Sérsic objects), see Section 3.3.2 for details. We also point out that the confusion rate of the model selection derived from the idealized simulations is a strictly lower limit. Realistic confusion rates for real galaxies, as calibrated against visual inspection, are quantified in Table 3 (Section 3.3.2) for the *r* band as well as Tables B1, B2, and B3 (Appendix B) for the *g*, *i*, and joint model selection.

6.4 Systematic uncertainties

Fig. 19 shows the results of our systematic error study, which we will now discuss in detail. Going from left to right we show three different plot types as labelled on the *x*-axis and described in more detail in the caption and below. Going from top to bottom, each plot type is shown for each single Sérsic parameter as labelled in the left-hand panels; and the colours of the lines in the plot indicate which sample was used according to the legend at the bottom of the figure. The left-hand panels show the distribution of absolute differences between the values fitted to both versions of a galaxy (in the indicated units);

the middle panels show the corresponding error distribution (errors added in quadrature for the two fits; in the same units as the fit); and the right-hand panels show the distribution of the absolute difference divided by its error (unit-less). The solid black lines labelled 'real data overlap' are using the overlap sample. The dashed dark blue lines labelled 'simulation overlap' show the same for the overlap sample in simulated galaxies, where we have run more simulations specifically to boost the number of simulated overlap galaxies to a similar value as we have for the real overlap sample ($\sim$700). The dashed light blue line labelled 'sim. overlap (t. PSF)' is the same as the dashed dark blue line, just that when fitting the galaxy, instead of the estimated PSF we passed in the true PSF (i.e. the one we used to convolve the model galaxy with originally). The dotted dark orange line labelled 'simulation fit versus true' shows the difference between the fitted value and the true value (instead of between the two fitted values in the overlap region) for the same sample of galaxies. Note the errors here now are just the errors of the fit since the true values do not have errors. The dotted light orange line ('sim. fit versus true (t. PSF)') is the same for the run that used the true PSFs.

All values are clipped to the plotting intervals (for plotting only). For scale parameters, all distributions are shown in logarithmic space (which the parameters were also fitted in). To make the scales comparable to the other parameters, the angle is shown in units of $30°$ (which it was also fitted in to make the MCMC step size comparable to the other parameters). For the magnitude and effective radius, we show both the fitted Sérsic values and the segment truncated values (cf. Section 3.3.3). Comparing rows 3 and 4 (Sérsic and segment magnitudes) or rows 5 and 6 (Sérsic and segment effective radii) against each other, it becomes clear immediately that the distributions for the segment values are narrower, i.e. limiting to segment sizes increases the stability and reduces the scatter in those parameters as already observed in many previous sections. Note, though, that our simulated galaxies do follow perfect single Sérsic profiles, so the differences between the segment and Sérsic values are generally expected to be much smaller in the simulations than in real data (see Section 3.3.3 for details).

For each distribution, we also give the median of the absolute difference and the errors and the 1σ-quantile (half of the range containing the central 68 per cent of data) of the normalized difference. These values (with errors) are also given in Table 6.

6.4.1 Overlap sample: real versus simulated

Focussing on the real data and simulation overlap samples (solid black and dashed dark blue lines), which are most directly comparable, it can be seen that the distributions of the parameter differences (left column) are broader for the real galaxies than for the simulated galaxies for all parameters; i.e. for two versions of the same galaxy in the overlap sample, the fitted values are on average closer to each other in the simulation than in the real data. This could be due to two reasons: either irregular galaxy features in combination with noise (i.e. perfect single Sérsic objects are just more easily constrained/less easily influenced by noise fluctuations); or differences in the KiDS data processing between tiles (e.g. inaccuracies in their background subtraction procedure) that affect the real galaxies but not the simulated ones since those were added later. In reality, it is probably a combination of these two effects (with the first one presumably dominating). All steps of our own analysis affect both the simulated and the real galaxies and hence on average will have the same effect on both.
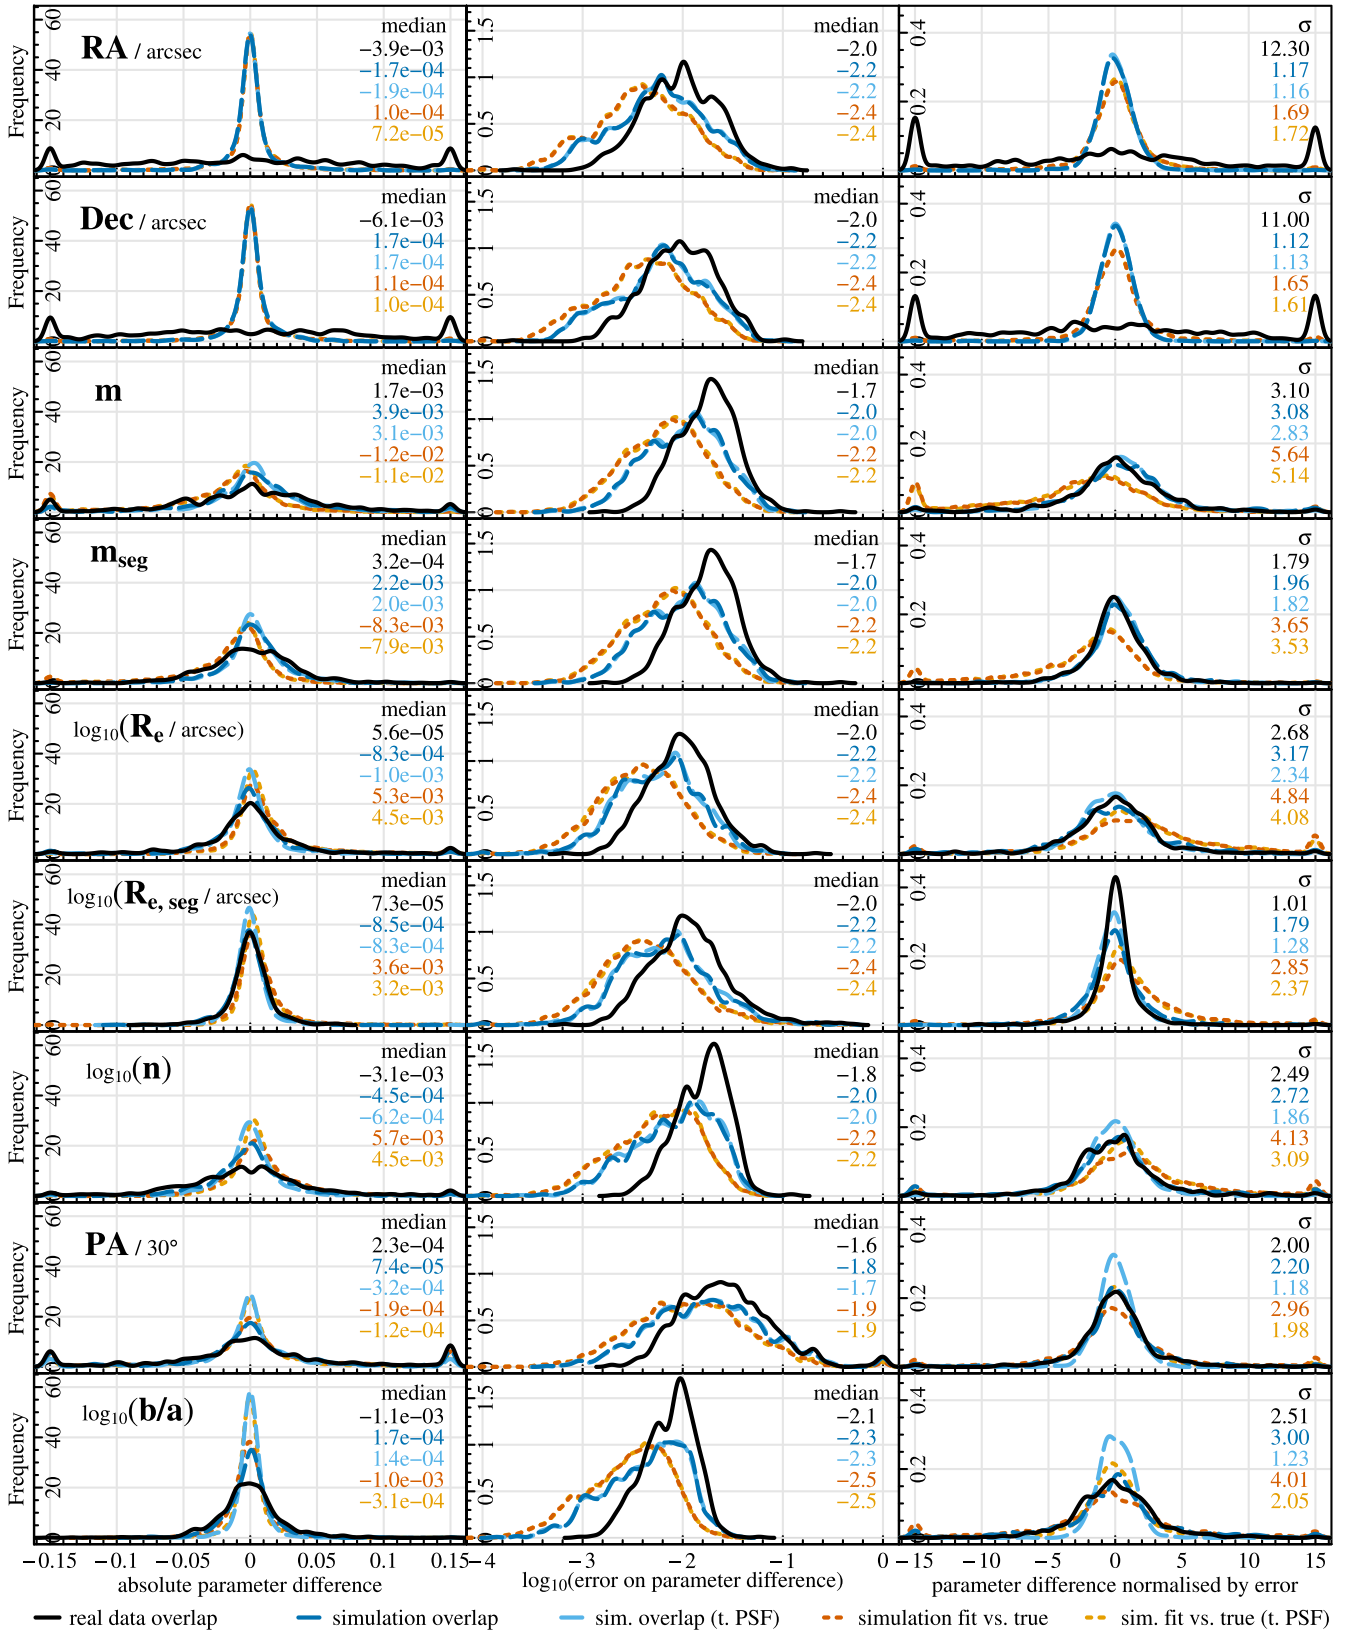
Figure 19. For all single Sérsic parameters as labelled top to bottom: **Left column:** The distribution of the absolute difference between the fitted and true values for simulated galaxies; or between the fitted values to two versions of the same (simulated or real) galaxy in the overlap sample. The legend at the bottom indicates which difference is shown; scale parameters are treated in logarithmic space throughout. See text for details. **Middle column:** The error on the parameter difference shown in the left column. **Right column:** The parameter difference normalized by its error (i.e. left column divided by middle column).

The errors (second column) do reflect this additional uncertainty in real galaxies in that they are larger by 0.2–0.3 dex for all parameters. In fact, the errors on the simulated galaxies seem to be more severely underestimated than those on the real galaxies, which becomes clear when looking at the parameter differences normalized by the respective errors (third column). In an ideal world, these would all be Gaussians centred on zero with a standard deviation of 1. As there will always be a few outliers due to interfering objects or image artefacts, instead of the mean and standard deviation we will consider their more robust equivalents, the median, and 1σ-quantile (shown in plots). Indeed, all overlap sample distributions (simulated and real, i.e. black, dark blue and light blue lines) are centred on zero, as already expected from the results shown in Fig. 14. However, it can be seen that for all parameters except segment effective radius, the 1σ-quantile is larger than 1 both for the simulated and the real sample: values generally range between 2 and 3 with simulated galaxies performing slightly worse due to the underestimated errors (which are most likely caused by the PSFs, see discussion below).

The exception to this is position (RA and Dec, top two rows), for which the normalized distribution is a factor of approximately 10 broader for real galaxies than for simulated ones. In fact, a considerable fraction of these distributions fall outside of the plotting range, such that the clipping to these intervals results in prominent peaks at the plot edges (top right two panels of Fig. 19). We believe this to be mainly due to the accuracy of the astrometric solution of the KiDS data, which shows a scatter of approximately 0.04 arcsec in DR4.0 in both RA and Dec (Kuijken et al. 2019; and we also confirmed this using the KiDS r-band source catalogues). This is a factor of 4 larger than the median MCMC error on position (top two panels in the middle row of Fig. 19). Accounting for this additional source of scatter between the tiles (which only affects real objects but not the simulations since those were inserted after the astrometric calibration) would bring the normalized error distribution for the real data overlap sample into much closer agreement with the simulated version. The remaining factor of $\sim$2–3 difference could also be due to the astrometry, considering that the overlap sample by definition sits at tile edges, where the astrometric solution is the most uncertain; or – as for all other parameters – due to irregular galaxy features in combination with noise (see discussion above). As a last point we would like to note that the absolute differences in position are usually still within 1 pixel (0.2 arcsec), i.e. although it stands out from the plot, this is a sub-pixel effect.

6.4.2 Simulated overlap: imperfect versus true PSFs

These distributions of the simulated overlap sample can be compared to their equivalent distributions using the true PSFs (dashed light blue lines in Fig. 19). This allows us to determine which parameters are affected by imperfect PSF estimates. However, we note that we can only make qualitative and relative statements here since we do not know how close to the truth our estimated PSFs for the real galaxies are. When simulating our galaxies, we convolve it with the model PSF fitted to the nearest real galaxy, i.e. this is the true PSF (cf. Section 6.2). When processing the simulated galaxy through our pipeline, the estimated PSF is then obtained by fitting nearby stars in the usual way. The nearest galaxy (which the true PSF is based on) is typically around 200 arcsec away, with the distribution ranging between $\sim$0 and $\sim$500 arcsec. This is close enough to provide a realistic PSF for the position of interest since KiDS tiles are much

larger than this (~ 1 deg^2) and the PSF varies only slowly across the tiles. However, it is further away than the stars used to obtain the estimated PSF, which are typically within $\sim$100–200 arcsec and can at the very maximum be $\sqrt{2} * 200$ arcsec away since the large cutouts used for PSF fitting are 400 arcsec on each side. This is expected since the density of stars is much higher than that of GAMA galaxies. However, it implies that the deviations of the estimated to the true PSFs in the simulations will on average be larger than for real data, hence leading to the errors on simulated galaxies being more severely underestimated as noted in the previous sections.

Comparing the simulations with true PSFs to the real data, it can be seen that the simulations now perform better than the real data for all parameters except segment effective radius (1σ-quantiles between 1 and 2.8). In addition, comparing the simulations with the true and the false PSFs against each other, we can assess which parameters are most affected (relatively speaking): position angle and axial ratio are most severely influenced; followed by Sérsic index, effective radius, and magnitude, while the position is nearly unaffected. This makes sense: the axial ratio and position angle are very sensitive to mistakes in the ellipticity and orientation of the PSF; while the fitted Sérsic index, effective radius, and magnitude depend on the concentration and FWHM of the PSF. The position is only very weakly affected since the PSF is always centred and symmetric.

Note that all 1σ-quantiles are still larger than 1 (ranging from 1.1 to 2.8) even for the simulations with true PSFs. This indicates that the error underestimates of these parameters are not exclusively caused by the effects studied so far (galaxy and/or image processing features not accounted for in the simulations and PSF uncertainties) but there is an additional contribution from features that are also present in the simulations such as nearby objects, noise fluctuations, background subtraction inaccuracies, or image artefacts.

6.4.3 Simulated sample: fitted versus true values

Finally, for the simulated samples we can compare the fitted values to the true values (instead of the overlap sample comparison), which is shown as dotted orange lines in Fig. 19. This allows us to detect biases, but is less directly comparable to the sample of real galaxies where the true value is unknown. Note the errors are generally slightly smaller compared to the overlap studies since for those, the errors for both fits were added in quadrature while the true values now obviously do not have errors. Correspondingly, the normalized distributions are slightly broader even though the absolute differences between parameters are comparable. Most notably, however, the median of the distribution is now shifted away from zero for magnitude, effective radius, and Sérsic index (see also Table 6). This is due to the bias caused by nearby objects already described in Section 6.2: magnitudes are too bright by ~ 0.01 mag; effective radii and Sérsic indices too large by approximately 1 per cent (always a bit better for segment values and/or simulations using the true PSFs). All other parameters still have their distributions centred on zero, i.e. do not show any bias – at least not one that we can test with our simulations. This makes sense since position, axial ratio, and angle will be influenced by nearby objects (and other effects) as well, but without any preferred direction and hence on average this leads to an error underestimate rather than an overall bias. Using the true PSFs (dotted light orange lines) narrows all distributions slightly as expected; but again there is only an error underestimate rather than an overall bias introduced by the wrong PSFs since they are 'randomly wrong'.

Table 6. Biases and error underestimates for all single Sérsic parameters derived from our systematic error studies (Section 6.4). The bias is additive (indicated with $\pm$) for those parameters that were treated in linear space and multiplicative (indicated with $\divideontimes$) for those treated in logarithmic space. Error underestimates are always multiplicative. The column 'bias/σ' gives the significance of each bias.

Param.	Bias (using true PSFs)	Bias/σ	Error underest.
RA	$\pm (7 \pm 18) \times 10^{-5}$ arcsec	0.39	12.27 ± 0.71
Dec	$\pm (10 \pm 19) \times 10^{-5}$ arcsec	0.56	10.98 ± 0.53
m	$\pm (-11.4 \pm 0.8) \times 10^{-3}$	-13.83	3.10 ± 0.16
m_{seg}	$\pm (-8 \pm 0.5) \times 10^{-3}$	-14.92	1.79 ± 0.09
R_e	$\divideontimes (1.0105 \pm 0.0009)$	11.61	2.68 ± 0.12
$R_{e, seg}$	$\divideontimes (1.0074 \pm 0.0005)$	14.39	1.01 ± 0.05
n	$\divideontimes (1.010 \pm 0.001)$	10.93	2.49 ± 0.10
PA	$\pm (-4 \pm 19) \times 10^{-3}$ deg	-0.19	2.00 ± 0.10
b/a	$\divideontimes (0.9993 \pm 0.0004)$	-2.07	2.51 ± 0.09

One source of potential bias that we cannot test with the simulations are galaxy features not accounted for in the models. If, for example, there is a large population of galaxies that have bars; and these bars lead to the bulge axial ratios being systematically underestimated, this is again a one-sided effect that could lead to an overall bias. In addition, such features could further increase the error underestimate because they will obviously tend to influence both fits to a galaxy in the overlap sample in the same way and hence are difficult to detect in the above analysis. If there are systematic one-sided deviations of the true PSFs from Moffat functions, these could lead to one-sided systematically wrong PSF estimates which could in turn also introduce an additional bias that cannot be tested by the simulations which use Moffat model PSFs.

6.4.4 Corrections for systematics and their validity

Table 6 summarizes the results of the systematic error studies: for all single Sérsic parameters (plus the segment magnitude and segment effective radius), we give the average bias and error underestimates with errors. The bias is estimated from the median of the offset between fitted and true values in the simulation using the true PSFs (light orange numbers in the first column of Fig. 19). The errors on the median are taken as the 1σ-quantiles of these distributions divided by the square-root of the number of data points (~ 2000). Logarithmic parameters are converted back into linear space to simplify bias correction in the catalogue. None the less, scale parameters have multiplicative correction factors while location parameters have additive corrections (in given units). In other words: to correct for the bias, subtract the values in Table 6 from the catalogue values for position, magnitude, and position angle; and divide by the given values for effective radius, Sérsic index, and axial ratio. Users should note, however, that due to the way these biases were estimated, they do not include all sources of potential bias (e.g. galaxy features such as bars and spiral arms; or systematically wrong PSFs). Also, we recommend to apply the bias correction only to statistically large and random samples; they are average values not applicable to individual galaxies as evident from Fig. 15.

The next column in Table 6 gives the significance of each bias, which is the deviation of the median from 0 (or 1 for scale parameters) divided by its error. It can be seen that position, position angle, and axial ratio are not biased (consistent with 0/1 within 2σ), while

magnitude, effective radius, and Sérsic index are biased (deviation from 0/1 of $>5\sigma$); as found and discussed before.

Finally, the last column in Table 6 gives the error underestimates estimated from the width of the distribution of the normalized difference between fits to two versions of the same galaxy in the overlap sample (black numbers in the last column of Fig. 19). The errors in this case were estimated by bootstrapping the distributions 1000 times each to get an estimate of the variation of the distribution width. Since these distributions were normalized (by the respective errors), there is no need to convert between linear and logarithmic space and the given values can hence directly be used as correction factors for the MCMC errors. We have applied the relevant correction to all quoted errors in the catalogue, but also give the original (purely random) errors for completeness. Also note that since these values are now based on real data, they do include PSF uncertainties (in contrast to the biases). Also, since the overlap sample is in many ways the worst in terms of data quality (sitting at tile edges), these are likely upper limits. However, they still do not include error underestimates caused by (galaxy) features not accounted for in the models such as bars, rings, spiral arms or similar, as well as nearby objects. Since these are physical (rather than related to the data taking or image processing), they will be present in both versions of the overlap sample galaxy and influence both fits in similar ways – leading to a (random) bias on individual galaxy fits; and hence an error underestimate for a large enough sample. These issues should be kept in mind when using the catalogue.

7 CONCLUSIONS

In this work we presented our pipeline for the single Sérsic fits and bulge-disc decompositions of 13 096 galaxies at redshifts $z < 0.08$ in the GAMA II equatorial fields in the KiDS g, r, and i bands. The galaxy modelling is done using PROFIT, the Bayesian two-dimensional surface profile fitting code of Robotham et al. (2017), fitting three models to each galaxy:

    (i) a single Sérsic component,
    (ii) a two-component model consisting of a Sérsic bulge plus exponential disc and
    (iii) a two-component model consisting of a point source bulge plus exponential disc (for unresolved bulges).

The preparatory work (image segmentation, background subtraction and obtaining initial parameter guesses) is carried out using the sister package PROFOUND (Robotham et al. 2018); with the PSF estimated by fitting nearby stars using a combination of PROFOUND and PROFIT. Segmentation maps are defined on joint *gri*-images, while the remaining analysis is performed individually in each band except for the model selection, for which we offer both a per-band and a joint version. The analysis is fully automated and self-contained with no dependency on additional tools.

In addition to the galaxy fitting, we perform a number of post-processing steps including flagging of bad fits and model selection. An overview of the number of galaxies successfully fitted in each band as well as the number classified in each category is given in Table 4 and Fig. 5. For our planned applications of the catalogue, which involves the statistical study of dust attenuation effects, we need fits that are most directly comparable to each other. Hence, we choose to model a maximum of two components for each galaxy even if more features may be present; and focus on achieving good fits in the high signal-to-noise regions of the galaxies by choosing relatively small segments for fitting. Consequently, we recommend

using truncated magnitudes and effective radii for all analyses instead of the Sérsic values which are extrapolated to infinity by definition. The quality of the fits is tested and ensured by visual inspection, comparing to previous works (Kelvin et al. 2012; Lange et al. 2015), studying independent fits of galaxies in the overlap regions of KiDS tiles and bespoke simulations. The latter two are also used for a detailed analysis of how systematic uncertainties affect our fits.

We find that the combination of PROFOUND and PROFIT is well suited to the automated analysis of large data sets. The fully Bayesian MCMC treatment enabled by PROFIT is able to overcome the main shortcomings of traditionally engaged downhill-gradient based optimizers, namely their susceptibility to initial guesses and their inability to easily derive realistic error estimates. The watershed de-blending algorithm used by PROFOUND is less prone to catastrophic segmentation failures and allows us to extract more complex object shapes than other commonly used algorithms based on elliptical apertures; while still preserving the total flux well. With its wealth of utility functions, it not only facilitates the robust segmentation of large sets of images but also provides sky background estimates and reasonable initial guesses for the MCMC fitting.

These characteristics, in combination with our own routines for quality assurance, led to results that are robust across a variety of galaxy types and image qualities and in reasonable agreement with previous studies given the different data, code, and focus of the analysis. The outlier rejection routine efficiently identifies objects for which none of our models is appropriate such as irregular galaxies or those compromised by masked areas. Model selection is based on a DIC cut and accurate to > 90 per cent compared to what could be achieved by visual inspection. There is a minimal bias in the fitted magnitude, effective radius, and Sérsic index of approximately 0.01 mag, 1 per cent, and 1 per cent, respectively (on average across the full sample), caused by excess flux from nearby other objects. The errors obtained from the MCMC chains are underestimated with respect to the true errors by factors typically between 2 and 3 (see Table 6) and can easily be corrected for statistically large samples of galaxies.

All results are integrated into the GAMA database as `v04` of the `BDDecomp` DMU. This consists of a total of eight catalogues giving the results of the preparatory work, the 2D surface brightness distribution fits and the post-processing of all 13 096 galaxies in our full sample ($z < 0.08$ in the GAMA II equatorial survey regions) in the KiDS *g*, *r*, and *i* bands (see Section 4 for details).

The full DMU is currently available to GAMA team members with a version restricted to SAMI galaxies available to the SAMI team. It will be made publicly available in one of the forthcoming GAMA data releases. Readers interested in using (parts of) the catalogue before it is publicly released are encouraged to contact the authors to explore the possibilities for a collaboration.[14]

We plan to extend this work to include the KiDS *u* and VIKING *Z*, *Y*, *J*, *H*, *K_s* bands, ideally making use of PROFIT's new multiframe and multiband fitting functionality. The decompositions will then be used to derive the stellar mass functions of bulges and discs and constrain the nature and distribution of dust in galaxy discs. This is achieved by comparing the distribution of bulges and discs in the luminosity–size plane to the dust radiative transfer models of Tuffs and Popescu (and preceding papers of this series Popescu et al. 2011); essentially expanding the work of Driver et al. (2007a) with more and better data and at several wavelengths.

---

[14] http://www.gama-survey.org/collaborate/

DATA AVAILABILITY

The data underlying this article are integrated into the GAMA database. For access options, see http://www.gama-survey.org/collaborate/.

REFERENCES

Akaike H., 1974, IEEE Trans. Autom. Control, 19, 716

Allen P. D., Driver S. P., Graham A. W., Cameron E., Liske J., de Propris R., 2006, MNRAS, 371, 2

Argyle J. J., Méndez-Abreu J., Wild V., Mortlock D. J., 2018, MNRAS, 479, 3076

Baldry I. K. et al., 2010, MNRAS, 404, 86

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

Baldry I. K. et al., 2014, MNRAS, 441, 2440

Barden M. et al., 2005, ApJ, 635, 959

Barsanti S. et al., 2021, ApJ, 911, 21

Bellstedt S. et al., 2020, MNRAS, 496, 3235

Benson A. J., Džanović D., Frenk C. S., Sharples R., 2007, MNRAS, 379, 841

Bertin E., Arnouts S., 1996, A&AS, 117, 393

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

Blanton M. R., Eisenstein D., Hogg D. W., Schlegel D. J., Brinkmann J., 2005, ApJ, 629, 143

Bluck A. F. L. et al., 2019, MNRAS, 485, 666

Bottrell C., Torrey P., Simard L., Ellison S. L., 2017, MNRAS, 467, 2879

Bottrell C., Simard L., Mendel J. T., Ellison S. L., 2019, MNRAS, 486, 390

Broyden C. G., 1970, J. Inst. Math. Appl., 6, 76

Bryant J. J. et al., 2015, MNRAS, 447, 2857

Cluver M. E. et al., 2020, ApJ, 898, 20

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

Cook M., Lapi A., Granato G. L., 2009, MNRAS, 397, 534

Cook R. H. W., Cortese L., Catinella B., Robotham A., 2019, MNRAS, 490, 4060

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

de Graaff A., 2022, MNRAS, 511, 2544

de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, Exp. Astron., 35, 25

de Jong J. T. A. et al., 2015, A&A, 582, A62
de Souza R. E., Gadotti D. A., dos Anjos S., 2004, ApJS, 153, 411
de Vaucouleurs G., 1948, Ann. d'Astrophysique, 11, 247
Dimauro P. et al., 2018, MNRAS, 478, 5410
Dimauro P. et al., 2019, MNRAS, 489, 4135
Driver S. P., Popescu C. C., Tuffs R. J., Liske J., Graham A. W., Allen P. D., de Propris R., 2007a, MNRAS, 379, 1022
Driver S. P., Allen P. D., Liske J., Graham A. W., 2007b, ApJ, 657, L85
Driver S. P., Popescu C. C., Tuffs R. J., Graham A. W., Liske J., Baldry I., 2008, ApJ, 678, L101
Driver S. P. et al., 2009, Astron. Geophys., 50, 5.12
Driver S. P. et al., 2011, MNRAS, 413, 971
Driver S. P., Robotham A. S. G., Bland-Hawthorn J., Brown M., Hopkins A., Liske J., Phillipps S., Wilkins S., 2013, MNRAS, 430, 2622
Driver S. P. et al., 2016, MNRAS, 455, 3911
Driver S. P. et al., 2022, MNRAS, 513, 439
Dutton A. A. et al., 2011, MNRAS, 410, 1660
Edge A., Sutherland W., Kuijken K., Driver S., McMahon R., Eales S., Emerson J. P., 2013, The Messenger, 154, 32
Erwin P., 2015, ApJ, 799, 226
Fernández Lorenzo M. et al., 2014, ApJ, 788, L39
Fletcher R., 1970, Comput. J., 13, 317
Gadotti D. A., 2009, MNRAS, 393, 1531
Gao H., Ho L. C., 2017, ApJ, 845, 114
Goldfarb D., 1970, Math. Comput., 24, 23
Graham A. W., 2013, in Oswalt T. D., Keel W. C., eds, Planets, Stars and Stellar Systems. Springer, Dordrecht, p. 91
Graham A. W., Driver S. P., 2005, PASA, 22, 118
Häussler B. et al., 2007, ApJS, 172, 615
Häußler B. et al., 2013, MNRAS, 430, 330
Häußler B. et al., 2022, A&A, 664, A92
Head J. T. C. G., Lucey J. R., Hudson M. J., Smith R. J., 2014, MNRAS, 440, 1690
Hilbe J. M., de Souza R. S., Ishida E. E. O., 2017, Bayesian Models for Astrophysical Data: Using R, JAGS, Python, and Stan. Cambridge Univ. Press, Cambridge
Hopkins A. M. et al., 2013, MNRAS, 430, 2047
Hyde J. B., Bernardi M., 2009, MNRAS, 394, 1978
Ivezić Ž. et al., 2019, ApJ, 873, 111
Jeffreys H., 1946, Proc. R. Soc., 186, 453
Johnson S. G., 2017, The NLopt Nonlinear-Optimization Package.. Available at http://ab-initio.mit.edu/nlopt
Kawinwanichakij L. et al., 2021, ApJ, 921, 38
Kelvin L. S. et al., 2012, MNRAS, 421, 1007
Kelvin L. S. et al., 2014, MNRAS, 444, 1647
Kennedy R. et al., 2016, MNRAS, 460, 3458
Kim K., Oh S., Jeong H., Aragón-Salamanca A., Smith R., Yi S. K., 2016, ApJS, 225, 6
Kormendy J., 1977, ApJ, 218, 333
Kruk S. J. et al., 2018, MNRAS, 473, 4731
Kuijken K. et al., 2019, A&A, 625, A2
La Barbera F., de Carvalho R. R., de La Rosa I. G., Lopes P. A. A., Kohl-Moreira J. L., Capelato H. V., 2010, MNRAS, 408, 1313
Lackner C. N., Gunn J. E., 2012, MNRAS, 421, 2277
Lagos C. d. P. et al., 2018, MNRAS, 473, 4956
Lange R. et al., 2015, MNRAS, 447, 2603
Lange R. et al., 2016, MNRAS, 462, 1470
Lawrence A. et al., 2007, MNRAS, 379, 1599
Liske J. et al., 2015, MNRAS, 452, 2087
Meert A., Vikram V., Bernardi M., 2015, MNRAS, 446, 3943
Meert A., Vikram V., Bernardi M., 2016, MNRAS, 455, 2440
Mendel J. T., Simard L., Palmer M., Ellison S. L., Patton D. R., 2014, ApJS, 210, 3
Méndez-Abreu J. et al., 2017, A&A, 598, A32
Moffett A. J. et al., 2016, MNRAS, 462, 4336
Nedkova K. V. et al., 2021, MNRAS, 506, 928
Oh S. et al., 2020, MNRAS, 495, 4638
Peng C. Y., Ho L. C., Impey C. D., Rix H.-W., 2010, AJ, 139, 2097

Pillepich A. et al., 2018, MNRAS, 473, 4077
Popescu C. C., Tuffs R. J., Dopita M. A., Fischera J., Kylafis N. D., Madore B. F., 2011, A&A, 527, A109
R Core Team, 2020, R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, https://www.R-project.org/
Robotham A. et al., 2010, PASA, 27, 76
Robotham A. S. G., Taranu D. S., Tobar R., Moffett A., Driver S. P., 2017, MNRAS, 466, 1513
Robotham A. S. G., Davies L. J. M., Driver S. P., Koushan S., Taranu D. S., Casura S., Liske J., 2018, MNRAS, 476, 3137
Robotham A. S. G., Bellstedt S., Driver S. P., 2022, MNRAS, 513, 2985
Rodriguez-Gomez V. et al., 2019, MNRAS, 483, 4140
Salo H. et al., 2015, ApJS, 219, 4
Sánchez-Janssen R. et al., 2016, ApJ, 820, 69
Schaye J. et al., 2015, MNRAS, 446, 521
Schwarz G., 1978, Ann. Stat., 6, 461
Sérsic J. L., 1963, Bol. Asociacion Argentina Astron. La Plata Argentina, 6, 41
Shanno D. F., 1970, Math. Comput., 24, 647
Shen S., Mo H. J., White S. D. M., Blanton M. R., Kauffmann G., Voges W., Brinkmann J., Csabai I., 2003, MNRAS, 343, 978
Shibuya T., Ouchi M., Harikane Y., 2015, ApJS, 219, 15
Simard L. et al., 2002, ApJS, 142, 1
Simard L., Mendel J. T., Patton D. R., Ellison S. L., McConnachie A. W., 2011, ApJS, 196, 11
Sivia D. S., Skilling J., 2006, Data Analysis: A Bayesian Tutorial, 2nd edn. Oxford Science Publications, Oxford Univ. Press, Oxford
Spiegelhalter D., Best N., Carlin B., Van Der Linde A., 2002, Journal of the Royal Statistical Society. Series B: Statistical Methodology, 64, 583
Statisticat LLC., 2018, LaplacesDemon: Complete Environment for Bayesian Inference. Available at https://web.archive.org/web/20150206004624/http://www.bayesian-inference.com/software
Taranu D. S., 2022, Astrophysics Source Code Library, record ascl:2201.005
Tasca L. A. M. et al., 2014, A&A, 564, L12
Taylor E. N. et al., 2011, MNRAS, 418, 1587
Trujillo I. et al., 2006, ApJ, 650, 18
van der Kruit P. C., Searle L., 1981, A&A, 95, 105
van der Wel A. et al., 2012, ApJS, 203, 24
van der Wel A. et al., 2014, ApJ, 788, 28
Vika M., Bamford S. P., Häußler B., Rojas A. L., 2014, MNRAS, 444, 3603
Vika M., Bamford S. P., Häußler B., Rojas A. L., Borch A., Nichol R. C., 2013, MNRAS, 435, 623
Vogelsberger M. et al., 2014, MNRAS, 444, 1518
York D. G. et al., 2000, AJ, 120, 1579

APPENDIX A: PSF ESTIMATION DETAILS

The selection of star candidates for PSF estimation proceeds in three steps.

From the segmentation statistics returned by PROFOUNDPRO-FOUND, we select relatively round and isolated objects as follows:

(i) objects that do not touch other segments, masked regions or image edges (edge fraction = 1)
(ii) objects with a regular boundary geometry (edge excess < 1)
(iii) objects with an axial ratio (minor/major axis) larger than 0.5
(iv) objects which were not flagged as possibly spurious

Of these relatively round and isolated objects, a given fraction (depending on the depth of the image and source extraction, 4 per cent for the *gri* stacks and a SKYCUT value of 2) are identified as star candidates via a joint cut in R50 (semi-major axis containing half the flux) and the concentration (R50/R90, where R90 is the semi-major axis containing 90 per cent of the total flux). This selection is based on the notion that we would expect stars to be small (i.e.

low R50) and highly concentrated (i.e. low R50/R90) and was calibrated empirically.

Around each of these star candidates, a smaller cutout is taken and a subsample selected:

(i) objects brighter than the 5σ point source detection limit and fainter than the saturation limit (both taken from the headers of the corresponding KiDS tile)

(ii) objects where less than 10 per cent of pixels in the cutout belong to other segments

(iii) objects where the star cutout does not overlap with the edge of the large cutout

(iv) objects with a positive sum of the cutout (excluding poorly background-subtracted and/or purely noise-dominated objects)

In the next step, the star candidate cutouts are normalized to a magnitude of 0, masked appropriately and fitted with a Moffat function using PROFIT (see fitting details in Section 3.1.5).

After fitting, suitable stars for PSF estimation are determined as follows:

(i) The fitted centre in x and y must be within ± 1 pixel of the centre of the cutout (fixed since the cutout was centred on the star)

(ii) The fitted magnitude must be within ± 0.1 mag of 0 (fixed since the cutout was normalized to a magnitude of zero according to the segmentation statistics, so any deviation indicates a difference between the magnitudes estimated by PROFOUND and PROFIT, which likely points to bad segmentation, additional objects in the segment, or a bad model fit)

(iii) The reduced chi-square (χ_ν^2) of the fit must be smaller than 3 (calibrated by inspection to exclude visually bad fits)

(iv) FWHM, concentration index, angle, and axial ratio must not be equal to the fit limits (except for the axial ratio, which is allowed to be exactly 1 although this is the upper limit of the fit).

(v) Outliers in any of FWHM, concentration index, angle, axial ratio or background are rejected via an iterative 2σ clip (in logarithmic space where appropriate).

The stars fulfilling these criteria are classified as suitable, from which the selection is made:

(i) The closest two from each quadrant (eight in total) are selected to make sure they are roughly evenly distributed around the galaxy.

(ii) If one or more quadrant contains less than two stars, the closest stars from any other quadrant (which are not already used) are taken instead to give eight stars in total.

(iii) If there are less than eight stars in total, all of them are used.

(iv) If there are no stars classified as suitable, the object is flagged as having a failed PSF and consequently skipped during the galaxy fitting.

The stars selected in this way are then used to create the model PSF as described in Section 3.1.5.

APPENDIX B: MODEL SELECTION ACCURACY

Tables B1, B2, and B3 show the confusion matrices for the model selection in the *g* and *i* bands and the joint model selection. The confusion matrix for the *r* band is already shown in Table 3.

In addition to this band-specific model selection, we perform a joint model selection for all three bands. For this, we sum the DIC values of all three bands for each model before computing the DIC differences. Then we perform the same optimization procedure as for the single bands (using all $\sim$2000 visually classified objects across the three bands) to obtain cuts in DIC difference which we

Table B1. The confusion matrix for our model selection based on a DIC difference cut compared against visual inspection for the *g* band. All values are in per cent of the total number of visually inspected *g*-band galaxies. Bold font highlights galaxies classified correctly, while grey shows those that were ignored during the calibration.

Visual classification	Number of components		
	1	1.5	2
'single'	**48.9**	0.4	1.0
'1.5'	2.1	**3.7**	0.6
'double'	1.8	0.1	**6.0**
'1.5 or double'	0.4	**1.5**	**1.9**
'unsure'	19.5	0.9	7.3
'unfittable'	1.3	0.3	2.1

Table B2. The confusion matrix for our model selection based on a DIC difference cut compared against visual inspection for the *i* band. All values are in per cent of the total number of visually inspected *i*-band galaxies. Bold font highlights galaxies classified correctly, while grey shows those that were ignored during the calibration.

Visual classification	Number of components		
	1	1.5	2
'single'	**51.8**	0.4	1.5
'1.5'	1.9	**3.0**	1.5
'double'	0.7	0	**8.9**
'1.5 or double'	0	**1.0**	**1.9**
'unsure'	10.0	0.7	12.4
'unfittable'	1.2	0.3	2.2

Table B3. The confusion matrix for our model selection based on a DIC difference cut compared against visual inspection for the joint model selection. All values are in per cent of the total number of visually inspected *g*, *r* and *i* band galaxies. Bold font highlights galaxies classified correctly, while grey shows those that were ignored during the calibration.

Visual classification	Number of components		
	1	1.5	2
'single'	**47.4**	0.5	1.5
'1.5'	2.2	**2.4**	1.4
'double'	2.2	0.4	**7.3**
'1.5 or double'	0.3	**1.1**	**2.1**
'unsure'	15.0	0.7	11.1
'unfittable'	1.0	0.4	2.4

subsequently apply for the model selection. Note that the model selected in this way is by necessity a compromise between the different bands, which have different depth and seeing. In this procedure, approximately 9 per cent of fits are classified wrongly across all bands compared to visual classification. The corresponding confusion matrix is shown in Table B3.

APPENDIX C: FIT EXAMPLES

Figs C1, C2, and C3 show examples of galaxies classified as single Sérsic, 1.5-component fit, and outlier (NCOMP $= 1$, 1.5, and -2), respectively. An example double component fit (NCOMP $= 2$) is already shown in Fig. 3. In addition, Fig. C4 shows an example of a very high B/T object where the bulge dominates both the centre and the wings of the object (see Section 4.5.2).
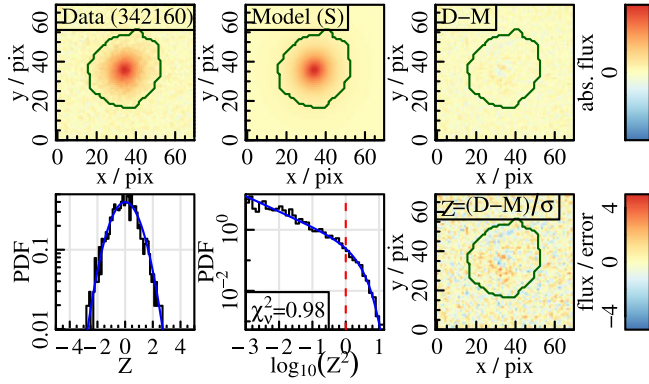
Figure C1. The single Sérsic fit to galaxy 342 160, classified as single component object (NCOMP = 1) in the KiDS *r* band. Panels are the same as the top two rows in Fig. 3.
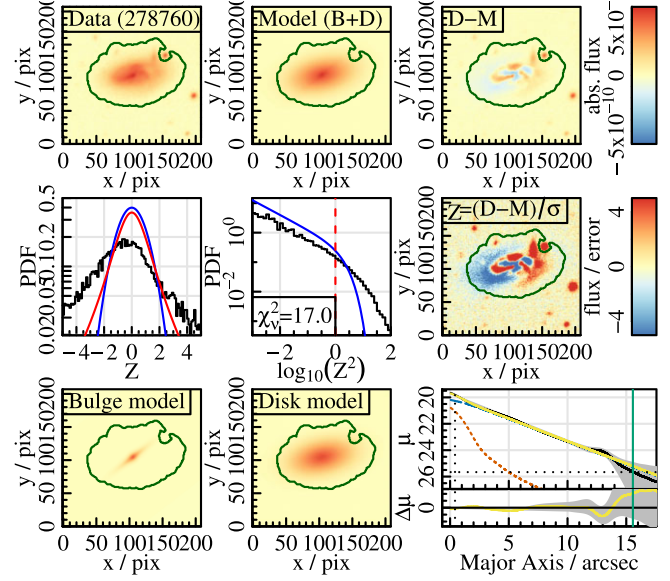

Figure C3. The double component fit to galaxy 278 760, classified as outlier (NCOMP = −2) in the KiDS *r* band. Panels are the same as those in Fig. 3.
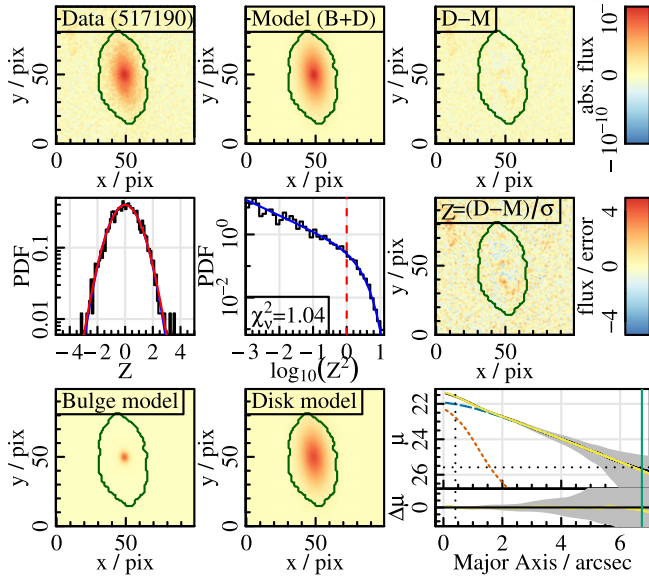

Figure C2. The 1.5-component fit to galaxy 517190, classified as 1.5-component object (NCOMP = 1.5) in the KiDS *r* band. Panels are the same as those in Fig. 3.
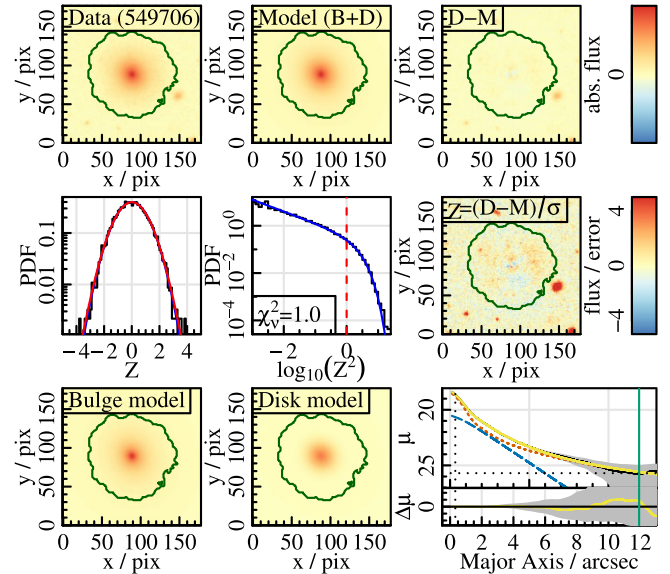

Figure C4. The double component fit to galaxy 549 706, classified as double component object but with a very high B/T ratio of 0.71 in the KiDS *r* band. Panels are the same as those in Fig. 3.

This paper has been typeset from a TeX/LaTeX file prepared by the author.