

Galaxy shear estimation from stacked images

Antony Lewis^{1*}

¹*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK.*

18 May 2009

ABSTRACT

Statistics of the weak lensing of galaxies can be used to constrain cosmology if the galaxy shear can be estimated accurately. In general this requires accurate modelling of unlensed galaxy shapes and the point spread function (PSF). I discuss suboptimal but potentially robust methods for estimating galaxy shear by stacking images such that the stacked image distribution is closely Gaussian by the central limit theorem. The shear can then be determined by radial fitting, requiring only an accurate model of the PSF rather than also needing to model each galaxy accurately. When noise is significant asymmetric errors in the centroid must be corrected, but the method may ultimately be able to give accurate un-biased results when there is a high galaxy density with constant shear. It provides a useful baseline for more optimal methods, and a test-case for estimating biases, though the method is not directly applicable to realistic data. I test stacking methods on the simple toy simulations with constant PSF and shear provided by the GREAT08 project, on which most other existing methods perform significantly more poorly, and briefly discuss generalizations to more realistic cases. In the appendix I discuss a simple analytic galaxy population model where stacking gives optimal errors in a perfect ideal case.

1 INTRODUCTION

Gravitational lensing of light from distant galaxies causes the shape of the galaxies to be distorted in a way that depends on the transverse gradients of the gravitational potential along the line of sight (Bartelmann & Schneider 2001). If the distortion can be measured accurately, it gives a constraint on the lensing potentials, and hence with large enough number of samples on the geometry and distribution of perturbations in the universe. Since the galaxy shapes vary greatly, this can only be done by analysing a very large number of galaxies, with galaxies that are sufficiently well separated that their intrinsic shape correlations can be modelled out or is small. The galaxies can then be assumed to be independent, so that any shape correlation is due entirely to lensing. The task is to find a way to estimate the lensing distortion, which can then be used to extract statistical results from an ensemble of galaxy images.

At leading order the main observable distortion is that of galaxy shear. As discussed further below, if we could observe the galaxies directly, fitting any sheared profile to each galaxy will give an unbiased estimator of this shear. The problem is however much more complicated, because in practice we can only measure the shape after convolution with the point spread function (PSF) of the instrument (e.g. due to atmospheric fluctuations and instrumental imperfections), and image pixelization. The levels of shear that are expected — a few percent — are comparable to those of typical PSFs, so the PSF must be modelled very accurately in order to isolate the cosmological signal. Since the PSF breaks the symmetries of the problem, in general this requires accurate modelling of both the unlensed galaxy shapes and the PSF. Finding methods of doing this that work to the required level of precision is an active area of current research. At the moment it unclear whether it is even possible to get useful high-precision shear constraints in the presence of

realistic ground-observation PSFs, or whether in fact there are unavoidable degeneracies with galaxy shapes and PSF modelling uncertainties. The correct statistical error on the shear measurement could also be too large for the number of observable galaxies to produce precision constraints.

In this paper I re-visit an old sub-optimal method of Kuijken (1999): stacking galaxy images. If the intrinsic galaxy shapes are uncorrelated, a stacked unlensed image should have circular symmetry. Since convolution is a linear operation, the observed stacked image should then be a PSF-convolved sheared version of a circularly symmetric average galaxy. If the PSF is known, the only modelling uncertainties are then in the averaged galaxy profile, which should be well determined by the data. Furthermore, under fairly general conditions a sum of independent samples should have a close-to-Gaussian distribution by the central limit theorem, so the statistics of the stacked image is known without needing to know anything about the distribution of individual galaxy shapes. Fitting a radial profile and shear to the data with a Gaussian error model gives an estimate of the shear that should be very independent of the actual distribution of galaxy shapes. The method therefore provides a useful baseline for comparing future more optimal methods that incorporate accurate modelling of individual galaxies.

In practice of course things are not so simple. To start with the shear and PSF are not expected to be constant, so any stacked galaxy image has to be interpreted with care. In addition, in the presence of noise, the process of stacking can itself produce biases since we cannot determine the centroid accurately: any shear- or PSF-correlated misalignments in the stacking procedure will introduce biases.

Given the complexity of the general problem, the lensing community has helpfully boiled the issues down into a series of much

simpler problems (Heymans et al. 2006; Massey et al. 2007; Bridle et al. 2008). Although existing methods perform adequately for current and near-future data, even in these highly simplified cases they are known to be inadequate for future surveys. I therefore focus on these simplified problems to try to isolate the important issues, in particular I shall assume the PSF is well measured from many low-noise star images and that shear is constant. If no methods works accurately even on this very simple toy problem, then that is clearly sufficient to show that ground-based weak lensing surveys with similar PSFs will be of no use for precision cosmology (i.e. future cosmological parameter constraints at the percent level or better). On the other hand if sufficiently accurate methods can be developed, the next task will be to make them applicable to more realistic situations where the PSF is likely to vary significantly and the shear has a realistic spatial correlation function. Space-based observations typically have rather different PSFs and would require a separate study.

I start by reviewing the case of shape estimation when there is no PSF, and then briefly explain why introducing a PSF qualitatively increases the complexity of the problem. I then move on to show that stacking can work well with low-noise simulations, and discuss various issues to do with pixel-scale stacking, centroid errors and non-constant PSFs. I test stacking methods on the GREAT08¹ simulations (Bridle et al. 2008) and show that it performs well compared to other existing methods, most of which involve modelling unlensed galaxy shape distributions with something that is known to be incorrect. Unlike other existing methods the stacking method is not directly applicable to more realistic data, but may be useful to motivate more general approaches.

2 FITTING GALAXIES AND SHEAR

2.1 Shear fitting with no PSF

At lowest order in the gravitational potentials weak lensing causes position \mathbf{x}_u on the unlensed image to be related to the corresponding position \mathbf{x}_l on the lensed image by

$$\mathbf{x}_u = \mathbf{S}\mathbf{x}_l \equiv \begin{pmatrix} 1 - g_1 & -g_2 \\ -g_2 & 1 + g_1 \end{pmatrix} \mathbf{x}_l, \quad (1)$$

where the components of the shear matrix g_1 and g_2 are the reduced shear in some coordinate system. For the purposes of this paper we can neglect the uniform convergence which is degenerate with the galaxy size and assume g_1, g_2 are constant across each galaxy image. For a thorough introduction to weak lensing see Bartelmann & Schneider (2001); Lewis & Challinor (2006).

Consider fitting a model $m(\mathbf{S}_m\mathbf{x}, \theta)$ to an unlensed perfect galaxy image $I_u(\mathbf{x})$, with model parameters θ and shear matrix \mathbf{S}_m . For example a simple least-squares fit would solve

$$\begin{aligned} \frac{\partial}{\partial \theta} \int d^2\mathbf{x} |I_u(\mathbf{x}) - m(\mathbf{S}_m\mathbf{x}, \theta)|^2 &= 0 \\ \frac{\partial}{\partial \mathbf{S}_m} \int d^2\mathbf{x} |I_u(\mathbf{x}) - m(\mathbf{S}_m\mathbf{x}, \theta)|^2 &= 0. \end{aligned} \quad (2)$$

Assuming there is a unique solution $\mathbf{S}_m = \mathbf{S}_0, \theta = \hat{\theta}$, the lensed image $I_l(\mathbf{x}) = I_u(\mathbf{S}\mathbf{x})$ would then be fit by $m(\hat{\mathbf{S}}\mathbf{x}, \hat{\theta})$ where $\hat{\mathbf{S}} = \mathbf{S}_0\mathbf{S}$. The best-fit unlensed shear matrix \mathbf{S}_0 is determined by the particular galaxy and model. The key assumption in galaxy

weak lensing is that the galaxy shapes are statistically isotropic, in other words that versions of each unlensed galaxy rotated by different angles (or flipped) are equally likely. So $\mathbf{R}^T\mathbf{S}_0\mathbf{R}$ is just as likely as \mathbf{S}_0 for some rotation matrix \mathbf{R} . Taking \mathbf{R} to be a rotation by 90° , on average over many galaxy orientations we have, $\langle \mathbf{S}_0 \rangle = \frac{1}{2}(\mathbf{S}_0 + \mathbf{R}^T\mathbf{S}_0\mathbf{R}) = \mathbf{I}$, and hence $\langle \hat{\mathbf{S}} \rangle = \mathbf{S}$: the shear matrix estimator is unbiased. Note that this is entirely independent of how well $m(\mathbf{S}_m\mathbf{x}, \theta)$ actually fits the unlensed galaxy, so in the idealized case we could fit any model we like to galaxy shapes and still on average get the correct answer. This will remain true for best-fits to more general log-likelihoods of the form

$$\chi^2 = \int d^2\mathbf{x} (I_l(\mathbf{x}) - m(\mathbf{S}_m\mathbf{x}, \theta))^T [N(I_l(\mathbf{x}), \mathbf{S}_m\mathbf{x})]^{-1} \times (I_l(\mathbf{x}) - m(\mathbf{S}_m\mathbf{x}, \theta)), \quad (3)$$

i.e. where the noise depends only on the lensed galaxy intensity or follows the alignment of the galaxy model. Similarly for generalizations with correlated noise.

2.2 Shear fitting with a PSF

Unfortunately we cannot observe lensed galaxies directly, but only after convolution with an instrumental point spread function and pixelization. Pixelization can be thought of as an additional contribution to the PSF, typically a convolution with a square window function, followed by sampling at the pixel centres. I shall discuss the PSF in this generalized sense, so that the observational data consists of a set of regularly-spaced samples of a PSF and pixel-convolved galaxy image. The noise-free observed value at position \mathbf{x} on the image plane is then

$$I_o(\mathbf{x}) = \int d^2\mathbf{y} P(\mathbf{x} - \mathbf{y}) I_l(\mathbf{y}), \quad (4)$$

where $P(\mathbf{x})$ is the total PSF, or simply $I_o = P \star I_l$. If we know the PSF function, we can fit a PSF-convolved galaxy model to the data; for example a least-squares solution would minimize

$$\begin{aligned} \chi^2 &= \int d^2\mathbf{x} \left(I_o(\mathbf{x}) - \int d^2\mathbf{y} P(\mathbf{x} - \mathbf{y}) m(\mathbf{S}_m\mathbf{y}, \theta) \right)^2 \\ &= \int d^2\mathbf{x} \left[\int d^2\mathbf{y} P(\mathbf{x} - \mathbf{y}) [I_u(\mathbf{S}\mathbf{y}) - m(\mathbf{S}_m\mathbf{y}, \theta)] \right]^2. \end{aligned} \quad (5)$$

If $\mathbf{S}_m = \mathbf{S}_0, \theta = \hat{\theta}$ is the best fit when there is no lensing, due to the position dependence of the PSF it is no longer in general the case that $\mathbf{S}_m = \mathbf{S}_0\mathbf{S}, \theta = \hat{\theta}$ is the best fit to the lensed image. Hence unlike in the case with no PSF, there is no longer any guarantee that fitting is giving an unbiased estimate of the shear. The only general exception is if the model fits the galaxy exactly, so the best fit has $I_u(\mathbf{S}\mathbf{y}) = m(\mathbf{S}_m\mathbf{y}, \theta)$, in which case fitting is giving the right answer on average independent of the known PSF.

Extracting unbiased shear constraints by model fitting in the presence of a PSF therefore in general requires modelling the galaxies accurately. This poses several significant problems. A large galaxy lensing survey will have most of its galaxies near the edge of its resolution, therefore there is typically only limited high-quality data to constrain the properties of the bulk of the galaxies in the selection function. A general Bayesian model could use information from some well-resolved galaxies, and fit a general model for uncertainties, but given the large variation in galaxy alignment with respect to the line of sight, and wide variations in the intrinsic shapes, a general model is likely to involve a large number of parameters and require many images to constrain well. The

¹ <http://www.great08challenge.info/>

galaxy model can also be constrained to some extent using all of the observed galaxies. But if the parameters can not all be well constrained by the data, it may be essential that the priors accurately represent the galaxy distributions in order to get unbiased answers. In addition any model with large numbers of parameters per galaxy is likely to become numerically time consuming. For an excellent discussion of many related issues see Bernstein & Jarvis (2002); Hirata & Seljak (2003), and a summary of other existing methods in Bridle et al. (2008). For promising recent results on Bayesian model fitting see Miller et al. (2007); Kitching et al. (2008), however the galaxy model used in this method is still unrealistic and results, though significantly better than many other methods, are still not good enough for high-precision cosmology (see Sec. 4).

3 SHEAR FITTING STACKED GALAXIES

General galaxy fitting should provide the best constraints on the shear. However given the problems outlined above, and given potential difficulties in knowing whether the modelling is accurate enough, it would be useful to have a simple less-optimal but more robust shear-estimation method that is more directly independent of the details of the galaxy distribution. In simple test cases this would be a useful cross-check, and provide a baseline for the levels of residual noise that better methods should be able to beat.

The method I shall focus on simply stacks the galaxies, and then fits a sheared average galaxy model to the stacked image, following Kuijken (1999). If the PSF is known, this should give unbiased results conditional only on being able to stack in an unbiased way, and being able to model the radial profile of the averaged unlensed galaxy accurately. A 1-dimensional radial model is clearly much easier to fit than a full 2D galaxy shape distribution, and since the average galaxy is expected to have a smooth radial profile only a modest number of parameters should be required. These parameters are likely to be well constrained with a reasonable number of galaxies (and hence the results fairly independent of the priors).

In order to stack galaxies images, we need to be able to define a rule for the relative galaxy alignment, e.g. by defining a centroid in each image and then stacking the images so that their centroids are aligned. Assuming this can be done, we then have an observed stack of N galaxy images

$$\hat{\mathbf{I}}_o(\mathbf{x}) \equiv \frac{1}{N} \sum_{i=1}^N \beta_i I_{o,i}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N [\beta_i P_i \star I_{i,i}](\mathbf{x}), \quad (6)$$

where β_i are some weights and P_i the PSF on galaxy i . Assuming the PSF is independent of the galaxy shape the expected value of the stacked image is

$$\langle \hat{\mathbf{I}}_o \rangle = \frac{1}{N} \sum_{i=1}^N [P_i \star \langle \beta_i I_{i,i} \rangle] = \left(\frac{1}{N} \sum_{i=1}^N P_i \right) \star \bar{I}_{i,\beta} = \bar{P} \star \bar{I}_{i,\beta}, \quad (7)$$

where $\bar{I}_{i,\beta}$ is the average of a weighted galaxy. By symmetry, taking \mathbf{x} to have origin at the centroid, in the unlensed case $\bar{I}_{u,\beta}(\mathbf{x}) = \bar{I}_{u,\beta}(|\mathbf{x}|)$. The average PSF \bar{P} — including pixelization — is precisely what is observed from a large statistically equivalent set of star images (assuming stars are point sources and have the same PSF as the galaxies).

Assuming the weights are independent of the shear and the shear is constant, the expectation of the stacked image is a sheared circularly-symmetric averaged galaxy, convolved with an average PSF. We can therefore proceed to fit a model to the observed stacked image, and if the radial profile can be fit accurately the

method should be unbiased. In the appendix I discuss a simple analytic galaxy population model in which, for the ideal noise-free case, stacking with an appropriate weighting is in fact optimal.

In the presence of noise, and with finite N so that there is dispersion about the expectation value, we need an error model. One benefit of using stacked images is that this is well defined: assuming each galaxy is independent, the distribution of $\hat{\mathbf{I}}_o$, a sum of many independent galaxy samples, should be nearly Gaussian by the central limit theorem.

In fact we can apply any linear function to $\hat{\mathbf{I}}_o$, and the distribution will still be Gaussian with expectation given by the equivalent linear function of the averaged convolved galaxy. This can be useful for data compression, e.g. to re-pixelize, or expand in moments, etc, anything that's likely to encapsulate most of the useful information in fewer numbers. If the stacked image is generated at much higher pixel sampling than the original image, there will be a large number of pixel values and hence a huge number of galaxies required for the covariance estimate to be accurate. Also since the noise is correlated on the scale of the original pixel size, the covariance would be singular, so applying some linear re-pixelization or other linear compression matrix \mathbf{M} can be useful. Writing the set of sampled \mathbf{x} values as a vector, for a data vector $\mathbf{X} = \mathbf{M}\hat{\mathbf{I}}_o$ the covariance $\mathbf{C}_X \equiv \langle \mathbf{X}\mathbf{X}^T \rangle$ can be estimated from N galaxy samples as

$$\hat{\mathbf{C}}_X = \frac{1}{N-2} \sum_i \mathbf{M}(\beta_i \mathbf{I}_{o,i} - \hat{\mathbf{I}}_o)(\beta_i \mathbf{I}_{o,i} - \hat{\mathbf{I}}_o)^T \mathbf{M}^T \quad (8)$$

(for large N , $N \gg \dim(\mathbf{X})$). The likelihood as a function of parameters θ and shear matrix \mathbf{S} can then be approximated as

$$-2 \ln \mathcal{L}(\mathbf{S}, \theta) \sim [\hat{\mathbf{I}}_o - \mathbf{m}_o(\mathbf{S}, \theta)]^T \mathbf{M}^T \hat{\mathbf{C}}_X^{-1} \mathbf{M} [\hat{\mathbf{I}}_o - \mathbf{m}_o(\mathbf{S}, \theta)] \quad (9)$$

where $\mathbf{m}_o(\mathbf{S}, \theta) = \bar{P} \star m(\mathbf{S}, \theta)$ is the model for the average PSF-convolved sheared circularly-symmetric galaxy. The simplest thing is to take \mathbf{M} to just re-pixelize, e.g. at the original pixel resolution. For simple PSFs it also loses little information to halve the number of points by taking \mathbf{M} to sum $I(\mathbf{x})$ and $I(-\mathbf{x})$ since this includes the shear and radial information, but ignores irrelevant dipole fluctuations. Note that \mathbf{C}_X includes variance due to both noise and 'shape noise' due to the differences in galaxy shapes. The latter is expected to be spatially correlated even if the noise is not, but in any case the central limit theorem result straightforwardly accounts for any correlated or non-Gaussian noise in individual images.

In the limit that the model fits the stacked image exactly, and the stated assumptions are met, the fitting procedure should be unbiased. However due to noise this will not quite be the case, so there is potentially a source of noise-bias through the PSF and shear dependence of the estimated covariance.

Note that even if the instrumental PSF is actually constant, the PSF for points in the stacked image plane varies from galaxy to galaxy due to the offset between the centres of the pixels in each image and the centre of the stacked image. If the stacked image is pixelized at higher resolution, then there are different PSFs for each non-equivalent high-resolution pixel centre. However the averaged PSF will be the same for each high-resolution pixel. The high resolution pixels are of course strongly correlated due to the pixelization of the individual images.

3.1 Centroid issues

Centroid errors on the galaxy image plane are harmless (other than increasing the error bars), since they merely effect the average

galaxy profile. For example we could consider defining the centroid in terms of a random displacement from the centre of light in the unlensed galaxy, which is perfectly legitimate. The problem is that in general the centroid error will depend on the galaxy shape, and hence also shear and PSF in non-trivial way. Since the centroid of a long thin shape is hard to determine in the long direction, the centroid error is typically strongly correlated with the shape of the galaxy; if the galaxies have a net ellipticity in one direction due to the PSF, the centroids will tend to have a net dispersion aligned with the PSF. This has the effect of making a naively stacked image give results biased in the direction of the PSF. There are similar effects due to shear. When the centroid error is not negligible compared to the galaxy sizes, the centroid error must be accounted for somehow in order to get unbiased results. In general this is difficult, though an approximate correction may be sufficient.

Two simple approaches immediately present themselves. We could simply attempt to model the effective centroid-error PSF and include it as part of the effective PSF on the stacked image. Or the centroid error could be modified to remove some of the sources of bias. The latter approach is likely to be more straightforward, if less optimal.

As a crude first attempt to remove the leading-order centroid bias I simply add Gaussian noise to each centroid so that the total centroid dispersion is approximately circularly symmetric. To do this I fit a 6-parameter Gaussian elliptical model to each observed (PSF-convolved) galaxy to determine the centroid, calculating the Hessian errors by numerical differentiation and then inverting to get an approximate centroid error matrix. Then to each estimated centroid I add a small Gaussian displacement in a direction chosen such that the total centroid error is then isotropic. If the estimate of the centroid error on each galaxy is fairly accurate, this should remove the correlation of centroid dispersion with galaxy alignment, and hence reduce the PSF bias. However the magnitude of the centroid error will still depend on the PSF-convolved sheared galaxy shape, and hence potentially lead to residual biases. Furthermore if the total centroid dispersion is accounted for by allowing the average galaxy profile to change, the centroid error really has to be sheared like the rest of the galaxy shape; a better approach could therefore use an approximate estimate of the shear to ensure that the total centroid error on the image plane is sheared approximately correctly².

The centroid determined by Gaussian model fitting that I use seems to have about 10% less dispersion than that obtained using adaptive moments (Bernstein & Jarvis 2002), however in the presence of more complicated PSF (e.g. with a dipole) the position of the centroid could be biased, so a more sophisticated method may be required. Unfortunately to get the centroid error correct in general requires modelling the shape of each galaxy correctly, which is just as hard as the original shape estimation problem. However as long as the centroid error is small compared to the size of the galaxies, an approximate correction may be sufficient. Indeed the output from a more realistic galaxy fitting code like *Lensfit* (Kitching et al. 2008) might be a good place to start trying to improve the crude Gaussian model used here. Simulations may also be reliable enough to find a fudge parameter to relate the estimated centroid error to the true centroid error to the required accuracy.

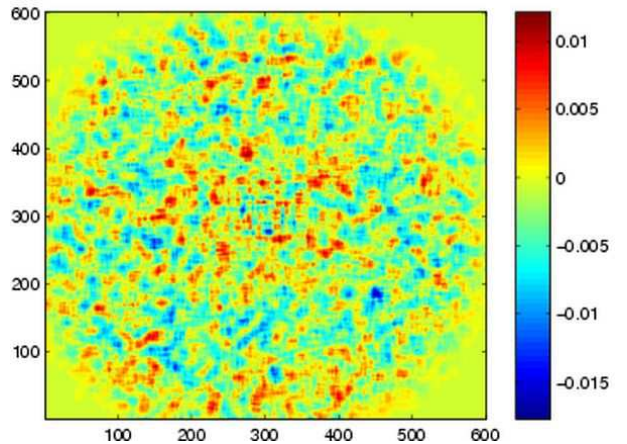


Figure 1. Typical residuals after fitting a unit-amplitude sheared circularly-symmetric galaxy to a stacked image. Note errors are correlated due to pixel-scale stacking correlations and correlated shape noise. Here pixels have been added at 15 sub-pixel resolution, less than needed for accurate results.

4 RESULTS WITH SIMULATIONS

I test the galaxy stacking method on simulations provided by the GREAT08 project (Bridle et al. 2008). These satisfy the required assumptions, in that the shear is constant over a larger number of galaxy images. The simulations also have constant (very simple) PSF, with a large number of low-noise star images so that the PSF can be determined essentially exactly. The PSF is anisotropic but has no dipole moment and the isophotes have the same shape at each radius; it is therefore a rather special case and is likely to be unrealistic in several important respects. Nonetheless, even with these radical simplifications from reality, most existing shear-estimation methods fail to produce results at an accuracy required for precision cosmology, so it makes an interested test case.

The disadvantage of non-optimal methods such as stacking is there are lots of free parameters, e.g. choice of β_i and radial fitting function, choice of the M reduction matrix, as well as resolution parameters governing the stacking. The weights β_i must be chosen in a shape-independent (or at least alignment-independent) manner, otherwise biases may be introduced. I take β_i to constant or inversely proportional to the integrated signal in each images (so that the average galaxy is then independent of the magnitude distribution of the galaxies); see the Appendix for a discussion of the optimal weighting in an idealized case. For noisy images this should probably be modified by an estimate of the signal to noise ratio to down-weight noise-dominated images. I chose M simply to re-pixelize the stacked image to the resolution of the original galaxies.

Parameterizing radial distribution of $m(S, \theta)$ using splines is convenient, so θ is a set of values at some radial spline nodes. Splines naturally have multiple resolution: e.g. we can do a quick fit with a few spline points, then increase the number of spline parameters to refine the result. This could be done in an adaptive way to make sure the data is fit but not over-fit. I simply choose to spline in the log of the radial amplitude, using 12 spline points over a radius range of 8 pixel units for fitting the stacked image, with stacked-image resolution 1/31 or 1/41 of the original pixel size (with initial fit at 1/9 resolution with 7 spline points to get close to the best-fit point quickly). The fitting could be done with MCMC to get accurate error bars, though at a quick look does not show evidence

² Thanks to Gary Bernstein for pointing out this issue.

of strong degeneracies or asymmetries in the error bars. So finding just the best fit is a reasonable first step, with errors approximated from a Hessian if required. To find the best-fit point I use the NEWUOA³ algorithm, which can be significantly faster than ‘AMOEBa’ downhill-simplex method (Nelder & Mead 1965) in many cases, though I need to be a bit careful to avoid local minima. The resulting reduced chi-squared is generally less than one, indicating that indeed the galaxy is being fit accurately, though values are hard to assess due to the non-realistic mirroring procedure used in the GREAT08 simulations to help reduce shape noise. Typical residuals are shown in Fig. 1.

For noisy images the centroid error needs to be corrected as discussed in the previous section. Comparison with the test simulation indicates that the centroid variance is underestimated by about a factor of a half, so I adopt a centroid-error fudge parameter $\alpha = 1.5$ (chosen to work with the test simulations), and assume that the actual centroid covariance on each galaxy is αC where C is estimated from the Hessian about the best-fit Gaussian model.

Accuracy of results for the purpose of GREAT08 is defined by a quality parameter Q (Bridle et al. 2008), so that the shear variance is

$$\langle (\hat{g}_1 - g_1)^2 + (\hat{g}_2 - g_2)^2 \rangle = \frac{10^{-4}}{Q}, \quad (10)$$

where \hat{g}_i is the shear estimated from a plate of 10000 galaxy images at constant shear and the same PSF, g_i is the true shear, and $\langle \hat{g}_i \rangle$ is estimated from an ensemble of different simulated plates with the same shear. For noisy simulations results are quoted for Q estimated from the expectation value from a set of simulations with different PSF and true shears. The target for future observations is $Q \sim \mathcal{O}(1000)$ (Amara & Refregier 2008), and current methods typically give $Q \lesssim \mathcal{O}(100)$. Biases on g_i therefore need to be below $\lesssim 3 \times 10^{-4}$ level, or a typical fractional shear error of less than about a percent. For low-noise images the definition is simply to take each plate separately

$$\hat{Q} = \frac{10^{-4}}{\langle (\hat{g}_1 - g_1)^2 + (\hat{g}_2 - g_2)^2 \rangle_{\text{plates}}}, \quad (11)$$

with the stacking method described here giving $\hat{Q} \sim 300$. The errors have a contribution from any systematics error and intrinsic shape noise (which may be significantly higher than possible due to the lossy nature of the stacking procedure). Most methods used with current data give $Q \lesssim 30$.

When the noise is significant the method is no longer strictly valid due to centroid issues, however using the centroid error correction described above still gives $\hat{Q} \sim 130$, which is as good as other existing methods at the time of this work, and within a factor of two of the best method eventually winning the GREAT08 challenge. However the stacking method is more reliant on the non-realistic constant-shear assumption than some other methods, so the main use may be as a baseline for simulation-based comparisons with better codes.

The fact that $Q \gtrsim 100$ can be obtained by this sub-optimal method, making essentially no assumptions about the galaxy distribution, is perhaps encouraging evidence that there will exist a better method that is good enough for precision cosmology using only modest assumptions about the galaxy distribution. There is some evidence for shear-calibration bias in the stacking results, with a

tendency for $|g|$ to be too large. More careful modelling of the centroid error, for example using better model fitting and an iterative shear estimate, could probably reduce the systematic error. Suitable time-consuming adjustment of the method parameters may also allow the method to perform significantly better.

5 CONCLUSIONS

I have revisited the simple shear-estimation stacking method of Kuijken (1999), and shown that it still makes a useful baseline that can compare favorably with currently used methods in idealized cases. Although it only works straightforwardly over regions with constant shear, it can be a useful test case, and help to understand possible sources of bias in other methods. Stacking has the advantage of giving results that are unbiased almost independent of the unknown distribution of unlensed galaxy shapes. Residual biases enter at a lower level, for example through correlations of the centroid error with galaxy shape. With low noise the method can produce accurate results, comparing favourably with methods that fit galaxy models that are known to be unrealistic. This should be unsurprising: Bayesian methods generally give the right results only if the correct model is used and priors truly reflect beliefs. Only in the very special case of no observational PSF does fitting any model to galaxy shapes give unbiased answers; a general PSF breaks all the symmetries, requiring accurate modelling of both the PSF and the unlensed galaxy shape distribution to get the right result.

Even if individual noisy galaxies are well fit by a simple galaxy model due to the large noise, if in reality galaxies have significant substructure or un-modelled shape variations, the combined high-precision shear estimate from fitting many galaxies separately may be biased due to the inconsistent shape modelling. It is possible the modelling bias is negligible, but unless carefully proven analytically or demonstrated numerically in realistic simulations it would be safer to assume otherwise. The noise-free stacking procedure is by construction linear in the galaxies, which is why substructure variations between galaxies effectively cancel. However fitting to individual galaxies is usually a non-linear procedure, and there is no reason to expect errors to cancel more generally. Future work may however be able to find fairly model-independent methods that can be applied to fitting individual galaxies, significantly improving on the stacking method both in terms of signal to noise, and in terms of application to more realistic cases. If not, stacking methods may still be useful. Future work could investigate how to apply stacking in more realistic cases where the shear varies from galaxy to galaxy, and the PSF can only be estimated locally with significant noise. At leading order, a fit to a stacked galaxy constructed over a region with small variations in shear should be probing the appropriately averaged shear. With a high galaxy density the corresponding suppression of small-scale power may be acceptable if it can be accurately modelled without significant bias.

6 ACKNOWLEDGEMENTS

I acknowledge a PPARC/STFC Advanced fellowship and thank members of the GREAT08 team, especially Sarah Bridle, for helpful discussions, Lance Miller for comments, Steven Gratton for matrix identity advice, and CITA for use of their old computer.

³ <http://www.inrialpes.fr/bipop/people/guilbert/newuoa/newuoa.html>

APPENDIX A: ANALYTIC EXAMPLE

Here we consider a very simple toy distribution of galaxies where we can attempt to calculate some things analytically. Consider the case where each galaxy has a Gaussian-shaped profile

$$I_u(\mathbf{x}|\mathbf{Q}_i) = A \frac{e^{-\mathbf{x}^T \mathbf{Q}_i^{-1} \mathbf{x}/2}}{|\mathbf{Q}_i|^{w/2}}, \quad (\text{A1})$$

where the distribution of the covariance \mathbf{Q}_i of each galaxy is drawn from a 2-dimensional inverse Wishart distribution (see e.g. Gupta & Nagar (1999) for review and results used below)

$$P(\mathbf{Q}) = \frac{|\Psi|^{(n-3)/2} e^{-\frac{1}{2} \text{Tr}(\Psi \mathbf{Q}^{-1})}}{2^{n-3} \pi^{1/2} \Gamma[(n-3)/2] \Gamma[(n-4)/2] |\mathbf{Q}|^{n/2}} \quad (\text{A2})$$

where $n > 6$. Since we assume the unlensed distribution is statistically isotropic $\Psi_0 = (n-6)\sigma_g^2 \mathbf{I}$ where σ_g is the average galaxy width. The parameter n determines how broad the galaxy shape distribution is, with $n \rightarrow \infty$ corresponding to a distribution of identical circular Gaussian galaxies. Typical galaxy ellipticities are $\mathcal{O}(n^{-1/2})$ with

$$\frac{\langle (Q_{11} - Q_{22})^2 \rangle}{\langle (Q_{11} + Q_{22})^2 \rangle} = \frac{\langle 4Q_{12}^2 \rangle}{\langle (Q_{11} + Q_{22})^2 \rangle} = \frac{1}{n-6}. \quad (\text{A3})$$

The parameter w governs how the magnitude varies, with $w = 0$ corresponding to all galaxies having the same peak amplitude, and $w = 1$ corresponds to them all having equal integrated light.

The averaged galaxy profile (with equal weight) is given by

$$\begin{aligned} \bar{I}_u(\mathbf{x}) &= \int d\mathbf{Q} P(\mathbf{Q}) I_u(\mathbf{x}|\mathbf{Q}) \\ &= \frac{\Gamma[n+w-4]}{\Gamma[n-4] |\Psi|^{w/2}} \frac{A}{(1 + \mathbf{x}^T \Psi^{-1} \mathbf{x})^{(n+w-3)/2}}. \end{aligned} \quad (\text{A4})$$

As expected this becomes the same as the individual galaxy shape as $n \rightarrow \infty$. The covariance can be calculated similarly as

$$\begin{aligned} \text{cov}(\mathbf{x}, \mathbf{y}) &= \int d\mathbf{Q} P(\mathbf{Q}) I_u(\mathbf{x}|\mathbf{Q}) I_u(\mathbf{y}|\mathbf{Q}) - \bar{I}_u(\mathbf{x}) \bar{I}_u(\mathbf{y}) \\ &= \frac{A^2}{\Gamma[n-4] |\Psi|^w} \left[\frac{\Gamma[n+2w-4]}{|\mathbf{I} + \Psi^{-1}(\mathbf{x}\mathbf{x}^T + \mathbf{y}\mathbf{y}^T)|^{(n+2w-3)/2}} \right. \\ &\quad \left. - \frac{\Gamma[n+w-4]^2}{\Gamma[n-4] [(1 + \mathbf{x}^T \Psi^{-1} \mathbf{x})(1 + \mathbf{y}^T \Psi^{-1} \mathbf{y})]^{(n+w-3)/2}} \right]. \end{aligned} \quad (\text{A5})$$

Note that

$$\begin{aligned} |\mathbf{I} + \Psi^{-1}(\mathbf{x}\mathbf{x}^T + \mathbf{y}\mathbf{y}^T)| &= \\ (1 + \mathbf{x}^T \Psi^{-1} \mathbf{x})(1 + \mathbf{y}^T \Psi^{-1} \mathbf{y}) - (\mathbf{x}^T \Psi^{-1} \mathbf{y})^2. \end{aligned} \quad (\text{A6})$$

The covariance is determined by the number of degrees of freedom governing the population, so that with a simple model the number of significantly non-zero eigenvalues is small. In the case here each galaxy shape is determined by the three independent numbers in \mathbf{Q} .

Using this analytic galaxy population model we can compare the errors (e.g. estimating the shear matrix \mathbf{S} such that $\Psi = \mathbf{S}^T \Psi_0 \mathbf{S}$) using stacking compared to what could be done using an optimal analysis. If \mathbf{Q}_i were simply measured directly from each galaxy (in the low-noise limit) then the optimal expected error is

$$\left\langle \sum_i \frac{\partial^2 \ln P(\mathbf{Q}_i)}{\partial g_a \partial g_b} \right\rangle_{g=0}^{-1} = \frac{\delta_{ab}}{4N(n-3)}, \quad (\text{A7})$$

where N is the number of galaxies. The corresponding error per

galaxy is $\sigma_1 = \sigma_2 = 1/(2\sqrt{n-3})$, where σ_i is the error on g_i . For $n = 7$ this corresponds to an error per galaxy $\sigma_i = 0.25$.

Using stacked galaxies with $w = 0$ in fact gives the same average error per galaxy, with the errors increasing only slightly for $w \sim \mathcal{O}(1)$. In this noise-free case with known distributions and centroids, the stacking method is close to optimal. To show that with $w = 0$ stacking gives optimal answers we only need to show that the full likelihood can be written in terms of the stacked image. Since

$$-2 \ln P(\Psi|\{\mathbf{Q}_i\}) = \sum_i [\text{Tr}(\Psi \mathbf{Q}_i^{-1}) - (n-3) \ln |\Psi|] + \text{const}, \quad (\text{A8})$$

where the last term is independent of Ψ , a sufficient statistic is $\sum_i \mathbf{Q}_i^{-1}$. However this can be measured by taking derivatives of the perfect stacked image

$$\hat{I}_u = \frac{A}{N} \sum_i e^{-\mathbf{x}^T \mathbf{Q}_i^{-1} \mathbf{x}} \quad (\text{A9})$$

at the origin, and hence stacking is lossless for measuring the shear in this ideal case. Since the number of degrees of freedom in the galaxy model is small, the stacked image does not actually need to be densely sampled to obtain close to optimal results.

In the zero-noise limit with infinite resolution, a known PSF can simply be deconvolved, so the above results also apply to PSF-smearred noise-free galaxies. Noise can be accounted for by adding an appropriate term to Eq. (A5) if the centroids are known, and will increase the expected error per galaxy. Analysing more realistic cases analytically is challenging.

REFERENCES

- Amara A., Refregier A., 2008, MNRAS, 391, 228, 0710.5171
 Bartelmann M., Schneider P., 2001, Phys. Rept., 340, 291, astro-ph/9912508
 Bernstein G. M., Jarvis M., 2002, AJ, 123, 583, astro-ph/0107431
 Bridle S., et al., 2008, 0802.1214
 Gupta A., Nagar D., 1999, Matrix Variate Distributions. Chapman & Hall, ISBN: 1584880465
 Heymans C., et al., 2006, MNRAS, 368, 1323, astro-ph/0506112
 Hirata C. M., Seljak U., 2003, MNRAS, 343, 459, astro-ph/0301054
 Kitching T. D., Miller L., Heymans C. E., van Waerbeke L., Heavens A. F., 2008, MNRAS, 390, 0802.1528
 Kuijken K., 1999, Astronomy and Astrophysics, 352, 355, astro-ph/9904418
 Lewis A., Challinor A., 2006, Phys. Rept., 429, 1, astro-ph/0601594
 Massey R., et al., 2007, MNRAS, 376, 13, astro-ph/0608643
 Miller L., Kitching T. D., Heymans C., Heavens A. F., van Waerbeke L., 2007, MNRAS, 382, 315, ADS, 0708.2340
 Nelder J. A., Mead R., 1965, The Computer Journal, 7, 308