

Bayesian hierarchical modelling of weak lensing: The golden goal

Alan Heavens*, Justin Alsing, Andrew Jaffe, Till Hoffmann

ICIC (Imperial Centre for Inference and Cosmology)

Imperial College, South Kensington

London SW7 2AZ, U.K.

**E-mail: a.heavens@imperial.ac.uk*

Alina Kiessling

Jet Propulsion Laboratory, California Institute of Technology,

4800 Oak Grove Drive,

Pasadena, CA 91109, USA

Benjamin Wandelt

Institut d'Astrophysique de Paris, UMR CNRS 7095,

Université Pierre et Marie Curie, 98bis boulevard Arago

75014 Paris, France

To accomplish correct Bayesian inference from weak lensing shear data requires a complete statistical description of the data. The natural framework to do this is a Bayesian Hierarchical Model, which divides the chain of reasoning into component steps. Starting with a catalogue of shear estimates in tomographic bins, we build a model that allows us to sample simultaneously from the the underlying tomographic shear fields and the relevant power spectra (*E*-mode, *B*-mode, and *EB*, for auto- and cross-power spectra). The procedure deals easily with masked data and intrinsic alignments. Using Gibbs sampling and messenger fields, we show with simulated data that the large (over 67000-) dimensional parameter space can be efficiently sampled and the full joint posterior probability density function for the parameters can feasibly be obtained. The method correctly recovers the underlying shear fields and all of the power spectra, including at levels well below the shot noise.

Keywords: Cosmology; statistics; gravitational lensing.

1. The risks of future experiments

Cosmology has often been described in recent years as moving into the *precision era*, where error bars at percent level are commonplace in the cosmic microwave background. In future, this precision will increase enormously in related fields, such as large-scale structure, weak gravitational lensing, and, in due course, 21cm studies, with ambitious observational programmes connected with Euclid, the Large Synoptic Survey Telescope, and the Square Kilometre Array. The surveys associated with these facilities will deliver exquisite precision, which will help to answer many of the major cosmological questions of the decade, such as the nature of Dark Energy, possible beyond-Einstein gravity, macroscopic extra dimensions, and the masses and hierarchy of neutrinos. However, with this precision comes risk: the small statistical error bars that come from the large datasets involved will lead to potentially far-reaching and erroneous conclusions if the *accuracy* of the results is not correspondingly high. To achieve this requires very good control of systematic effects

at all stages of the experiment, including in the statistical analysis of the data. In this report, we focus on weak gravitational lensing, and propose a rigorous statistical framework that can encapsulate essentially everything the data tell us, allowing for accurate and suitably precise inferences to be made, with full propagation of errors.

2. The goal of scientific inference

Let us focus on the scientific method. A very general description of an experiment is that there is some prior information I , some new data d , and one or more models M with parameters θ . We assume we wish to infer the parameters of a given model in the light of the new data collected by the experiment. From a Bayesian perspective, everything we know is encapsulated in the posterior probability,

$$p(\theta|d, I) \propto p(d|\theta, I)p(\theta|I)$$

where we have used Bayes' theorem to write it in terms of the likelihood and the prior, and the dependence on the model has been suppressed. This is it; one can if desired form *estimates* of the parameters from this, but that is not necessarily useful, and one can from this form Bayesian credible intervals for the parameters. The posterior is the desired outcome of the experiment — it is in a sense the ‘golden goal’ of statistical analysis.

2.1. Computing the posterior

Inevitably, for anything other than a very simple experiment, the posterior probability is not an analytic function, and cannot be computed directly, and the most common general technique is to draw *samples* from the posterior (or sometimes the likelihood). Asymptotically, methods such as Markov Chain Monte Carlo will sample the target distribution with a density of samples that is proportional to the target, so with a sufficient number of samples, the chain characterises the target to whatever accuracy is required, although there may be computational constraints.

So, is there any reason not to try to construct the posterior? The answer is basically no — if you can do it, you should. However, for some situations this may be computationally a very demanding task. However, recent theoretical advances and the increase in computational power mean that many problems for which previously this golden goal was unachievable are now soluble. The posterior itself is complicated, but it may be broken down into a hierarchy of elements, each of which we do understand and can sample from. This forms what is called a ‘Bayesian Hierarchical Model’ (BHM). The state-of-the-art has reached a stage where one of the most demanding cosmological probes is now accessible to a rigorous treatment — weak gravitational lensing on a cosmic scale.

3. Weak gravitational lensing

Weak gravitational lensing⁴ distorts the shapes and sizes of galaxy images, and the measured shape of an individual image can be used as a very noisy estimate of

the spin-weight 2 shear field induced by lensing. A Bayesian Hierarchical Model for lensing could in principle be built from the raw image pixel data all the way to inferences about cosmology, but here we restrict ourselves to a subset of the analysis process, taking as a starting point estimates of the shear at the locations of individual galaxies. This does not propagate all errors fully, but is a major first step forward. We also assume that we have photometric redshifts for the galaxies, and bin them in both photo- z and pixels on the sky, in sufficient numbers that we can assume the errors (which are dominated by intrinsic ellipticity dispersion) are gaussian. Thus we have pixelised noisy shear maps in some number of tomographic bins, represented by the data, written in terms of the true shear maps s and the noise n (characterised by a covariance matrix N) as $d = s + n$ (d represents the complete data vector, including pixels in all tomographic bins). The statistical properties of the maps are set by various power spectra (auto- [within a bin], and cross-, in both E - and B -mode, and EB cross-powers). These power spectra are controlled in a physical model by cosmological parameters, and we have a choice of performing a cosmology-independent characterisation of the data, or inferring cosmological parameters. Thus we have as unknown parameters either:

- True shear maps s , plus power spectra $C_{\alpha\beta}^{XY}$; $X, Y = E$ or B ; $\alpha, \beta = \text{bin}$,
- or true shear maps s , plus cosmological parameters θ .

There are advantages to doing both, but let us concentrate initially on the first. The power spectra all enter into $C = \langle s s^T \rangle$, we may parametrise C by band powers. We have a relatively large number of parameters about C to infer, and a very large number ($\sim 10^5 - 10^7$) of true pixelised shear values. Thus it is a very high-dimensional problem. Statistically we are interested most in $p(C|d)$, so would marginalise over s , but if one is interested in map-making, then one can obtain $p(s|d)$ by marginalising over C .

3.1. Weak lensing Bayesian Hierarchical Model

A weak lensing BHM is shown in Fig. 1; Alsing et al.¹ give further details. The generative model for the data consists of drawing samples of C , from which shear maps can be drawn, since the conditional probability $p(s|C)$ is known, and noise be added since the conditional probability $p(d|s, N)$ is also known. BHMs lend themselves naturally to Gibbs samplers, which sample alternately from the power spectra and the map:

$$\begin{aligned} C^{i+1} &\leftarrow P(C|s^i) \\ s^{i+1} &\leftarrow P(s|C^i, d, N). \end{aligned} \quad (1)$$

Interestingly, for gaussian fields, the distributions (conditioned on the observed data d) can be written down: $p(C|s)$ is an inverse-Wishart distribution; $p(s|C, d, N)$ is a gaussian with a mean given by the Wiener-filtered map, $d_{\text{WF}} = (C^{-1} + N^{-1})^{-1} N^{-1} d$ with covariance $(C^{-1} + N^{-1})^{-1}$. However, there is a difficulty in that these matrices

are huge (up to say $10^7 \times 10^7$ elements). Unless they are diagonal, they cannot be handled at all. Now, C is diagonal in the harmonic domain, and N is diagonal in the pixel domain, but there is no basis in which both of them are diagonal. This problem has been elegantly solved² by the introduction of a *messenger field*.




Fig. 1. Simple Bayesian Hierarchical Model for weak lensing. C represents the lensing power spectra, s the true shear fields, d the data, and N the noise covariance.

3.2. Results from simulations

We applied the algorithm¹ sketched above to simulated noisy, masked 2-bin tomographic shear maps generated using SUNGLASS simulations³, inferring 1940 power spectrum parameters. The shear maps contain 128×128 pixels, bringing the total number of parameters in the inference task to 67476.

We ran three Gibbs chains with independent starting points, obtaining 3.6M samples, each taking 0.5s on a desktop. Convergence was determined with a Gelman-Rubin statistic $r < 1.1$ for the marginal distributions of all parameters. Fig. 2 shows the results in the map domain. Fig. 3 shows the recovered E -mode tomographic (auto and cross) power spectra for the two tomographic bins.

4. Cosmological parameter inference and intrinsic alignments

An alternative to using the power spectra as parameters is to use cosmological parameters, with a Boltzmann code to generate the lensing power spectra. This has some advantages, in that the number of parameters is reduced, and some sampling issues in the very low S/N regime may be avoided. An additional advantage is that it may be possible to include non-gaussianity in the conditional distributions. The

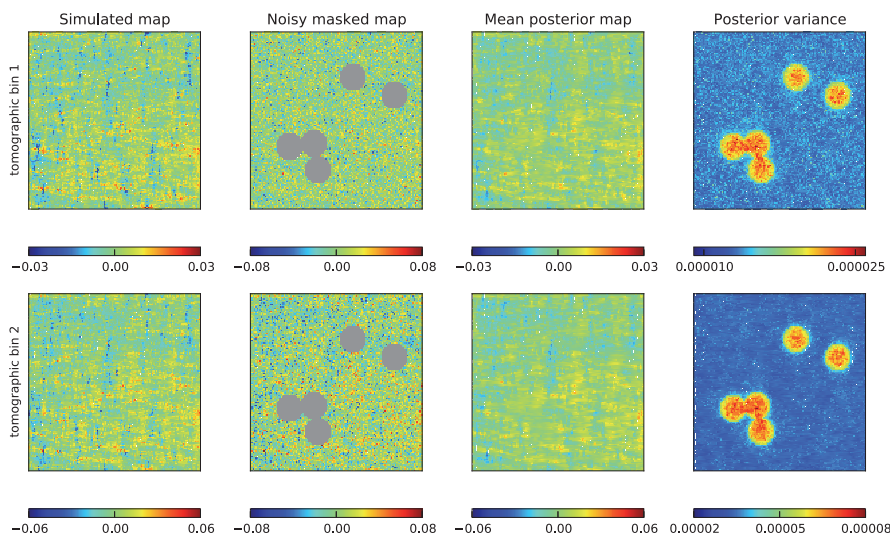


Fig. 2. Tomographic maps of the γ_1 component of the shear for the noiseless simulated shear maps (far left), noisy masked simulated maps (second from left), mean posterior maps (third from left) and the posterior variance (far right). The mean posterior maps recover most of the structure from the simulated shear maps. Note that inference is made about the field in masked regions, but the posterior variance in those regions is significantly higher than in unmasked regions.

disadvantage is that the analysis becomes model-dependent, and the power spectrum approach gives a neat encapsulation of the statistical properties of the data at the 2-point level, regardless of the origin of the signal. The other advantage of using the power spectrum approach is that it is in principle possible to do the cosmological parameter inference as a second step, using the power-spectrum chains. This way, different models for intrinsic alignments of galaxies (a potential source of significant systematic error in weak lensing) may be investigated without rerunning chains. They can be included in a post-processing step as an addition to the theoretical power spectrum (rather than as a source of correlated noise, which would be much harder to deal with). Using power spectra, one could in principle include higher-order, non-gaussian terms (bispectrum, trispectrum), but in the general case this is unfeasible, as there is an extremely large number of parameters unless they are regularised by imposing a physical model. In this case the correct procedure will be to use a gaussian likelihood, as it is the maximum entropy distribution given a mean and covariance, and is thus in some sense the least informative and most conservative assumption. Both approaches should be undertaken.




Fig. 3. Recovered E -mode tomographic shear power spectra: the orange (grey) bands indicate the 68% (95%) credible regions, the black lines show the posterior means, the red lines show the estimated band powers from the noiseless, mask-less simulated shear catalogue and the horizontal blue lines show indicate the mean ellipticity-noise level.

5. Conclusion

We have shown here how a Bayesian Hierarchical Model can be used to generate samples efficiently from the posterior distribution — the ‘golden goal’ of statistical analysis — for tomographic weak lensing shear maps, and a multitude of lensing power spectra, inferring parameters in a $\sim 10^5$ -dimensional space. It solves many of the awkward problems of cosmic shear, such as how to treat the mask (pixel variances are simply set to infinity), and how to include intrinsic alignments (simply as an additional signal power contribution).

References

1. Alsing J., Heavens A., Jaffe A., Kiessling A., Wandelt B., Hoffmann T., *MNRAS* **455**, 4452 (2016).
2. Elsner F., Wandelt B. D., *A&A* **549**, A111 (2013).
3. Kiessling A., Heavens A., Taylor A., Joachimi B., *MNRAS* **414**, 2235 (2011).
4. Munshi D., Valageas P., van Waerbeke L., Heavens A., *Phys. Rep.* **462**, 67 (2008).