

Bayesian bulge–disc decomposition of galaxy images

J. J. Argyle,¹★ J. Méndez-Abreu,^{1,2,3} V. Wild¹ and D. J. Mortlock^{4,5,6}

¹*School of Physics and Astronomy, University of St Andrews, North Haugh, St Andrews KY16 9SS, UK*

²*Instituto de Astrofísica de Canarias, Calle Via Lactea s/n, E-38205 La Laguna, Tenerife, Spain*

³*Departamento de Astrofísica, Universidad de La Laguna, E-38206 La Laguna, Spain*

⁴*Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK*

⁵*Statistics Section, Department of Mathematics, Imperial College London, London SW7 2AZ, UK*

⁶*Department of Astronomy, Stockholm University, Albanova, SE-10691 Stockholm, Sweden*

Accepted 2018 June 25. Received 2018 June 22; in original form 2017 August 22

ABSTRACT

We introduce `PHI`, a fully Bayesian Markov chain Monte Carlo algorithm designed for the structural decomposition of galaxy images. `PHI` uses a triple layer approach to effectively and efficiently explore the complex parameter space. Combining this with the use of priors to prevent non-physical models, `PHI` offers a number of significant advantages for estimating surface brightness profile parameters over traditional optimization algorithms. We apply `PHI` to a sample of synthetic galaxies with Sloan Digital Sky Survey (SDSS)-like image properties to investigate the effect of galaxy properties on our ability to recover unbiased and well-constrained structural parameters. In two-component bulge+disc galaxies, we find that the bulge structural parameters are recovered less well than those of the disc, particularly when the bulge contributes a lower fraction to the luminosity, or is barely resolved with respect to the pixel scale or point spread function (PSF). There are few systematic biases, apart from for bulge+disc galaxies with large bulge Sérsic parameter, n . On application to SDSS images, we find good agreement with other codes, when run on the same images with the same masks, weights, and PSF. Again, we find that bulge parameters are the most difficult to constrain robustly. Finally, we explore the use of a Bayesian information criterion method for deciding whether a galaxy has one or two components.

Key words: methods: data analysis – methods: statistical – techniques: image processing – techniques: photometric – galaxies: photometry – galaxies: structure.

1 INTRODUCTION

Galaxy morphologies are complex, arising from many different processes that dictate the formation and evolution of the galaxy as a whole. Accurately characterizing galaxy structure, i.e. bulges, discs, bars, and further complex components, is crucial for furthering our understanding.

The human brain is extremely adept at pattern recognition and the classification of galaxy images began with the ‘Hubble tuning fork’ (Hubble 1936; Sandage 1961). Today visual classification is still widely used and has been recently revitalized by the Galaxy Zoo project enabling amateur galaxy classifiers from across the globe to process vast amounts of galaxy structural information (Lintott et al. 2008, 2011; Willett et al. 2013; Simmons et al. 2017). Furthermore, the introduction of new techniques to mimic how the human brain captures the full, complex distribution of light has further advanced

the usefulness of visual classification (Huertas-Company et al. 2008, 2015; Dieleman, Willett & Dambre 2015).

An alternative and complementary approach is through the quantitative description of galaxy structures, using either parametric or non-parametric methods. Examples of non-parametric classifiers are concentration, clumpiness, and asymmetry (Conselice 2003; Lotz, Primack & Madau 2004; Pawlik et al. 2016). Parametric methods include Sérsic profile fitting (Sersic 1968) and multi-Gaussian expansion (Bendinelli 1991; Emsellem 1998; Fasano & Filippi 1998; Cappellari 2002; Odewahn et al. 2002). Sérsic profile fitting has become increasingly popular in recent decades due to its ability to reproduce the basic structures of many nearby and distant galaxies with typically only one or two (bulge and disc) components. Initially ellipticity-averaged 1D surface brightness profiles were used to fit the photometric components. However, this was shown to lead to systematic errors, as it does not account for the intrinsic shapes or position angle of the bulges (Kormendy 1977; Boroson 1981), and most modern studies now fit the 2D images pixel-by-pixel.

* E-mail: jairomendezabreu@gmail.com

There are a large number of 2D fitting algorithms: GIM2D (Simard 1998), GALFIT (Peng et al. 2002, 2010), BUDDA (de Souza, Gadotti & dos Anjos 2004), GASP2D (Méndez-Abreu et al. 2008, 2017), IMFIT (Erwin 2015), GALPHAT (Yoon, Weinberg & Katz 2011), and PROFIT (Robotham et al. 2017). Some of these codes (e.g. GALFIT and GASP2D) use minimization algorithms to efficiently search for the best solution using the gradient of the model with respect to the parameters. Although these frequentist algorithms have commonly been employed to solve multicomponent, non-linear fits, they suffer from some important drawbacks when faced with a problem as complex as photometric decomposition. Lange et al. (2016) listed five commonly occurring factors that lead to failure of the Levenberg–Marquardt (LM) fitting algorithm: (i) local minima trapping; (ii) unrealistic solutions; (iii) reversal of components (Allen et al. 2006); (iv) indecisiveness as to which model to use; and (v) bad representation of final errors. To avoid some of these issues, they advocated the use of a grid of starting values combined with a convergence test to obtain robust parameter values. Typically 20–30 per cent of automatic fits are deemed physically unrealistic; previous studies have often employed logical filters to weed these out (e.g. Allen et al. 2006; Simard et al. 2011; Meert, Vikram & Bernardi 2015; Méndez-Abreu et al. 2017).

To circumvent these difficulties, a more modern approach is to embed the galaxy morphology analysis into the broader context of inference and hypothesis testing with the use of Bayesian inference. The above problems can then be solved in turn: the exploration of parameter space can overcome runs that become trapped in local minima; initial priors can prevent unrealistic solutions and the reversal of components; model comparison tests can help determine the most probable morphology; and the posterior distribution gives a proper description of the parameter uncertainties. GALPHAT, PROFIT, and IMFIT (version 1.4) offer a Markov chain Monte Carlo (MCMC) approach to help overcome these problems.

This paper introduces a new adaptive Bayesian MCMC algorithm, which has been purpose designed to obtain robust galaxy morphologies from galaxy images. We demonstrate its use on two-component bulge–disc decomposition of both synthetic and real galaxy images. The aim is to provide a flexible, open source code in which it is simple for users to define their own models, point spread functions (PSFs), priors, and likelihoods. The code is available for download in IDL (Interactive Data Language, <https://github.com/SEDMORPH/PHI>). A PYTHON version is also under development.

In Section 2, we describe the basic formalism of the inference methodology, including an overview of Bayesian statistics and details of the priors used. In Section 3, we apply the method to an ensemble of synthetic galaxy images and discuss the interpretation of the outputs. In Section 4, the algorithm is applied to real galaxies within the Sloan Digital Sky Survey (SDSS). In Section 5, we investigate a method to formally compare the one- and two-component model fits. Finally, in Section 6, we discuss and summarize the results of this paper.

2 INFERENCE METHODOLOGY

In this section we describe the main attributes of the code PHI (2D PHotometric decompositions using Bayesian Inference). The key steps of our inference method to perform 2D photometric decompositions, illustrated in Fig. 1, are as follows.

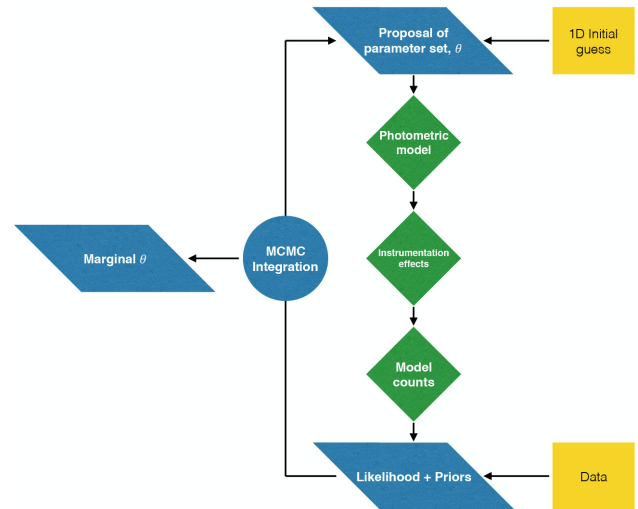


Figure 1. The flow chart of the code PHI (2D PHotometric decompositions using Bayesian Inference).

(i) PHI reads in the flux and error maps of an image in FITS format. A PSF must either be provided as an image or specified in functional form.

(ii) The user provides functions describing the components they wish to fit (Section 2.1) and defines the priors. Initial guesses for the model parameters may be provided, but are not essential for the algorithm to function correctly.

(iii) PHI simulates the galaxy image with the chosen model and initial model parameters, and uses a fast Fourier transform (FFT) to convolve the simulated image with the PSF.

(iv) The likelihood and posterior probability are calculated for the model and data.

(v) The MCMC engine commences by iteratively updating the model parameters and repeating steps (iii)–(iv) until the full posterior distribution has been mapped, or the user-defined maximum iteration number is reached.

2.1 2D photometric model functions

PHI has been designed to easily fit any parametric model to a galaxy image. However, for simplicity, in this paper we focus on the two main observed components of galaxies: spheroids and discs. In general, discs are well described by an exponential profile (Freeman 1970), with the intensity I changing with radius R as

$$I(R) = I_0 \exp\left(-\frac{R}{h}\right), \quad (1)$$

where I_0 is the central intensity and h is the disc scale length. The spheroid component of galaxies can be modelled by Sérsic’s (1968) generalization of de Vaucouleurs (1948) $R^{1/4}$ function to give the $R^{1/n}$ surface density profile,

$$I(R) = I_e \exp\left\{-b_n \left[\left(\frac{R}{R_e}\right)^{1/n} - 1\right]\right\}, \quad (2)$$

where I_e is the intensity at the effective radius R_e that encloses half of the total light from the model, and n , the Sérsic index, describes the concentration of the light profile. The final parameter, b_n , is specified by n . When $n = 1$ the model follows an exponential surface intensity profile and $n = 4$ reproduces the de Vaucouleurs’ model; thus the Sérsic profile can describe the two main observed

Table 1. Details of the prior distributions set for each of the model parameters used in this paper. $U[a, b]$ specifies a uniform distribution between lower and upper limits a and b , and $\delta(a)$ specifies a Kronecker delta function with probability of 1 at a and 0 otherwise. S_{image} is the size of the image.

Parameter	Symbol	Prior	Range
<i>Individual parameters</i>			
Effective intensity	I_e	Uniform in $\log(I_e)$	$\log(I_e/\text{counts}) \in U[0.01, 10]$
Effective radius	R_e	Uniform in $\log(R_e)$	$\log(R_e/\text{pixels}) \in U[0.01, 10]$
Sérsic index	n	Uniform in n	$n \in U[0.4, 8]$
Central intensity	I_0	Uniform in $\log(I_0)$	$\log(I_0/\text{counts}) \in U[0.01, 10]$
Scale length	h	Uniform in $\log(h)$	$\log(h/\text{pixels}) \in U[0.01, 10]$
Axial ratio	q	Uniform in q	$q \in U[0.2, 1]$
Position angle	PA	Uniform in PA	$\text{PA}^\circ \in U[-360, 360]$
Central coordinates	x_0, y_0	Uniform in x_0 and y_0	x_0 and $y_0 \in U[0, S_{\text{image}}]$
<i>Combined parameters</i>			
Effective radius/scale length	R_e/h	Uniform in R_e/h	$R_e/h \in U[0.05, 1.678]$
Bulge-to-total flux ratio	B/T	Uniform in B/T	$B/T \in U[0.01, 1]$
Bulge-to-disc flux ratio ($< R_e$)	$B/D(< R_e)$	Uniform in $B/D(< R_e)$	$B/D(< R_e) \in U[1, -]$
Number of crossing points	N_x	$\delta(N_x) = 1$ for $N_x = 1$	

components of galaxies. A detailed review of the Sérsic profile and associated quantities is given by Graham & Driver (2005).

We characterize the intensity profile of a galaxy by concentric elliptical isophotes with position angle θ_{PA} in degrees counterclockwise from the vertical axis of the image, and ellipticity $\epsilon = 1 - q$, where $q = b/a$ is the ratio between the semiminor and semimajor axis of the ellipse. The projected radius is given by

$$r = x_p^2 + \frac{y_p^2}{q^2}, \quad (3)$$

where x_p and y_p are coordinates in the reference frame centred on the image centre (x_0, y_0) and rotated to the position angle relative to the image x -axis ($\text{PA} = \theta_{\text{PA}} + 90^\circ$),

$$\begin{aligned} x_p &= (x - x_0) \cos(\text{PA}) + (y - y_0) \sin(\text{PA}), \\ y_p &= -(x - x_0) \sin(\text{PA}) + (y - y_0) \cos(\text{PA}). \end{aligned} \quad (4)$$

Many previous studies have shown that a careful analysis of the PSF is needed to perform robust photometric decompositions. Méndez-Abreu et al. (2008) found that errors of ~ 2 per cent in the PSF full width at half-maximum (FWHM) led to errors of up to 10 per cent in the R_e and n of the bulge. Gadotti (2009) found that to reliably retrieve the structural properties of bulges, the effective radius must be larger than ~ 80 per cent of the PSF FWHM. It is therefore crucial that an accurate model for the PSF is provided to `PHI`, and that `PHI` then treats the PSF correctly. Given the large number of models that must be built for comparison with the data during the running of `PHI`, we explored a variety of methods for convolving the model image with the PSF. We found that the FFT returns the required accuracy for the decomposition of images.

2.2 Bayesian framework

Bayesian methods combine prior knowledge about a model with data to obtain a probabilistic description of the model. This is described by Bayes' theorem,

$$p(\theta|D) = \frac{L(D|\theta)p(\theta)}{\int L(D|\theta)p(\theta)d\theta}, \quad (5)$$

where the model is characterized by the parameter vector θ , $p(\theta|D)$ is the posterior probability of a set of parameters θ given the data D , $L(D|\theta)$ is the likelihood function or the probability of the data given θ , and $p(\theta)$ is the prior probability of the parameter vector θ .

The denominator ensures that the probability is unity when summed over all possible models.

2.2.1 Likelihood

For a large number of photons detected in each independent CCD pixel, the measurement errors can be considered to be Gaussian with mean zero and no covariance. The likelihood for a given pixel i is then given by

$$p(d_i|\theta) = \frac{1}{(2\pi\sigma_i^2)^{1/2}} \exp\left\{-\frac{1}{2}\left[\frac{d_i - f(x_i;\theta)}{\sigma_i}\right]^2\right\}, \quad (6)$$

where $f(x_i;\theta)$ is the model function consisting of known quantities x (i.e. constants, control variables, etc.) and the unknown parameters θ . Since each pixel is considered to be an independent measurement, the combined $-2 \ln$ likelihood of all the N pixels is

$$-2 \ln p(D|\theta) = \chi^2 + \sum_{i=1}^N \ln \sigma_i^2 + N \ln 2\pi, \quad (7)$$

where

$$\chi^2 = \sum_{i=1}^N \frac{[d_i - f(x_i;\theta)]^2}{\sigma_i^2}. \quad (8)$$

2.2.2 Priors

The prior distribution expresses our knowledge and prejudices about certain values of the parameters, or relationships between them. Where we know some values to be more probable than others, a carefully selected prior distribution can encode this knowledge. However, biases may arise if the prior distribution is *informative*, e.g. a Gaussian probability distribution function (PDF) with a narrow width. For the purpose of galaxy image analysis, where the parameter space is known to have many local minima, it is only advisable to use an informative prior where there is a clear justification. `PHI` will accept any valid function for the prior distributions defined by the user. In this paper we demonstrate the use of `PHI` with maximally *uninformative* priors, i.e. a uniform distribution between certain parameter limits.

The top section of Table 1 lists the prior functions used in this paper for each parameter of the model in Section 2.1. In the case of photometric decompositions of galaxy images there are some

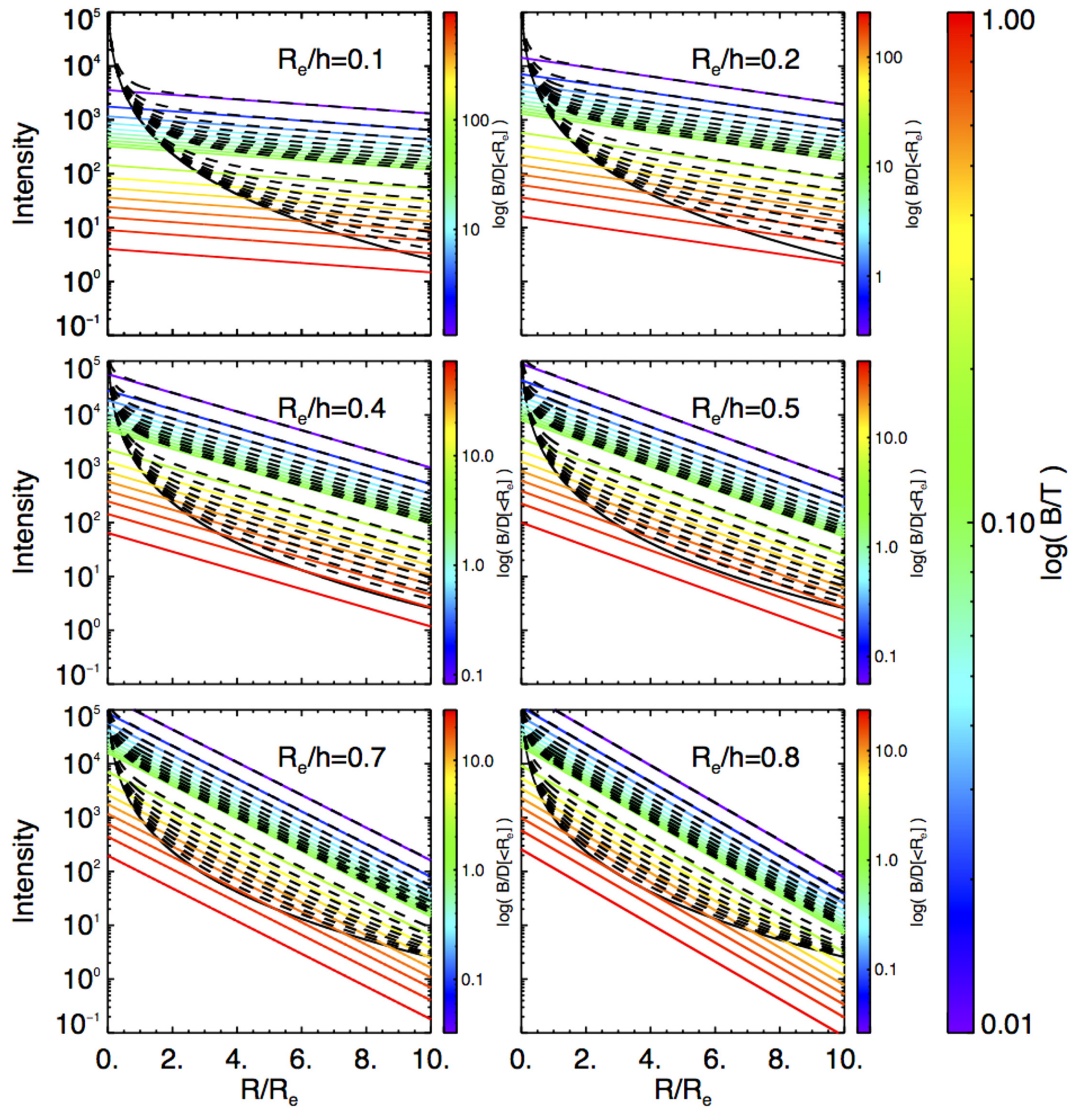


Figure 2. The intensity profiles for two-component models, where the radius has been normalized by the effective radius of the bulge. The solid black line shows the bulge profile in the case of a Sérsic model with $n = 4$. The coloured lines show the exponential disc component, with different bulge-to-total flux ratios (B/T) as given in the large colour bar, or bulge-to-disc flux ratios (B/D) within one effective radius as shown by the individual colour bars. The black dashed lines show the combined (total) intensity profiles. Each panel shows models with different ratios of bulge effective radius to disc scale length (R_e/h) as given in the legend.

clearly physically motivated limits on the parameters. For example, negative or very large values for the radius are unphysical. We allow values for the Sérsic index, n , to take any value within the range $0.5 < n < 8$. Larger values produce unphysical concentrations of light in the centre.

One advantage of the Bayesian inference framework for the photometric decomposition of *two-component* galaxy images is that the accidental fitting of a one-component model, or the reversal of bulge and disc components, can be explicitly avoided through the use of combined priors. To illustrate the use of combined priors, we restrict our models to fit a bulge defined to be an excess of light over the inner extrapolation of an exponential disc. We note that this particular model may not be appropriate for all science goals, for example in the study of embedded discs. To do this we must constrain our two-component galaxies to be a combination of an *inner* Sérsic profile (the bulge) and *outer* exponential profile (the disc). Unfortunately, this is less straightforward than it sounds: for a bulge

profile with Sérsic parameter $n > 1$ and an exponential disc profile, at some (large) radius the inner component will again dominate over the outer. To explore where the reversal of components could occur in practice, Fig. 2 shows the effect on the intensity profiles of varying each of the model parameters presented in Section 2.1, for a bulge with $n = 4$. The bulge component dominates at large radii when the ratio of bulge effective radius to disc scale length (R_e/h) is large and the bulge-to-total flux ratio (B/T) is large. In this example, we employ combined priors to enforce our prejudice about the relative strength and positioning of the two components in three ways: (i) preventing the disc effective radius ($R_{\text{disc}} = 1.678h$) becoming smaller than the bulge R_e ; (ii) ensuring the bulge-to-disc flux ratio (B/D) within one effective bulge radius is always larger than unity; and (iii) ensuring that the number of crossing points in the bulge and disc 1D light profiles (N_c) is unity. This third prior ensures that two components are fit, yet prevents the light profile of the bulge becoming dominant in the outer edges. It is implemented




Figure 3. The MCMC trace plots for *some* of the profile parameters in a two-component model. The left-hand column shows the bulge (Sérsic) profile parameters: the effective intensity (I_e), effective radius (R_e), and Sérsic index (n). The right-hand column shows the disc (exponential) profile parameters: the central intensity (I_0) and scale length (h). The red, green, and blue correspond to the three levels of the algorithm. The orange shows the transition between first and second levels, and the light blue shows the burn-in part of the third level that is discarded.

via a Newton–Raphson algorithm, which is run until the total model intensity falls below the mean of the sky background; after this a second crossing point can occur (and is inevitable for galaxies with $n > 1$). These combined priors are summarized in the second half of Table 1, where we again use maximally uninformative priors for simplicity. More complex priors may be included trivially by the user, depending on their science goals.

2.3 Model comparison

The aim of model fitting is to construct probabilistic models that represent, or sufficiently approximate, the data. Once the simplest model has been fit, we can increase the complexity of the model by adding extra parameters. It is important to then test whether each additional parameter is justified on the grounds of a significantly improved fit, given the increased number of degrees of freedom.

In this paper we investigate the use of the Bayesian information criterion (BIC; Schwarz 1978), which compares the maximum likelihood of each model $L(D|\theta_{ML})$,

$$\text{BIC} = -2 \ln(p(D|\theta_{ML})) + m \ln(N), \quad (9)$$

where N is the number of data points and m is the number of free parameters in the model (θ_{ML} is the corresponding maximum

likelihood parameter vector). The difference between the BIC of two models (A and B) is

$$\Delta \text{BIC} = -2 \ln \left(\frac{p(D|\theta_A)}{p(D|\theta_B)} \right) + (m_A - m_B) \ln(N). \quad (10)$$

The actual calculation of the BIC requires the Bayes factor, which is not provided by the PH1 algorithm. We therefore approximate ΔBIC by fitting both models to the data and taking the posterior medians for each set of fitted parameters. Unfortunately, the appropriate demarcation to distinguish between the models is somewhat dependent on the problem in hand. We therefore carry out simulations in Section 5 to determine the appropriate values for our data set.

2.4 The MCMC engine

Exact Bayesian analysis is restricted by the need to perform integrations analytically. For a simple model with two or three parameters, one can obtain a good estimate of the posterior probability by exploring all possible parameter values on a grid. For a high-dimensional parameter space, such as the models used to describe the surface brightness distribution of galaxies, the characterization of the posterior distribution becomes increasingly difficult, and the use of a grid is prohibitively time consuming. Sampling-based methods allow the exploration of highly multidimensional and complex parameter spaces. One example of these methods is MCMC: by generating repeated states by a first-order Markovian process, MCMC asymptotically converges to the posterior distribution.

The purpose-built MCMC algorithm used in PH1 consists of three levels that aim to achieve an efficient convergence and accurate estimation of the posterior distribution. Fig. 3 shows a typical run of the entire algorithm with different colours depicting the three levels and transitions between them. In the following subsections we will address the intricacies of each level individually.

2.4.1 Level one: blocked adaptive Metropolis

PH1 begins with a variation on the adaptive-Metropolis-within-Gibbs algorithm introduced in Roberts & Rosenthal (2009). The purpose of this level is to obtain an estimate of the scale of each parameter in the Markov chain. By knowing how each parameter scales PH1 can efficiently sample from the parameter space, which overcomes problems with poorly chosen initial parameter values, as well as complex probability distributions with many local minima. Given a current value in the Markov chain, X_i , a new value or set of values Y is proposed, where i denotes the i th step in the Markov chain. The new values Y are either accepted as a valid move so the next starting location is $X_{i+1} = Y$, or are rejected and $X_{i+1} = X_i$, according to the criteria,

$$\begin{cases} X_{i+1} = Y & \text{if } U < \min[1, \pi(Y)/\pi(X_i)], \\ X_{i+1} = X_i & \text{if } U \geq \min[1, \pi(Y)/\pi(X_i)], \end{cases} \quad (11)$$

where U is a uniformly chosen random number $U \sim U(0, 1)$ and $\pi(\cdot)$ is the target distribution (i.e. the combination of likelihood and prior distributions in equation 5). This process is repeated for every parameter sequentially.

The proposed new parameters are drawn from a Gaussian function $Y \sim N(X_{i,j}, \sigma_{i,j}^2)$, where $\sigma_{i,j}$ represents the size of the step the algorithm makes when choosing the proposed values at each iteration i and for each parameter j . The correct value for $\sigma_{i,j}$ will provide a compromise between being able to jump from one region of parameter space to another quickly, and being able to explore in detail the target distribution. If $\sigma_{i,j}$ is too large we see a drop in the




Figure 4. Example marginal posterior distributions for a two-component synthetic SDSS *i*-band galaxy image. The marginal distribution for each of the structural parameters is shown on the diagonal. Joint marginal distributions for pairs of parameters are shown in the off-diagonal panels. The three colour contours represent the 68, 95, and 99 per cent confidence levels. The solid grey line shows the true value of this synthetic galaxy, and the dashed line indicates the median of the posterior distribution.

acceptance rate, as we are drawing from a region of parameter space with low probability. For a $\sigma_{i,j}$ that is too small we will accept values at almost every iteration. The optimal acceptance rate is 0.44 for a one-dimensional Markov chain and 0.23 for dimensions greater than one (Roberts & Rosenthal 2001). The adaptive-Metropolis-within-Gibbs algorithm uses information from past iterations to adapt $\sigma_{i,j}$ until the desired acceptance rate is achieved.

To accomplish diminishing adaptation we initially calculate the average acceptance rate for the past n_{step} (default $n_{\text{step}} = 100$) of iterations and allow an update to the $\sigma_{i,j}$ by adding or subtracting 5 per cent of $\sigma_{i,j}$ to adjust the acceptance rate closer to the optimal value (Roberts & Rosenthal 2001). Once the acceptance rate falls within 0.15 and 0.32, the average acceptance rate is calculated over the last $2n_{\text{step}}$ until every parameter (or blocked set of parameters) again has an average acceptance rate within 0.15 and 0.32. We then monitor the acceptance rate for a further $4n_{\text{step}}$ iterations, and adjust

$\sigma_{i,j}$ until the acceptance rate falls between 0.15 and 0.32. At that point, adaptation is stopped and the final $\sigma_{i,j}$ values are saved.

It is important that the final Markov chain closely matches the target distribution, so the chain is continued without any further adaptation of the σ_j until the chain's gradient tends to zero. This is done by calculating the average parameter value after every 200 iterations, and once five averages are obtained a line is fit and the gradient of this line determined. When this gradient is close to zero the chain is converging to the target distribution and the algorithm can move onto Level two.

2.4.2 Level two: adaptive Metropolis

The aim of this level is to obtain a similar covariance structure for the proposal distribution (Y) to that of the target distribution, which leads to greater success rates for the proposal distribution (Haario,

Table 2. Ranges for structural and geometric parameters used to simulate synthetic elliptical and bulge+disc galaxies. Geometric parameters are independently determined for each component.

Parameter	[min, max]
<i>Synthetic elliptical</i>	
Sérsic magnitude ($m_{S, i}$)	[14, 17]
Effective radius (R_e in arcsec)	[1.5, 6]
Sérsic index (n)	[1.9, 7.5]
<i>Synthetic bulge+disc</i>	
Sérsic magnitude ($m_{S, i}$)	[15, 21]
Effective radius (R_e in arcsec)	[0.4, 2.24]
Sérsic index (n)	[0.5, 7]
Exponential magnitude ($m_{E, i}$)	[15, 18]
Scale length (h in arcsec)	[1.3, 7]
<i>Geometric parameters of bulge or disc</i>	
Axial ratio (q)	[0.6, 1]
Position angle (PA in $^\circ$)	[-360, 360]




Figure 5. The posterior fractional error distribution for the entire ensemble of synthetic elliptical galaxies. The marginalized error distributions for each structural parameter are shown on the diagonal. Joint marginal error distributions are shown in the off-diagonal panels. The colour contours represent the 68, 95, and 99 per cent confidence levels. The solid grey lines indicate the true values and the dashed lines represent the median of the posterior error distributions. A positive (negative) value indicates the fitted value is larger (smaller) than the true value.

Saksman & Tamminen 2001; Roberts & Rosenthal 2009). Y is drawn in a similar way to before: $Y \sim N(X_i, c\Sigma_i)$, again where X_i is the current state of the chain, and the same accept/reject Metropolis rule is used as in Level one. Σ_n is the covariance matrix of all the previously generated values of the chain since the adaptation of Level one finished, and c is a constant that is included to yield an optimal acceptance rate: $c = 2.382^2/d$, where d is the number of parameter dimensions (see Haario et al. 2001; Roberts & Rosenthal 2009).

To establish if further adaptation will improve the chain the algorithm tests that the covariance structure of the target distribution has been correctly identified. This can be determined directly from

the past iterations of the chain. Every N_{L2} (user input) iterations the mean squared difference between each successive iteration ($X_{i-1} - X_i$)² for each parameter is calculated, and after $5N_{L2}$ a linear model is fit. If the gradient of the mean squared differences appears to have an increasing or decreasing gradient, then the algorithm continues to adapt; if the gradient is close to zero, then adaptation stops and the algorithm moves to the final level.

2.4.3 Level three and chain convergence

The final level of PHI involves a symmetric random walk Metropolis algorithm drawing the proposed values from $Y \sim N(X_i, c\Sigma_{L2})$, where Σ_{L2} is the last covariance matrix calculated before adaptation stopped in Level 2 and c is again the constant to help achieve the target acceptance rate. In PHI the default method to test for convergence is to run multiple chains simultaneously, and then to use a Gelman–Rubin diagnostic (Gelman & Rubin 1992). Alternatively, a Geweke diagnostic (Geweke 1992) can be used to determine whether a single Markov chain has converged. Once the Markov chains have converged, the chains are combined to form the final sample distribution that will be used in the analysis stage.

2.5 Run time

In a typical run the IDL version of PHI requires between 10^4 and $10^{4.5}$ iterations for three simultaneously running chains to converge. The median total generation time for a 250×250 pixels image is $t_{\text{total}} = 0.029$ s for a single Sérsic model and $t_{\text{total}} = 0.041$ s for a Sérsic+exponential model. The wall clock time for a complete run on a 2.5 GHz Intel core i5 CPU is ~ 10 and ~ 20 min for a single Sérsic and a Sérsic+exponential model, respectively. Run times are similar for real and mock galaxies. The quoted times are for convolution with a 50×50 pixels PSF.

3 APPLICATION TO SYNTHETIC GALAXY IMAGES

In this section, we use synthetic galaxy images to test the accuracy and robustness of PHI. Synthetic galaxies lack the complexity present in real galaxies, but allow us to check for any systematic errors inherent in the method. The images were made to mimic SDSS (Strauss et al. 2002) i -band images as closely as possible, with a pixel scale of 0.396 arcsec pixel⁻¹, CCD gain of $4.86e$ /ADU and read-out noise of $5.76e$. The appropriate noise level was estimated from the SDSS data frames by removing all objects and fitting a Poisson distribution to the residual counts.

Fig. 4 shows the posterior distribution produced by PHI for fitting an example two-component synthetic galaxy image with an inner Sérsic and outer exponential component as described in Section 2.1 and with priors as described in Section 2.2.2. The medians of the posterior distributions have a fractional error of at most 2 per cent in relation to the true values. We also clearly see that there is a strong covariance between parameters within the individual components, i.e. I_e , R_e , and n for the Sérsic profile and I_0 and h for the exponential component. There is also some covariance between the two components, i.e. n versus h . The quantification of these covariances is important, as it may cause correlations between physical parameters (e.g. scaling relations) to appear stronger than they are in reality. The entire posterior distribution for a galaxy can be used when testing hypotheses about galaxy populations and this will be explored in a subsequent paper.




Figure 6. Same as Fig. 5, but for the synthetic two-component bulge+disc galaxies.

Table 3. The 16th, 50th, and 84th percentiles of the posterior fractional error distributions for the synthetic elliptical and bulge+disc galaxies.

Parameter	Median	16%	84%
<i>Synthetic elliptical</i>			
$\Delta I_e/I_{e, \text{in}}$	0.003	−0.022	0.049
$\Delta R_e/R_{e, \text{in}}$	0.004	−0.021	0.018
$\Delta n/n_{\text{in}}$	−0.004	−0.033	0.012
<i>Synthetic bulge+disc</i>			
$\Delta I_e/I_{\text{in}}$	0.005	−0.25	0.22
$\Delta R_e/R_{e, \text{in}}$	−0.001	−0.15	0.14
$\Delta n/n_{\text{in}}$	−0.014	−0.19	0.05
$\Delta I_0/I_{0, \text{in}}$	−0.002	−0.09	0.04
$\Delta h/h_{\text{in}}$	0.0001	−0.02	0.03

3.1 Population of synthetic galaxies

To test PH1 on a realistic range of galaxy structural parameters, we take the parameters of 260 elliptical and 380 bulge+disc galaxies in the SDSS fitted by Gadotti (2009) with single Sérsic and

Sérsic+exponential profiles, respectively. Further details about this sample are given in Section 4. Table 2 summarizes the range of parameters tested in this way. In particular, we test a large range in the bulge-to-total flux ratio of bulge+disc galaxies with $0.01 < B/T < 0.8$. For each synthetic galaxy we also assume a Gaussian PSF with FWHM as provided by Gadotti (2009). According to this paper, these values were largely taken from the SDSS Data Release 2 (DR2) imaging headers, although erroneous values were corrected for by fitting nearby stars. The mean FWHM of the sample is 1.5 arcsec, with a range between 0.4 and 2.6 arcsec.

3.2 Fractional errors on structural parameters

To visualize the errors over a large sample of galaxies, we calculate the fractional error distribution on each parameter as $(x_{\text{out}} - x_{\text{in}})/x_{\text{in}}$, where x_{in} is the input parameter value and x_{out} are the values given in the MCMC output. The stacked posterior fractional error distributions of all the synthetic galaxies are shown in Fig. 5 for elliptical galaxies and in Fig. 6 for bulge+disc galaxies. The clear covariances between the model parameters seen in Fig. 4 lead to a high

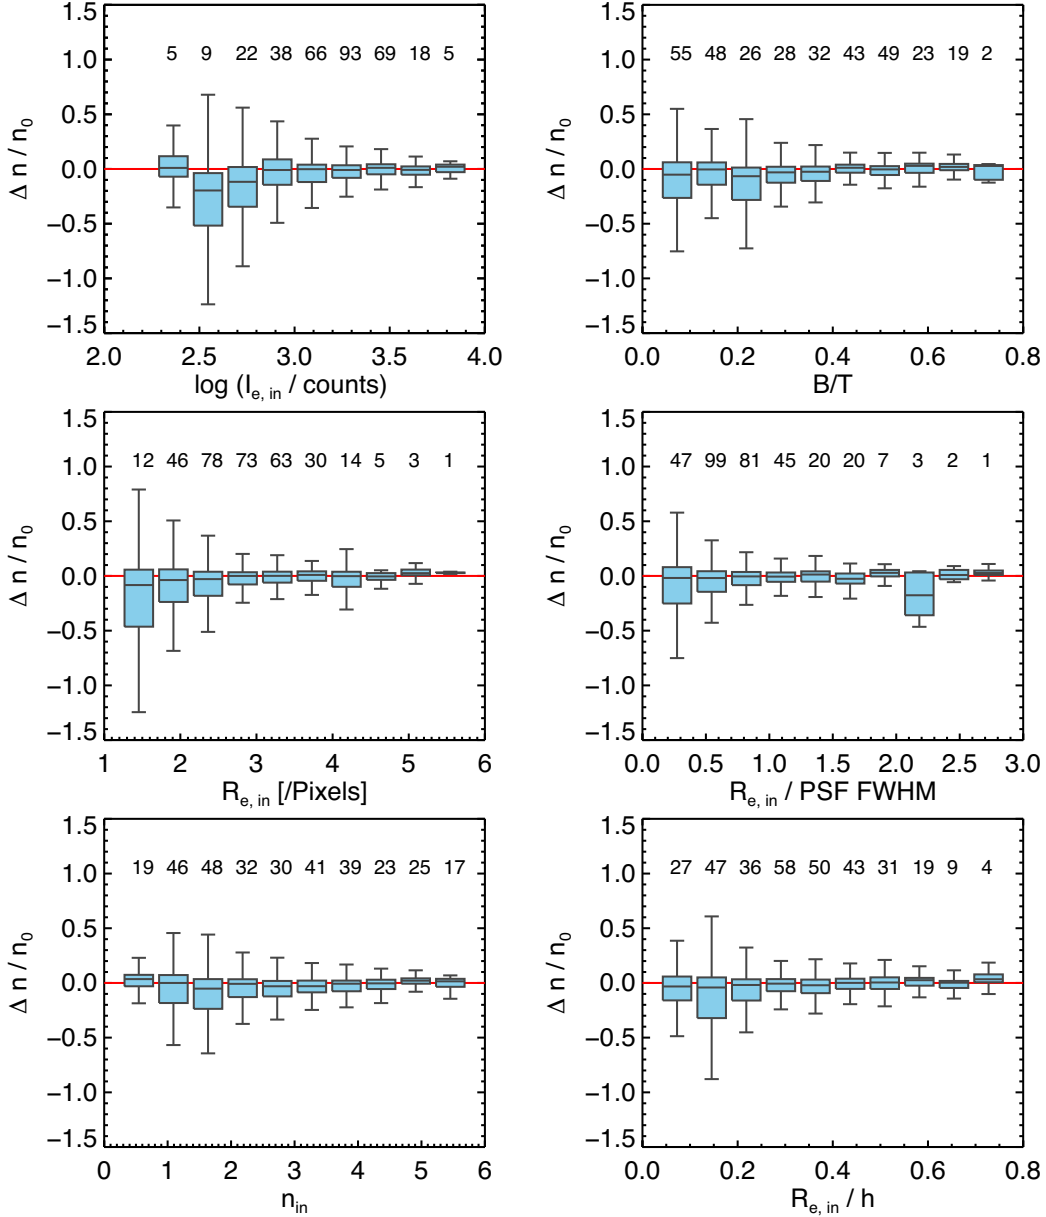


Figure 7. Box plots representing the fractional error on n for bulge+disc synthetic galaxies as a function of input parameters: I_e , R_e , n , B/T , R_e/h , as well as the ratio between input R_e and the PSF FWHM. The box limits represent the 16th and the 84th percentiles of the posterior fractional error distributions and the horizontal line shows the 50th percentile. The whiskers show the total extent of the distributions. The value above each box gives the number of galaxies in each bin.

degree of correlation in the fractional errors. These figures clearly show the well-known degeneracies inherent in Sérsic fits: (i) if the effective intensity is overestimated (underestimated), the effective radius will be underestimated (overestimated) to compensate; (2) if the effective radius is overestimated (underestimated), the Sérsic index will also be overestimated (underestimated) to increase the concentration of intensity within a now larger effective radius.

Table 3 presents the statistics of the posterior fractional error distributions. It can be seen that the systematic bias on parameters is minimal for both elliptical and bulge+disc galaxies, and the 1σ errors (16th and 84th percentiles) are usually <20 per cent. In the following subsections we look in more detail at the impact of degeneracies between parameters on obtaining unbiased estimates for the structural parameters and B/T .

3.2.1 Impact on bulge+disc structural parameters

Fig. 7 shows the fractional error on n as a function of a selection of structural parameters for the two-component bulge+disc synthetic galaxies, as well as the ratio between R_e and the PSF FWHM. The blue box indicates the 16th and 84th percentiles of the marginalized posterior fractional error distribution, with the horizontal line showing the median, and the whiskers indicating the maximum extent of the posteriors. As we saw in Fig 6 the fractional errors are typically <20 per cent and the median shows no overall bias. However, this figure shows how the errors on n increase for models with smaller I_e , smaller B/T , smaller R_e , smaller ratio of R_e to PSF FWHM, and smaller R_e/h . The errors appear to be largest for $n \sim 2$ bulges.

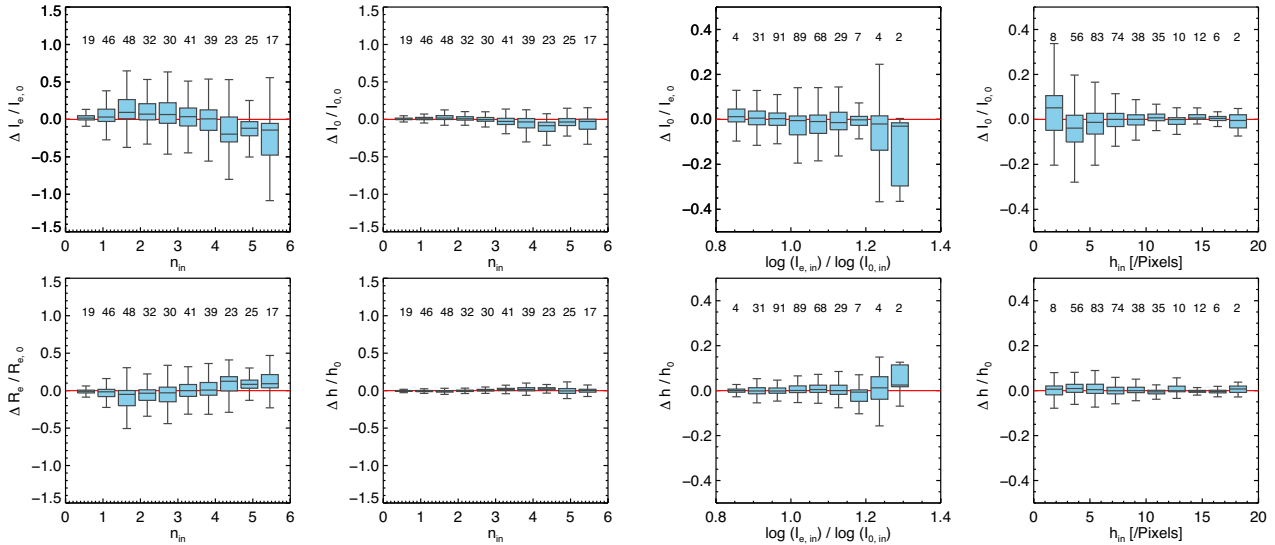


Figure 8. The fractional errors on I_e , I_0 , R_e , and h as a function of input n (left) and I_0 and h as a function of I_e/I_0 and h (right). See the caption of Fig. 7 for more details.

Similarly, the two left-hand panels of Fig. 8 show the fractional error on the other structural parameters as a function of input n . As before we see that there is typically no overall bias in the estimation of parameters, with the median sitting close to zero. The exception is at large n where a bias of ~ 12 and ~ 10 per cent is found on R_e and I_e , respectively, for $n = 6$. For high values of Sérsic index the algorithm also takes longer to converge, suggesting difficulties in this region. Further investigation suggests these difficulties are caused by a flattening in the likelihood space for galaxies with higher n values ($n \geq 4.5$). Subsequent changes in the $n > 4.5$ region therefore result in little variation in the likelihood values, thus decreasing the accuracy and precision. We explore this specific case of high- n bulges in more detail in Appendix A, finding that they are particularly susceptible to strong degeneracies between R_e and n when B/T is low.

3.2.2 Impact on B/T

Galaxy morphology is commonly parametrized by the bulge-to-total flux ratio (B/T), it is therefore important to assess the biases on this particular parameter caused by correlations in the parameter errors observed in Fig. 6. Fig. 9 shows the fractional error on the measured B/T as a function of input parameters I_e , B/T , R_e , I_0 , n , and h . The median values are remarkably stable for the whole range of models tested here, indicating that B/T should provide a robust parametrization of galaxy morphology.

3.2.3 Impact on the disc parameters

The two right-hand panels of Fig. 8 show the fractional errors on the disc structural parameters I_0 and h , as a function of input I_e/I_0 and h . Additionally, the centre left-hand column of Fig. 8 shows how the Sérsic index affects for the same parameters. The fractional errors on the disc scale length h are the smallest of all structural parameters, this is because the disc is well resolved in the images. There is no clear bias in the measured disc parameters, although as expected the fractional errors increase where the disc becomes less dominant (as traced by the I_e/I_0 ratio) and the disc scale length becomes smaller.

4 APPLICATION TO DATA

The next step is to apply PH1 to a sample of real galaxy images to assess the functionality and robustness of the method in a fully realistic scenario. It is first useful to visualize the fits from PH1. To do this, we collapse the 2D data and model images into 1D using ellipse averaging. In order to represent the full posterior uncertainty on the fitted parameters, and therefore the profiles, we draw randomly from the parameter posterior distributions to estimate the median surface brightness and 1σ errors as a function of radius.

Fig. 10 shows the 1D surface brightness profile of a bulge+disc galaxy in the Gadotti (2009) sample to which we have fitted both a one-component Sérsic model (left-hand panel) and a two-component Sérsic+exponential model (right-hand panel). The black dots and grey region show the image data and root-mean-square of the flux values in a given ellipse. The green band shows the ensemble of model fits generated from drawing parameters randomly from the posterior distribution and convolving the resulting model with the PSF. Residuals from the model generated from the median posterior parameters are shown in the lower panels. In the case of the one-component model, the central region is well fitted, but the model deviates significantly from the data beyond 3 arcsec. This is improved by the two-component model, with an acceptable fit out to 8 arcsec, although the residuals beyond this radius suggest that this galaxy has a truncated disc.

Fig. 11 shows the 2D surface brightness profile of the same bulge+disc galaxy presented in Fig. 10. The top row shows the image, the fitted one-component Sérsic model, and the residual. The middle and bottom rows show an exponential-only and the Sérsic+exponential models, respectively, with their residuals. The exponential-only model is a one-component Sérsic model with fixed $n = 1$; this model is not used in this paper, but presented here to aid understanding. The models were generated from the posterior medians from the MCMC outputs. As in Fig. 10 we see that the one-component Sérsic model fits the central region well, which cannot be fit with a pure $n = 1$ profile (middle panel). The two-component model provides the lowest residuals in both the central and outer regions, leaving a clear signal from the spiral arms that are not included in the model.

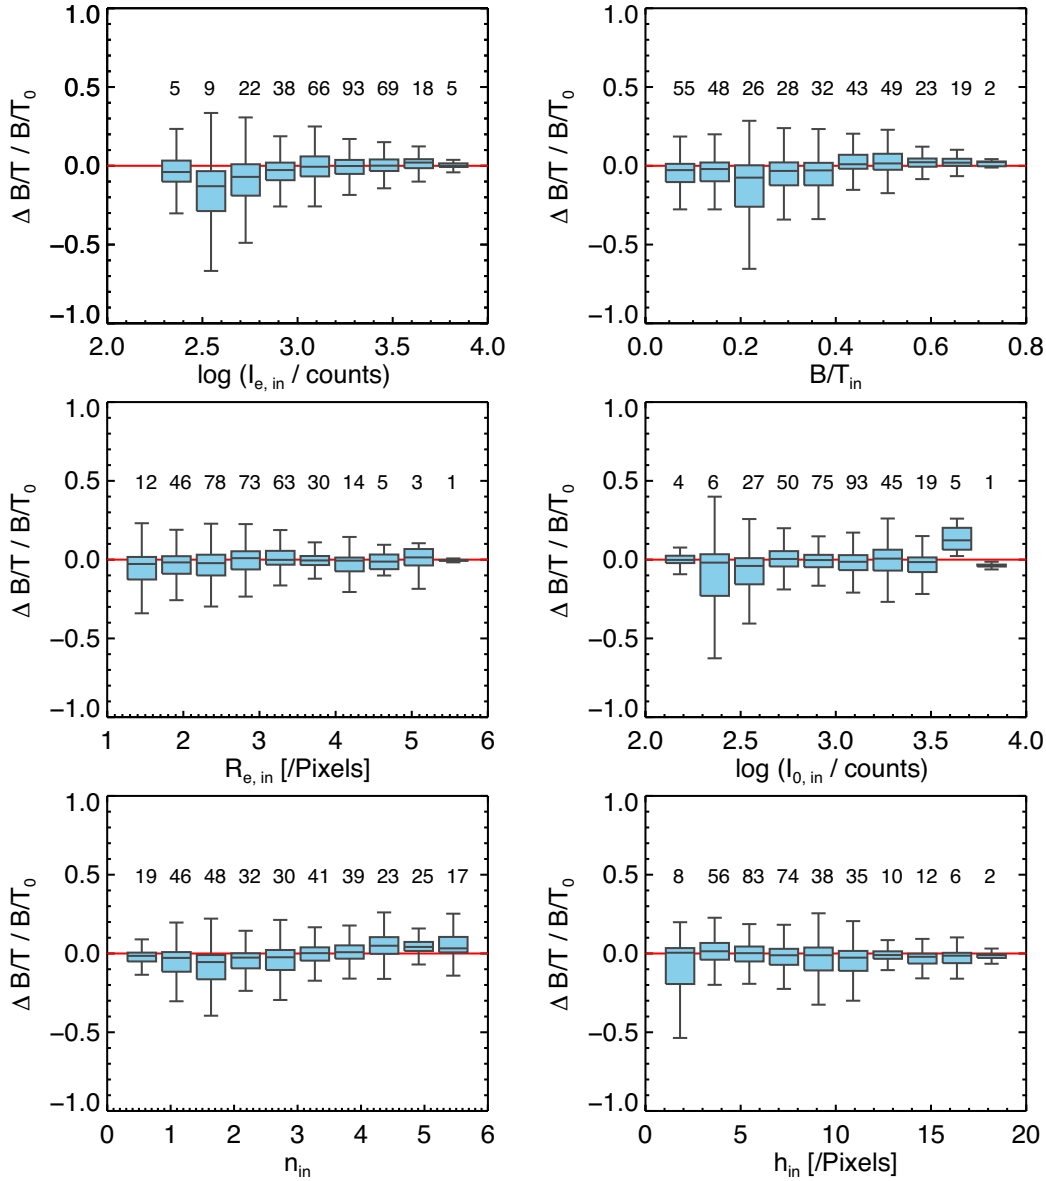


Figure 9. Same as Fig. 7, but showing the fractional errors on B/T as a function of input I_e , B/T , R_e , I_0 , n , and h .

Although the graphical and visual representation are useful tools to ensure the code is working as it should, the quantitative model comparison presented below in Section 5 is required in order to make statistical claims about which model is best.

4.1 Analysis of SDSS images

We study SDSS i -band images of galaxies with stellar masses $>10^{10}M_{\odot}$, $0.04 \leq z \leq 0.06$, and $q \geq 0.9$ that were previously analysed by Gadotti (2009) using the BUDDA code (de Souza et al. 2004) to perform bulge/disc/bar 2D photometric decompositions. Gadotti (2009) selected 1000 galaxies and separated them into elliptical and bulge+disc based on the i -band Petrosian concentration index, C , as given in the SDSS data base, defining ellipticals to have $C > 3$, disc galaxies to have $C < 2.5$, and bulge+disc galaxies to lie in between. Despite the constraints imposed on axis ratio and concentration parameter, the sample is considered to be a fair representation of the galaxy population in the local Universe. We select

only those galaxies that were classified as elliptical or bulge+disc, and remove from the sample any galaxies that were found to have a bar (visually identified from residual fits). As noted above, for this code presentation paper we have chosen to focus on one- and two-component galaxies, although the code is able to fit any model specified by the user. Barred galaxies will be studied with PH1 in a future publication. This leaves us with 260 elliptical and 380 bulge+disc galaxies.

In Gadotti (2009) the imaging used to classify the sample and perform the 2D photometric decompositions was from the SDSS DR2 (Abazajian et al. 2004). For this study we use the SDSS Data Release 7 (DR7; Abazajian et al. 2004) images. Moffat PSFs were obtained for each individual galaxy by fitting a Moffat function to 5–10 stars nearby to each galaxy and using the median value for each parameter. Segmentation maps were created following a similar approach to that used by SExtractor (Bertin & Arnouts 1996). There is no concern about overlapping sources as the galaxies were originally selected to be isolated. We took the gain, read noise, and




Figure 10. Ellipse-averaged radial profile of the surface brightness of SDSS galaxy J084149.16+504711.1 (black dots) with the root-mean-square error of the flux values in a given ellipse (grey region). The green band shows the synthetic galaxies generated from random draws from the posterior parameter distributions output by *PHI*, with the resulting models convolved with the PSF. The lower panels show the residuals between the data and the model created with the posterior median parameters. Left: for a one-component Sérsic model. Right: for a two-component Sérsic+exponential model. The blue and red bands show the random draws for the exponential and Sérsic components, respectively, without PSF convolution. The ΔBIC for this galaxy is 137, indicating that an improved fit is provided by the two-component model (see Section 5).

sky values from the SDSS image headers, and combined these with the galaxy shot noise to compute the weight maps in the standard way.

PHI was run on a square cut-out image, typically 250×250 pixels in size, following removal of unwanted sources using the segmentation map. The size of the cut-out was selected to include the entire galaxy, although slightly smaller cut-outs were used in a few cases where the galaxy fell close to the edge of the image. We ran *PHI* with three simultaneous chains for both one-component Sérsic and two-component Sérsic+exponential models on every galaxy present in both samples, in order to perform a model comparison in Section 5. We additionally analysed the same images using *GASP2D* (see Méndez-Abreu et al. 2014, 2017), with the same PSF, weight maps, and masks. Finally, we compared our results to those presented in Gadotti (2009) that use SDSS DR2 images, and primarily takes the PSF values from the SDSS image headers. *GASP2D* uses the LM and *BUDDA* uses the Nelder–Mead simplex minimization method so we can investigate whether any differences that arise are due to different codes or different images/treatment of the images.

4.2 Comparison of elliptical galaxies

Fig. 12 compares the measured Sérsic profile parameters for the galaxies classified as elliptical in Gadotti (2009). The blue circles show the difference between *PHI* and *GASP2D* and the red diamonds the difference between *PHI* and Gadotti (2009). We removed any catastrophic failures that occurred when running *GASP2D*, leaving 250 elliptical and 350 bulge+disc galaxies for the final sample.

We focus first on the comparison between *PHI* and *GASP2D*, where the same images are fit with the same PSFs, weight maps, and segmentation maps. The results are well correlated and agree on average. The standard deviations for the parameter differences are $\sigma_{\log(I_c)} = 0.10$ counts, $\sigma_{R_e} = 2.0$ pixels, and $\sigma_n = 0.49$. There is a subtle deviation from an exact one-to-one match at larger values of R_e and n . In *GASP2D*, a run is determined to have reached the global minimum when the deviations between the χ^2 of two consecutive iterations are lower than a given threshold. This threshold cut results in changing errors with parameter values: due to the exponential nature of the Sérsic profile, changes to n where $n \leq 2$ have a greater impact on the surface brightness profile than at larger values of n . This effect is not seen in *PHI* due to the efficient exploration of parameter space and adaptable step sizes.

When we compare the median posterior parameter values measured by *PHI* to the best-fitting values obtained by Gadotti (2009) we see significant differences for all the parameters. Most notably, the values fit by *PHI* span larger ranges than found by Gadotti (2009). For example, the distribution of n in Gadotti (2009) has a mean of 3.8 and standard deviation of 0.9, compared to a mean of 4.1 and standard deviation of 1.2 found by *PHI*. This difference in range results in the visible trend between the size of the offset and the value of the fitted parameter.

4.3 Comparison of bulge+disc galaxies

Fig. 13 compares the results between *PHI*, *GASP2D*, and Gadotti (2009) for the sample of bulge+disc galaxies. The results found by *GASP2D*

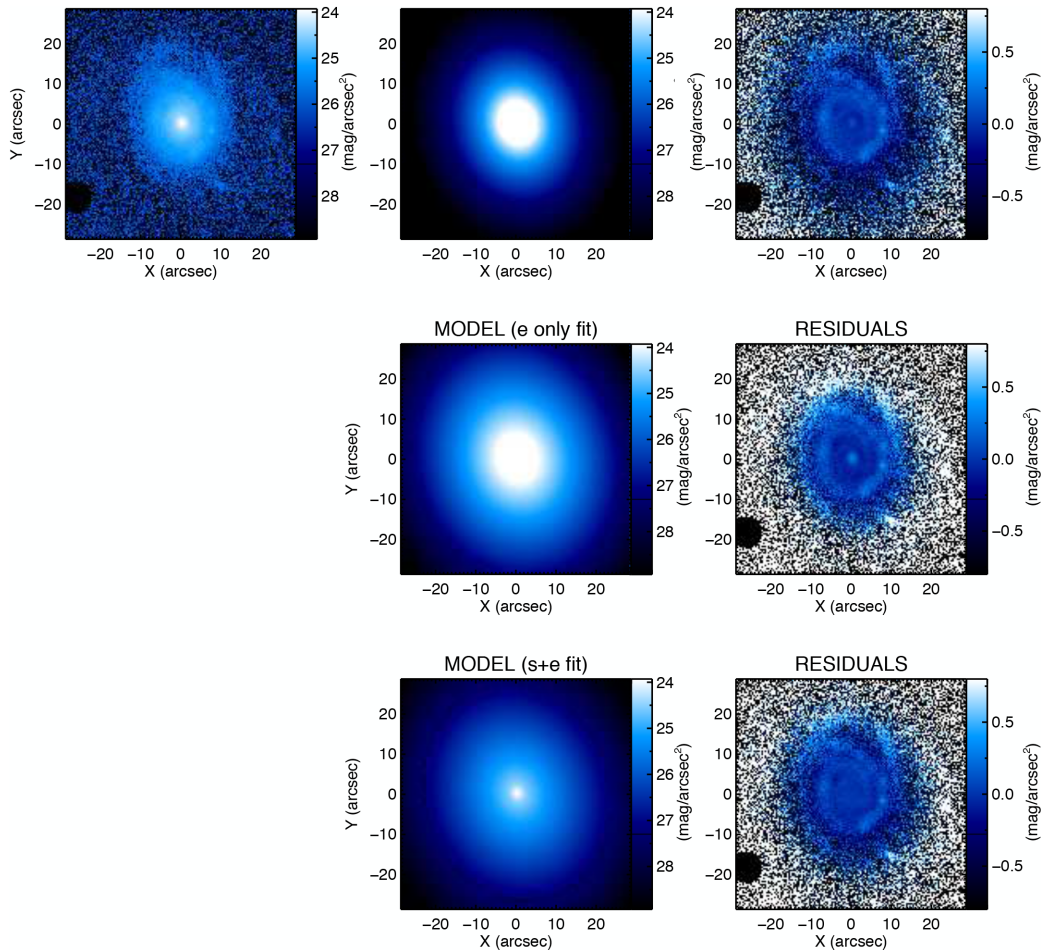


Figure 11. A 2D representation for the SDSS galaxy presented in Fig. 10. The top row shows the data (left), the Sérsic-only model fit (middle), and the residual (right). The second row shows an exponential only fit (i.e. with fixed $n = 1$) and corresponding residual. The bottom row shows the bulge–disc model fit and corresponding residual. All the models were made using the medians from the posterior distributions.

and PHI are broadly consistent, with no obvious bias. Standard deviations for the parameter differences are $\sigma_{\log(I_e)} = 0.4$ counts, $\sigma_{R_c} = 2.0$ pixels, $\sigma_n = 1.9$, $\sigma_{\log(I_0)} = 0.1$ counts, and $\sigma_h = 1.5$ pixels.

Comparing between the posterior median values found by PHI and the best-fitting values found by Gadotti (2009), we see that the bulge components are systematically different, as a function of parameter value. This is identical to the pattern seen for the elliptical galaxies and is caused by the much smaller range of parameter values that is fitted by Gadotti (2009) compared to PHI . On the other hand, the values found for the disc parameters are much more consistent, which agrees with our analysis of synthetic galaxies where disc parameters show much smaller fractional errors. It is clear from our analysis of both single Sérsic and Sérsic+exponential galaxies with BUDDA, GASP2D, and PHI that significant disagreements appear in the fitted bulge parameters when different code, images (DR2 versus DR7), weights, and PSFs are used. This highlights a fundamental limitation of bulge+disc decomposition, in that the estimation of bulge parameters will always be susceptible to biases and systematics when they are barely resolved in comparison to the PSF. We explore this issue with the synthetic galaxies in Section 3.2.1 and Appendix A. Using a fully Bayesian code such as PHI allows one to explore potential biases, errors, and covariances with ease, but ultimately a full code comparison study is clearly required to understand the limitations in more detail.

5 MODEL COMPARISON

In this section we test the use of the ΔBIC introduced in Section 2 to formally distinguish between one- and two-component galaxies. We define the $\Delta\text{BIC} = \text{BIC}_{1c} - \text{BIC}_{2c}$, so a larger value indicates that a two-component model is preferred. In simple and ideal situations a $\Delta\text{BIC} > 10$ is typically taken to be decisive evidence that a more complex model is preferred over a simpler one. However, in the case of image decomposition simulations are required to inform the choice of boundaries.

For every synthetic galaxy in Section 3 we have performed a one- and two-component fit; the distributions in ΔBIC are shown in the top panel of Fig. 14 with statistics of the distributions summarized in Table 4. The one-component/elliptical synthetic galaxies show a very tight ΔBIC distribution centred on zero, while the two-component/bulge+disc synthetic galaxies have larger ΔBIC values. The results from the synthetic galaxies suggest that a minimum ΔBIC value of 13.4 could be used to differentiate one- and two-component galaxies. This would incorrectly classify 16 per cent of one-component galaxies as two-component (i.e. contaminant level), but would identify 82 per cent of two-component galaxies correctly (i.e. high completeness).

The middle panel shows the ΔBIC distribution for the galaxies studied in Section 4 and classified by their concentration index

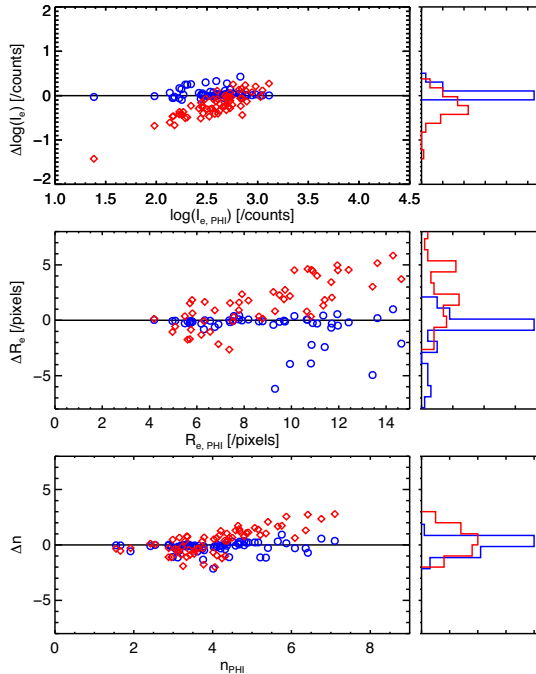


Figure 12. The differences between parameter estimates for elliptical galaxy images fit with a one-component Sérsic profile. From top to bottom the parameters are: the effective intensity (I_e), effective radius (R_e), and Sérsic index (n). The blue circles and red diamonds are the posterior medians given by PHI compared to the best fit from GASP2D and Gadotti (2009), respectively (i.e. $\text{PHI} - \text{GASP2D}$ and $\text{PHI} - \text{GAD09}$). The projected distributions are shown to the right, with the blue and red histogram comparing PHI to GASP2D and Gadotti (2009), respectively.

as either elliptical or bulge+disc. Table 4 shows that the low concentration index galaxies have a significantly higher median $\Delta\text{BIC} = 99.43$ than the high concentration index galaxies with $\Delta\text{BIC} = 2.78$. However, there is no clear differentiating line between the two samples. 22 per cent of galaxies with $C > 3$ are classified as two-component systems by this method, and 38 per cent of galaxies with $C < 2.5$ are classified as one-component. This shows that classifying galaxies by concentration index is not equivalent to classifying them by the results of 2D photometric bulge–disc decomposition.

Finally, we compare with a machine learning morphological classification method by Huertas-Company et al. (2011), based on support vector machines. They assign a probability to each galaxy that it is an elliptical, S0, SAB, or SCD. The algorithm was trained on visual classifications from the Galaxy Zoo first release catalogue (Lintott et al. 2008, 2011). We classify as elliptical any galaxy with $p(E) > 0.5$, and plot the distributions of ΔBIC in the lower panel of Fig. 14 for galaxies above and below this cut. We see a similar result to the case of classification by concentration index, with galaxies with higher ΔBIC more likely to have a disc, however, there is no clean demarcation between the two samples. 36 per cent of galaxies with $p(E) > 0.5$ are classified as two-component systems by their ΔBIC , and 38 per cent of galaxies with $p(E) < 0.5$ are classified as one-component. We note that this will include one-component disc galaxies, so their classification may in fact agree if we were to look in detail at the fitted parameters.

In the synthetic galaxies the ΔBIC can clearly be used as a classification method to separate one- and two-component galaxies. However, for real galaxies the lack of any significant demarcation between galaxies classified by other methods suggests that

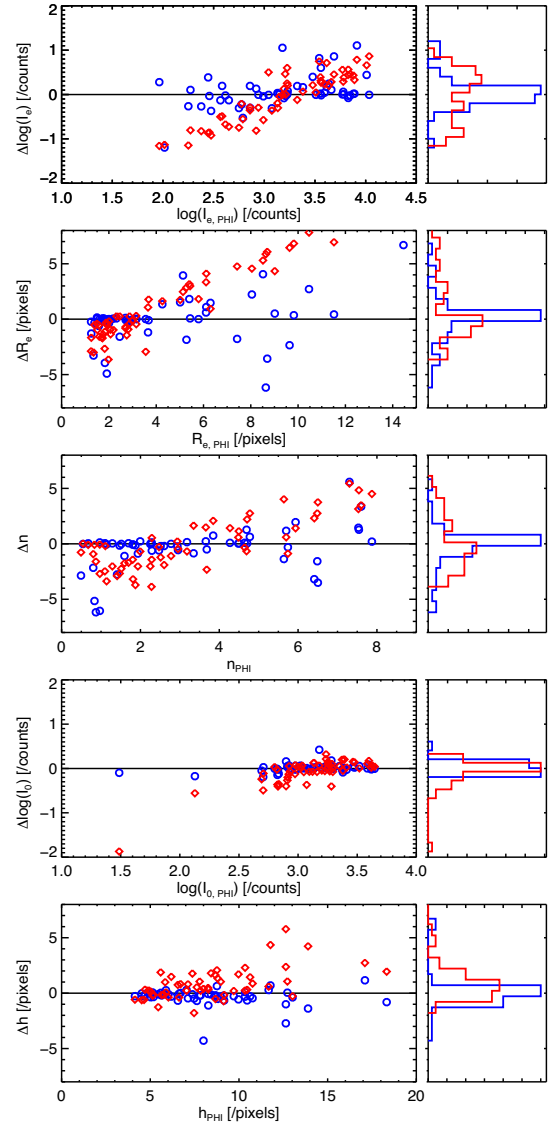


Figure 13. The same as Fig. 12, but for the sample of bulge+disc galaxies. From top to bottom we have the effective intensity (I_e), the effective radius (R_e), the Sérsic index (n), the central intensity (I_0), and the scale length (h).

the complex structure of real galaxies limits the usefulness of the ΔBIC approach, certainly for a simple bulge+disc model as studied here. While galaxies with higher values of ΔBIC will have a higher probability of having a disc, we advocate that the ΔBIC should be used in combination with other methods to determine the number of structural components in a galaxy.

6 SUMMARY

We have used a new fitting algorithm (PHI) to perform 2D photometric decompositions of galaxy images from a Bayesian perspective. PHI offers a number of significant advantages for estimating surface brightness profile parameters over traditional downhill optimization algorithms.

(i) PHI uses a triple layer approach to effectively and efficiently explore the complex parameter space. The first layer uses a blocked adaptive Metropolis algorithm to obtain an estimate of the scale for each parameter in the chain. The second layer uses an




Figure 14. Histograms and cumulative distributions showing the $\Delta\text{BIC} = \text{BIC}_{1c} - \text{BIC}_{2c}$ distributions. A positive value shows that the two-component model is preferred. Top: synthetic one- and two-component galaxies shown as red and blue lines, respectively. Middle: SDSS galaxies classified by their concentration index as elliptical (red line) or bulge+disc (blue line). Bottom: SDSS galaxies classified as elliptical (red line) or discy (blue line) by the machine learning algorithm of Huertas-Company et al. (2011).

adaptive Metropolis algorithm with the purpose of estimating the target covariance matrix. The final level uses this estimated covariance matrix to quickly and effectively explore the parameter space. This reduces the chances of local minima trapping.

(ii) The algorithm naturally and explicitly incorporates priors that force the parameters to be realistic and physical, e.g. positive in the case of the dimensions and intensities. These priors replace the need for filtering processes to remove non-physical parameter outcomes.

(iii) Priors on parameters can be combined to further strengthen the model in an explicit way. In this paper, to prevent the reversal of components (i.e. the desired inner component profile switching to

Table 4. The statistics of the ΔBIC distributions for the synthetic and SDSS samples. Where C is the i -band concentration parameter provided in the SDSS catalogue, and $p(E)$ is the probability that the galaxy is an elliptical from the machine learning classification of Huertas-Company et al. (2011).

	Mean	Median	16%	84%
<i>Synthetic ΔBIC</i>				
Elliptical	-427.47	-31.09	-149.74	13.36
Bulge+disc	1218.51	341.64	35.13	2272.50
<i>SDSS ΔBIC</i>				
$C > 3$	-89.70	2.78	-25.30	89.63
$C < 2.5$	225.67	99.43	-4.29	480.77
$p(E) > 0.5$	-16.88	18.82	-21.91	152.45
$p(E) < 0.5$	280.88	134.56	3.69	551.71

fit the outer and vice versa) we use a prior combination that specifies that the bulges of galaxies are better modelled by a Sérsic profile and the discs are described by an exponential profile. We do this via a Newton–Raphson algorithm to determine the crossing points in the total light profile, as well as ensuring that the bulge is the dominant component in the central region.

(iv) PH1 gives the full posterior probability distribution for a set of model parameters. This is a powerful description of the model uncertainties that can be used in further analyses of galaxy structures.

We used a sample of synthetic galaxies with SDSS-like image properties to ensure that there are no internal systematics due to the code, and to investigate the effect of galaxy properties on our ability to recover unbiased and well-constrained structural parameters. In bulge+disc galaxies we find that the bulge structural parameters are recovered less well than those of the disc, particularly when the bulge contributes a lower fraction to the luminosity, or is barely resolved with respect to the pixel scale or PSF. The only systematic biases occur in bulge+disc galaxies with high bulge Sérsic index ($n > 5$), where the code fits a bulge with an effective radius that is too large by 50 per cent and a central intensity that is too small by 20 per cent. No bias is found in the bulge-to-total luminosity ratio, which is important given the popularity of this quantity for parametrizing galaxy morphologies.

We have also applied PH1 to a sample of SDSS galaxies to compare with previous algorithms. Under the same image conditions, i.e. images with the same masks, weights, and PSF, PH1 achieves consistent results with a standard minimization code, with a low level of scatter. This validates both algorithms and approaches when assessing galaxy structures in the nearby Universe. However, we found less consistency when comparing to results from a previous analysis performed on different images, with different image conditions. The bulge structural parameters were the most affected, which we believe is due to the limited resolving power of SDSS images for local galaxy bulges.

Finally, we investigated the BIC as a method for deciding whether a galaxy has one or two components. In synthetic images the BIC cleanly separates the two types of galaxies, however, for real galaxies there was a less clean demarcation between galaxies classified as elliptical or bulge+disc by other methods in the literature. This suggests that the complexities of real galaxies are preventing a clean statistical test, and the BIC may need to be used in tandem with other methods to ensure that the correct model is selected.

For future large-area, deep optical surveys such as *Euclid* and Large Synoptic Survey Telescope (LSST), a full Bayesian

analysis of local galaxy morphologies will be essential for unlocking the remaining unanswered questions about galaxy structures. Both fast, non-parametric approaches and full Bayesian fitting methods will need to be employed to quantify galaxy structures and successfully link them to the assembly history of galaxies. In the era of massive cosmological simulations taking galaxy evolution into a quantitative comparative science, observers must be careful to account for degeneracies between structural parameters when scaling relations are calculated. The addition of Φ_{HI} into the 2D photometric decomposition toolbox will help advance our future understanding of galaxy properties.

ACKNOWLEDGEMENTS

We thank the anonymous referee for many valuable comments that helped to improve this paper. JJA, JMA, and VW acknowledge support of the European Research Council via the award of a starting grant (SEDMorph; PI: V. Wild). JMA acknowledges support from the Spanish Ministerio de Economía y Competitividad (MINECO) by the grant AYA2013-43188-P. Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org>.

REFERENCES

- Abazajian K., Adelman-McCarthy J. K., Agüeros M. A., Allam S. S., 2004, *AJ*, 128, 502
- Allen P. D., Driver S. P., Graham A. W., Cameron E., Liske J., de Propris R., 2006, *MNRAS*, 371, 2
- Bendinelli O., 1991, *ApJ*, 366, 599
- Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
- Boroson T., 1981, *ApJS*, 46, 177
- Cappellari M., 2002, *MNRAS*, 333, 400
- Conselice C. J., 2003, *ApJS*, 147, 1
- de Souza R. E., Gadotti D. A., dos Anjos S., 2004, *ApJS*, 153, 411
- de Vaucouleurs G., 1948, *Ann. d’Astrophys.*, 11, 247
- Dieleman S., Willett K. W., Dambre J., 2015, *MNRAS*, 450, 1441
- Emsellem E., 1998, in Sofue Y., ed., *Proc. IAU Symp.* 184, *The Central Regions of the Galaxy and Galaxies*. Kluwer, Dordrecht, p. 397
- Erwin P., 2015, *ApJ*, 799, 226
- Fasano G., Filippi M., 1998, *A&AS*, 129, 583
- Freeman K. C., 1970, *ApJ*, 160, 811
- Gadotti D. A., 2009, *MNRAS*, 393, 1531
- Gelman A., Rubin D. B., 1992, *Stat. Sci.*, 7, 457
- Geweke J., 1992, in Bernardo J. M., Berger J. O., Dawid A. P., Smith A. F. M., eds, *Bayesian Statistics*. Oxford Univ. Press, Oxford, Vol. 4, p. 169
- Graham A. W., Driver S. P., 2005, *Publ. Astron. Soc. Aust.*, 22, 118
- Haario H., Saksman E., Tamminen J., 2001, *Bernoulli*, 7, 223
- Hubble E. P., 1936, *Realm of the Nebulae*. Yale Univ. Press, New Haven, CT
- Huertas-Company M., Rouan D., Tasca L., Soucail G., Le Fèvre O., 2008, *A&A*, 478, 971
- Huertas-Company M., Aguerri J. A. L., Bernardi M., Mei S., Sánchez Almeida J., 2011, *A&A*, 525, A157
- Huertas-Company M. et al., 2015, *ApJS*, 221, 8
- Kormendy J., 1977, *ApJ*, 217, 406
- Lange R. et al., 2016, *MNRAS*, 462, 1470
- Laurikainen E., Salo H., Buta R., Knapen J. H., Comerón S., 2010, *MNRAS*, 405, 1089
- Lintott C. J. et al., 2008, *MNRAS*, 389, 1179
- Lintott C. et al., 2011, *MNRAS*, 410, 166
- Lotz J. M., Primack J., Madau P., 2004, *AJ*, 128, 163
- Meert A., Vikram V., Bernardi M., 2015, *MNRAS*, 446, 3943

- Méndez-Abreu J., Aguerri J. A. L., Corsini E. M., Simonneau E., 2008, *A&A*, 478, 353
- Méndez-Abreu J., Debattista V. P., Corsini E. M., Aguerri J. A. L., 2014, *A&A*, 572, A25
- Méndez-Abreu J. et al., 2017, *A&A*, 598, A32
- Odehahn S. C., Cohen S. H., Windhorst R. A., Philip N. S., 2002, *ApJ*, 568, 539
- Pawlik M. M., Wild V., Walcher C. J., Johansson P. H., Villforth C., Rowlands K., Mendez-Abreu J., Hewlett T., 2016, *MNRAS*, 456, 3032
- Peng C. Y., Ho L. C., Impey C. D., Rix H.-W., 2002, *AJ*, 124, 266
- Peng C. Y., Ho L. C., Impey C. D., Rix H.-W., 2010, *AJ*, 139, 2097
- Roberts G. O., Rosenthal J. S., 2001, *Stat. Sci.*, 16, 351
- Roberts G. O., Rosenthal J. S., 2009, *J. Comput. Graphical Stat.*, 18, 349
- Robotham A., Taranu D. S., Tobar R., Moffett A., Driver S., 2017, *MNRAS*, 466, 1513
- Sandage A., 1961, *The Hubble Atlas of Galaxies*. Carnegie Institution of Washington, Washington, DC
- Schwarz G., 1978, *Ann. Stat.*, 6, 461
- Sersic J. L., 1968, *Atlas de Galaxias Australes*. Observatorio Astronomico, Cordoba, Argentina
- Simard L., 1998, in Albrecht R., Hook R. N., Bushouse H. A., eds, *ASP Conf. Ser. Vol. 145, Astronomical Data Analysis Software and Systems VII*. Astron. Soc. Pac., San Francisco, p. 108
- Simard L., Mendel J. T., Patton D. R., Ellison S. L., McConnachie A. W., 2011, *ApJS*, 196, 11
- Simmons B. D. et al., 2017, *MNRAS*, 464, 4420
- Strauss M. A. et al., 2002, *AJ*, 124, 1810
- Willett K. W. et al., 2013, *MNRAS*, 435, 2835
- Yoon I., Weinberg M. D., Katz N., 2011, *MNRAS*, 414, 1625

APPENDIX A: LIMITATIONS TO MODELLING GALAXY BULGES

Here we use a library of synthetic galaxies to investigate how the final parameter distributions are dependent on the resolution of the components, i.e. the total number of pixels that make up each component. It is clear that the ability to resolve the inner component (the bulge) will diminish as R_e tends to the Nyquist limit. A decrease in resolution will affect the precision of other parameters of interest, which we demonstrate here. For these synthetic galaxies we fix $h = 10$ pixels and keeping $I_0 = 10^{2.8}$ counts to ensure the disc is well resolved. These values are typical for the SDSS sample studied in Section 4.

We define the extent of the bulge region, R_{BD} , by subtracting the modelled disc component from the modelled bulge component and fitting an ellipse to the positive central pixels. This provides a good description of where the bulge is dominant in the bulge+disc galaxy model. Fig. A1 shows R_{BD} as a function of B/T for models with a range of values of n (coloured lines) and R_e/h ratio (different panels).

The figure clearly illustrates the highly non-linear relation between the structural parameters and the extent of the bulge region. We see that galaxies with $B/T > 0.7$, $n > 2$, and $R_e/h > 0.2$ will have the greatest number of data points available to fit. Understandably, bulges with $n = 1$ will be more difficult to fit, as will bulges with low R_e/h or B/T . Interestingly, galaxies with high B/T values will still have small bulge extents where R_e/h is low.

In Fig. A2 we focus on galaxies with bulges with high n but low B/T , i.e. small bulges with profiles highly distinct from a disc-like structure. We show the posterior fractional error distribution in n versus R_e for nine synthetic galaxies with different values of n and B/T . The figure shows the strong degeneracy between the fitted n and R_e , resulting in very large errors on the median values, when




Figure A1. Bulge extent R_{BD} (see text) as a function of bulge-to-total luminosity ratio (B/T). Different colour lines represent different Sérsic indices and each panel shows a different R_e/h ratio. For these synthetic galaxies $h = 10$ pixels and $I_0 = 10^{2.8}$ counts.

B/T is very small. This degeneracy can lead to biases in population fits, as found in Section 3.2.1 and Fig. 8.

This is relevant, because it is commonly believed that galaxies with a dominant bulge (high B/T) component are more likely to have a higher n , while lower B/T galaxies tend to have low n values (e.g. this claim is demonstrated in fig. 11 of Gadotti 2009, and also in fig. 6 of Laurikainen et al. 2010, who use more complex decompositions of higher resolution data). One implication of Figs A1 and A2 is that

both these parameter combinations are easier to model, compared to galaxies with low B/T and high n , or high B/T and low n . This is a clear example of where studies need to be aware of degeneracies between parameters, and the potential for systematic biases. Future studies should be aware that the B/T - n relation may be influenced by the resolution effects described here.




Figure A2. The posterior fractional error distribution between R_e and n for nine synthetic galaxies. Each column (left to right) represents synthetic galaxies with true n values of $n = 4, 5,$ and 6 , while each row (top to bottom) has a different B/T , $B/T = 0.1, 0.3,$ and 0.5 . The contours are the 68, 95, and 99 per cent confidence regions. The dashed line shows the median of the distribution.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.