

Mapping gravitational lensing of the CMB using local likelihoods

Ethan Anderes

Statistics Department, University of California, Davis, California 95616, USA

Lloyd Knox

Physics Department, University of California, Davis, California 95616, USA

Alexander van Engelen

Physics Department, McGill University, Montréal, Quebec H3A 2T8, Canada

(Received 8 December 2010; published 24 February 2011; publisher error corrected 28 February 2011)

We present a new estimation method for mapping the gravitational lensing potential from observed CMB intensity and polarization fields. Our method uses Bayesian techniques to estimate the average curvature of the potential over small local regions. These local curvatures are then used to construct an estimate of a low pass filter of the gravitational potential. By utilizing Bayesian/likelihood methods one can easily overcome problems with missing and/or nonuniform pixels and problems with partial sky observations (E- and B-mode mixing, for example). Moreover, our methods are local in nature, which allow us to easily model spatially varying beams, and are highly parallelizable. We note that our estimates do not rely on the typical Taylor approximation which is used to construct estimates of the gravitational potential by Fourier coupling. We present our methodology with a flat sky simulation under nearly ideal experimental conditions with a noise level of $1 \mu\text{K-arcmin}$ for the temperature field, $\sqrt{2} \mu\text{K-arcmin}$ for the polarization fields, with an instrumental beam full width at half maximum (FWHM) of 0.25 arcmin .

DOI: [10.1103/PhysRevD.83.043523](https://doi.org/10.1103/PhysRevD.83.043523)

PACS numbers: 98.70.Vc, 98.62.Sb

I. INTRODUCTION

Over the past decade the cosmic microwave background (CMB) has emerged as a fundamental probe of cosmology and astrophysics. In addition to the primary fluctuations of the early Universe, the CMB contains signatures of the gravitational bending of CMB photon trajectories due to matter, called gravitational lensing. Mapping this gravitational lensing is important for a number of reasons including, but not limited to, understanding cosmic structure, constraining cosmological parameters [1,2] and detecting gravity waves [3–5]. In this paper we present a local Bayesian estimate that can accurately map the gravitational lens in high resolution, low noise measurements of the CMB temperature and polarization fields.

There is extensive literature on estimating the lensing of the CMB (classic references include [6–8]) and some recent observational detections [9,10]. The current estimators in the literature can be loosely characterized into two types. The first type was initiated in [8] (see also [11,12]) and utilizes quadratic combinations of the CMB and its gradient to infer lens structure. The optimal quadratic combinations were then discovered by [7,13,14] and are generally referred to as “the quadratic estimator”. This is arguably the most popular estimate of the gravitational potential and uses a first order Taylor approximation to establish mode coupling in the Fourier domain which can be estimated to recover the gravitational potential (real space analogs to these estimators can be found in [15,16]). The second type is an approximate global maximum likelihood estimate and was developed in [6,17].

Our method, in contrast, locally approximates a quadratic form for the gravitational potential and estimates the coefficients locally using Bayesian methods. The locally estimated coefficients are then globally *stitched together* to construct an estimate of a low pass filter of the gravitational potential. The local analysis allows us to avoid using the typical first order Taylor expansion for the quadratic estimator and avoids the likelihood approximations used in global estimates. Moreover, we are able to easily handle missing pixels, problems with partial sky observations (E- and B-mode mixing, for example), and spatially varying or asymmetric beams. The motivation for developing this estimate stems, in part, from current speculation that likelihood methods will allow superior mapping of the lensing structure (compared to the quadratic estimator) under low noise levels, and that global likelihood methods can be prohibitively computational intensive—indeed intractable—without significant approximation.

We illustrate our mapping methodology on a high resolution, low noise simulation of the CMB temperature and polarization field on a $17^\circ \times 17^\circ$ patch of the flat sky. This simulation is used throughout the paper to demonstrate findings and techniques. To get an overview of the performance of our method, Fig. 1 shows the estimated potential (left) from the simulated lensed CMB temperature and polarization field (observational noise levels are set at $1 \mu\text{K-arcmin}$ for the temperature field, $\sqrt{2} \mu\text{K-arcmin}$ for the polarization fields, with a beam FWHM of 0.25 arcmin). The input gravitational potential is shown in the right diagram in Fig. 1. The details of the simulation procedure




FIG. 1 (color online). *Left*: Estimated gravitational potential on a $17^\circ \times 17^\circ$ patch of the simulated flat sky. *Right*: Input gravitational potential used in the simulation. See Section I and Appendix A for the simulation details.

can be found in Appendix A. It is clear from Fig. 1 that the mapping accurately traces the true, unknown gravitational potential. To get an idea of the noise of this reconstruction for different realizations of the CMB + noise we present Fig. 2 which shows the different estimates of the projected matter power spectrum using the estimated projected mass—with the local likelihood approach—for 10 different CMB + noise realizations (dashed lines) while keeping




FIG. 2 (color online). Plot of the projected mass power spectrum (red solid line) along with the estimated power spectrum using the true, but unknown, projected mass (blue dotted line). The dashed lines correspond to different estimates of the power spectrum using the estimated projected mass—with the local likelihood approach—for different CMB realizations but the same lensing potential realization. See Section I and Appendix A for the simulation details.

the gravitational potential in Fig. 1 fixed. The dotted line (blue) shows the estimated projected mass power spectrum if one had access to the true gravitational potential used in our simulations. Finally we plot the theoretical ensemble average projected mass power spectrum in red to get an idea of the magnitude of the errors in the mass reconstruction.

II. LOCAL MAXIMUM A POSTERIORI ESTIMATES OF SHEAR AND CONVERGENCE

The CMB radiation in the flat sky limit can be expressed in term of the Stokes parameters T , Q , U which measure total intensity $T(\mathbf{x})$, and linear polarization $Q(\mathbf{x})$ and $U(\mathbf{x})$ with respect to some coordinate frame $\mathbf{x} = (x, y) \in \mathbb{R}^2$. Instead of directly observing T , Q , U we observe a remapping of the CMB due to the gravitational effect of intervening matter. This lensed CMB can be written $T(\mathbf{x} + \nabla\phi(\mathbf{x}))$, $Q(\mathbf{x} + \nabla\phi(\mathbf{x}))$ and $U(\mathbf{x} + \nabla\phi(\mathbf{x}))$ where ϕ denotes the gravitational potential (see [18], for example).

To describe our estimate of the gravitational potential, ϕ , first consider a small circular observation patch with diameter δ in the flat sky centered at some point \mathbf{x}_0 , denoted $\mathcal{N}_\delta(\mathbf{x}_0) \subset \mathbb{R}^2$. Over this small region we decompose ϕ into an overall local quadratic fit and error term

$$\phi = q^\phi + \epsilon$$

where q^ϕ is a local quadratic approximation of the potential ϕ with error term $\epsilon \equiv \phi - q^\phi$. In what follows we estimate q^ϕ , denoted \hat{q}^ϕ , and associate this estimate with the neighborhood midpoint \mathbf{x}_0 . Then we repeat this procedure for other local midpoints \mathbf{x}_0 throughout the observation window. After a shrinkage adjustment is made to the local estimates we show, in Section II D, how to *stitch* the

estimates together to produce the final estimated potential $\hat{\phi}$ shown in Fig. 1.

Notice that as $\delta \rightarrow 0$ the expected magnitude of the error ϵ approaches zero. This has the effect of improving the following Taylor approximation

$$T(\mathbf{x} + \nabla\phi(\mathbf{x})) = T(\tilde{\mathbf{x}}) + \nabla\epsilon(\mathbf{x}) \cdot \nabla T(\tilde{\mathbf{x}}) + \dots \quad (1)$$

for $\mathbf{x} \in \mathcal{N}_\delta(\mathbf{x}_0)$, where we use the notation $\tilde{\mathbf{x}} \equiv \mathbf{x} + \nabla q^\phi(\mathbf{x})$ (with a similar Taylor expansion for both $Q(\mathbf{x} + \nabla\phi(\mathbf{x}))$ and $U(\mathbf{x} + \nabla\phi(\mathbf{x}))$). Notice that $\tilde{\mathbf{x}}$ depends not only on \mathbf{x} but also the unknown coefficients of the quadratic term q^ϕ . We briefly mention that these are related to the convergence κ and shear $\gamma = \gamma_1 + i\gamma_2$ of the gravitational lens by

$$\begin{aligned} \kappa &\approx -(q_{xx}^\phi + q_{yy}^\phi)/2 \\ \gamma_1 &\approx -(q_{xx}^\phi - q_{yy}^\phi)/2 \\ \gamma_2 &\approx -q_{xy}^\phi \end{aligned}$$

using the shear notation given in [8]. Now when δ is sufficiently small we can truncate the expansion in (1) to get

$$\begin{bmatrix} T(\mathbf{x} + \nabla\phi(\mathbf{x})) \\ Q(\mathbf{x} + \nabla\phi(\mathbf{x})) \\ U(\mathbf{x} + \nabla\phi(\mathbf{x})) \end{bmatrix} \approx \begin{bmatrix} T(\mathbf{x} + \nabla q^\phi(\mathbf{x})) \\ Q(\mathbf{x} + \nabla q^\phi(\mathbf{x})) \\ U(\mathbf{x} + \nabla q^\phi(\mathbf{x})) \end{bmatrix} \quad (2)$$

on the local neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$. By regarding q^ϕ as unknown we can use the right hand side of (2) to develop a likelihood for estimating the coefficients of q^ϕ . Nominally q^ϕ has 6 unknown coefficients for which to estimate. However, we can ignore the linear terms in q^ϕ since the CMB temperature and the polarization are statistically invariant under the resulting translation in ∇q^ϕ . Therefore, one can write q^ϕ as $c_1(x - x_0)^2/2 + c_2(x - x_0)(y - y_0) + c_3(y - y_0)^2/2$ for unknown coefficients $c_1 = q_{xx}^\phi$, $c_2 = q_{xy}^\phi$, $c_3 = q_{yy}^\phi$.

An important probe of gravitational lensing from the CMB polarization is the creation of a curl-like B-mode from the lensing [19,20]. We remark that a local quadratic approximation in (2) still has the power to detect this B-mode power so that the local procedure is not blind to this information source. To see this notice that a quadratic lensing potential remaps the coordinates by

$$\tilde{\mathbf{x}} = \mathbf{x} + \begin{bmatrix} q_{xx}^\phi & q_{xy}^\phi \\ q_{xy}^\phi & q_{yy}^\phi \end{bmatrix} (\mathbf{x} - \mathbf{x}_0).$$

If we assume the original polarization ($Q(\mathbf{x}), U(\mathbf{x})$) is curl free then the lensed polarization has curl given by

$$\text{curl}(Q(\tilde{\mathbf{x}}), U(\tilde{\mathbf{x}})) = -2\gamma_1 U_x(\tilde{\mathbf{x}}) + \gamma_2 [Q_x(\tilde{\mathbf{x}}) - U_y(\tilde{\mathbf{x}})].$$

Therefore the shear parameter γ , and not the convergence κ , is what creates local B-mode power. The dominant source of information for B-mode power is in the cross correlation between the lensed Stokes parameters $Q(\tilde{\mathbf{x}})$ and

$U(\tilde{\mathbf{x}})$. This agrees with [7] that the E-B cross estimator provides optimal signal to noise under nearly ideal experimental conditions.

We finish this section with a remark on the accuracy of the Taylor approximation (1). As the signal to noise ratio increases and the pixel resolution improves one can shrink the local neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$ so the term ϵ becomes smaller (which improves the Taylor approximation). However, as $\delta \rightarrow 0$, the fields T , Q and U become nearly linear and one may expect some loss of information from the shrinking power in T , Q and U at frequencies with wavelengths smaller than the neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$. It therefore may be statistically advantageous to artificially increase the neighborhood size while simultaneously increasing the order of the local polynomial fit q^ϕ . Then, instead of recording the full polynomial fit at each midpoint \mathbf{x}_0 , one can retain the second order derivatives $q_{xx}^\phi(\mathbf{x}_0)$, $q_{xy}^\phi(\mathbf{x}_0)$, $q_{yy}^\phi(\mathbf{x}_0)$ for estimates of κ and γ . It is yet to be seen, however, what δ and what polynomial order will be optimal for a given noise and resolution level. In Section III we present an information metric for choosing the neighborhood size δ for the simulation specifics and for a quadratic polynomial q^ϕ .

A. The local posterior

Using the Gaussian approximation of the CMB along with the quadratic potential approximation given by (2) we describe how to construct the likelihood as a function of the unknown quadratic coefficients in q^ϕ . Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ denote the observation locations of the CMB within the local neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$ centered at \mathbf{x}_0 . Using approximation (2), the CMB observables in this local neighborhood are well modeled by white noise corruption of a convolved (by the beam) lensed intensity and polarization field. Let \mathbf{t} , \mathbf{q} , \mathbf{u} denote n -vectors of observed CMB values at the corresponding pixel locations in $\mathcal{N}_\delta(\mathbf{x}_0)$ for the intensity T and Stokes parameters Q , U , respectively. Using Gaussianity of the full vector of CMB observables, $\mathbf{z} = (\mathbf{t}^\dagger, \mathbf{q}^\dagger, \mathbf{u}^\dagger)^\dagger$, the log likelihood (up to a constant) as a function of the quadratic fit q^ϕ can be written

$$\mathcal{L}(q^\phi | \mathbf{z}) = -\frac{1}{2} \mathbf{z}^\dagger (\Sigma_{q^\phi} + N)^{-1} \mathbf{z} - \frac{1}{2} \text{ln det}(\Sigma_{q^\phi} + N) \quad (3)$$

where $\Sigma_{q^\phi} + N$ is the covariance matrix of the observation vector \mathbf{z} (we use the subscript to emphasize the dependence on the unknown quadratic q^ϕ), $N = \text{diag}(\sigma_T^2 I, \sigma_Q^2 I, \sigma_U^2 I)$ is the noise covariance structure and I is the $n \times n$ identity matrix. Notice that the noise structure does not depend on the unknown quadratic q^ϕ . In the next section we will derive the exact form of the prior distribution on q^ϕ , denoted $\pi(q^\phi)$, but briefly mention that the posterior distribution on q^ϕ , which we maximize to estimate q^ϕ , is

$$p(q^\phi | \mathbf{z}) \propto e^{\mathcal{L}(q^\phi | \mathbf{z})} \pi(q^\phi). \quad (4)$$

The entries of $\Sigma_{q^\phi} + N$ contain the covariances $\langle t_k t_j \rangle_{\text{CMB}}$, $\langle q_k q_j \rangle_{\text{CMB}}$, $\langle u_k u_j \rangle_{\text{CMB}}$ and all cross-covariances among \mathbf{t} , \mathbf{q} , \mathbf{u} (we use t_k to denote the k th entry of \mathbf{t} , for example). Let φ denote the instrumental beam and σ_T , σ_Q , σ_U denote the noise standard deviations of T , Q , U so that, for example, the k th entry of \mathbf{t} is modeled as

$$t_k \equiv \int_{\mathbb{R}^2} d^2 \mathbf{x} \varphi(\mathbf{x}) T(\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}) + \sigma_T n_k \quad (5)$$

where the n_k 's are independent standard Gaussian random variables, $\tilde{\mathbf{x}}_k = \mathbf{x}_k + \nabla q^\phi(\mathbf{x}_k)$ and $\tilde{\mathbf{x}} = \mathbf{x} + \nabla q^\phi(\mathbf{x})$. Note that this is an approximate model for t_k based on (2). In actuality, the k th temperature measurement is $\int_{\mathbb{R}^2} d^2 \mathbf{x} \varphi(\mathbf{x}) T(\mathbf{x}_k - \mathbf{x} + \nabla \phi(\mathbf{x}_k - \mathbf{x})) + \sigma_T n_k$, but the linearity of ∇q^ϕ allows us to write $\mathbf{x}_k - \mathbf{x} + \nabla \phi(\mathbf{x}_k - \mathbf{x}) \approx \text{constant} + \tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}$ on the small neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$. Under the assumption of zero B-mode, the spectral densities associated with Q , U can be written

$$C_\ell^Q = C_\ell^E \cos^2(2\varphi_\ell) \quad (6)$$

$$C_\ell^U = C_\ell^E \sin^2(2\varphi_\ell) \quad (7)$$

$$C_\ell^{QU} = C_\ell^E \cos(2\varphi_\ell) \sin(2\varphi_\ell) \quad (8)$$

where $\tan(\varphi_\ell) = \ell_2/\ell_1$ and $\ell = (\ell_1, \ell_2) \in \mathbb{R}^2$. Since one can write $\mathbf{x} + \nabla q^\phi(\mathbf{x}) = M\mathbf{x}$ where the M is a 2×2 real matrix, the sheared Stokes parameters $T(\tilde{\mathbf{x}})$, $Q(\tilde{\mathbf{x}})$ and $U(\tilde{\mathbf{x}})$ are stationary random fields with spectral densities given by $C_{M^{-1}\ell}^T \det M^{-1}$, $C_{M^{-1}\ell}^Q \det M^{-1}$ and $C_{M^{-1}\ell}^U \det M^{-1}$, respectively. After adjusting for the beam (which is applied after lensing) the covariance between the observations in \mathbf{t} can be written

$$\langle t_k t_j \rangle_{\text{CMB}} = \sigma_T^2 \delta_{ij} + \int_{\mathbb{R}^2} \frac{d^2 \ell}{(2\pi)^2} e^{i\ell \cdot (\mathbf{x}_k - \mathbf{x}_j)} |\varphi(\ell)|^2 \frac{C_{M^{-1}\ell}^T}{\det M}. \quad (9)$$

The computations are similar to complete the entries of covariance matrix $\Sigma_{q^\phi} + N$. At face value the above integral seems too computationally intensive for every pair $\mathbf{x}_k - \mathbf{x}_j$. Moreover, to apply Newton type algorithms for maximizing the posterior (4) one needs to compute the derivatives of $\langle t_k t_j \rangle_{\text{CMB}}$ with respect to elements of M . In Appendix B we show that some of these computational challenges can be overcome by utilizing a single fast-Fourier-transform algorithm (FFT) to quickly compute the above integral for sufficient resolution in the argument $\mathbf{x}_k - \mathbf{x}_j$ to recover $\langle t_k t_j \rangle_{\text{CMB}}$ for all pairs k, j .

B. Taylor truncation bias

The quadratic function q^ϕ is defined as the best least square fit of ϕ over the neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$. The residual

$\epsilon = \phi - q^\phi$, defined over $\mathcal{N}_\delta(\mathbf{x}_0)$, is nonstationary and will therefore not have a spectral density that diagonalizes the covariance structure. However, stationarity is a good approximation for order of magnitude calculations on the truncation error in (1). We approximate the spectral density of ϵ as an attenuated version of C_ℓ^ϕ by arguing that the quadratic fit effectively removes the spectral power at wavelengths greater than 2δ . Reasoning similarly we expect the quadratic fit to have negligible impact on the spectral power at wavelengths smaller than δ . By assuming the spectral power grows linearly in the intermediary spectral range, from zero at $\ell = \pi/\delta$ to $C_{2\pi/\delta}^\phi$ at $\ell = 2\pi/\delta$, we obtain an approximate model for the spectral density of ϵ

$$C_\ell^\epsilon \approx \min\left\{1, \left[\frac{\delta}{\pi} |\ell| - 1\right]^+\right\}^2 C_\ell^\phi$$

where x^+ denotes the positive part of the real number x . Notice that the attenuation happens on the realizations of ϕ , hence requiring the square on the low pass filter in the spectral density. This implies that the second term in the Taylor expansion (1) has approximate spectral density

$$C_\ell^{\nabla T(\tilde{\mathbf{x}}) \cdot \nabla \epsilon(\mathbf{x})} \approx \int \frac{d^2 \ell'}{(2\pi)^2} (M^{-1}\ell' \cdot (\ell - \ell'))^2 C_{\ell - \ell'}^\epsilon \frac{C_{M^{-1}\ell'}^T}{\det M}.$$

In our simulation we use a neighborhood diameter of $\delta = 0.006$ radians (20.6 arcmin). This diameter was chosen using the information criterion developed in Section III. The corresponding approximate rms of $\nabla T(\tilde{\mathbf{x}}) \cdot \nabla \epsilon(\mathbf{x})$ is $\sim 2.3 \mu\text{K}$ with an order of magnitude reduction for the polarization field. Brute force simulation of $\langle_{\mathbf{x}_k \in \mathcal{N}_\delta(\mathbf{x}_0)}^{\text{Mean}} \{T(\mathbf{x}_k + \nabla \phi(\mathbf{x}_k)) - T(\tilde{\mathbf{x}}_k)\}_{\text{CMB}}^2 \rangle^{1/2}$ yields a value closer to $\sim 3.6 \mu\text{K}$, suggesting a reasonable stationary approximation to ϵ . These approximations show that the polarization truncation error is smaller (by an order of magnitude) than the simulation noise level $\sqrt{2} \mu\text{K}$ -arcmin. However, the temperature truncation error is greater than the temperature noise level $1 \mu\text{K}$. A consequence is that the likelihood explains the additional high frequency power in the observations (from the error term) by adjusting the estimate of q^ϕ to artificially magnify the convergence κ estimates. Indeed, this bias seems relatively constant and can be clearly seen in Fig. 3 in the top blue points. Each blue point corresponds to a local neighborhood: the x -coordinate representing the true $\nabla^2 q^\phi$ associated with that neighborhood; the y -coordinate representing the estimated local value shifted up by 0.2, i.e. $\nabla^2 \hat{q}^\phi + 0.2$. The bias of nearly ~ 0.1 above the top dashed blue line $y = x + 0.2$, shows the effect of the additional high frequency power of the error term $\nabla T(\tilde{\mathbf{x}}) \cdot \nabla \epsilon(\mathbf{x})$. To adjust this, we subtract the overall mean of the local estimates, reasoning that the observation window is large enough at $17^\circ \times 17^\circ$ so that the overall mean of the true values q_{xx}^ϕ , q_{xy}^ϕ , q_{yy}^ϕ is close to zero. For smaller observation windows it

may be possible to estimate an overall quadratic fit to correct for this bias but we do not investigate that here.

C. The prior $\pi(q^\phi)$

The stationary approximation for ϵ also yields an approximation for the prior distribution of the local quadratic fit q^ϕ using the identity $q^\phi = \phi - \epsilon$. Since ϵ is well modeled by a high pass filter of ϕ , the quadratic function q^ϕ can be modeled by the corresponding low pass filter

$$q^\phi(\mathbf{x}) \approx \int \frac{d^2\ell}{2\pi} e^{i\mathbf{x}\cdot\ell} \phi^{\text{lp}}(\ell)$$

over $\mathbf{x} \in \mathcal{N}_\delta(\mathbf{x}_0)$, where $\phi^{\text{lp}}(\ell) \equiv \min\{1, [2 - \frac{\delta}{\pi}|\ell|]^+\} \phi(\ell)$ which has spectral density $\min\{1, [2 - \frac{\delta}{\pi}|\ell|]^+\}^2 C_\ell^\phi$. Therefore a natural candidate for the prior on the coefficients of q^ϕ are the random variables $\frac{\partial^2 \phi^{\text{lp}}(0)}{\partial x_i \partial x_j}$ which are mean zero and Gaussian with variances obtained by the corresponding spectral moments of ϕ^{lp} . This prior is used on each local neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$ to derive the local maximum *a posteriori* estimate. For the simulation used in this paper, the neighborhood width was set to $\delta = 0.006$ radians (20.6 arcmin) which gives prior variances 0.0023, 0.0008, 0.0023 for q_{xx}^ϕ , q_{xy}^ϕ and q_{yy}^ϕ , respectively, (the only nonzero cross covariance is between q_{xx}^ϕ and q_{yy}^ϕ and is 0.0008).

D. Reconstructing ϕ from \hat{q}^ϕ

When observing the full sky, the estimates of κ will allow one to recover the gravitational potential ϕ by solving the poisson equation $\nabla^2 \phi = -2\kappa$ (up to a constant). With partial sky observations, however, the shear is needed to break ambiguity corresponding to different boundary conditions. We do this in two stages, first using \hat{q}_{xx}^ϕ , \hat{q}_{xy}^ϕ and \hat{q}_{yy}^ϕ (regarded as functions of the local neighborhood midpoint \mathbf{x}_0) to recover the estimated displacement field $(\hat{\phi}_x, \hat{\phi}_y)$, then using this displacement field to recover the estimated potential $\hat{\phi}$. To handle this, we adopt the method of [21,22] and define $\hat{\phi}_x$, $\hat{\phi}_y$ as minimizers of functionals F_1 and F_2 defined as

$$F_1(\phi_x) \equiv \int dx dy [(\phi_{xx} - \hat{q}_{xx}^\phi)^2 + (\phi_{xy} - \hat{q}_{xy}^\phi)^2]$$

$$F_2(\phi_y) \equiv \int dx dy [(\phi_{xy} - \hat{q}_{xy}^\phi)^2 + (\phi_{yy} - \hat{q}_{yy}^\phi)^2].$$

In particular, $\hat{\phi}_x$ satisfies $F_1(\hat{\phi}_x) = \min_{\phi_x} F_1(\phi_x)$ and similarly for $\hat{\phi}_y$. See [23] for details of the minimization algorithm. Now we use the estimated displacements $(\hat{\phi}_x, \hat{\phi}_y)$ to define the estimated potential $\hat{\phi}$ as the minimizer of the functional F_3 defined as

$$F_3(\phi) \equiv \int dx dy [(\phi_x - \hat{\phi}_x)^2 + (\phi_y - \hat{\phi}_y)^2].$$

The minimization is needed to account for the fact that our estimates are noisy versions of the truth and therefore may not correspond to an integral vector field for which a potential exists. A consequence is that the estimate $\hat{\phi}$ is “shrunk” towards zero when the algorithm fits a gradient to a vector field which may have nonvanishing curl. This shrinking can be seen in Fig. 3 looking at the scatter plot of green points. These points show $(\nabla^2 \phi(\mathbf{x}_0), \nabla^2 \hat{\phi}(\mathbf{x}_0))$ for each local neighborhood $\mathcal{N}_\delta(\mathbf{x}_0)$. One can clearly see the shrinkage effect by noticing the slope of the trend in the green points is less than one. We undo this shrinkage effect by multiplying $\hat{\phi}$ by a factor that undoes this shrinkage. The multiplication factor, denoted c , is determined by matching the variance of the raw estimates $\nabla^2 \hat{q}^\phi$ with $c \nabla^2 \hat{\phi}$. The result of this correction factor is seen in the scatter plot of the red points, in Fig. 3, which show the local




FIG. 3 (color online). Estimated values of $\nabla^2 \phi(\mathbf{x}_0)$, for each local neighborhood midpoint \mathbf{x}_0 , plotted against the simulation truth at different stages of the algorithm. The blue points (top) correspond to the raw estimates at each local neighborhood; The green points (middle) to the estimates after fitting a gravitational potential; The red points (bottom) after a shrinkage correction. The y coordinates of the blue points (top) are shifted up by 0.2 and the red points (bottom) are shifted down by 0.2 to fit on the same diagram. See Sections II B and II D for discussion.




FIG. 4 (color online). The right diagram shows $\hat{\phi}_x$, where $\nabla\phi \equiv (\phi_x, \phi_y)$ is the true gravitational displacement field used in the simulation. The left diagram shows the estimate $\hat{\phi}_x$ which is derived from the local quadratic estimates using the methodology described in Section II D.

convergence estimates versus truth after the correction factor ($\nabla^2\phi, c\nabla^2\hat{\phi} - 0.2$).

The estimated $\hat{\phi}_x$ (after correcting for the shrinkage) along with the true displacement ϕ_x (used in the simulation) are shown in Fig. 4. The estimated $\hat{\phi}$ along with the true gravitational potential ϕ are shown in Fig. 1. These two figures demonstrate accurate reconstruction of both the gravitational potential and the displacement field. In addition, by differentiating the estimated potential, $\hat{\phi}$, one obtains smoothed estimates of convergence and shear (smoothed from the fitting of $\hat{\phi}$). In Fig. 5 we plot the estimate $\hat{\phi}_{xy}$ (which corresponds to minus the imaginary part of the shear γ), along with ϕ_{xy} (bottom right) and the low pass filter ϕ_{xy}^{lp} (bottom left) defined in Section II C. Notice that the estimate $\hat{\phi}_{xy}$ tracks the derivatives of the low pass filter ϕ_{xy}^{lp} , whereas the additional high frequency in ϕ_{xy} is not accurately estimated from $\hat{\phi}_{xy}$. This is presumably due to the local fitting of a quadratic potential over the neighborhoods $\mathcal{N}_\delta(\mathbf{x}_0)$.

III. NEIGHBORHOOD SIZE AND STRUCTURE

We define the following measure of information which is used as a metric for choosing the width of the neighborhood and other parameters of our estimation method:

$$\text{Information for } q^\phi \equiv \frac{\text{variance of the prior on } q^\phi}{\text{expected variance of the posterior on } q^\phi}.$$

The above information metric is essentially a measure of signal to noise ratio (squared). The variance of the prior corresponds to the squared magnitude of the signal,

whereas the expected variance of the posterior is a proxy for the squared magnitude of the noise. We use simulations to estimate this information (while using the hessian of the posterior density at \hat{q}^ϕ to approximate posterior variance) and use it for guidance when choosing the tuning parameters for our estimation algorithm. Note: we avoided a lengthy and rigorous simulation study to choose global optimal tuning parameters, opting for a less rigorous simulation study which yields, potentially, suboptimal but reasonable algorithmic parameters.

The main parameter that needs tuning is the local neighborhood size δ . Notice that our information measure attempts to balance two competing quantities when choosing a neighborhood size, the larger the neighborhood the smaller the signal q^ϕ (from the low pass filter). On the other hand, larger neighborhoods correspond to more data when the resolution is fixed. Using this metric, $\delta = 0.006$ radians (20.6 arcmin) emerges as a good neighborhood size when the beam FWHM is 0.25 arcmin and the noise levels are $\sqrt{2}$ and 1 $\mu\text{K-arcmin}$ pixels for Q , U and T , respectively.

Because of computational limitations associated with larger neighborhoods we found it necessary to down-sample the local neighborhoods by discarding pixels. Using the information metric we were able to isolate that randomly sampling the pixels seemed preferable to evenly down-sampling to a coarser grid. Moreover, we found that using different randomly selected pixels for T , Q and U was preferable to using the same random pixels for all the Stokes fields. Therefore, for each local neighborhood we selected 300 random pixels in $\mathcal{N}_\delta(\mathbf{x}_0)$ for the T observations, then randomly selected 300 pixels from those remaining for Q and finally 300 pixels from the remaining unselected pixels for U (allowing overlaps when the local neighborhood size had fewer than 900 pixels).




FIG. 5 (color online). The top diagram shows the estimate of ϕ_{xy} (which corresponds to minus the imaginary part of the shear γ) where ϕ denotes the gravitational potential. The bottom two diagrams show the simulation truth: bottom left shows ϕ_{xy}^{lp} where ϕ^{lp} denotes the low pass filter $\phi^{\text{lp}}(\ell) \equiv \min\{1, [2 - \frac{0.006}{\pi}|\ell|]^+\}\phi(\ell)$ (see Section II C for a discussion); bottom right shows ϕ_{xy} . Notice that the estimate of ϕ_{xy} tracks the low pass filter ϕ_{xy}^{lp} and does not have the high frequency behavior seen in the simulation truth ϕ_{xy} .

IV. DISCUSSION

We have demonstrated the feasibility of using a local Bayesian estimate to accurately map the gravitational potential and displacement fields under low noise, small beam experimental conditions. The motivation for developing this estimate stems, in part, from speculation that likelihood methods will allow superior mapping of the lensing structure (compared to the quadratic estimator) under low noise levels. The main difference between the global estimates of [6,17] and the local estimate presented here is the nature of the likelihood approximation. In [6,17] the global likelihood is defined as a functional on the unknown gravitational potential ϕ and approximations are made to this functional. Our method, in contrast, uses a nearly exact likelihood—exact up to approximation (B1) in Appendix B—but under a local modeling approximation

that assumes a quadratic ϕ . One advantage is the added precision available to model instrumental and foreground characteristics. For example, the local analysis models the beam convolved CMB rather than the deconvolved CMB. Deconvolution induces spatial correlation in the additive instrumental noise which is potentially nonstationary if the beam spatially varies. Since this noise is not invariant under warping it complicates the global likelihood. Another advantage is that the local estimates are relatively easy to implement and parallelize. In addition, the local estimate automatically uses the highest signal to noise combinations of Q , U and T so there is no need to rederive the optimal quadratic combinations for different experimental conditions.

The local analysis is not free from disadvantages however. A global analysis is presumably much better suited for estimating long wavelengths in the gravitational

potential and wavelengths that are shorter than the local neighborhood size. Moreover, since our estimates are defined implicitly—as the maximum of the posterior density—it is difficult to derive expected error magnitudes. However, the results presented here show that under some experimental conditions the advantages overcome the disadvantages. Moreover our local estimate uses an approximation that is inherently different from the Taylor approximation used to derive the quadratic estimator. This leaves open the possibility that the local estimate may have different bias and error characteristics which could compliment the quadratic estimator, rather than replace it.

ACKNOWLEDGMENTS

E. Anderes received support through NSF Grant No. 1007480. L. Knox received support through NSF Grant No. 0709498.

APPENDIX A: SIMULATION DETAILS

The fiducial cosmology used for the simulations is based on a flat, power law Λ CDM cosmological model, with baryon density $\Omega_b = 0.044$; cold dark matter density $\Omega_{\text{cdm}} = 0.21$; cosmological constant density $\Omega_\Lambda = 0.74$; Hubble parameter $h = 0.71$ in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$; primordial scalar fluctuation amplitude $A_s(k = 0.002 \text{ Mpc}^{-1}) = 2.45 \times 10^{-9}$; scalar spectral index $n_s(k = 0.002 \text{ Mpc}^{-1}) = 0.96$; primordial helium abundance $Y_p = 0.24$; and reionization optical depth $\tau_r = 0.088$. The code for anisotropies in the microwave background is used to generate the theoretical power spectra [24].

We start by simulating maps of the unlensed CMB Stokes parameters T , Q , U . The following Riemann sum approximation is used for the random fields T , Q , U

$$T(x) \approx \sum_{\ell} \frac{Z_{\ell}^T \sqrt{\Delta \ell_1 \Delta \ell_2}}{2\pi} e^{ix \cdot \ell} \sqrt{C_{\ell}^T} \quad (\text{A1})$$

$$Q(x) \approx \sum_{\ell} \frac{Z_{\ell}^E \sqrt{\Delta \ell_1 \Delta \ell_2}}{2\pi} e^{ix \cdot \ell} \cos(2\varphi_{\ell}) \sqrt{C_{\ell}^E} \quad (\text{A2})$$

$$U(x) \approx \sum_{\ell} \frac{Z_{\ell}^E \sqrt{\Delta \ell_1 \Delta \ell_2}}{2\pi} e^{ix \cdot \ell} \sin(2\varphi_{\ell}) \sqrt{C_{\ell}^E} \quad (\text{A3})$$

where $\varphi_{\ell} = \tan^{-1}(\ell_2/\ell_1)$; $\Delta \ell_1$, $\Delta \ell_2$ are the frequency spacing in the two coordinate directions; for each ℓ , Z_{ℓ}^T and Z_{ℓ}^E are mean zero complex Gaussian random variables such that $\langle Z_{\ell}^T Z_{\ell'}^{T*} \rangle = \langle Z_{\ell}^E Z_{\ell'}^{E*} \rangle = \delta_{\ell-\ell'}$, $\langle Z_{\ell}^T Z_{\ell'}^{E*} \rangle = \frac{C_{\ell}^{TE}}{\sqrt{C_{\ell}^T} \sqrt{C_{\ell}^E}} \delta_{\ell-\ell'}$, $Z_{-\ell}^T = Z_{\ell}^{T*}$ and $Z_{-\ell}^E = Z_{\ell}^{E*}$. To enforce the proper cross correlation between Z_{ℓ}^T and Z_{ℓ}^E we set

$$\begin{bmatrix} Z_{\ell}^T \\ Z_{\ell}^E \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} -\sqrt{1-\rho} & \sqrt{1+\rho} \\ \sqrt{1-\rho} & \sqrt{1+\rho} \end{bmatrix} \begin{bmatrix} W_{\ell}^1 \\ W_{\ell}^2 \end{bmatrix} \quad (\text{A4})$$

where $\rho \equiv \frac{C_{\ell}^{TE}}{\sqrt{C_{\ell}^T} \sqrt{C_{\ell}^E}}$, and for each ℓ , W_{ℓ}^1 , W_{ℓ}^2 are mean zero complex Gaussian random variables such that $\langle W_{\ell}^1 W_{\ell'}^{1*} \rangle = \delta_{\ell-\ell'}$, $\langle W_{\ell}^2 W_{\ell'}^{2*} \rangle = \delta_{\ell-\ell'}$, $\langle W_{\ell}^1 W_{\ell'}^{2*} \rangle = 0$, $W_{-\ell}^1 = W_{\ell}^{1*}$ and $W_{-\ell}^2 = W_{\ell}^{2*}$.

In our simulation, the above sums—Eqs. (A1)–(A3)—are taken over frequencies $\ell \in \{\frac{2\pi}{L} \mathbf{k} : \mathbf{k} \in \{-N/2, \dots, N/2 - 1\}^2\}$ where $L = 0.2967$ radians so that T will be periodic on $[-L/2, L/2]^2$. The limit $N = L/\Delta_x$ is chosen to match the resolution in pixel space, denoted Δ_x , so that FFT can be used to compute the sums (A1)–(A3) which, after simplification, becomes

$$T(j\Delta_x) \approx L^{-1} \sum_{\mathbf{k}} Z_{(2\pi/L)\mathbf{k}}^T e^{i2\pi\mathbf{k} \cdot \mathbf{j}/N} \sqrt{C_{(2\pi/L)\mathbf{k}}^T} \quad (\text{A5})$$

$$Q(j\Delta_x) \approx L^{-1} \sum_{\mathbf{k}} Z_{(2\pi/L)\mathbf{k}}^E e^{i2\pi\mathbf{k} \cdot \mathbf{j}/N} \cos(2\varphi_{\ell}) \sqrt{C_{(2\pi/L)\mathbf{k}}^E} \quad (\text{A6})$$

$$U(j\Delta_x) \approx L^{-1} \sum_{\mathbf{k}} Z_{(2\pi/L)\mathbf{k}}^E e^{i2\pi\mathbf{k} \cdot \mathbf{j}/N} \sin(2\varphi_{\ell}) \sqrt{C_{(2\pi/L)\mathbf{k}}^E} \quad (\text{A7})$$

for each $\mathbf{j} \in \{N/2, \dots, N/2 - 1\}^2$ where the sums range over $\mathbf{k} \in \{-N/2, \dots, N/2 - 1\}^2$. The matrix of values $[T(j\Delta_x)]_{\mathbf{j} \in \{-N/2, \dots, N/2 - 1\}^2}$, for example, can then be simulated by a two dimensional FFT of the matrix $[L^{-1} Z_{(2\pi/L)\mathbf{k}}^T e^{i2\pi\mathbf{k} \cdot \mathbf{j}/N} \sqrt{C_{(2\pi/L)\mathbf{k}}^T}]_{\mathbf{k} \in \{-N/2, \dots, N/2 - 1\}^2}$. The identities $W_{-\ell}^1 = W_{\ell}^{1*}$ and $W_{-\ell}^2 = W_{\ell}^{2*}$ are enforced using a two dimensional FFT of two $N \times N$ matrices with independent standard Gaussian random entries.

Remark: Typically the above method suffers from an aliasing error when truncating to a finite sum in (A1)–(A3). We avoid any such complication by setting the power spectrum in C_{ℓ}^T , C_{ℓ}^Q and C_{ℓ}^U to zero for all frequencies beyond $|\ell| = 6000$. We justify this truncation since both diffusion damping and the beam FWHM of $0.25'$ combine to produce negligible amplitude in the CMB Stokes parameters at frequencies $|\ell| \geq 6000$ compared to the noise level.

Remark: Since the full sky Stokes parameters T , Q , U are defined on the sphere, the theoretical power spectrum for C_{ℓ}^T , C_{ℓ}^Q , C_{ℓ}^U are only defined on integers ℓ . Our flat sky approximation is obtained by extending C_{ℓ}^T , C_{ℓ}^Q and C_{ℓ}^U to $\ell \in \mathbb{R}^2$ by rounding the magnitude $|\ell|$ to the nearest integer. See Appendix C in [25] for a derivation of this flat sky approximation.

To get a realization of the lensed CMB Stokes parameters T , Q , U we use the above method to generate a high resolution simulation of T , Q , U and the gravitational

potential ϕ on a $17^\circ \times 17^\circ$ patch of the flat sky with 0.25 arcmin pixels. The lensing operation is performed by taking the numerical gradient of ϕ , then using linear interpolation to obtain the values $T(\mathbf{x} + \nabla\phi(\mathbf{x}))$, $Q(\mathbf{x} + \nabla\phi(\mathbf{x}))$, $U(\mathbf{x} + \nabla\phi(\mathbf{x}))$. The beam effect is then performed in Fourier space using FFT of the lensed fields. Finally, we down-sample the lensed fields, every 4th pixel, to obtain the desired arcmin pixel resolution for the simulation output.

APPENDIX B: NEWTON'S METHOD FOR MAXIMIZING THE LOCAL POSTERIOR

In this section we discuss our numerical procedure for maximizing the local posterior given by (4). We remark that calculations need to be fast since they will be performed on each local neighborhood for which a shear and convergence estimate is required. We discuss how the FFT can be used to compute the covariance matrix, denoted $\Sigma_{q^\phi} + N$ in Section II A, and the corresponding derivatives with respect to the unknown coefficients of q^ϕ . We let M be the symmetric 2×2 matrix defined as

$$M \equiv \begin{pmatrix} 1 + q_{xx}^\phi & q_{xy}^\phi \\ q_{xy}^\phi & 1 + q_{yy}^\phi \end{pmatrix}$$

so that $\mathbf{x} + \nabla q^\phi(\mathbf{x}) = M\mathbf{x}$. The matrix M is regarded as the unknown which will be estimated from the data in the local neighborhood $\mathcal{N}_\delta(x_0)$. Let $T_{b,M}$ denote a sheared temperature field, convolved with a Gaussian beam (with standard deviation σ_b) so that

$$T_{b,M}(\mathbf{x}) = \int_{\mathbb{R}^2} d^2\mathbf{y} T(M\mathbf{x} - M\mathbf{y}) e^{-|\mathbf{y}|^2/(2\sigma_b^2)} (\sigma_b^2 2\pi)^{-1}.$$

To compute the covariance matrix of the T observations $\mathbf{t} = (t_1, \dots, t_n)^\dagger$ in $\mathcal{N}_\delta(x_0)$ (see Eq. (5)) one needs to evaluate the following covariance function for a given test shear matrix M at all vector lags $\mathbf{h} = \mathbf{x}_j - \mathbf{x}_k$

$$\begin{aligned} C_{T_{b,M}}(\mathbf{h}) &\equiv \langle T_{b,M}(\mathbf{x} + \mathbf{h}) T_{b,M}(\mathbf{x}) \rangle_{\text{CMB}} \\ &= \int_{\mathbb{R}^2} \frac{d^2\ell}{(2\pi)^2} e^{i\ell \cdot \mathbf{h}} e^{-\sigma_b^2 |\ell|^2} \frac{C_{M^{-1}\ell}^T}{\det M}. \end{aligned}$$

All these calculations can be approximated using a FFT by noticing

$$C_{T_{b,M}}(\mathbf{j}\Delta_x) \approx \sum_{\mathbf{k}} \frac{\Delta_\ell^2}{(2\pi)^2} e^{i2\pi\mathbf{k} \cdot \mathbf{j}/N} e^{-\sigma_b^2 |\Delta_\ell \mathbf{k}|^2} \frac{C_{\Delta_\ell M^{-1}\mathbf{k}}^T}{\det M} \quad (\text{B1})$$

where the sum ranges over $\mathbf{k} \in \{-N/2, \dots, N/2 - 1\}^2$, Δ_x is the pixel spacing, $\mathbf{j} \in \{-N/2, \dots, N/2 - 1\}^2$, $\Delta_\ell = 2\pi/L$ and $L = N\Delta_x$. Then to compute the covariance between t_j and t_k we simply select the entry of the matrix $[C_{T_{b,M}}(\mathbf{j}\Delta_x) + \sigma_T^2 \delta_{j_1 j_2}]_{\mathbf{j} \in \{-N/2, \dots, N/2 - 1\}^2}$ such that

$\mathbf{j}\Delta_x = \mathbf{x}_j - \mathbf{x}_k$ (which was obtained by a single FFT). A similar technique can be used to compute all other covariance and cross-covariances among T , Q and U to construct the covariance matrix Σ_{q^ϕ} . We remark that to speed up the computations we choose a smaller N than the one used in the simulations ($N = 4096$ in the simulation but $N = 256$ for the approximation of $C_{T_{b,M}}(\mathbf{j}\Delta_x)$).

Once the covariance matrix $\Sigma_{q^\phi} + N$ is constructed using the approximation (B1) (and the analogous approximations for Q , U and all cross correlations) the posterior is easily computed as $p(q^\phi | \mathbf{z}) \propto e^{\mathcal{L}(q^\phi | \mathbf{z})} \pi(q^\phi)$ where \mathcal{L} denotes the log likelihood (3) and π is the prior distribution derived in Section II C. In principle, one can now simply use preexisting minimization algorithms for maximizing the posterior $p(q^\phi | \mathbf{z})$ with respect to q^ϕ . If one desires a more sophisticated Newton type algorithm for maximizing the posterior one often needs to compute the gradient and hessian of the posterior. Using the techniques of *automatic differentiation* (see [26], for example) one can easily compute such derivatives if one can compute the rates of change of the covariance Σ_{q^ϕ} with respect to the elements of M .

We finish this Appendix by noticing that the FFT can be used to approximate the derivatives of Σ_{q^ϕ} with respect to the elements of the matrix M , denoted $M_{k,j}$ for $k, j \in \{1, 2\}$. For illustration we focus on the covariance structure of the temperature field T and mention that the extension to Q , U is similar. First notice that by transforming variables $\ell' = M^{-1}\ell$ one gets

$$\begin{aligned} \frac{dC_{T_{b,M}}(\mathbf{h})}{dM_{k,j}} &= \frac{d}{dM_{k,j}} \int \frac{d^2\ell'}{(2\pi)^2} e^{i(M\ell') \cdot \mathbf{h} - \sigma_b^2 |M\ell'|^2} C_{\ell'}^T \\ &= \int \frac{d^2\ell'}{(2\pi)^2} e^{i(M\ell') \cdot \mathbf{h} - \sigma_b^2 |M\ell'|^2} C_{\ell'}^T \\ &\quad \times \frac{d}{dM_{k,j}} [i(M\ell') \cdot \mathbf{h} - \sigma_b^2 |M\ell'|^2]. \end{aligned}$$

Now $d[i(M\ell') \cdot \mathbf{h} - \sigma_b^2 |M\ell'|^2]/dM_{k,j}$ can be written as a sum $\sum_{\mathbf{k}} c_{\mathbf{k}}(\mathbf{h}) g_{\mathbf{k}}(M, \ell')$ so that by retransforming variables to $\ell = M\ell'$ one gets

$$\frac{dC_{T_{b,M}}(\mathbf{h})}{dM_{k,j}} = \sum_{\mathbf{k}} c_{\mathbf{k}}(\mathbf{h}) \int \frac{d^2\ell}{(2\pi)^2} e^{i\ell \cdot \mathbf{h} - \sigma_b^2 |\ell|^2} g_{\mathbf{k}}(M, M^{-1}\ell) \frac{C_{M^{-1}\ell}^T}{\det M}.$$

The point is that the above integrals can now be approximated using the FFT to approximate the matrix of values $[\frac{dC_{T_{b,M}}(\mathbf{j}\Delta_x)}{dM_{k,j}}]_{\mathbf{j} \in \{-N/2, \dots, N/2 - 1\}^2}$. The same method applies to approximate all higher order derivatives of the covariance matrix. These derivatives can then be used in a Newton type algorithm for finding the maximum *a-posteriori* estimates \hat{q}^ϕ .

- [1] M. Kaplinghat, L. Knox, and Y. Song, *Phys. Rev. Lett.* **91**, 241301 (2003).
- [2] K. Smith, W. Hu, and K. Manoj, *Phys. Rev. D* **74**, 123002 (2006).
- [3] M. Kesden, A. Cooray, and M. Kamionkowski, *Phys. Rev. Lett.* **89**, 011304 (2002).
- [4] L. Knox and Y. Song, *Phys. Rev. Lett.* **89**, 011303 (2002).
- [5] U. Seljak and C. Hirata, *Phys. Rev. D* **69**, 043005 (2004).
- [6] C. Hirata and U. Seljak, *Phys. Rev. D* **68**, 083002 (2003).
- [7] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
- [8] M. Zaldarriaga and U. Seljak, *Phys. Rev. D* **59**, 123507 (1999).
- [9] C. M. Hirata, S. Ho, N. Padmanabhan, U. Seljak, and N. A. Bahcall, *Phys. Rev. D* **78**, 043520 (2008).
- [10] K. Smith, O. Zahn, and O. Dore, *Phys. Rev. D* **76**, 043510 (2007).
- [11] J. Guzik, U. Seljak, and M. Zaldarriaga, *Phys. Rev. D* **62**, 043517 (2000).
- [12] U. Seljak and M. Zaldarriaga, *Phys. Rev. Lett.* **82**, 2636 (1999).
- [13] W. Hu, *Astrophys. J.* **557**, L79 (2001).
- [14] T. Okamoto and W. Hu, *Phys. Rev. D* **67**, 083002 (2003).
- [15] M. Bucher, C. S. Carvalho, K. Moodley, and M. Remazeilles, [arXiv:1004.3285](https://arxiv.org/abs/1004.3285).
- [16] C. S. Carvalho and K. Moodley, *Phys. Rev. D* **81**, 123010 (2010).
- [17] C. Hirata and U. Seljak, *Phys. Rev. D* **67**, 043001 (2003).
- [18] S. Dodelson, *Modern Cosmology* (Academic Press, San Diego, 2003).
- [19] M. Kamionkowski, A. Kosowsky, and A. Stebbins, *Phys. Rev. D* **55**, 7368 (1997).
- [20] U. Seljak and M. Zaldarriaga, in *Fundamental Parameters in Cosmology: Proceedings of the XXXIIIrd Rencontres de Moriond 1998* (Editions Frontieres, Paris, 1998).
- [21] B. Horn, *Int. J. Comput. Vis.* **5**, 37 (1990).
- [22] T. Simchony, R. Chellappa, and M. Shao, *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 435 (1990).
- [23] A. Agrawal, R. Raskar, and R. Chellappa, in *Lecture Notes in Computer Science* (Springer, Berlin/Heidelberg, 2006), 3951, pp. 578–591.
- [24] A. Lewis and A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000).
- [25] W. Hu, *Phys. Rev. D* **62**, 043007 (2000).
- [26] R. Neidinger, *SIAM Rev.* **52**, 545 (2010).