Bayesian galaxy shape measurement for weak lensing surveys - I. Methodology and a fast-fitting algorithm

L. Miller, T. D. Kitching, C. Heymans, A. F. Heavens, L. van Waerbeke

▶ **To cite this version:**

L. Miller, T. D. Kitching, C. Heymans, A. F. Heavens, L. van Waerbeke. Bayesian galaxy shape measurement for weak lensing surveys - I. Methodology and a fast-fitting algorithm. Monthly Notices of the Royal Astronomical Society, 2007, 382, pp.315-324. 10.1111/j.1365-2966.2007.12363.x . hal-03646534

HAL Id: hal-03646534
https://hal.science/hal-03646534

Submitted on 30 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Mon. Not. R. Astron. Soc. **382,** 315–324 (2007) doi:10.1111/j.1365-2966.2007.12363.x

Bayesian galaxy shape measurement for weak lensing surveys – I. Methodology and a fast-fitting algorithm

L. Miller,[1⋆] T. D. Kitching,[1] C. Heymans,[2,3] A. F. Heavens[4] and L. Van Waerbeke[2]

[1]*Department of Physics, Oxford University, Keble Road, Oxford OX1 3RH*
[2]*Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC, Canada V6T 1Z1*
[3]*Institut d'Astrophysique de Paris, UMR7095 CNRS, 98 bis bd Arago, 75014 Paris, France*
[4]*SUPA†, Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ*

Accepted 2007 August 17. Received 2007 August 14; in original form 2007 June 1

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

ABSTRACT

The principles of measuring the shapes of galaxies by a model-fitting approach are discussed in the context of shape measurement for surveys of weak gravitational lensing. It is argued that such an approach should be optimal, allowing measurement with maximal signal-to-noise ratio, coupled with estimation of measurement errors. The distinction between likelihood-based and Bayesian methods is discussed. Systematic biases in the Bayesian method may be evaluated as part of the fitting process, and overall such an approach should yield unbiased shear estimation without requiring external calibration from simulations. The principal disadvantage of model fitting for large surveys is the computational time required, but here an algorithm is presented that enables large surveys to be analysed in feasible computation times. The method and algorithm is tested on simulated galaxies from the Shear TEsting Programme (STEP).

Key words: gravitational lensing – methods: data analysis – methods: statistical – techniques: miscellaneous.

1 INTRODUCTION

Measurement of the effects of weak gravitational lensing has become a key technique in the arsenal of methods used to measure the distribution of matter, both associated with individual objects such as galaxy clusters or individual galaxies, and on large scales through the measurement of 'cosmic shear'. A key advantage of such measurement is that it directly measures the total matter distribution, generally dominated by the dark matter component, which may then be related directly to theory without needing to understand the uncertain effects of the physics of baryons in galaxies, provided one avoids the highly non-linear regime (White 2004; Zhan & Knox 2004; Jing et al. 2006). Through the use of photometric redshifts, 3D analyses (Hu 1999; Bacon & Taylor 2003; Heavens 2003) can be used to further measure both the cosmological growth of structure and the values of cosmological parameters (Heavens, Kitching & Taylor 2006; Kitching et al. 2007; Massey et al. 2007a; Taylor et al. 2007). Until recently such surveys have been of limited size, but even so the results obtained provided useful constraints on cosmological parameters and an important test of the values deduced from other methods. One long-standing puzzle has been that the range of values for the power-spectrum normalization parameter σ_8 found by weak lensing analyses has tended to be higher than found by some other methods (see the discussion in Spergel et al. 2007), an effect that persists at some level in the latest studies. For the best-fitting 3-yr *Wilkinson Microwave Anisotropy Probe* (*WMAP*) value of the matter density parameter $\Omega_0 = 0.24$, the 3D analysis of Massey et al. (2007b) finds the value $\sigma_8 = 0.96^{+.09}_{-.07}$ and the 2D analysis of Benjamin et al. (2007) finds $\sigma_8 = 0.84 \pm 0.07$. These results can be compared with the 3-yr *WMAP* value $\sigma_8 = 0.76 \pm 0.05$ (Spergel et al. 2007).

Measurement of the effect of weak gravitational lensing requires the statistical analysis of large samples and is sensitive to any systematic errors in measured quantities. Possible systematic errors in lensing signals introduced by uncertainty in photometric redshifts has been discussed by Edmondson, Miller & Wolf (2006). Another fundamental concern with the method is whether the shapes of galaxies, which are used to deduce the signal, may be measured in an unbiased manner. The problem of shape measurement in optical imaging data is that galaxy images are convolved with a possibly varying point spread function (PSF) which must be accurately corrected for when deducing galaxy shape. Convolution with the PSF tends to make galaxy images appear rounder (for reasonably circularly symmetric PSFs) whereas addition of photon shot noise has the systematic effect of tending to make round galaxies appear less round. These two observational effects thus tend to work in opposite senses, and are independent of each other, so that both accurate PSF correction and calibration to remove the effects of noise on shape are required. Following the seminal paper by Kaiser, Squires & Broadhurst (1995) there have been many suggestions for possible measurement processes, which are discussed by Heymans et al. (2006) and Massey et al. (2007b) as part of the 'Shear TEsting

⋆E-mail: l.miller1@physics.oxford.ac.uk
†The Scottish Universities Physics Alliance.

© 2007 The Authors. Journal compilation © 2007 RAS

Programme' (STEP). Those papers discuss 18 published methods for shear measurement. The existence of so many suggested methods implies that no consensus has yet emerged on the best way to measure weak lensing signals, and therefore naturally leads us to ask whether there might in fact be one method that may be regarded as being optimal. In this paper we investigate whether a model-fitting approach to galaxy shape measurement can both achieve this aim of optimal measurement and also be constructed such that it is computationally feasible for large surveys.

In the following we shall suppose that galaxies may be characterized by a measurement of their ellipticity ($\boldsymbol{e}$) and that a weak lensing signal, such as cosmic shear, may be inferred either from the mean ellipticity or from some form of cross-correlation of the ellipticities of different galaxies.

We can stipulate a number of requirements that a weak lensing measurement technique should satisfy.

(i) Optimal measurement of lensing signal, in the sense of maximum signal-to-noise ratio (S/N).

(ii) Unbiased measurement of lensing signal.

(iii) Ability to calculate the statistical uncertainties of the measurement.

A standard approach that in principle allows us to meet these criteria is that of model fitting, which is the method discussed in this paper. We first discuss some general principles, including whether we should use a frequentist or Bayesian approach and how shear may be measured in an unbiased way from a Bayesian posterior probability distribution. However, the principal disadvantage of a model-fitting approach is that it might be computationally prohibitive for very large surveys. In Section 3 we discuss a novel galaxy shape model-fitting algorithm that allows good estimation of the likelihood surface in a usefully short computational time. We also discuss the evaluation of shear sensitivity within the Bayesian framework, which allows individual galaxy contributions to be assessed and unbiased estimation of shear to be made, fulfilling the second criterion above. Some initial results and further considerations are then discussed. More detailed results from applying the algorithm to the STEP simulations are given in a companion paper (Kitching et al., in preparation).

2 A MODEL-FITTING APPROACH TO SHAPE MEASUREMENT

2.1 General considerations

The basic rationale for fitting a model of a galaxy's surface brightness distribution is that, if the family of models is a good representation of the true surface brightness profile, the highest possible S/N of the resulting parameters should be obtained. When model and data agree the model encapsulates the full information content of the data. Although this has been recognized previously in weak lensing shape measurement (Bernstein & Jarvis 2002), no implementation of weak lensing shape measurement methods published to date has this property, because the methods usually adopt some simplification of the surface brightness profile, such as assuming that second moments entirely characterize the profile (e.g. Tyson, Wenk & Valdes 1990; Kaiser et al. 1995) or equivalently assuming Gaussian profiles or weights (e.g. Kuijken 1999; Bernstein & Jarvis 2002; Bridle et al. 2002; Bardeau et al. 2005). Model fitting has been used for some time for detailed determination of galaxy surface brightness profiles and shapes (e.g. Peng et al. 2003). Kuijken (1999) proposed model fitting to averaged galaxy images specifically for weak lensing measurement, and Bridle et al. (2002) proposed a method of measuring shear by fitting galaxies and PSFs with multiple Gaussian components. The latter method has been applied to surveys of weak lensing around galaxy clusters by Bardeau et al. (2005, 2007) and Kneib et al. (2003), among others. A Monte Carlo method is used to find best-fitting galaxy model parameters for each individual galaxy, where Gaussian surface brightness profiles, or combinations of two Gaussian profiles, are assumed for both galaxy and PSF. Shear measurement and the computational time required for that model-fitting method has been evaluated by Heymans et al. (2006). Recently, sets of basis functions known as 'shapelets' have been used to describe surface brightness profiles (Refregier 2003; Refregier & Bacon 2003 and the related work of Bernstein & Jarvis 2002) but there is no requirement for the individual shapelet functions to match real galaxy profiles. Moving to a pure model-fitting approach allows us to choose whichever brightness profiles we like, and for galaxies it clearly makes most sense to choose either exponential or de Vaucouleurs surface brightness profiles. Naturally, the above statements are qualitative, we do not know *how much* improvement one obtains by fitting a profile that is closer to the actual profile, but the principle at least is a sound one, which we expect to satisfy the first of our criteria from the Introduction.

Either frequentist model fitting, based on determining the likelihood function, or Bayesian model fitting that determines the posterior probability distribution of model parameters, allow error estimates to be made, satisfying the third of our criteria. This is not the case for early versions of weak lensing shear estimators, although error estimates have been made in some recent methods (Bernstein & Jarvis 2002; Bridle et al. 2002; Bardeau et al. 2005; Kuijken 2006).

Finally, we should address the question of whether a method can be determined to be unbiased. This is a serious issue for weak lensing studies, where the signal is so small that even a small systematic bias can have a devastating effect. Evaluation of existing methods by STEP demonstrate that they are indeed biased, with significant magnitude-dependent biases that need to be corrected empirically from comparison with simulations (Heymans et al. 2006; Massey et al. 2007b). We discuss in the next section why in principle a Bayesian method should be unbiased provided a correct choice of prior is made, but note that realistic implementations result in a quantifiable bias that may be corrected for.

2.2 Bayesian estimation of the sample ellipticity distribution

We have previously mentioned the problem of shape measurement, which not only is the shape changed by convolution with the PSF, but also noise-biases the measured shape, and in general tends to make nearly circular objects systematically appear more elliptical. We discuss in this section how a Bayesian method may be formulated that precisely corrects for this phenomenon, provided we make a correct choice of prior.

Consider a set of observations of N galaxies that yields the surface brightness distribution for each galaxy denoted by a vector of pixel values $\boldsymbol{y}$. The shape of each galaxy may be characterized by its two-component ellipticity $\boldsymbol{e}$: the particular definition we choose for $\boldsymbol{e}$ in this paper is given in Section 2.4 but what follows below applies to any shape estimator that we may choose. If the sample of galaxies has a probability distribution of intrinsic ellipticities (i.e. the value of ellipticity that would be measured by the observer in the absence of degradation by the PSF or by noise) $f(\boldsymbol{e})$, then the probability

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

distribution of $\boldsymbol{y}$ is

$$n(\boldsymbol{y}) = \int f(\boldsymbol{e})\epsilon(\boldsymbol{y}\,|\,\boldsymbol{e})\mathrm{d}\boldsymbol{e},$$

where $\epsilon(\boldsymbol{y}|\boldsymbol{e})$ is the probability distribution for $\boldsymbol{y}$ given $\boldsymbol{e}$.

For each of these galaxies we can generate a Bayesian posterior probability distribution for its ellipticity

$$p_i(\boldsymbol{e}\,|\,\boldsymbol{y}_i) = \frac{\mathcal{P}(\boldsymbol{e})\mathcal{L}(\boldsymbol{y}_i\,|\,\boldsymbol{e})}{\int \mathcal{P}(\boldsymbol{e}')\mathcal{L}(\boldsymbol{y}_i|\boldsymbol{e}')\,\mathrm{d}\boldsymbol{e}'},$$

where $\mathcal{P}(\boldsymbol{e})$ is the ellipticity prior probability distribution and $\mathcal{L}(\boldsymbol{y}_i|\boldsymbol{e})$ is the likelihood of obtaining the ith set of data values $\boldsymbol{y}_i$ given an intrinsic ellipticity $\boldsymbol{e}$.

We would hope that the true distribution of intrinsic ellipticities can be obtained from the data by considering the summation over the data:

$$\left\langle \frac{1}{N}\sum_i p_i(\boldsymbol{e}\,|\,\boldsymbol{y}_i)\right\rangle = \int \mathrm{d}\boldsymbol{y}\,\frac{\mathcal{P}(\boldsymbol{e})\mathcal{L}(\boldsymbol{y}\,|\,\boldsymbol{e})}{\int \mathcal{P}(\boldsymbol{e}')\mathcal{L}(\boldsymbol{y}\,|\,\boldsymbol{e}')\mathrm{d}\boldsymbol{e}'}\int f(\boldsymbol{e}'')\epsilon(\boldsymbol{y}\,|\,\boldsymbol{e}'')\,\mathrm{d}\boldsymbol{e}'',$$

where on the right-hand side we are integrating over the probability distributions to obtain the expectation value of the summed posterior probability distribution for the sample. We can see that this will be achieved if both $\epsilon(\boldsymbol{y}\,|\,\boldsymbol{e}) = \mathcal{L}(\boldsymbol{y}\,|\,\boldsymbol{e})$ and $\mathcal{P}(\boldsymbol{e}) = f(\boldsymbol{e})$, assuming the likelihood is normalized, $\int \mathcal{L}(\boldsymbol{y}\,|\,\boldsymbol{e})\mathrm{d}\boldsymbol{y} = 1$, from which we obtain

$$\left\langle \frac{1}{N}\sum_i p_i(\boldsymbol{e}\,|\,\boldsymbol{y})\right\rangle = \mathcal{P}(\boldsymbol{e}) = f(\boldsymbol{e}).$$

The strength of this result is that we can in principle recover statistically knowledge of the intrinsic distribution of shapes independently of assumptions about the shapes of the likelihood surfaces: in particular the likelihood surfaces for ellipticity measurement must be non-Gaussian, being bounded at $|e| < 1$, but this has no effect on the results we expect. This result parallels the analogous result discussed by Edmondson et al. (2006) for the case of Bayesian photometric redshift estimation. It says that we must know the mechanism by which data values are generated in order to construct the likelihood function, and that we must know the expected distribution of intrinsic ellipticities, in which case the summed posterior probability distribution will recover that intrinsic distribution. It might be thought that a Bayesian approach then has a non-useful requirement that we need to know the answer before we start, but the point of course is that with the correct choice of prior we then expect the posterior probability distribution for each *individual* galaxy to yield an unbiased estimate of ellipticity, and those sets of individual posterior probability distributions may then be used to infer the spatially varying shear arising from gravitational lensing. We discuss in Section 5.1 one possible method for creating the correct prior.

2.3 Frequentist or Bayesian measurement?

So far the framework has been described in a purely Bayesian context, but we can also ask whether there is a frequentist equivalent of the above formalism: can weak lensing shear be measured using likelihood functions alone? Conversely, are there any disadvantages to using a Bayesian method?

It is important to recognize that likelihood and Bayesian estimators measure different things. We can illustrate this by considering a sample of galaxies (say) with some intrinsic property x that we wish to determine from fitting to some measurements y. We shall look at the results obtained with either a Bayesian estimator, $\hat{x}_{\mathcal{B}} = \int xp(x|y)\mathrm{d}x$, or a likelihood estimator, $\hat{x}_{\mathcal{L}} = \int x\mathcal{L}(y|\,x)\mathrm{d}x$ (the general considerations discussed here apply also to maximum likelihood estimators). Suppose the intrinsic distribution of x has a normal distribution of variance a^2, and that for each x drawn from this distribution the measurement process causes a normally distributed uncertainty of variance b^2. Fig. 1 shows the results obtained in a Monte Carlo realization for the illustrative case $a = 0.3$, $b = 0.4$.

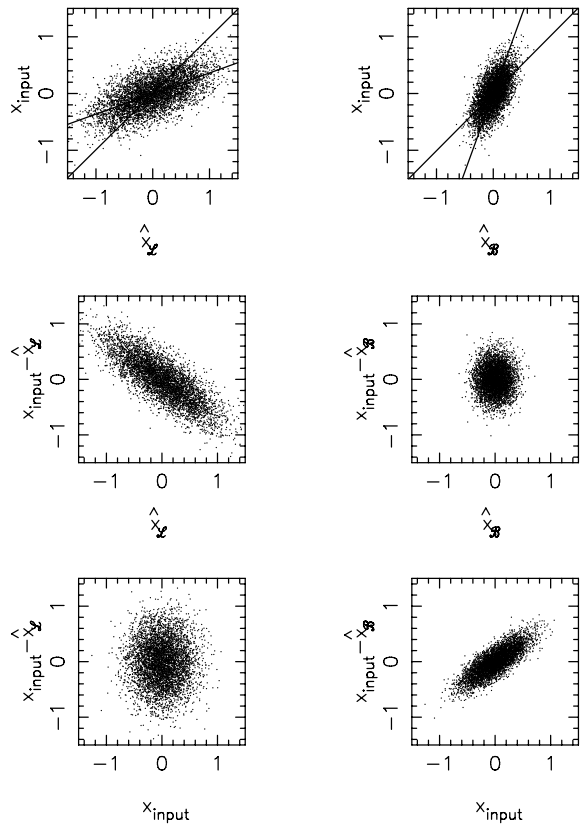

Figure 1. Illustration of the properties of an ideal likelihood estimator $\hat{x}_{\mathcal{L}}$ (left-hand panels) and ideal Bayesian estimator $\hat{x}_{\mathcal{B}}$ (right-hand panels) for the Gaussian example described in the text. The top pair of graphs shows the correlation between the input and deduced values. Two regression lines are shown on each, one being the regression of input on estimated value, the other being the regression of estimated on input value. The next two pairs show the distribution of the difference between input and estimated values compared with either the estimated values (centre) or the input values (bottom). Note the graph x-axes differ between the centre and bottom panels. For a given input value, the likelihood estimator yields an unbiased estimate (regression slope unity) whereas the Bayesian estimator appears biased (regression slope 2.75). However, for a given estimated value, the likelihood estimator is biased (regression slope 0.36) and the Bayesian estimator is unbiased. The Bayesian estimator returns the best estimate of the input value for a given measurement.

The likelihood estimator is based on the function $\mathcal{L}(y|x)$ and in Fig. 1 is unbiased in the regression of input on estimated values. The Bayesian estimator is based on the function $p(x|y)$ and is unbiased in the regression of estimated on input values: i.e. for a given set of measurements the Bayesian estimator yields the best estimate of the input values. The likelihood estimator yields a distribution of measured values that is broader than the intrinsic distribution. Note, however, that the Bayesian estimator yields a distribution that is

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

narrower, despite the result of Section 2.2 that the summed posterior probability distribution yields the intrinsic distribution if the correct prior is chosen. This apparent paradox is resolved by realizing that each of the individual estimated values is associated with its own posterior probability distribution, so that the sum of the distributions is broader than the distribution of expectation values. While this might seem undesirable, i.e. inevitable in any noisy measurement process.

In the case of ellipticity shape measurement, we expect there to be other reasons why a likelihood estimator might be biased, in particular at large $\boldsymbol{e}$ values or for low S/N, where the boundary $|\boldsymbol{e}| < 1$ renders the likelihood function asymmetric and highly non-Gaussian. With no prior, the strong degeneracy between size and ellipticity in the likelihood fitting can create regions of high likelihood at extreme values of ellipticity, and given the hard bound $|\boldsymbol{e}| < 1$ such an estimator cannot be unbiased.

In the frequentist approach, it is also possible to estimate errors for the individual galaxies, and this is important to establish the contribution to the signal from each galaxy. As in the Bayesian approach, we expect the information on shear to decrease as the S/N decreases: even with an unbiased estimator $\hat{\boldsymbol{e}}$ it would be important to quantify this effect and allow for it in the shear estimation.

A final consideration is that in weak lensing surveys we are not simply interested in measuring the shapes of individual galaxies, but rather in measuring the systematic lensing shear in a sample. Usually this is measured from the mean ellipticity: thus it may be possible to have an ellipticity estimator that is biased but where the shear estimate from a sample is unbiased, or vice versa. If any bias were isotropic, corresponding to a bias in the value of $|\boldsymbol{e}|$ but not in orientation, then we might hope that the bias would average out. However, even in this case we should not *assume* that the shear estimator $\langle \boldsymbol{e} \rangle$ is unbiased, since the likelihood functions for $\boldsymbol{e}$ measurement must be non-Gaussian and $\boldsymbol{e}$ dependent, any shift $\boldsymbol{g}$ in the distribution of $\boldsymbol{e}$ caused by lensing would lead to bias in the estimated shear.

However, even the Bayesian method is not immune to the problem of bias, particularly in a realistic implementation of the Bayesian method where we are forced to assume a zero-shear prior, as discussed below. But the bias can be quantified and the method provides a self-contained framework within which we can work out all the required quantities. This is the framework that we return to in the remainder of this paper.

2.4 Bayesian shear estimation and the shear sensitivity

Following Heymans et al. (2006) we assume observed galaxy ellipticity $\boldsymbol{e}$ is related to the intrinsic galaxy ellipticity $\boldsymbol{e}^s$ in the weak lensing regime via

$$\boldsymbol{e} = \frac{\boldsymbol{e}^s + \boldsymbol{g}}{1 + \boldsymbol{g}^\star \boldsymbol{e}^s}$$

from Schramm & Kayser (1995), Seitz & Schneider (1997), where $\boldsymbol{e}$ is represented as a complex variable and $\boldsymbol{g}$, $\boldsymbol{g}^*$ are the reduced shear and its complex conjugate, respectively. $\boldsymbol{e}$ is defined in terms of the major and minor axes and orientation a, b, θ, respectively, as $\boldsymbol{e} = (a - b)/(a + b) \exp(2\mathrm{i}\theta)$. In this formalism, we expect

$$\langle \boldsymbol{e} \rangle = \boldsymbol{g} \tag{1}$$

for an unbiased sample where $\langle \boldsymbol{e}^s \rangle = 0$, and so $\langle \boldsymbol{e} \rangle$ for a sample of galaxies is adopted as our estimator of $\boldsymbol{g}$. Note that this result differs from the other commonly used formalism where ellipticity is instead defined as $\boldsymbol{e} = (a^2 - b^2)/(a^2 + b^2) \exp(2\mathrm{i}\theta)$.

For a population of galaxies, $\langle \boldsymbol{e} \rangle = \int \boldsymbol{e}\, f(\boldsymbol{e}) \mathrm{d}\boldsymbol{e}$ where $f(\boldsymbol{e})$ is the ellipticity probability distribution for the sample. But in the Bayesian formalism we can write a similar expression for an individual galaxy if we know its Bayesian posterior probability distribution, $\langle \boldsymbol{e} \rangle_i = \int \boldsymbol{e}\, p(\boldsymbol{e}|\boldsymbol{y}_i) \mathrm{d}\boldsymbol{e}$. Hence for a sample of N galaxies we can evaluate the sample mean as

$$\langle \boldsymbol{e} \rangle = \frac{1}{N} \sum_i \int \boldsymbol{e}\, p_i(\boldsymbol{e}|\boldsymbol{y}_i) \mathrm{d}\boldsymbol{e} = \frac{1}{N} \int \boldsymbol{e} \sum_i p_i(\boldsymbol{e}|\boldsymbol{y}_i)\, \mathrm{d}\boldsymbol{e}.$$

In practice we shall use the first of these expressions, as estimation of ellipticities for individual galaxies allows error estimates to be made for each galaxy, and its contribution to the signal to be evaluated.

However, in measuring shear we cannot know in advance the correct prior to apply, even if we know the intrinsic unsheared ellipticity prior distribution, because the amount of shear varies over the sky in a way that we are attempting to measure. We must therefore use a prior that contains zero shear. The effect of this is that as S/N decreases, the measured ellipticity distribution tends to the prior, and in the limit of zero S/N no shear signal is recoverable. This is precisely what we should expect of course: no measurement method can extract a measured shear value from data with zero S/N, and a Bayesian method is no different in that respect. A Bayesian method does however allow us to estimate the magnitude of this effect for each individual galaxy. Consider the Bayesian estimate of ellipticity $\langle \boldsymbol{e} \rangle_i$, defined above, which is measured for the ith galaxy, and express its dependence on each component of shear $\boldsymbol{g}$ as a Taylor series. For component e_1,

$$\langle e_1 \rangle_i \simeq e^s_{1i} + g_1 \partial \langle e_1 \rangle_i / \partial g_1 + g_2 \partial \langle e_1 \rangle_i / \partial g_2 + \cdots \tag{2}$$

and similarly for component e_2, where numeric subscripts indicate the components of $\boldsymbol{e}$ and $\boldsymbol{g}$. In the weak lensing limit the cross-terms vanish. If we sum over N galaxies in an unbiased sample we find

$$\sum_i^N \langle e_1 \rangle_i \simeq g_1 \sum_i^N \partial \langle e_1 \rangle_i / \partial g_1.$$

We may optionally multiply both sides in equation (2) by a statistical weight for each galaxy, w_i. Provided w_i is uncorrelated with $\boldsymbol{e}^s_i$ we may then define a weighted estimate of shear for the sample:

$$\hat{g_\mu} \equiv \frac{\sum_i^N w_i \langle e_\mu \rangle_i}{\sum_i^N w_i \partial \langle e_\mu \rangle_i / \partial g_\mu} \tag{3}$$

for $\mu = 1, 2$. We shall call $\partial \langle e_\mu \rangle_i / \partial g_\mu$ the *shear sensitivity*. It is a measure of how much each Bayesian estimate is biased by the use of the zero-shear prior, and it takes values in the range $0 < \partial \langle e_\mu \rangle_i / \partial g_\mu \leqslant 1$, where the lower bound is expected in the limit of zero S/N. The upper bound would be attained in the case of ideal measurement at high S/N, where we expect no bias: in this case the Bayesian measure is a good measure of the true ellipticity, regardless of which prior is assumed, and differentiating equation (1) yields unity for the shear sensitivity. The weights w_i could in principle be tuned for optimal S/N in the measurement(s) being made, such as the values of cosmological parameters. Care should be taken that any weights that are a function of ellipticity do not introduce bias into shear measurement. Since the shear sensitivity is effectively a measure of how much information about the effect of lensing is carried by each galaxy, the weights should themselves also include a dependence on $\partial \langle e_\mu \rangle_i / \partial g_\mu$, as well as on the measurement error and on the redshift-dependent cosmological effect of lensing on each galaxy.

The shear sensitivity may be estimated for each galaxy and for the survey as a whole without recourse to external calibration from

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

simulations, as described below. Kaiser et al. (1995), Luppino & Kaiser (1997), Kaiser (2000) and Bernstein & Jarvis (2002) have emphasized the utility of knowing the shear 'polarizability' or 'responsivity' for individual galaxies, as this not only allows accurate optimized shear measurement but also allows future surveys to be planned and optimized.

The estimator $\hat{g}_\mu$ is appropriate for a survey where the shear is uniform over some region (and this is assumed in the STEP simulations discussed below), but in the more general case we instead infer the shear correlation function or some related quantity such as shear variance from a measurement such as $\langle \boldsymbol{e}_i \boldsymbol{e}_j \rangle$. In this case we can compute the analogous estimator

$$\widehat{\langle g_\mu g_\mu \rangle} = \frac{\sum_{i,j} w_i w_j \langle e_\mu \rangle_i \langle e_\mu \rangle_j}{\sum_{i,j} w_i w_j (\partial \langle e_\mu \rangle_i / \partial g_\mu)(\partial \langle e_\mu \rangle_j / \partial g_\mu)}.$$

We now discuss possible approaches to calculating the shear sensitivity, first for normal prior and likelihood distributions, then for the more general case where the shear sensitivity may be evaluated numerically from the measured likelihood surfaces of individual galaxies.

2.5 Calculation of shear sensitivity

As an illustration of the calculation of shear sensitivity, suppose the prior is described by a normal distribution $\mathcal{P}(\boldsymbol{e})$ of variance a^2 and $\langle \boldsymbol{e} \rangle = 0$, and that the likelihood $\mathcal{L}(\boldsymbol{e})$ for a particular galaxy also has a normal distribution of variance b^2 centred on some value $\boldsymbol{e}_0$. It is straightforward then to show that the Bayesian posterior probability $p(\boldsymbol{e}|\boldsymbol{y})$ also has a normal distribution of variance $a^2 b^2/(a^2 + b^2)$ and expectation value $\langle \boldsymbol{e} \rangle = \boldsymbol{e}_0 a^2/(a^2 + b^2)$. For perfect measurement of ellipticity ($b^2 \ll a^2$) we expect equation (1) to hold, so for this galaxy $\partial \langle e_\mu \rangle_i / \partial g_\mu = \partial e_{0\mu} / \partial\, g_\mu = 1$. For more noisy measurement, we expect

$$\frac{\partial \langle e_\mu \rangle_i}{\partial g_\mu} = \frac{a^2}{a^2 + b^2} \frac{\partial e_{0\mu}}{\partial g_\mu} = \frac{a^2}{a^2 + b^2}.$$

The shear sensitivity decreases as the measurement error increases. The value of the shear sensitivity is also given by the inverse of the slope of the regression of the intrinsic ellipticity on estimated ellipticity illustrated in Fig. 1.

The above illustration indicates that it is straightforward to calculate the shear sensitivity; however, in general it would not be safe to assume normal distributions: not least because $\boldsymbol{e}$ is defined such that $|\boldsymbol{e}| < 1$, so when the measurement error becomes large $\mathcal{L}(\boldsymbol{e})$ cannot be normally distributed. We discuss here one method of calculating the shear sensitivity numerically. We should emphasize that this can be done entirely internally to the fitting process, with no need to calibrate shear sensitivity externally from simulations.

Consider first the response of the posterior probability distribution to a small amount of shear. The prior probability does not depend on the shear in our implementation. Let us assume that applying a weak lensing shear shifts the likelihood function by some small amount, $\mathcal{L}(\boldsymbol{e} - \boldsymbol{e}^s) \rightarrow \mathcal{L}(\boldsymbol{e} - \boldsymbol{e}^s - \boldsymbol{g})$ and expand as a Taylor series:

$$\mathcal{L}(\boldsymbol{e} - \boldsymbol{e}^s - g_\mu) \simeq \mathcal{L}(\boldsymbol{e} - \boldsymbol{e}^s) - g_\mu \frac{\partial \mathcal{L}}{\partial e_\mu} + \cdots.$$

Then, substituting into

$$\langle \boldsymbol{e} \rangle = \frac{\int \boldsymbol{e} \mathcal{P}(\boldsymbol{e}) \mathcal{L}(\boldsymbol{e}) \mathrm{d}\boldsymbol{e}}{\int \mathcal{P}(\boldsymbol{e}) \mathcal{L}(\boldsymbol{e}) \mathrm{d}\boldsymbol{e}}$$

and differentiating with respect to $\boldsymbol{g}$ we find

$$\frac{\partial \langle e_\mu \rangle}{\partial g_\mu} \simeq \frac{\int (\langle \boldsymbol{e} \rangle - \boldsymbol{e})\, \mathcal{P}(\boldsymbol{e}) (\partial \mathcal{L}/\partial e_\mu) \mathrm{d}\boldsymbol{e}}{\int \mathcal{P}(\boldsymbol{e}) \mathcal{L}(\boldsymbol{e}) \mathrm{d}\boldsymbol{e}}$$

as an estimate of weak lensing shear sensitivity. This expression is cast in terms of the derivatives of the likelihood surface multiplied by the prior: it may also be expressed in terms of derivatives of the prior multiplied by the likelihood

$$\frac{\partial \langle e_\mu \rangle}{\partial g_\mu} \simeq 1 - \left[\frac{\int (\langle \boldsymbol{e} \rangle - \boldsymbol{e})\, \mathcal{L}(\boldsymbol{e}) (\partial \mathcal{P}/\partial e_\mu) \mathrm{d}\boldsymbol{e}}{\int \mathcal{P}(\boldsymbol{e}) \mathcal{L}(\boldsymbol{e}) \mathrm{d}\boldsymbol{e}} \right].$$

This may be evaluated numerically from the posterior probability surface for each galaxy, and is preferred over the preceding expression in the case where the derivative of the prior is known analytically. For the case of normal distributions of $\mathcal{P}(\boldsymbol{e})$ and $\mathcal{L}(\boldsymbol{e})$ the expression yields the analytic result above.

3 FAST SHAPE MEASUREMENT

3.1 The algorithm

The technique we adopt for measuring $\langle \boldsymbol{e} \rangle$ and its uncertainty is to fit model galaxy surface brightness profiles to the data for individual galaxy images. The simplest galaxy model has six free parameters if the form of the surface brightness profile is fixed: central surface brightness, size, ellipticity and celestial position. The problem of fitting six parameters to large samples of galaxies is that this could be a time-consuming task, probably prohibitively so. However, we can greatly speed up the process if we can marginalize over any parameters that are not of interest to the weak lensing measurement. It turns out that *for isolated galaxies* it is straightforward to marginalize over three of the parameters, central surface brightness and position, if the model fitting is treated in Fourier space, as described below. And because there exist fast Fourier transform algorithms this approach can be done in a short amount of computational time.

We can start by writing the statistic

$$\begin{aligned} \chi^2 &= \sum_i \left(\frac{y_i - C y_i^m}{\sigma_i} \right)^2 \\ &= \sum_i \frac{y_i^2}{\sigma_i^2} + A(C - B)^2 - AB^2, \end{aligned}$$

where y_i is the data value in pixel i, σ_i is the statistical uncertainty of that data value, y_i^m is a model value for that pixel, C is the model amplitude and where

$$A = \sum_i \left(\frac{y_i^m}{\sigma_i} \right)^2, \qquad B = \sum_i \frac{y_i y_i^m}{\sigma_i^2} \Bigg/ \sum_i \left(\frac{y_i^m}{\sigma_i} \right)^2.$$

We assume the pixel noise is stationary and uncorrelated, which is appropriate for shot noise in CCD detectors in the sky noise limit. Bright galaxies may make a significant contribution to photon shot noise, but in this case the non-stationarity of the noise makes it not possible to work in Fourier space. Hence this algorithm is appropriate for model fitting to faint galaxies in the sky-limited regime. The method can be generalized to the case where the noise is stationary but correlated between pixels, and the example of radio interferometer observations is discussed qualitatively in Section 5.4.

Then if we adopt a prior $\mathcal{P}(C)$ for the model amplitude we can marginalize the likelihood $\mathcal{L} = \mathrm{e}^{-\chi^2/2}$ over C:

$$\mathcal{L} = \mathrm{e}^{-\sum y_i^2/2\sigma^2} \mathrm{e}^{AB^2/2} \int_{C_{\min}}^{C_{\max}} \mathrm{e}^{-A(C-B)^2/2} \mathcal{P}(C)\, \mathrm{d}C.$$

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

We shall adopt a uniform prior for $\mathcal{P}(C)$ in the range $C_{\min} \leqslant C \leqslant C_{\max}$. We expect $C > 0$, but if the galaxy is significantly detected the Gaussian form of the likelihood causes the value of the prior to become unimportant at both large and small C, and we can simplify the calculation by allowing $C_{\min} \to -\infty$ and $C_{\max} \to \infty$, so that

$$\mathcal{L} \simeq \sqrt{\frac{2\pi}{A}} \mathrm{e}^{-\sum y_i^2/2\sigma_i^2} \mathrm{e}^{AB^2/2}.$$

However, although we have eliminated the amplitude C, $\mathcal{L}$ still depends on the model second moment, $A = \sum (y_i^m/\sigma_i)^2$. Thus we need to introduce the model constraint A = constant, achieved by renormalizing each model appropriately. Since for a given data set $\sum y_i^2/\sigma_i^2$ is also fixed, we can write $\mathcal{L} \propto \mathrm{e}^{AB^2/2}$ when maximizing.

We can also rapidly calculate the marginalization over galaxy celestial position if we work in Fourier space, writing

$$y_i = \sum_k y_k \mathrm{e}^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}_i}, \qquad y_i^m = \sum y_k^m \mathrm{e}^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}_i}.$$

We can simplify the various summations by assuming that we are dealing with faint galaxies in weak lensing measurement, such that σ_i is dominated by the background photon shot noise and is constant for all pixels. And since the model y_i^m is real, $y_k^m = y^{m*}_k$ and $\sum_i y_i y_i^m = \sum_k y_k y_k^{m\star}$. If we introduce a shift $\boldsymbol{X}$ into the model position, the new model becomes

$$y_i^{m\prime} = \sum_k y_k^m \mathrm{e}^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}_i} \mathrm{e}^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{X}}$$

and

$$\sum_i y_i y_i^{m\prime} = \sum_k y_k y_k^{m\star} \mathrm{e}^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{X}} = h(\boldsymbol{X})$$

where $h(\boldsymbol{X})$ is the cross-correlation of the data y_i with the model y_i^m. So the likelihood becomes

$$\mathcal{L} \propto \exp\left[\frac{|h(\boldsymbol{X})|^2}{2\sigma^2 \sum y_i^{m\,2}}\right].$$

To marginalize over $\boldsymbol{X}$ we need to adopt a prior $\mathcal{P}(\boldsymbol{X})$, but in this case it cannot be uniform as $\mathcal{L} \to$ constant as $|\boldsymbol{X}| \to \infty$ and the marginalized likelihood would not be finite. This problem arises because, no matter how large a pixel value, it always has a finite chance of being due to random noise, with the true galaxy being positioned elsewhere. We shall adopt a prior which is centred on some assumed galaxy position that has been previously estimated and which falls off to zero at large distances: this is equivalent to assuming that a galaxy does indeed exist somewhere near the location we have chosen. We shall assume a prior which is symmetric and centred on the nominal galaxy position, which for convenience is at the coordinate origin, such as

$$\mathcal{P}(\boldsymbol{X})\mathrm{d}^2\boldsymbol{X} = \frac{1}{2\pi b^2} \mathrm{e}^{-|\boldsymbol{X}|^2/2b^2} \mathrm{d}^2\boldsymbol{X}.$$

The process of model fitting is seen from the above to be one of cross-correlating the data with a model. Galaxies generally have smooth centrally concentrated surface brightness distributions which are convolved with near-Gaussian PSFs in an observed image. The model is also smooth, centrally concentrated and convolved with the same PSF. From the central limit theorem such a cross-correlation should be well represented by a 2D Gaussian distribution,

$$h(\boldsymbol{X}) = h_0 \exp\left[-(\boldsymbol{X} - \boldsymbol{X}_0)\mathbf{C}^{-1}(\boldsymbol{X} - \boldsymbol{X}_0)^{\mathrm{T}}\right],$$

where $\mathbf{C}^{-1}$ is the inverse covariance matrix and some shift $\boldsymbol{X}_0$ of the maximum from the origin is allowed. In what follows we assume circular symmetry for simplicity, although this assumption may be removed without affecting the final result (a 2D Gaussian distribution may always be transformed to a circularly symmetric distribution by a coordinate transformation). If the cross-correlation function has the form $h = h_0 \exp\left[-|\boldsymbol{X} - \boldsymbol{X}_0|^2/s^2\right]$ then

$$\mathcal{L} \propto \frac{1}{2\pi b^2} \int_0^\infty \exp\left[\beta \mathrm{e}^{-|\boldsymbol{X}-\boldsymbol{X}_0|^2/s^2}\right] \mathrm{e}^{-|\boldsymbol{X}|^2/2b^2} \mathrm{d}^2\boldsymbol{X},$$

where

$$\beta = \frac{h_0^2}{2\sigma^2 \sum y_i^{m\,2}}.$$

We could evaluate this by, for example, expanding the first exponential as a Taylor series and hence obtaining a series solution for the marginalized likelihood. We could also evaluate it purely numerically, but this would require evaluation of the cross-correlation function on an extremely fine grid in order to achieve adequate accuracy. Either of these approaches would be computationally expensive, and an alternative is to find an approximate value of the integral by writing

$$\mathcal{L} \propto \frac{1}{2\pi b^2} \int_0^\infty \left\{\exp\left[\beta \mathrm{e}^{-|\boldsymbol{X}-\boldsymbol{X}_0|^2/s^2}\right] - 1 + 1\right\} \mathrm{e}^{-|\boldsymbol{X}|^2/2b^2} \mathrm{d}^2\boldsymbol{X}.$$

If $b \gg s$,

$$\mathcal{L} \propto 1 + \frac{1}{2\pi b^2} \mathrm{e}^{-|\boldsymbol{X}_0|^2/2b^2} \int_0^\infty \left\{\exp\left[\beta \mathrm{e}^{-|\boldsymbol{X}-\boldsymbol{X}_0|^2/s^2}\right] - 1\right\} \mathrm{d}^2\boldsymbol{X}$$

and changing variables to a polar system centred on $\boldsymbol{X}_0$,

$$\begin{aligned}\mathcal{L} &\propto 1 + \frac{s^2}{b^2} \mathrm{e}^{-|\boldsymbol{X}_0|^2/2b^2} \int_0^\infty \left[\exp\left(\beta \mathrm{e}^{-r^2}\right) - 1\right] r\,\mathrm{d}r \\ &\propto \frac{s^2 \mathrm{e}^{\beta}}{2\beta b^2} \mathrm{e}^{-|\boldsymbol{X}_0|^2/2b^2} \qquad \beta \gg 1 \end{aligned} \tag{4}$$

approximately, where the constant of proportionality has no model dependency provided A is held invariant (we could obtain a similar result more exactly if we were to adopt a top-hat prior for the galaxy position). If the width s, amplitude h_0 and centroid $\boldsymbol{X}_0$ of the cross-correlation function can be measured, the marginalized likelihood may be estimated from equation (4).[1] In the more general case, where the cross-correlation function is approximated by a bivariate Gaussian with widths s_1 and s_2 in the two principal directions, then equation (4) is modified by $s^2 \to s_1 s_2$.

3.2 Implementation

Using the above, the likelihood may be estimated for a given set of model parameters, marginalized over the 'uninteresting' position and brightness of the galaxy. If we choose either an exponential disc or a de Vaucouleurs model for the surface brightness, the free parameters are the scalelength and two ellipticity values $\boldsymbol{e}$, or equivalently the scalelength, the axial ratio and the orientation. The values of $\boldsymbol{e}$ are restricted to lie in the range $|\boldsymbol{e}| < 1$, and for faint galaxies the probability $p(\boldsymbol{e})$ is broad, making a grid search in $\boldsymbol{e}$ an easy and not-too-expensive approach. The resulting likelihood may be numerically marginalized over the galaxy scalelength, which is also 'uninteresting' for weak lensing measurement, to obtain a likelihood surface, which is a function of ellipticity alone.

To evaluate the cross-correlation function $h(\boldsymbol{X})$ (returning now to the general case where we do not assume circular symmetry) we can use the fast Fourier transform method, which proceeds as follows.

[1] It may seem that the requirement for a prior on position may be removed by allowing $b \to \infty$. However, this is an artefact of the approximation. There is no clear way of identifying a value for b, but it should be set sufficiently small that confusion from other nearby galaxies is eliminated.

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

(i) Generate a series of 2D galaxy surface brightness models on a 3D grid in parameter space of scalelength and ellipticity. These models can be discrete Fourier transformed and those transforms stored for use with all the galaxies. The choice of a grid of models allows a considerable multiplex gain to be realized: the models can be pre-generated on that grid and the same set used for fitting to every galaxy.

(ii) Estimate the surface brightness profile of the PSF on the same pixel scale as the models. Usually this would be done by stacking images of stars from the region of an image as the galaxies being measured. The PSF can also be Fourier transformed and stored. If the PSF varies over an image or between images, the image may be divided into zones over which the PSF is approximately invariant, and the Fourier transform of the PSF for each zone stored separately. If a mathematical model for the varying PSF is known this may also be used to generate a smoothly varying PSF (e.g. Rhodes et al. 2007).

(iii) Estimate the rms noise in each pixel from the entire image.

(iv) Identify a set of nominal galaxy positions to be measured, most likely from a separate image analysis tool such as SE XTRACTOR (Bertin & Arnouts 1996).

(v) In turn for each galaxy, extract a subimage centred on that galaxy, Fourier transform it, and temporarily store the result.

(vi) For this galaxy, take each possible model in turn, multiply by the transposed PSF and model transforms to carry out the cross-correlation, measure the amplitude, width and position of the maximum of the resulting cross-correlation, and hence evaluate the likelihood for this model and galaxy. Repeat for all models on the grid (or for a subset of models if a more intelligent maximum likelihood or MCMC search algorithm is being employed).

(vii) Numerically marginalize over the scalelength parameter. In the implementation described here we adopt a uniform prior for the distribution of galaxy scalelength. This could be replaced by a prior close to the actual distribution of galaxy sizes, although such a prior would need to be magnitude dependent.

(viii) Discard the extracted data when all models have been explored, and repeat for the next galaxy.

The result is a grid of likelihood values in ellipticity parameter space which thus defines the probability surface $p(\boldsymbol{e})$. The reduced shear may then be directly estimated from $\langle \boldsymbol{e} \rangle$, and the uncertainty in individual $\boldsymbol{e}$ values may be estimated from the width of the likelihood surface.

There is a significant multiplex gain obtained by Fourier transforming the models, the PSFs and the data and storing the results. The time-consuming step then is the cross-correlation, which comprises some multiplications and a single inverse Fourier transform to obtain the cross-correlation function. It is this multiplex gain, combined with the elimination of three parameters by marginalization, which yields a fast-fitting algorithm.

The algorithm is approximate, in the sense that we require the cross-correlation amplitude to be high enough that $\beta \gg s^2/2b^2$, and also in that we assume the core of the cross-correlation function can be adequately modelled as a Gaussian, and we have assumed that the pixel noise is invariant. This latter constraint may impose a maximum brightness limit on galaxies that may be fitted, as the pixel noise is not invariant in the case where the galaxy itself makes a significant contribution to the noise. For a fixed size of extracted region around each galaxy, there is also a maximum galaxy size that can be adequately measured. Larger sizes are possible at the expense of greater computation time.

We note that, in this method, the final PSF used is itself a convolution of PSF components arising from the atmosphere, telescope and the pixel response of the detector. We do not need to distinguish the origin of the final PSF that is used, the method takes a galaxy model and convolves that with an estimate of the final PSF in order to cross-correlate with the data. Ultimately however this, and all shape measurement methods, are limited by the extent to which the sampled data fully encapsulate the information in the sky: the effect of sampling is to alias spatial frequencies higher than the Nyquist sampling frequency. This affects both the creation of the stacked PSF and the model fitting itself. If astronomical observations were band limited this would not be a problem, but in reality some aliasing is inevitable. Poorly sampled observations should ideally be 'dithered' in order to reduce such aliasing effects.

4 RESULTS

4.1 Tests on simulated galaxy images

The algorithm has been implemented and tested on simulations provided for the STEP (Heymans et al. 2006; Massey et al. 2007b). Images of galaxies were simulated for the Canada–France–Hawaii Telescope with pixel scale 0.206 arcsec. The simulations used here to demonstrate basic shape measurement are those with an isotropic PSF of FWHM 0.9 arcsec and zero lensing shear (tests of shear measurement in the companion paper will cover all the simulated PSF shapes and shear values). Simulated galaxies with a mixture of bulge/disc components were used but all were fitted with a single exponential surface brightness profile.

As here we are testing the Bayesian method, and not our ability to locate galaxies, we use as input galaxy positions those that were used when making the simulations. We also adopt as the prior $\mathcal{P}(\boldsymbol{e})$ the input ellipticity distribution used in the simulations.

For these tests the size of each subimage was 32 pixels. The choice of subimage size is a compromise between (i) having the subimage large enough that the galaxy surface brightness distribution is not unduly truncated and (ii) not allowing the computation time to become excessively long. In our initial implementation we have also required that only a single galaxy should occupy each subimage, thereby eliminating close pairs. The choice of 32 pixels for the STEP galaxies ensured that the subimage was larger than the half-light diameter in every case. In principle, the subimage size could be a function of galaxy size or brightness, but this sophistication would introduce some complexity into the code and has not been tested here.

The likelihood was evaluated on a Cartesian grid in $\boldsymbol{e}$, sampling at intervals of 0.1 out to a maximum axial ratio $a/b = 10$. The resulting galaxy shapes were found to be consistent for ellipticity grid intervals less than 0.1 for these STEP galaxies. The choice of grid interval may need to be adjusted for different surveys.

The PSF was created by stacking stars from the simulation, allowing subpixel registration using sinc-function interpolation. Ultimately any shear measurement survey will be limited by the accuracy to which the PSF is known. Systematic PSF errors will of course cause a systematic error in estimated shear, and if the PSF varies on some angular scale within a survey this will imprint a signal on that scale on the shear power spectrum. This concern is common to all methods of shape and shear measurement, and we do not specifically address this problem here.

An assumption of the fast-fitting algorithm is that we are fitting to individual galaxies, and hence close pairs of galaxies cannot be fitted with this algorithm. In practice one could identify such close

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

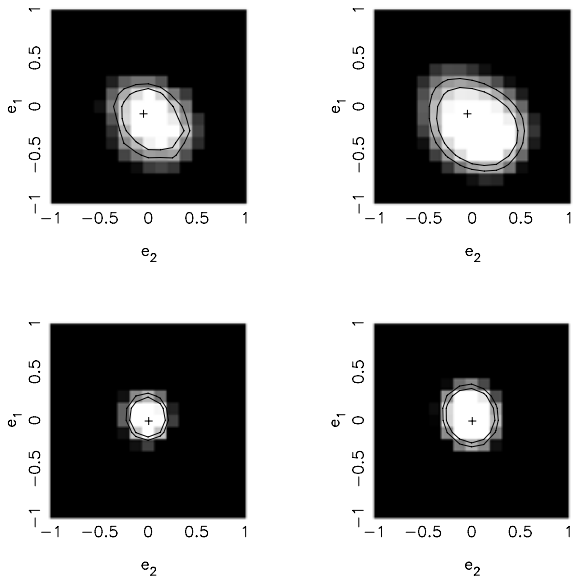

Figure 2. Comparison of Bayesian posterior probability $p(\boldsymbol{e})$ (left-hand panels) and likelihood $\mathcal{L}(\boldsymbol{e})$ (right-hand panels) surfaces for two individual galaxies. The grey-scale is logarithmic showing a range of 5 in $\Delta \log \mathcal{L}$ below the maximum value (shown as white) in each case. The upper panel shows results from fitting to a magnitude 24.17 simulated STEP galaxy, the lower panel a magnitude 23.15 galaxy. Solid lines show the two parameter 1σ and 2σ contours. The cross shows the input ellipticity value.

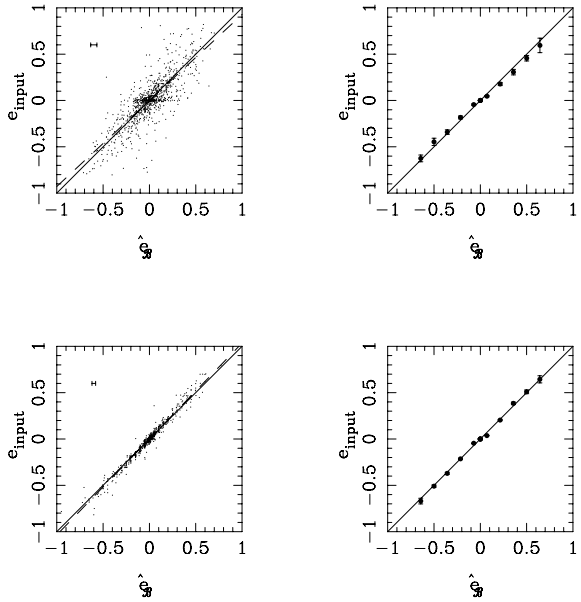

Figure 3. Tests on the STEP 1 simulated galaxy sample, as a function of galaxy apparent magnitude. Each graph shows the expectation value of the Bayesian estimate of component e_1 (x-axis) plotted against the input value (y-axis). Results for component e_2 are similar and are not shown. Left-hand panels show individual simulated galaxies, right-hand panels show results binned in intervals of the measured ellipticity. Two magnitude ranges are shown, $m > 22$ (upper panels) and $m \leqslant 22$ (lower panels). The solid lines have a slope of unity, the dashed lines on the left-hand panels show the least-squares regression of input values on estimated values. The mean error on individual measured ellipticities is shown on the left-hand panels. Vertical error bars on the right-hand panels indicate the error in the mean input values in each interval of measured values.

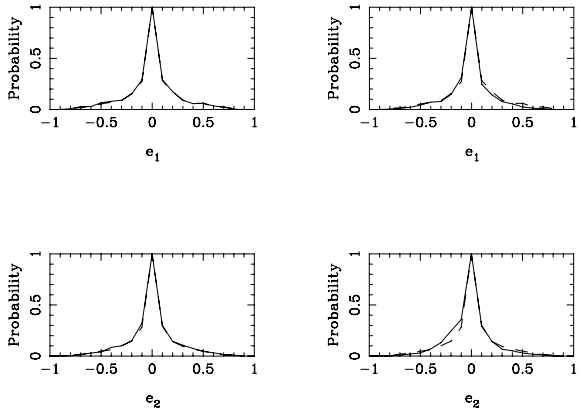

Figure 4. The summed posterior probability distribution of measured ellipticity values e_1 (top) and e_2 (bottom) as a function of apparent magnitude. The prior $\mathcal{P}(\boldsymbol{e})$ is also shown for comparison as a dashed line. The magnitude ranges of the simulated galaxies are $m \leqslant 22$ (left-hand panels) and $m > 22$ (right-hand panels).

pairs in the data at the galaxy-detection stage, and on those galaxies we could use a fitting algorithm that fits multiple components. In this case the full six parameters per galaxy would need to be fitted, with marginalization over uninteresting parameters being carried out post-fitting. There would still be a significant speed advantage to be gained by using the fast-fitting algorithm on the more isolated galaxies however. In this paper we focus on testing the Bayesian method and the fast-fitting algorithm, and hence in the results presented here we exclude cases where multiple objects are identified within a single galaxy subimage. This procedure excludes 13 per cent of galaxies in the STEP simulations. Some rejection of close pairs is required in most other methods of shape estimation also: in future, development of a fast multiple-component fitting algorithm might allow this constraint to be relaxed. We also test the fit returned by the fast-fitting algorithm to determine whether the fitted centroid of a galaxy is within a reasonable range of the nominal position, given the prior on the galaxy position: this would identify some of the cases of multiple galaxies. The criterion adopted is that the fitted galaxy position should lie within 3σ of the nominal position, where σ^2 is the prior position variance, and this excludes 9 per cent of the initial simulated galaxy sample but no others that are excluded by the 'close pairs' criterion.

Fig. 2 shows the posterior probability surfaces that result from fitting to two of the simulated galaxies. For completeness we also show the likelihood surfaces, which are broader and more biased away from the nominal value of ellipticity.

Fig. 3 shows the results for each galaxy in the simulations (only the first component of ellipticity, e_1, is shown, similar results are obtained for e_2). At bright magnitudes there is good correspondence between input and measured ellipticity values. The slope appears slightly steeper than unity, but with a value for the slope of 1.04 ± 0.08 the departure from unity is not very significant.

At fainter magnitudes, as the S/N decreases, an increasing fraction of galaxies with a given value of the Bayesian measure are drawn from a wider range of input ellipticities, as expected from the earlier discussion. The slope of the relation between input and measured values is again close to unity, with value 0.93 ± 0.11. At all magnitudes the summed posterior probability distribution is a faithful reproduction of the distribution of the input prior distribution (Fig. 4) as expected from Section 2.2. There is also no detectable

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

correlation between estimated values of e_1 and e_2 in this simulated galaxy sample.

We can also investigate the effect of measurement uncertainty in the prior positions of the galaxies. A random position uncertainty drawn from a normal distribution was introduced to each galaxy and the shapes remeasured. The prior assumed in the fitting was a normal distribution of rms 3 pixels throughout. In the results, no change was found in the slope of Fig. 3 for rms position uncertainties as large as 10 pixel. The scatter about the mean relation did not change for rms uncertainties less than 3 pixels and increased by 4 per cent for rms uncertainties as large as 10 pixels. This test indicates that the results are not sensitive to uncertainties in galaxy position measurement. We would recommend that the position prior that is chosen should match the actual position uncertainty for the faintest galaxies that are reliably fitted.

Tests of the algorithm, again using the full suite of STEP simulations specifically to measure the shear values recovered, are made in the companion paper (Kitching et al., in preparation).

4.2 Speed

The algorithm has been implemented in the C programming language[2] for use on desktop computing systems, with discrete fast Fourier transforms being supplied by the FFTW library (Frigo & Stevens 2005).[3] The computational speed per galaxy obviously depends on the computing system being used as well as on issues such as the extent to which the multiplex advantage of having many galaxies per PSF function can be exploited. In the simulations described above, using readily available 2-GHz desktop PCs in 2007 and evaluating the likelihood surfaces on a grid of sampling interval 0.1 in $\boldsymbol{e}$, we found computation times around 1.0 s per galaxy, implying that a survey of 10^7 galaxies could be analysed in a few months on a single standard desktop PC. The computation time-scales inversely with the square of the sampling in ellipticity and increases approximately as $m^2 \log m$ for subimages of size $m \times m$.

5 FURTHER CONSIDERATIONS

5.1 The ellipticity prior

A number of studies have been made of the distribution of galaxy ellipticities (e.g. Lambas, Maddox & Loveday 1992; Brainerd, Blandford & Smail 1996; Ebbels, Kneib & Ellis 1999; Bernstein & Jarvis 2002). These studies find a wide variation in distribution of axial ratios, which appears strongly dependent on apparent magnitude, presumably largely as a result of the changing mix of galaxies with brightness and redshift. The distribution of ellipticities at the faint magnitudes probed by ongoing and future weak lensing surveys is even less well known, and the best estimate would come from the lensing data itself. For a sufficiently large survey the prior estimate could also be allowed to be a function of galaxy brightness, redshift or colour, if that information were available. One way to estimate the ellipticity prior may be to adopt an iterative approach: evaluate the summed posterior probability distribution starting from an initial guess of the prior distribution; then iteratively adjust the assumed prior until the summed posterior and prior distributions agree. We would expect this to be a stable iteration in the absence of sampling noise, because if the prior is initially assumed distributed to values that are smaller than are required to explain the data, the next iteration will adjust the prior to be distributed to large values, and vice versa. Such an approach might however be unstable with small surveys where sampling noise might be important.

In the case of lensing shear estimation, the ellipticity prior should also include the shear effect, and should not just be the intrinsic pre-sheared distribution. As the shear varies on relatively small scales, and we are unlikely to have sufficient number of galaxies to measure accurately the ellipticity distribution in small regions, we suggest that correct generation of the prior should be to force the prior to be circularly symmetric, centred on $\langle e \rangle = 0$, and to be obtained from the large numbers of galaxies that comprise the full survey. In this way 'false' shear variation arising from noise on the prior would be avoided, but the resulting shear values would be slightly biased to low values, in a magnitude-dependent way. This bias has already been discussed in Section 2.4 and a method of correcting for the bias using the shear sensitivity has been described.

5.2 Choice of model surface brightness profile

A key advantage of the model-fitting approach over other methods is that a surface brightness profile may be chosen that accurately represents the actual profiles of galaxies. Two obvious choices of profile are exponential or de Vaucouleurs. In fact, it is notoriously difficult to choose between these profiles when fitting to faint galaxies, so we do not expect the accuracy of the weak lensing measurement to depend strongly on which of these profiles is chosen. Some greater freedom in profile could be allowed by adding the Sérsic index as a free parameter, allowing exponential and de Vaucouleurs models to be treated as special cases of this generalized profile (e.g. Dunlop et al. 2003); however, it is unlikely that the addition of an extra parameter can be justified on evidence grounds. A similar consideration is that galaxies generally are composed of bulge and disc components, which when viewed at an inclined angle may present differing ellipticities: accurate modelling of galaxies requires both components to be fitted, but again, for generic shape measurement of faint galaxies where information at this level of sophistication is not present in the data, this seems unwarranted.

5.3 Addition of multiple images or wavebands

It may be that a weak lensing survey comprises multiple images of the same region of sky, but taken at different times and hence with differing PSFs, and possibly in different wavebands. The latter is likely if broad-band optical photometric redshifts are also being estimated from the same data that are being used for shape measurement. Clearly one would like to optimally estimate galaxy shape from the combination of all these data, but it would not be optimal simply to co-add the data, because of the differing PSFs in each image. The model-fitting algorithm described above allows a natural way to optimally include all the data, since all we need do is add the likelihoods for the models fitted to each galaxy. In doing this, we should take care that the nominal galaxy positions are the same in each image, so the optimal way to proceed would be to co-add images for the purpose of detecting and measuring nominal galaxy positions only, and then fitting each individual image with models convolved with the appropriate PSF and adding the resulting likelihoods. Images with a mixture of seeing qualities are thus optimally combined for the shape measurement.

[2] The code LENSFIT is available on request from the authors: modification to the data input stages is likely to be required for any particular survey.
[3] http://www.fftw.org.

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022

5.4 Weak lensing from radio interferometer data

It is clear that in large future optical surveys systematic uncertainties in PSF correction will be a dominating concern, indeed this is a significant factor in the case for space-based weak lensing missions. Ground-based optical PSFs vary temporally and very often on spatial scales comparable to those on which the cosmic shear signal is detectable. Even HST lensing studies suffer significantly from PSF variation (Rhodes et al. 2007; Schrabback et al. 2007). In principle radio interferometers have precisely known PSFs, being determined by the antenna positions (note that full 3D knowledge of antenna positions is required, to allow for curvature of the Earth and natural height variations). The PSF varies with hour angle and declination, but in a completely deterministic way. Other effects such as bandwidth and sampling-time smearing can also be precisely computed and incorporated into the shape measurement process (Chang, Refregier & Helfand 2004). Because interferometer measurement are made in the Fourier domain, and because the noise also originates in that domain (being associated with individual antennas) it makes sense to measure galaxy shapes in Fourier space (in the image plane the noise is correlated between pixels, effectively being also convolved with the PSF). Chang & Refregier (2002) and Chang et al. (2004) have already shown how a shapelets (Refregier 2003; Refregier & Bacon 2003) based approach can be extended to the Fourier domain. The Bayesian algorithm presented in this paper already operates in the Fourier domain, so it should be easily adapted for radio interferometer data, which will be particularly relevant for future deep radio surveys such as those proposed for the Square Kilometre Array.

6 CONCLUSIONS

We have argued that a model-fitting approach to galaxy shape measurement should provide an optimum approach to shape measurement for large weak lensing surveys, with the advantages that the S/N of the shape measurement should be optimized and random measurement errors can be estimated. We have further argued that a Bayesian estimation process allows unbiased shape estimation to be made, although even in a realistic implementation of a Bayesian method there is a bias in recovered shear values introduced by the presence of the prior probability distribution. This bias may be calculated from the measured likelihood surfaces, however, and in this paper we have spent some time discussing the calculation of the 'shear sensitivity'. Overall this approach to shape measurement should provide a framework for shear measurement that does not need external calibration by comparison with simulations.

A traditional disadvantage of model fitting is that it may be computationally time consuming, and in this paper we present a fast algorithm for measuring the shapes of individual galaxies. The algorithm makes use of analytic marginalization over surface brightness amplitude, and by working in Fourier space enables rapid marginalization over galaxy position. The algorithm has been tested and has an adequate speed on current generations of computers for use with large ongoing and planned weak lensing surveys. Close pairs of galaxies are not treated by the algorithm, but provided such close pairs can be identified in the data a separate fitting process may be applied to those.

The Bayesian method and fast-fitting algorithm have been tested on simulated galaxies created for the STEP (Heymans et al. 2006; Massey et al. 2007b) and promising results on the measurement of individual galaxy ellipticities have been obtained. A companion paper (Kitching et al., in preparation) will test the measurement of weak lensing shear in the STEP simulations.

ACKNOWLEDGMENTS

CH acknowledges the support of a European Commission Programme 6th framework Marie Curie Outgoing International Fellowship under contract M01F-CT-2006-21891, and a CITA National Fellowship. LVW is supported by NSERC, CIfAR and CFI. TDK is supported by the Science and Technology Facilities Council, research grant number E001114.

REFERENCES

Bacon D. J., Taylor A. N., 2003, MNRAS, 344, 1307
Bardeau S., Kneib J.-P., Czoske O., Soucail G., Smail I., Ebeling H., Smith G. P., 2005, A&A, 434, 433
Bardeau S., Soucail G., Kneib J.-P., Czoske O., Ebeling H., Hudelot P., Smail I., Smith G. P., 2007, A&A, 470, 449
Benjamin J. et al. 2007, MNRAS, preprint (astro-ph/0703570)
Bernstein G. M., Jarvis M., 2002, AJ, 123, 583
Bertin E., Arnouts S., 1996, A&AS, 117, 393
Brainerd T. G., Blandford R. D., Smail I., 1996, ApJ, 466, 623
Bridle S., Kneib J.-P., Bardeau S., Gull S., 2002, in Natarajan P., ed., Proceedings of the Yale Cosmology Workshop, 'The Shapes of Galaxies and Their Dark Matter Halos', New Haven, Connecticut, USA, 28–30 May 2001. World Scientific Press, Singapore
Chang T.-C., Refregier A., 2002, ApJ, 570, 447
Chang T.-C., Refregier A., Helfand D. J., 2004, ApJ, 617, 794
Dunlop J. S., McLure R. J., Kukula M. J., Baum S. A., O'Dea C. P., Hughes D. H., 2003, MNRAS, 340, 1095
Ebbels T., Kneib J.-P., Ellis R. S., 1999, in Sato K., ed., Proceedings of the 183rd symposium of the International Astronomical Union held in Kyoto, Japan, August 18–22, 1997, Cosmological Parameters and the Evolution of the Universe. Kluwer, Dordrecht, Boston, p. 247
Edmondson, E. M., Miller, L., Wolf, C., 2006, MNRAS, 371, 1639
Frigo M., Johnson S. G. 2005, Proc. IEEE, 93, 216
Heavens A., 2003, MNRAS, 343, 1327
Heavens A. F., Kitching T. D., Taylor A. N., 2006, MNRAS, 373, 105
Heymans C. et al. 2006, MNRAS, 368, 1323
Hu W., 1999, ApJ, 522, 21
Jing Y. P., Zhang P., Lin W. P., Gao L., Springer V., 2006, ApJ, 640, L119
Kaiser N., 2000, ApJ, 537, 555
Kaiser N., Squires G., Broadhurst T., 1995, ApJ, 449, 460
Kitching T. D., Heavens A. F., Taylor A. N., Brown M. L., Meisenheimer K., Wolf C., Gray M. E., Bacon D. J., 2007, MNRAS, 376, 771
Kneib J.-P. et al., 2003, ApJ, 598, 804
Kuijken K., 1999, A&A, 352, 355
Kuijken K., 2006, A&A, 456, 827
Lambas D. G., Maddox S. J., Loveday J., MNRAS, 258, 404
Luppino, G. A., Kaiser, N., 1997, ApJ, 475, 20
Massey R. et al., 2007a, ApJS, 172, 239
Massey R. et al., 2007b, MNRAS, 376, 13
Peng C. Y., Ho L. C., Impey C. D., Rix H.-W., 2003, AJ, 124, 266
Refregier A., 2003, MNRAS, 338, 35
Refregier A., Bacon, D., 2003, MNRAS, 338, 48
Rhodes J. D. et al., 2007, ApJ, ApJS, 172, 203
Schrabback T. et al., 2007, A&A, 468, 823
Schramm T., Kayser R., 1995, A&A, 299, 1
Seitz C., Schneider P., 1997, A&A, 318, 687
Spergel D. et al., 2007, ApJS, 170, 377
Taylor A. N., Kitching T. D., Bacon D. J., Heavens A. F., 2007, MNRAS, 374, 1377
Tyson J. A., Wenk R. A., Valdes F., 1990, ApJ, 349, L1
White M., 2004, Astropart. Phys., 22, 211
Zhan H., Knox L., 2004, ApJ, 616, L75

This paper has been typeset from a TEX/LATEX file prepared by the author.

Downloaded from https://academic.oup.com/mnras/article/382/1/315/984384 by guest on 30 April 2022