

Autodifferentiable likelihood pipeline for the cross-correlation of CMB and large-scale structure due to the kinetic Sunyaev-Zeldovich effect

Yurii Kvasiuk^{*} and Moritz Münchmeyer

Department of Physics, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA



(Received 5 June 2023; accepted 28 February 2024; published 15 April 2024)

We develop an optimization-based maximum likelihood approach to analyze the cross-correlation of the cosmic microwave background (CMB) and large-scale structure induced by the kinetic Sunyaev-Zeldovich (kSZ) effect. Our main goal is to reconstruct the radial velocity field of the Universe. While the existing quadratic estimator (QE) is statistically optimal for current and near-term experiments, the likelihood can extract more signal-to-noise in the future. Our likelihood formulation has further advantages over the QE, such as the possibility of jointly fitting cosmological and astrophysical parameters and the possibility of unifying several different kSZ analyses. We implement an autodifferentiable likelihood pipeline in JAX, which is computationally tractable for a realistic survey size and resolution, and evaluate it on the Agora simulation. We also implement a machine-learning-based estimate of the electron density given an observed galaxy distribution, which can increase the signal-to-noise for both the QE and the likelihood method.

DOI: [10.1103/PhysRevD.109.083515](https://doi.org/10.1103/PhysRevD.109.083515)

I. INTRODUCTION

The kinetic Sunyaev-Zeldovich (kSZ) effect is the dominant contribution to the total blackbody cosmic microwave background (CMB) anisotropies at small scales, overcoming CMB lensing around $\ell \simeq 4000$. It was predicted in the early 1970s [1] and first detected with the Atacama Cosmology Telescope in 2012 [2]. The kSZ effect is a Doppler shifting of CMB photons due to the Thompson scattering on free electrons on the line of sight. The kSZ signal is sensitive to bulk velocities as well as to galaxy and cluster astrophysics. Measurements of the kSZ are able to constrain different dark energy models [3–5], neutrino masses [6,7], modified gravity [8,9], and compensated isocurvature perturbations [10]. Further, the kSZ signal can be used to tightly constrain inflationary models that predict so-called local primordial non-Gaussianity f_{NL} [11,12].

To use the kSZ signal for cosmological constraints, it needs to be combined with a tracer of the electron density field, for example, using a galaxy survey. By cross-correlating the temperature measurement from a CMB experiment with the galaxy field from a large-scale structure survey, one can reconstruct the radial velocity (or momentum) field v_r on large scales [13]. One can then convert this measurement to a measurement of the radial matter overdensity using the continuity equation $\delta_m(k) \sim k^2 v_r(k)$ and obtain an extremely low noise measurement of this field [14].

The canonical way to estimate v_r from the kSZ effect is to construct an analytic quadratic estimator (QE) in spherical harmonics or Fourier space [13–16]. Studies of this method on simulations were presented in [17,18]. Prior to the QE approach, other kSZ estimators that are sensitive to the velocity field were developed also in [19–21]. In the present work, we develop an alternative maximum likelihood approach to kSZ velocity reconstruction. The possible estimators for kSZ velocity reconstruction are similar to those of gravitational lensing in the CMB. In fact, the QE was first introduced in [22] for CMB lensing. In CMB lensing, the QE is optimal at low signal-to-noise, but at higher signal-to-noise (in particular where the total CMB anisotropy is dominated by lensing) it is substantially suboptimal. For this reason, several authors have developed a maximum likelihood approach to CMB lensing, which is optimal at all noise levels, but comes at the cost of a computationally expensive optimization process. The first maximum likelihood estimator for lensing was developed in [23]. A closely related method and code implementation (LENSIT) to find the maximum *a posteriori* (MAP) iteratively was presented in [24]. A different lensing likelihood was presented in [25,26], with associated code CMBLENSING.JL, and recently applied to South Pole Telescope data using Hamiltonian Monte Carlo (HMC) to sample the likelihood [27]. For the small high-resolution sky patch analyzed in [27], it was shown that the likelihood outperforms the QE on real data. For CMB S4, it is expected that lensing likelihood methods can improve the signal-to-noise over the QE by up to a factor of 2 [28].

^{*}Corresponding author: kvasiuk@wisc.edu

The goal of our present work is to develop a similar likelihood approach to kSZ velocity reconstruction or, more generally, to the cross-correlation of the kSZ signal with a large-scale structure survey. As we will see, at high enough signal-to-noise, the likelihood again outperforms the quadratic estimator substantially. There are, however, some important differences compared to the CMB lensing likelihood. CMB lensing is a nonlinear remapping of the primary CMB temperature, while the kSZ, at leading order, is linearly added to the primary CMB. This should make it easier to optimize the likelihood. Further, one needs to model the relation of the galaxy distribution and the electron distribution in the likelihood, which can only be approximated. Recently, a different maximum likelihood kSZ estimator was proposed in [29]. However, this estimator is not optimization-based and instead makes some analytic assumptions that allow for an analytic solution. We will comment below in more detail on the relation to this prior work. Another related work is [30], which developed a pipeline to detect the kSZ signal using a forward-modeling-based velocity template and a kSZ model at the cluster level.

A number of autodifferentiable forward models for cosmology have recently been proposed in other contexts. In [31] an autodifferentiable likelihood is used to analyze gravitational lensing of the CMB and patchy screening during the epoch of reionization. For the study of large scale structure (LSS) formation, pmwd [32], a differentiable cosmological particle mesh, implemented with JAX, has been presented. In the context of weak gravitational galaxy lensing, recent developments include MADLENS [33], a package for autodifferentiable calculation of non-Gaussian convergence maps, and DLL [34], a GPU-based differentiable weak lensing forward model which is differentiable with respect to all cosmological parameters.

In the present work we develop an autodifferentiable approach to kSZ analysis and aim to answer the following main questions:

- (i) What likelihood is the most useful for our purpose, and what choices can be made? We perform a separation of scales between the large-scale (linear) velocity field and the small-scale galaxy/electron distribution, which avoids the need for nonlinear forward modeling of large-scale structures. We also discuss different assumptions for the statistical relation of galaxies and electrons.
- (ii) Can we consider a realistically sized dataset (angular resolution and redshift bins) and find the MAP effectively on current hardware? To achieve this, we implement the likelihood using the autodifferentiable language JAX, which provides a highly optimized GPU implementation.
- (iii) How does the signal-to-noise in the likelihood compare with the quadratic estimator as a function of experimental parameters? As we discuss below,

for realistic current and next-generation data the kSZ likelihood can probably not outperform the QE (see below for details); however, we show that for more futuristic experiments the improvement can be substantial, despite the presence of high redshift kSZ from reionization.

- (iv) Can machine learning be used to raise the signal-to-noise by learning an improved template of the electron distribution given observed galaxies, by training on simulations? While a full exploration of this idea requires a study on hydrodynamic simulations, we show encouraging initial results using the matter distribution that could already improve the velocity reconstruction for Simons Observatory.

We make forecasts for a number of idealized experimental configurations. Our simulation study is not fully realistic, since we do not take into account effects such as photo- z errors (which would degrade forecasts) or halo mass tracers (which could improve them). For the upcoming Simons Observatory or CMB-S4 combined with the Rubin Observatory, we do not predict a signal-to-noise improvement of the likelihood over the QE, but our results are not fully conclusive because we have not analyzed a simulation with the full halo density of the Rubin Observatory. However, even in the case where the QE is statistically optimal, an optimization-based formulation has several advantages over the QE, some of which we will study in future work:

- (i) The likelihood can include both cosmological and astrophysical parameters, and it becomes thus possible to jointly fit both of these and obtain combined constraints with covariances. In the present work, we do not yet implement such a joint fit, but this is a main goal for future work. This could also include a model of the reionization kSZ.
- (ii) A likelihood approach can naturally take into account systematic uncertainties, such as calibration issues, because it allows one to fit a joint model to the data that takes into account all known experimental effects. This was recently demonstrated for CMB lensing in [26].
- (iii) The kSZ likelihood can include several different estimators. For example, the likelihood directly provides us with a partially kSZ-cleaned primary CMB map. So-called de-kSZing has recently been studied in [35] using a template method. The signal extracted by squared kSZ statistics (projected fields) [36], which can probe the baryon distribution, is also included in the likelihood formulation because the likelihood includes a complete probabilistic model of all relevant fields.
- (iv) It is conceptually straightforward (but practically and computationally difficult) to combine several CMB secondary effects into an overall $\text{CMB} \times \text{LSS}$

likelihood. A joint lensing-kSZ likelihood could estimate both the lensing potential and the radial velocities, include the integrated Sachs-Wolfe effect (e.g., [37]) and the moving lens effect (e.g., [38]), include lensing of the kSZ, and include nonblack-body contributions such as the thermal Sunyaev-Zeldovich effect.

- (v) In the likelihood formulation, we can use non-Gaussian priors that capture the non-Gaussian small-scale statistics of the galaxy and electron density. A non-Gaussian prior could be implemented, for example, with a normalizing flow [39,40], trained on hydrodynamic simulations, and can increase the signal-to-noise over a Gaussian prior, since these fields are substantially non-Gaussian.
- (vi) The kSZ likelihood can be included in a more general forward modeling setup. In forward modeling of large-scale structure, one starts with the Gaussian initial conditions and then maps them to late-time observables with a differentiable forward model of structure formation (see, e.g., [41,42]). Radial velocities were recently used for reconstruction of the initial conditions in [43].
- (vii) Once a likelihood is available, one can, in principle not only find the MAP but also perform a fully Bayesian analysis, for example, using variants of HMC or variational inference. It seems unlikely that this is computationally tractable at the full resolution of the data, but it may be possible to develop useful approximations, for example, by approximately integrating out the small-scale fields.

The paper is organized as follows. In Sec. II we set up our notation and review the conventional quadratic estimator formulation of kSZ velocity reconstruction in our coordinates. In Sec. III we discuss the kSZ likelihood, its relation to the QE, its implementation, and how to find the MAP of field parameters and scalar parameters. In Sec. IV we apply our method to simulations and evaluate the improvement factor over the QE. In Sec. V we discuss a machine learning method, which can improve the velocity reconstruction for either the QE or the likelihood. Our conclusions and future work are summarized in Sec. VI.

II. REVIEW OF THE QUADRATIC ESTIMATOR APPROACH

We start by reviewing the standard kSZ velocity reconstruction. For computational simplicity, we will work in binned flat-sky coordinates. This is appropriate, in particular, for a photometric survey such as Rubin Observatory. Redshift binning does not lose significant information compared to a full three-dimensional analysis if the photometric error bars are larger than the redshift binning.

A. Notation

We adopt the following Fourier conventions for scalar fields in a 2D flat-sky Cartesian basis:

$$f(\mathbf{x}) = \frac{1}{(2\pi)^2} \int e^{i\mathbf{l}\cdot\mathbf{x}} f(\mathbf{l}) d^2l \quad f(\mathbf{l}) = \int e^{-i\mathbf{l}\cdot\mathbf{x}} f(\mathbf{x}) d^2x. \quad (1)$$

The 2D power spectrum $P^f(l)^1$ is defined as

$$\langle f(\mathbf{l}_1) f^*(\mathbf{l}_2) \rangle = (2\pi)^2 \delta^{(2)}(\mathbf{l}_1 - \mathbf{l}_2) P^f(l_1). \quad (2)$$

In this work, we will use $\frac{\Delta T}{T} \equiv \theta$ for the CMB anisotropy; matter (galaxy, electron) overdensity is denoted by $\delta^{m,(g,e)} \equiv \rho^{m,(g,e)} / \bar{\rho}^{m,(g,e)} - 1$ and should not be confused with 2D Dirac $\delta^{2D}(\mathbf{x})$ or Kronecker symbol $\delta_{\alpha\beta}^{(K)}$. The radial velocity is denoted by v_r and sometimes we will omit the superscript r for notational simplicity. The data required in this work are the CMB temperature fluctuation $\theta(\mathbf{l})$ and the binned galaxy field $\delta_g^\alpha(\mathbf{l})$.

B. Binned kSZ signal

The temperature anisotropy generated by the kSZ is proportional to the electron density and the velocity and is given by the following line-of-sight integral:

$$\theta(\hat{n})|_{\text{kSZ}} = -\sigma_T \int d\chi a n_e(\hat{n}, \chi) v_r(\hat{n}, \chi), \quad (3)$$

over comoving distance χ . In this equation, σ_T is the Thomson scattering cross section, a is the scale factor [$a(z=0)=1$], and n_e is the electron number density. We now assume that we can probe the electron density n_e with a finite radial resolution, given roughly by the redshift error of the experiment (or below by the radial binning of the available simulations). Given a set of radial bins indexed by α (60 bins for $0.5 < z < 3$ in our main analysis), we can write the binned kSZ as [13]

$$\theta(\mathbf{x})|_{\text{kSZ}}^{\text{binned}} = \sum_{\alpha=1}^{N_{\text{bins}}} \tau^\alpha(\mathbf{x}) v_r^\alpha(\mathbf{x}). \quad (4)$$

Here $\tau^\alpha(\mathbf{x})$ and $v_r^\alpha(\mathbf{x})$ are the integrated optical depth and radial velocity in the bin α as functions of 2D Cartesian coordinate \mathbf{x} . We further assume that

$$\tau^\alpha(\mathbf{x}) = f^\alpha(1 + \delta_e^\alpha(\mathbf{x})), \quad (5)$$

where δ_e is the electron overdensity integrated over the bin. Our binning is defined with a top-hat window function with the width of the bin Δ_α ,

¹Relations between 2D flat-sky, 2D angular, and 3D power spectra are reviewed in Appendix A 1 and A 2.

$$\delta_m^\alpha = \frac{1}{\Delta_\alpha} \int_{z_\alpha^* - \Delta_\alpha/2}^{z_\alpha^* + \Delta_\alpha/2} dz \delta_m(z). \quad (6)$$

The binned kSZ from the radially averaged τ^α and v_r^α do not constitute the total kSZ that an experiment will see. Small-scale fluctuations within the bin will contribute additional kSZ, which however cannot be described by the binned data. To have a realistic kSZ power in the map, including contributions from smaller radial scales than the bin width, occasionally we will add an additional “uncorrelated” kSZ component. This component also includes kSZ from redshifts where we have no galaxy data, including reionization. Thus we model the total kSZ temperature anisotropy as

$$\theta(\mathbf{x})|_{\text{kSZ}}^{\text{total}} = \theta(\mathbf{x})^{\text{kSZ,binned}} + \theta(\mathbf{x})^{\text{kSZ,uncorr}}. \quad (7)$$

We model $\theta(\mathbf{x})|_{\text{kSZ}}^{\text{uncorr}}$ as a Gaussian random field, which is uncorrelated with the binned galaxy and velocity fields, and adjust its power spectrum to have a realistic total kSZ amplitude in our maps. The total measured CMB temperature is thus

$$\theta_{\text{obs}}(\mathbf{l}) = \theta(\mathbf{l})^{\text{pLCMB}} + \theta(\mathbf{l})^{\text{kSZ,binned}} + \theta(\mathbf{l})^{\text{kSZ,uncorr}} + n_\theta(\mathbf{l}), \quad (8)$$

which includes contributions from lensed primary CMB, kSZ effect, and noise. We write $\tilde{P}^\theta(l)$ for its corresponding power spectrum.

C. Flat-sky quadratic estimator for the velocity field

We now derive the quadratic estimator in binned flat-sky coordinates (see Appendix A), which closely follows the binned spherical harmonics derivation in [13]. The kSZ-induced correlation between total CMB temperature perturbation field $\theta_{\text{obs}}(\mathbf{l})$ and galaxy overdensity bin $\delta_{\text{obs}}^{g,\alpha}(\mathbf{l}')$ (the tildes here mean including instrumental noise) is

$$\begin{aligned} & \langle \theta_{\text{obs}}(\mathbf{l}_1) \delta_{\text{obs}}^{g,\alpha}(\mathbf{l}_2) \rangle \\ &= \sum_\beta \frac{1}{(2\pi)^4} \int d^2l' d^2l'' (2\pi)^2 \delta^{2D}(\mathbf{l}_1 - \mathbf{l}' - \mathbf{l}'') v_\beta(\mathbf{l}'') \\ & \quad \times \langle \tau_\beta(\mathbf{l}') \delta^{g,\alpha}(\mathbf{l}_2) \rangle \\ &= \int d^2l' d^2l'' \delta^{2D}(\mathbf{l}_1 - \mathbf{l}' - \mathbf{l}'') v_\beta(\mathbf{l}'') P_\alpha^{\tau g}(\mathbf{l}') \delta^{2D}(\mathbf{l}_2 + \mathbf{l}') \\ &= \int d^2l \delta^{2D}(\mathbf{l}_1 + \mathbf{l}_2 - \mathbf{l}) v_\alpha(\mathbf{l}) P_\alpha^{\tau g}(\mathbf{l}_2), \end{aligned} \quad (9)$$

where, under the assumption of independent bins, we defined

$$\langle \tau_\beta(\mathbf{l}') \delta_\alpha^g(\mathbf{l}) \rangle = (2\pi)^2 P_\alpha^{\tau g}(\mathbf{l}) \delta_{\alpha\beta}^{(K)} \delta^{2D}(\mathbf{l} + \mathbf{l}'). \quad (10)$$

The general quadratic estimator is then of the form

$$\hat{v}_\alpha(\mathbf{l}) = \int d^2l_1 d^2l_2 W_\alpha(\mathbf{l}_1, \mathbf{l}_2) \theta_{\text{obs}}(\mathbf{l}_1) \delta_{\text{obs}}^{g,\alpha}(\mathbf{l}_2) \delta^{2D}(\mathbf{l} - \mathbf{l}_1 - \mathbf{l}_2). \quad (11)$$

We demand that $\langle \hat{v}_\alpha(\mathbf{l}) \rangle = v_\alpha(\mathbf{l})$ and look for $\hat{v}_\alpha(\mathbf{l})$ that has the smallest possible variance under the assumption that fields are Gaussian. That yields

$$W_\alpha(\mathbf{l}_1, \mathbf{l}_2) = \lambda(l) \frac{P_\alpha^{\tau g}(\mathbf{l}_2)}{\tilde{P}^\theta(\mathbf{l}_1) \tilde{P}_\alpha^g(\mathbf{l}_2)}, \quad (12)$$

where $\lambda(l)$ is a normalization,

$$\lambda^{-1}(l) = \int d^2l_1 d^2l_2 \delta^{2D}(\mathbf{l} - \mathbf{l}_1 - \mathbf{l}_2) \frac{(P_\alpha^{\tau g}(\mathbf{l}_2))^2}{\tilde{P}^\theta(\mathbf{l}_1) \tilde{P}_\alpha^g(\mathbf{l}_2)}. \quad (13)$$

Assuming that the contribution to the kSZ signal comes only from small scales $|\mathbf{l}_2| \gg |\mathbf{l}|$, we can simplify this expression to get

$$\lambda^{-1}(l) \approx 2\pi \int d^2l_2 \frac{(P_\alpha^{\tau g}(\mathbf{l}_2))^2}{\tilde{P}^\theta(\mathbf{l}_2) \tilde{P}_\alpha^g(\mathbf{l}_2)}. \quad (14)$$

When calculating the integral for $v_\alpha(\mathbf{l})$, we notice that weights neatly factorize: $W_\alpha(\mathbf{l}_1, \mathbf{l}_2) = W_\theta(\mathbf{l}_1) W_{g,\alpha}(\mathbf{l}_2)$. We thus find the quadratic estimator of the radial velocity field,²

$$\hat{v}_\alpha(\mathbf{l}) = (2\pi)^2 \lambda(l) \int d^2x [\theta_{\text{obs}}(\mathbf{x})]_{W_\theta(\mathbf{l})} [\delta_{\text{obs}}^{g,\alpha}(\mathbf{x})]_{W_{g,\alpha}(\mathbf{l})}, \quad (16)$$

²In [15] it was discussed that it is somewhat ambiguous what the true velocity field v_r is, which \hat{v}_r is supposed to reconstruct. In the case of N-body simulations, which are a collection of particles, one most directly obtains the sum over particles

$$\mathbf{q}(\mathbf{x}) = \frac{1}{n_p} \sum_i \mathbf{v}_i \delta^3(\mathbf{x} - \mathbf{x}_i), \quad (15)$$

where n_p is the number density average over the whole volume. This quantity is the particle or mass-weighted velocity field, i.e., the “momentum field.” From the momentum field, one can define the “velocity field” by choosing a smoothing scale and defining the velocity to be the smoothed momentum, divided by the smoothed density (appropriately regulated to avoid dividing by zero in voids). In [15] it was shown that the momentum field correlates somewhat better with the estimator \hat{v}_r than the velocity field defined in this way. However, on the theoretical side, it is easier to work with velocities rather than momenta since they are first order in the perturbations. For this reason, in this work, we use the velocity field rather than the momentum field. On large scales, the difference between the two quantities is small, because the anisotropy in the radial momentum field $q_r = (1 + \delta)v_r$ is dominated by the larger fluctuations in v_r .

where

$$[f(\mathbf{x})]_{W(l)} \equiv \frac{1}{(2\pi)^2} \int d^2 l f(l) W(l) e^{i l \mathbf{x}}; \quad W_\theta(l) = \frac{1}{\tilde{P}^\theta(l)};$$

$$W_{g,\alpha}(l) = \frac{P_\alpha^{g\tau}(l)}{\tilde{P}_\alpha^g(l)}. \quad (17)$$

The prefactor $(2\pi)^2 \lambda(l)$ is the noise of the quadratic estimator. We can split the correlation function of QE velocities into two terms: $\langle \hat{v}_\alpha(l) \hat{v}_\beta(l') \rangle = (N^{\text{QE}} + S^{\text{QE}}) (2\pi)^2 \delta^{2D}(l + l')$. Then the expression for noise is as follows:

$$N^{\text{QE}} = (2\pi)^2 \left(\int d^2 l_1 d^2 l_2 \delta^{2D}(l - l_1 - l_2) \frac{(P_\alpha^{g\tau}(l_2))^2}{\tilde{P}^\theta(l_1) \tilde{P}_\alpha^g(l_2)} \right)^{-1}. \quad (18)$$

The noise thus depends on the CMB noise $\theta(l)^N$ as well as the noise of the large-scale structure tracer. As a large-scale structure tracer, we use the halo (or galaxy) field $\delta_{\text{obs}}^{g,\alpha}(l) = \delta_\alpha^g(l) + n^g(l)$. Here the stochastic shot noise term $n^g(l)$ comes from discrete sampling and has a power spectrum $N^g(l) = n(z)^{-1}$, where $n(z)$ is number of halos per redshift interval per steradian. In this derivation, we assumed that $\langle \tau_\alpha \delta_\beta^g \rangle \propto \delta_{\alpha,\beta}^{(K)}$, i.e., that small-scale matter fields are uncorrelated between different redshift bins. In [16] it was shown how to include redshift bin correlations of the small-scale fields in the QE. However, we found that, for our study with 60 redshift bins, this correlation is small and we neglect it both in our QE implementation and our likelihood.

III. LIKELIHOOD APPROACH

We now discuss the different elements of our novel likelihood formulation. We discuss several different likelihoods with different properties, in increasing order of complexity.

A. Likelihood for the velocity field from a kSZ observation assuming a known matter field

We can write a maximum likelihood estimator for the velocities by writing down the Gaussian likelihood of CMB perturbations including the kSZ,

$$-2 \ln \mathcal{L}(\theta^{\text{obs}} | v_r) = (\theta^{\text{obs}} - \theta^{\text{kSZ}}(\delta_e, v_r))^T \times P_\theta^{-1} (\theta^{\text{obs}} - \theta^{\text{kSZ}}(\delta_e, v_r)) + \text{const.} \quad (19)$$

Here $\theta^{\text{obs}} = \theta^{\text{pCMB}} + \theta^{\text{kSZ}} + n_\theta$ is the lensed primary CMB with added noise and $P = S + N$ is the covariance of $\theta^{\text{pCMB}} + n_\theta$ and θ^{pCMB} can also include uncorrelated kSZ

such as from reionization (however, assuming Gaussianity thereof). We use a basis-independent notation where θ are 2D fields and δ_e and v_r are redshift binned 2D fields, both represented as a vector.

The kSZ is given in position space by the radial integral

$$\theta^{\text{kSZ}}(n) = f_\tau \int dr \delta_e(r, n) v_r(r, n) \quad (20)$$

and we can write it as basis independent in index notation as

$$\theta_l^{\text{kSZ}} = K_l^{x,y} \delta_x v_y, \quad (21)$$

where $K_l^{x,y}$ is the kSZ projection matrix.

This likelihood is the same as in [29]. Assuming that one has an estimate of δ_e , for example, from a galaxy survey (see below), the maximum likelihood estimator for the velocities can, in principle, be found analytically by evaluating

$$\frac{\delta \mathcal{L}}{\delta v} \stackrel{!}{=} 0. \quad (22)$$

However, in the case of multiple redshift bins, this estimator is ill defined, and even for a single bin it gives a poor reconstruction in terms of residuals with the truth. In [29] this problem was circumvented by introducing a coarse-graining procedure. Here we follow a more Bayesian approach and include a Gaussian prior on the velocities, which is physically appropriate at large scales. The posterior is then

$$-2 \ln \mathcal{P}(\theta^{\text{obs}} | v_r) = (\theta^{\text{obs}} - \theta^{\text{kSZ}}(\delta_e, v_r))^T \times P_\theta^{-1} (\theta^{\text{obs}} - \theta^{\text{kSZ}}(\delta_e, v_r)) + v^T P_v^{-1} v + \text{const}, \quad (23)$$

and the MAP is given by

$$\frac{\delta \mathcal{P}}{\delta v} \stackrel{!}{=} 0. \quad (24)$$

We will derive an analytic solution for this case in the next section, however, in practice we maximize the posterior numerically. The simple posterior discussed in this section has the advantage that there is only a single unknown field, the velocities v . Below we will jointly fit the matter and velocity fields as probabilistic degrees of freedom and we will not use this likelihood in our main implementation. However, in Appendix B we use the posterior in Eq. (23) to demonstrate that we can obtain correct error bars from its Hessian.

B. Analytic MAP estimator for the velocity field

While ultimately we find MAP solutions numerically, it is interesting to consider whether there is an analytic MAP estimator. As we shall see, this is possible in the case of a known τ field (or one that has been reconstructed from galaxies independent of the kSZ data). The estimator we develop here is different from the estimator in [29] in that we include a Gaussian prior on the large-scale velocity field, thus obtaining a maximum *a posteriori* (rather than maximum likelihood) estimator.

We assume that we observe some θ_o which is sourced by θ_p , primary CMB anisotropy, and θ_{kSZ} with Gaussian noise with known covariance N_θ . In this section, we write the binned electron field in terms of τ rather than in terms of δ_e for compactness of notation. The two quantities are related by Eq. (5). Following our discretization of the radial kSZ integral (4), in position space we have $\theta_{\text{kSZ}} = \sum_{i=1}^{N_{\text{bins}}} v_i \tau_i$, with both v_i and τ_i being 2D fields. We include Gaussian priors on velocities and primary CMB and assume that we know the optical depth τ . Then, we can write a posterior,

$$\mathcal{P}(v_i, \theta_p | \theta_o, \tau_i) \simeq e^{-\frac{1}{2}[(\theta_o - \theta_p - \tau_i v^i)^T N_\theta^{-1} (\theta_o - \theta_p - \tau_i v^i) + (v^T)^i (C_v^{-1})_{ij} v^j + \theta_p^T C_\theta^{-1} \theta_p]}. \quad (25)$$

Weak equality \simeq here means equality up to a constant, multiplicative normalization (or additive constant if log quantities are considered) of $\mathcal{N}^{-1} = \sqrt{\det C_\theta} \times \sqrt{\det C_v} \times \sqrt{\det C_\theta}$. In the expression of a type $\theta_p^T N_\theta^{-1} \theta_p$, a sum over some discretization (pixels) of a field is implied, i.e., $\theta^T N_\theta^{-1} \theta \equiv \sum_{p_i, p_j}^{N_{\text{pix}}} (\theta)_{p_i} (\theta)_{p_j} (N_\theta^{-1})_{p_i, p_j}$. We are interested in the MAP for v_i , hence we are going to neglect the normalization (this is possible since v_i does not appear in covariance) and marginalize over the primary CMB θ_p ,

$$\mathcal{P}(v_i | \theta_o, \tau_i) = \int D\theta_p \mathcal{P}(v_i, \theta_p | \theta_o, \tau_i). \quad (26)$$

Let us consider the θ_p part,

$$(-2 \ln \mathcal{P}(v_i, \theta_p | \theta_o, \tau_i))_{\theta_p} \simeq [\theta_p^T C_\theta^{-1} \theta_p + \theta_p^T N_\theta^{-1} \theta_p - (\theta_o - \tau_i v^i)^T N_\theta^{-1} \theta_p - \theta_p^T N_\theta^{-1} (\theta_o - \tau_i v^i)]. \quad (27)$$

We define $\tilde{C}_\theta^{-1} \equiv C_\theta^{-1} + N_\theta^{-1}$ and $\tilde{\theta} \equiv \tilde{C}_\theta N^{-1} (\theta_o - \tau_i v^i)$. Then we can rewrite

$$(-2 \ln \mathcal{P}(v_i, \theta_p | \theta_o, \tau_i))_{\theta_p} \simeq \theta_p^T \tilde{C}_\theta^{-1} \theta_p - \tilde{\theta}^T \tilde{C}_\theta^{-1} \theta_p - \theta_p^T \tilde{C}_\theta^{-1} \tilde{\theta} = (\theta_p - \tilde{\theta})^T \tilde{C}_\theta^{-1} (\theta_p - \tilde{\theta}) - \tilde{\theta}^T \tilde{C}_\theta^{-1} \tilde{\theta}. \quad (28)$$

Now we can perform marginalization over θ_p easily. Since $\int D\theta_p e^{-\frac{1}{2}(\theta_p - \tilde{\theta})^T \tilde{C}_\theta^{-1} (\theta_p - \tilde{\theta})} \simeq 1$, we are left with

$$-2 \ln \mathcal{P}(v_i | \theta_o, \tau_i) \simeq (\theta_o - \tau_i v^i)^T N_\theta^{-1} (\theta_o - \tau_i v^i) - (\theta_o - \tau_i v^i)^T N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1} (\theta_o - \tau_i v^i) + (v^T)^i (C_v^{-1})_{ij} v^j. \quad (29)$$

We can try to rewrite this expression as we did before for θ_p . The velocity part is as follows:

$$(-2 \ln \mathcal{P}(v_i | \theta_o, \tau_i))_v \simeq (v^T)^i ((C_v^{-1})_{ij} + (\tau^T)_i (N_\theta^{-1} - N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1}) \tau_j) v^j \quad (30)$$

$$- (v^T)^i (\tau^T)_i (N_\theta^{-1} - N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1}) \theta_o - \theta_o^T (N_\theta^{-1} - N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1}) \tau_j v^j. \quad (31)$$

Assuming that the inverse exists, we can define a τ -dependent covariance,

$$(\tilde{C}_v^{-1})_{ij} \equiv (C_v^{-1})_{ij} + (\tau^T)_i (N_\theta^{-1} - N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1}) \tau_j. \quad (32)$$

Then \tilde{v}_i , which is the MAP estimator we were looking for, is given by

$$\tilde{v}_i = (\tilde{C}_v^{-1})_i^j (\tau^T)_j (N_\theta^{-1} - N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1}) \theta_o. \quad (33)$$

Now the log-likelihood takes the following form:

$$-2 \ln \mathcal{P}(v_i | \theta_o, \tau_i) \simeq (v^T - \tilde{v}^T)^i (\tilde{C}_v^{-1})_{ij} (v - \tilde{v})_j + (\tilde{v}^T)^i (\tilde{C}_v^{-1})_{ij} (\tilde{v})_j + \theta_o^T (N_\theta^{-1} - N_\theta^{-1} \tilde{C}_\theta N_\theta^{-1}) \theta_o. \quad (34)$$

It is worth noting that had we instead treated τ_i as unknown, we could have added $-2 \ln \mathcal{P}(\tau_i | \tau_o) = (\tau_o^T - \tau^T)_i (N_\tau^{-1})^{ij} (\tau_o - \tau)_j + (\tau^T)_i (C_\tau^{-1})^{ij} (\tau)_j$ and, after marginalizing over θ_p , marginalized over τ . That procedure would result in a posterior for velocity $\mathcal{P}(v_i | \theta_o, \tau_o)$ that depends only on observed data. However, finding an analytical MAP would not be

possible in that case, which is consistent with the fact that it is not possible to further analytically marginalize the posterior that we obtained over τ_i . We also note that our MAP depends on an inverse of the operator that is explicitly τ_i dependent. It is worth adding that one way to proceed with an analytic MAP for the velocity, in this case, is to find the MAP of τ_i alone from galaxy survey data and use those values to evaluate the covariance operator [Eq. (32)].

Our result is consistent with the results of [29], where the inverse was computed via the “coarse-graining” procedure. Indeed, Eq. (14) in [29], which is the maximum likelihood estimator (MLE) condition there, reads as follows:

$$(\tau^\alpha)^\dagger (C^{\text{pCMB}})^{-1} \Theta = (\tau^\alpha)^\dagger (C^{\text{pCMB}})^{-1} (\tau^\beta) v_\beta. \quad (35)$$

That is equivalent to our case if we neglect the prior and consider that we observe θ^p . As was mentioned in the same work, QE is obtained if we approximate $(\tau^\alpha)^\dagger (C^{\text{pCMB}})^{-1} (\tau^\beta)$ with $\langle (\tau^\alpha)^\dagger (C^{\text{pCMB}})^{-1} (\tau^\beta) \rangle_\tau$. In that case, the covariance is independent of τ and the resulting expression for MAP v is quadratic in the observed fields.

Our analytic MAP estimator given in Eq. (33) is not easy to evaluate in practice because it includes several inverses of large matrices. For this reason, even for the velocity field likelihood presented in this section, we find the MAP by gradient descent rather than from the analytic expression.

C. Joint likelihood for CMB and matter

We now describe a likelihood where the fields $v, \delta_m, \theta^{\text{CMB}}$ are treated probabilistically. In this section, we assume that $\delta_e = \delta_m$ and that we can directly observe δ_m with some Gaussian noise (which we will take to have a flat shot noise power spectrum). We also assume that both the galaxy survey and the CMB experiment have mutually uncorrelated Gaussian noise. The likelihood is then given by

$$\begin{aligned} & -2 \ln \mathcal{L}(\theta^{\text{obs}}, \delta_m^{\text{obs}} | \theta^{\text{pCMB}}, \delta_m, v_r) \\ &= (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_m, v_r))^T N_\theta^{-1} \\ & \quad \times (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_m, v_r)) \\ & \quad + (\delta_m^{\text{obs}} - \delta_m)^T N_m^{-1} (\delta_m^{\text{obs}} - \delta_m) + \text{const.} \end{aligned} \quad (36)$$

We again use a basis-independent notation where θ is a 2D field and δ_m and v_r are redshift binned 2D fields, both represented as a vector.

To get the MAP estimate, we maximize the product of the likelihood times priors for the unobserved quantities. By Bayes theorem,

$$\begin{aligned} \mathcal{P}(\theta^{\text{CMB}}, \delta_m, v_r | \theta^{\text{obs}}, \delta_m^{\text{obs}}) &\propto \mathcal{L}(\theta^{\text{obs}}, \delta_m^{\text{obs}} | \theta^{\text{CMB}}, \delta_m, v_r) \\ &\quad \times \mathcal{P}(T^{\text{CMB}}, \delta_m, v_r). \end{aligned} \quad (37)$$

If we assume Gaussian priors for the fields, we get

$$\begin{aligned} -2 \ln \mathcal{P}(\theta^{\text{CMB}}, \delta_m, v_r) &= (\theta^{\text{CMB}})^T P_\theta^{-1} \theta^{\text{CMB}} + v^T P_v^{-1} v \\ &\quad + \delta_m^T P_m^{-1} \delta_m + \text{const}, \end{aligned} \quad (38)$$

where P_x are the signal covariance matrices (power spectra) for fiducial cosmological parameters. In this prior we have assumed a separation of scales, i.e., we draw the velocities v_r^L from a Gaussian field on large scales, and we draw the matter field δ_m^S on small scales only. In our implementation, we split these scales at $\ell_{\text{split}} = 700$. The analysis is not sensitive to the precise choice of ℓ_{split} because kSZ is induced by significantly smaller-scale fluctuations of matter overdensity that are modulated by substantially larger-scale velocity fluctuations. In addition, we can only observe kSZ at much higher ℓ (where it becomes a dominant anisotropy) and we can only reconstruct velocities at much lower ℓ (due to signal-to-noise). Our likelihood thus treats the velocities on large scales v^L independent from the matter field on small scales δ_m^S . This neglects the gravitation-induced correlation between v^L and δ_m^S , such as the bispectrum of form $\langle v^L \delta_m^S \delta_m^S \rangle$. However, these squeezed limit correlations are small and can be neglected. For some applications of parameter fitting, one would also like to include information from the large-scale galaxy field and enforce a relation between matter density and velocity (i.e., the continuity equation on large scales). However, for our present purpose of reconstructing the large-scale velocity field, such an extended likelihood is not required.

We now want to maximize the posterior with respect to $v_r, \delta_m, \theta^{\text{CMB}}$. Unlike the case in Sec. III B, there is no analytic solution to the MAP. We thus use optimization to find the MAP of the fields. In particular, δ_m is now constrained from both the observed matter data and the observed kSZ data (for example, a cluster is visible both to a galaxy survey and a high-resolution CMB experiment, so they contain mutual information). Our likelihood formulation is optimal if the noise covariances N and signal covariances (in the prior) are correct, including bin-to-bin correlations in the signal. This assumes Gaussianity in all fields. We will discuss how to relax this assumption with machine learning below.

D. Joint likelihood for CMB and galaxies

In realistic experiments, our datasets are the observed CMB temperature anisotropy θ^{obs} and the observed galaxy overdensity δ_g^{obs} (rather than δ_m or δ_e). We assume that these quantities are determined by the following underlying unobserved fields: true radial velocity v_r , true galaxy overdensity δ_g (without stochastic shot noise), true electron density δ_e , and true lensed primary CMB temperature θ^{pCMB} . We also assume that both the galaxy survey and the CMB experiment have mutually uncorrelated Gaussian noise. In this case the log-likelihood of seeing θ^{obs} and δ_g^{obs} for given $v_r, \delta_g, \delta_e, \theta^{\text{pCMB}}$ is the following:

$$\begin{aligned}
& -2 \ln \mathcal{L}(\theta^{\text{obs}}, \delta_g^{\text{obs}}) \\
& = (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_e, v_r))^T \\
& \quad \times N_T^{-1} (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_e, v_r)) \\
& \quad + (\delta_g^{\text{obs}} - \delta_g)^T N_g^{-1} (\delta_g^{\text{obs}} - \delta_g) + \text{const}, \quad (39)
\end{aligned}$$

where N_T is the noise covariance of the CMB experiment and N_g is the shot noise of the galaxy survey due to the finite number of observed galaxies. A crucial question is now how to relate the galaxy density δ_g and the unobserved electron density δ_e . There are two general options how to accomplish this, deterministic or stochastic, as we now discuss. In our numerical tests below we will only implement the deterministic approach.

In the expression above we have assumed that the noise of the CMB field and the noise of the galaxy field are uncorrelated. This holds exactly true in our main analysis (Sec. IV C) because we approximate the galaxy noise by a constant pixel-independent shot noise, which is not correlated with the primary CMB, CMB instrumental noise, or kSZ. In reality, however, a part of the kSZ, as well as foreground residuals which we do not model, come from the halos that the galaxy survey resolves, which leads to a correlation between CMB and galaxy noise residuals. To take into account a noise correlation we define $n_t = \theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_e, v_r)$ and $n_g = \delta_g^{\text{obs}} - \delta_g$. Then we can define a vector $n^T = (n_t, n_g)$ and a covariance matrix C_N ,

$$C_N = \begin{pmatrix} N_T & \text{Cov}(n_t, n_g) \\ \text{Cov}(n_g, n_t) & N_g \end{pmatrix}. \quad (40)$$

Then the likelihood equation takes the usual correlated Gaussian form,

$$-2 \ln \mathcal{L}(\theta^{\text{obs}}, \delta_g^{\text{obs}} | \theta^{\text{pCMB}}, \delta_g, \delta_e, v_r) = n^T (C_N)^{-1} n + \text{const}. \quad (41)$$

The correlated noise $\text{Cov}(n_t, n_g)$ could be estimated from simulations or perhaps be calculated in the halo model. We also note that the noise PDF on small scales is not guaranteed to be Gaussian and could be learned with a machine learning approach. In the following, we do not take into account noise correlation and instead work with Eq. (39), which, as we discussed, holds exactly for our Poisson noise setup (Sec. IV C) and also performs well for the halo analysis (Sec. IV D).

1. Using an estimator for δ_e given δ_g

The most straightforward way to estimate δ_e from δ_g is the following template estimator (e.g., [14,29]):

$$\hat{\delta}_e(k) = \frac{P_{eg}}{P_{gg}} \delta_g(k), \quad (42)$$

where the small-scale power spectra P_{eg} and P_{gg} can be calculated in the halo model (see Ref. [14]) or estimated from simulations (which we do here).

We can then write the likelihood as

$$\begin{aligned}
& -2 \ln \mathcal{L}(\Theta^{\text{obs}}, \delta_g^{\text{obs}} | \theta^{\text{pCMB}}, \delta_g, v_r) \\
& = (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\hat{\delta}_e(\delta_g), v_r))^T N_T^{-1} \\
& \quad \times (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\hat{\delta}_e(\delta_g), v_r)) \\
& \quad + (\delta_g^{\text{obs}} - \delta_g)^T N_g^{-1} (\delta_g^{\text{obs}} - \delta_g) + \text{const}. \quad (43)
\end{aligned}$$

The kSZ calculated from $\hat{\delta}_e$ will only include the resolved kSZ, i.e., the one that can be explained by the galaxy distribution. The estimator in Eq. (42) is not guaranteed to be the optimal estimator since these small-scale fields are non-Gaussian. Below in Sec. V we will explore a machine learning method, using a convolutional neural network to learn a better template $\hat{\delta}_e(\delta_g)$. If we assume Gaussianity for the fields we also get the prior

$$\begin{aligned}
-2 \ln \mathcal{P}(\theta^{\text{pCMB}}, \delta_g, v_r) & = (\theta^{\text{pCMB}})^T P_T^{-1} \theta^{\text{pCMB}} + v^T P_v^{-1} v \\
& \quad + \delta_g^T P_g^{-1} \delta_g + \text{const}, \quad (44)
\end{aligned}$$

where P_x are the signal covariance matrices for fiducial cosmological parameters.

2. Coupling δ_e and δ_g with a correlated prior

A second, more Bayesian, option is to impose a stochastic relation between δ_e and δ_g . In this case, the likelihood is

$$\begin{aligned}
& -2 \ln \mathcal{L}(\theta^{\text{obs}}, \delta_g^{\text{obs}} | \theta^{\text{pCMB}}, \delta_g, \delta_e, v_r) \\
& = (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_e, v_r))^T N_T^{-1} \\
& \quad \times (\theta^{\text{obs}} - \theta^{\text{pCMB}} - \theta^{\text{kSZ}}(\delta_e, v_r)) \\
& \quad + (\delta_g^{\text{obs}} - \delta_g)^T N_g^{-1} (\delta_g^{\text{obs}} - \delta_g) + \text{const}. \quad (45)
\end{aligned}$$

To impose a relation between δ_e and δ_g , we need a correlated prior. If we assume the Gaussianity of the small-scale fields, the prior is

$$-2 \ln \mathcal{P}(\delta_e, \delta_g) = (\delta_e \ \delta_g) \begin{pmatrix} P_{ee} & P_{ge} \\ P_{ge} & P_{gg} \end{pmatrix}^{-1} \begin{pmatrix} \delta_e \\ \delta_g \end{pmatrix} + \text{const}, \quad (46)$$

where the covariance matrix must be invertible, i.e., the two fields cannot be fully correlated (as is the case if one naively assumes $\delta_g \propto \delta_e$). The posterior is then

$$\begin{aligned}
& \mathcal{P}(\theta^{\text{pCMB}}, \delta_g, v_r, \delta_e | T^{\text{obs}}, \delta_g^{\text{obs}}) \\
& \propto \mathcal{L}(T^{\text{obs}}, \delta_g^{\text{obs}} | T^{\text{pCMB}}, \delta_g, \delta_e, v_r) \times \mathcal{P}(T^{\text{pCMB}}, \delta_e, \delta_g, v_r). \quad (47)
\end{aligned}$$

In this approach, again the assumption of Gaussianity could be circumvented by machine learning. A non-Gaussian prior $\mathcal{P}(\delta_e, \delta_g)$ can, in principle, be learned from simulations using the technique of normalizing flows [39,40]. We are planning to explore this idea in the future.

E. Outlook: Jointly estimating cosmological and astrophysical parameters

So far we have discussed how to estimate fields, assuming that the priors are known, i.e., they have been evaluated for known astrophysical and cosmological parameters. We now generalize the posterior to include such parameters. We are splitting this discussion into two types of parameters, according to our separation of scales. We defer an application of such a fit to simulations to future work.

Astrophysical parameters affect the small-scale δ_e and δ_g fields, in the deeply nonlinear regime where most of the observed kSZ signal is generated. For example, one can calculate P_{ge} in the halo model as a function of the slope of the radial electron profile (see Ref. [14]). P_{ge} appears in the likelihood above in Eqs. (42) and (46). For example, if we assume Gaussian fields, the prior in Eq. (46), written in terms of the vector $\delta_{ge} = \begin{pmatrix} \delta_e \\ \delta_g \end{pmatrix}$ becomes

$$-2 \ln \mathcal{P}(\delta_e, \delta_g | \Lambda_S) = \delta_{ge}^T C_{ge}^{-1}(\Lambda_S) \delta_{ge} + \log \det C_{ge}(\Lambda_S). \quad (48)$$

Note that now we need to keep track of the determinant factor since it depends on the Λ_S parameters for which we optimize. We can use a prior on the astrophysical parameter Λ_S to express $\mathcal{P}(\delta_e, \delta_g, \Lambda_S) = \mathcal{P}(\delta_e, \delta_g | \Lambda_S) \mathcal{P}(\Lambda_S)$.

Cosmological parameters affect v_r and δ_g at large scales. To avoid complications due to nonlinear evolution (which could be included with a differentiable forward model of structure formation), we can probe cosmological parameters on linear scales $k \ll k_{NL}$. In particular, we can constrain local non-Gaussianity f_{NL} [11], galaxy bias b_g , kSZ optical depth bias b_v [29], and the Hubble constant H . To include constraining power from the galaxy field on large scales, we need to add a field likelihood for the large-scale galaxy field $\delta_{g,L}$ given by

$$-2 \ln \mathcal{L}(\delta_{g,L}^{\text{obs}} | \delta_{g,L}) = (\delta_{g,L}^{\text{obs}} - \delta_{g,L})^T N_g^{-1} (\delta_{g,L}^{\text{obs}} - \delta_{g,L}). \quad (49)$$

The joint prior $\mathcal{P}(\delta_{g,L}, v_r)$ of the fields $\delta_{g,L}$ and v_r on large (linear) scales as a function of cosmological parameters Λ_L is Gaussian. Written in terms of the vector $\delta_{gv} = \begin{pmatrix} \delta_{g,L} \\ v_r \end{pmatrix}$ it is given by

$$-2 \ln \mathcal{P}(\delta_{g,L}, v_r | \Lambda_L) = \delta_{gv}^T C_{gv}^{-1}(\Lambda_L) \delta_{gv} + \log \det C_{gv}(\Lambda_L). \quad (50)$$

Analytic expressions for the covariance matrix C_{gv} on linear scales are given for example in [11].

Including all these parameters, in full generality, the posterior is

$$\begin{aligned} & \mathcal{P}(\theta^{\text{pCMB}}, \delta_g, \delta_e, v_r, \Lambda_S, \Lambda_L | \theta^{\text{obs}}, \delta_g^{\text{obs}}) \\ & \propto \mathcal{L}(\theta^{\text{obs}}, \delta_g^{\text{obs}} | \theta^{\text{pCMB}}, \delta_g, \delta_e, v_r) \\ & \times \mathcal{P}(\theta^{\text{pCMB}}, \delta_e, \delta_g, v_r, \Lambda_S, \Lambda_L). \end{aligned} \quad (51)$$

Physical constraints are then various joint or marginal distributions of this complicated posterior. Finding the MAP and marginalizing over fields will be challenging and is deferred to future work.

F. Implementation and optimization of the likelihood

We implement the likelihood equations in JAX—a PYTHON package with a versatile and highly optimized autodifferentiation. In this way, we can take analytic derivatives with respect to all fields and scalar parameters. We then find the MAP, i.e., the parameter and field configuration which maximizes the posterior probability density. We now give a detailed description of our PYTHON implementation, which is based on the JAX library.

Data dimensionality. Given n redshift bins and for an angular side length N_{side} ,³ we have in total $N_{\text{side}}^2 \times (2n + 1)$ parameters to solve for (n velocities, n electron densities, and the primary CMB).

Masking the data. Since we have nonperiodic simulation data, we need to apodize the data with a smooth mask at the boundary. An exact treatment would deal with the mask by Wiener filtering with an inhomogeneous noise covariance matrix. However, Wiener filtering at high signal-to-noise converges very slowly, so instead we use the common but somewhat suboptimal apodization technique (as implemented, e.g., in [44]).

Data representation. The CMB likelihood term [e.g., Eq. (36)] is naturally represented in pixel space (x space) because the kSZ is obtained with a simple pixelwise multiplication and the noise covariance is diagonal in pixel space (even for inhomogeneous uncorrelated noise). The prior terms, however, are naturally represented in Fourier space (l space), since covariance matrices are diagonal there. We thus use both representations and FFT between the two.

Optimizer. We minimize the negative log posterior numerically for parameters of interest, for example, we minimize Eq. (37) with respect to $\{\delta_m, v_r, \theta^{\text{pCMB}}\}$. We

³ N_{side} here represents the number of pixels per side of a square patch. In our simulation results presented below, it also coincides with the N_{side} parameter for the input HEALPIX maps, since the total number of pixels in a HEALPIX map with given N_{side} is $N_{\text{pix}} = 12 \times N_{\text{side}}^2$ and we analyze exactly 1/12 of the full-sky HEALPIX map, projected to a flat patch.

freely switch from x to l space and vice versa utilizing JAX FFT routines. We found that the popular Adam optimization algorithm, which is widely used for neural network training, works well for our setup. We used the implementation of Adam from the Optax library. Convergence of the optimization depends on the noise levels and required some tuning. We found that convergence can be improved by setting different initial learning rates for each field (for example, $\sim 10^{-3}$ for velocity, $\sim 10^{-2}$ for density, and $\sim 10^{-5}$ for primary CMB anisotropy worked in many cases). We also used the exponential decline of learning rates during the optimization (using a decay factor of 0.83 every 150 optimization steps).

Initialization. We initialize $\{\delta_g = \delta_g^{\text{obs}}, v_r = \hat{v}_r^{\text{OE}}, \theta^{\text{pCMB}} = \theta^{\text{obs}}\}$. Here, \hat{v}_r^{OE} is the quadratic estimator for v_r . Although all the results in this paper were obtained with $v_r = \hat{v}_r^{\text{OE}}$ initialization, we noticed that this is not essential and one can start with $v_r = 0$ as well.

Multi-GPU parallelization. Because of GPU memory limitations, there is a trade-off between resolution (both pixel- and binwise) and physical volume. However, we also found it possible to make use of several GPU devices in parallel to increase the size of the redshift interval of our analysis. We kept different groups of redshift bins on different GPUs and calculated kSZ temperature anisotropy after each optimization step via collective communication.

Hardware and convergence time. We can include up to 23 redshift bins with maps of 2048^2 pixels on a single 16 GB RTX-A4000. We also note that, in addition to parameters, one also needs to keep various constants in memory, such as the observed fields $\delta_g^{\text{obs}}, \theta^{\text{obs}}$ and theoretical covariances of the parameters. The optimization process on one GPU takes of order 10 min. With the help of three GPUs we are able to include exactly 3 times more bins. For three GPUs, the optimization time increased to approximately one hour, which includes inter-GPU and GPU-CPU communication. In our main analysis presented in Sec. IV, we use 60 redshift bins spread over three GPUs.

G. Joint posterior, marginal posterior, and error bars

The full posterior, including cosmological and astrophysical parameters, which we keep fixed in our results below, is given by

$$\mathcal{P}(\theta^{\text{pCMB}}, \delta_g, \delta_e, v_r, \Lambda_S, \Lambda_L | \theta^{\text{obs}}, \delta_g^{\text{obs}}). \quad (52)$$

We will refer to this posterior as the “joint posterior” (following the naming of lensing posteriors in [26]). From this posterior one can, at least in principle, obtain various “marginalized posteriors.” In particular, if one is only interested in reconstructing the velocity field, one may want to marginalize over the small-scale fields, i.e.,

$$\begin{aligned} & \mathcal{P}(v_r, \Lambda_S, \Lambda_L | \theta^{\text{obs}}, \delta_g^{\text{obs}}) \\ &= \int d\theta^{\text{pCMB}} d\delta_g d\delta_e \mathcal{P}(\theta^{\text{pCMB}}, \delta_g, \delta_e, v_r, \Lambda_S, \Lambda_L | \theta^{\text{obs}}, \delta_g^{\text{obs}}). \end{aligned} \quad (53)$$

In the CMB lensing likelihood analysis, Refs. [23,24] maximize the marginal posterior, while [25,26] maximize the joint posterior. In the lensing case, it is possible to analytically marginalize over the CMB fields, but maximizing the marginal posterior comes at the cost of repeatedly evaluating a computationally expensive determinant. In our present work, we maximize the joint posterior of the fields, but physically both the joint and marginal posteriors are interesting. To make unbiased estimates of cosmological or astrophysical parameters, it is, in principle, necessary to marginalize over the field variables. How this can be achieved in the case of CMB lensing was recently shown in [45], and we defer the kSZ case to future work.

In addition to finding the MAP, we need to give error bars for the modes of the reconstructed velocity field v_r (and ultimately also for cosmological and astrophysical parameters Λ_S and Λ_L). This can be done by considering the curvature of the posterior (or, at least in principle, by sampling from the posterior). Assuming that the posterior is not multimodal, it can be expanded around the MAP to second order and error bars can be given in terms of the second derivatives of the posterior. As usual, we define the Hessian of the posterior,

$$F_{ij} = - \left. \frac{\partial^2 \ln \mathcal{P}}{\partial \lambda_i \partial \lambda_j} \right|_{\lambda=\hat{\lambda}}, \quad (54)$$

where λ are all parameters of the posterior and $\hat{\lambda}$ is their MAP. One can obtain their unmarginalized error bars from the diagonal part as $\sigma_i = \frac{1}{\sqrt{F_{ii}}}$ and the marginalized ones as $\sigma_i = (F^{-1})_{ii}$ if this inversion is tractable. Unfortunately, inverting the full Hessian for the high-dimensional posterior with all fields is not computationally tractable. Here we restrict ourselves to a simpler setup. In Appendix B we show that for the simplest likelihood, given in Eq. (23), we can obtain error bars on the reconstructed velocity field from the Hessian of the posterior, which match the residuals with the truth (which are available in simulations but not in real data). For this simpler likelihood, it is even possible to invert the full Hessian. At least in this simple case, we can thus set error bars on the reconstructed velocity field without the need to run large numbers of simulations, by using autodifferentiation to calculate the likelihood curvature at the MAP. We will investigate error bars in the full likelihood formulation, including methods to integrate out the small-scale fields either analytically or numerically, in future work.

IV. APPLICATION TO SIMULATIONS

In this section, we apply our likelihoods to a simulation. We explore under which conditions the MAP outperforms the QE and find that both a low CMB noise and a low large-scale structure noise are required for a significant improvement factor.

A. Agora simulation

There are currently not many available simulations that have high-resolution halos, light-cone coordinates, include CMB, and cover a large sky fraction. The simulation requirements for kSZ velocity reconstruction are particularly challenging because one both needs the large velocity scales and the small scales where the kSZ is visible, and here we also require a high halo density. One suitable simulation is the recently published Agora simulation, which we use here. Agora [46,47] is a multicomponent simulation on a light cone that uses halos and particles from Multidark-Planck2 N-body simulation and also models CMB primary and secondary anisotropies including lensing, kSZ, thermal Sunyaev-Zeldovich effect (tSZ), and cosmic infrared background (CIB). While here we only include the kSZ effect, other secondary anisotropies may be helpful in the future.

With our hardware setup of three 16 GB memory RTX-A4000 GPUs, we found the following resolution to be tractable. We downgrade the original full-sky maps from HEALPIX $N_{\text{side}} = 8192$ to $N_{\text{side}} = 2048$, where the latter corresponds to the resolution of 1.72 arcmin. We then apply flat-sky approximation, keeping 1/12 of the maps and treating them as fields on a 2D Cartesian grid of resolution 2048^2 pixels (see footnote 2 above). For this purpose, we found the `reproject` function from the `PIXELL` library useful. We consider the redshift interval of $z = 0.5\text{--}3$ and use $N = 60$ equidistant comoving radial bins so that each bin has a width of 50 Mpc. This binning merges the original 120 bins of Agora into pairs of two. Our optimization for three-dimensional fields thus uses arrays of size [60, 2048, 2048]. In the future, when we apply our method to real data (CMB-S4), it would be useful to raise the angular resolution further, which would require a larger bank of GPUs. For our present goal of testing our likelihood method, the current resolution is sufficient.

B. Velocity field, galaxy density, and CMB map

The first required simulation product in our analysis is the radial velocity field. Since the Agora simulation is patched together by repeating the same base simulation several times, this could potentially lead to wrong powers on large scales. In Fig. 1 we show the radial velocity field and the matter field of the Agora simulation for the radial bin at $z = 1.15$, compared to the expectation from CAMB. The power spectra here and below were estimated from the simulation after taking a 1/12th of the sky flat-sky projection and apodizing the mask.

We find generally good agreement with theory at all redshifts, sufficient for our purpose.

The second required simulation product is the halo/galaxy catalog. In Fig. 2 (left) we show the differential number density of the halo catalog from the Agora⁴ simulation, as well as the expected galaxy number density of the Rubin Observatory Legacy Survey of Space and Time (LSST) Y10 gold sample [48]. We note that we used the light-cone halo catalog from Agora which is not tailored for a specific survey. Agora does provide an LSST-Y1 galaxy catalog, but we are not using it in this analysis. Given the densities in Fig. 2 (left), we thus expect weaker constraints with Agora halos (due to the mass threshold of the Agora halo catalog of 10^{11} solar masses) than can be obtained with Rubin Observatory LSST-Y10. Unfortunately, as we shall find, to increase the signal-to-noise with the likelihood, a larger number density than Agora is required. For this reason, in addition to analyzing the actual halo catalog in Agora, we also test our analysis on Poisson sampled biased matter density maps for various shot noise levels. More precisely, we emulate higher-density halo maps by multiplying Agora matter densities by halo biases and adding Poisson shot noise n_g^{-1} corresponding to the target density. While Poisson shot noise is certainly not a perfect approximation of small-scale halo formation, this allows us to approximately extend our analysis to very high galaxy densities. Figure 2 (right) shows the power spectrum of both the real Agora halo map as well as the emulated one with Rubin Observatory density. In all of our results below we do not include photo- z redshift errors; however, the redshift resolution is limited by the choice of 60 radial bins.

The third required simulation product is the observed CMB map. While Agora comes with its own simulated CMB map, which includes properly correlated CMB lensing and various CMB secondary anisotropies, for our present purpose it is more convenient to generate our own CMB map. We generate primary CMB from the lensed CMB power spectrum of CAMB. We then add kSZ generated according to the binned kSZ model described in Sec. II B, by multiplying the binned radial velocity maps with the binned electron density maps (assuming $\delta_e = \delta_m$) at full resolution and adding all 60 bins up to redshift 3. This is the kSZ generated by our simulation volume. The real data will include some additional kSZ. First, some kSZ will be generated by radial distances below our binning resolution (see Ref. [16]). Second, kSZ from higher redshifts including reionization kSZ will also contribute. To approximate both of these effects, we add uncorrelated kSZ to the CMB map, as described in Eq. (8). Below we both show results for the simulation kSZ alone, as well as for an added uncorrelated kSZ with the same power spectrum and amplitude as the simulation kSZ. We made this choice

⁴<https://yomori.github.io/agora/index.html>.




FIG. 1. Example of power spectra of radial velocity (left) and matter density (right) from Agora simulations (solid blue) compared to theoretical values calculated with CAMB (dashed black) at $z = 1.15$. The power spectra were estimated from the apodized flat-sky projection of 1/12 of the full-sky AGORA simulation.




FIG. 2. Left: halo density (Agora) and galaxy density (LSST Y-10) per sterad per redshift as a function of redshift. Right: power spectra of halo overdensity from Agora and emulated halo overdensity (i.e., Poisson sampled biased matter overdensity) for LSST-Y10 gold sample galaxy density level. At the kSZ length scales of $\ell > 4000$ the Agora map is noise dominated, while the emulated LSST map is signal dominated.

because high- z kSZ (here $z > 3$ including reionization) may have a similar total amplitude as low- z kSZ [35,49].

In our analysis, we consider three different CMB noise levels, as shown in Fig. 3. CMB-S4 noise corresponds to the CMB-S4 ultra-deep internal linear combination (ILC) curve taken from [50]. Finally, the curve named “1,1” was generated from the power spectrum $N(l) = N_T^2 e^{\frac{l(l+1)\theta_{\text{FWHM}}^2}{8 \ln 2}}$ with $N_T = 1 \mu\text{K arcmin}$ and $\theta_{\text{FWHM}} = 1 \text{ arcmin}$ (without ILC foreground cleaning) to be comparable to the forecast in [11]. The CMB-HD noise curve was generated according to [51] by coadding instrumental noises from the frequency channels of 150 and 90 GHz. The level of foregrounds in 150 and 90 GHz, which is left after foreground cleaning, is shown with a blue line in Fig. 1 of [51]. We note that it is dominated by the reionization kSZ signal, the influence of which is discussed later in this work. The level of CIB foreground that is left is much less than the instrumental noise. Whether CMB-HD can reach such low foreground levels (using other frequency channels for foreground cleaning) is currently being investigated by the CMB-HD Collaboration. In this work, we do not include correlated foreground residuals, which need to be studied in more detail even in the case of the standard quadratic




FIG. 3. Lensed CMB power spectrum along with three different instrumental noise levels and generated kSZ signal (simulation volume only). For CMB-S4, the kSZ (shown only for our simulation volume $0.5 < z < 3$) is below the noise level, while for the futuristic experiments, it is above the noise level. Here the CMB-S4 noise curve includes ILC foreground deprojection, while the 1, 1 noise means white noise with $N_T = 1 \mu\text{K arcmin}$ and $\theta_{\text{FWHM}} = 1 \text{ arcmin}$ (without ILC). For either of these experiments, our analysis uses only scales up to $\ell \simeq 6000$ to limit the GPU requirements of this theoretical study.

estimator method. We show these noise spectra as well as the kSZ from our simulation volume in Fig. 3.

C. Results for observed matter δ_m with Poisson shot noise

We now test our likelihood pipeline on the Agora simulations, assuming that we observe the Poisson noise-corrupted biased matter field (see Fig. 2 right) in addition to the CMB. We thus find the MAP to the posterior, Eq. (37), which includes the likelihood, Eq. (36). In the next section, we will then analyze the Agora halo catalog instead of the matter field.

We first consider the shot noise that corresponds to the n_g prediction for Rubin Observatory (Fig. 2 left) and the $N_{1,1}^{\text{pCMB}} = N_{1,1}^{\text{pCMB}}$ CMB noise as described in Sec. IV B. Note that we do not take into account photo- z errors (except by redshift binning), so the results for real Rubin Observatory galaxies would be somewhat weaker. Figure 4 shows the cross-correlation coefficient and reconstruction noise for QE and MAP estimator for three different redshifts $z = 0.7, 1, 2$. The reconstruction noise is defined as $N(\ell) = \langle |v^{\text{true}} - b\hat{v}|^2 \rangle$ and the cross-correlation coefficient is $r(\ell) = \frac{\langle \hat{v} \hat{v} \rangle}{\sqrt{\langle \hat{v} \hat{v} \rangle \langle v^{\text{true}} v^{\text{true}} \rangle}}$. For this configuration the MAP estimator performs significantly better than QE, especially at lower and intermediate redshifts, giving lower reconstruction noise and higher cross-correlation coefficient.

Let us discuss the curves in Fig. 4 in somewhat more detail. First, the quadratic estimator has lower noise at intermediate redshift $z = 1$ than at both lower and higher redshift bins. This is expected because the volume of the $z = 1$ bin is larger than at $z = 0.7$, and the number density is larger at $z = 1$ than at $z = 2$ (see Fig. 2), resulting in the lowest noise at intermediate redshift. Second, cosmic variance fluctuations (due to having only a single simulation) between the QE and the MAP are similar, as expected. Third, the noise of the MAP (as well as the amplitude of the reconstructed velocity signal) goes down at large ℓ , while the noise of the QE stays flat. This is the expected behavior because the MAP includes inverse variance weighting so that noisy scales are being suppressed by $S(S + N)^{-1}$ (where S is the signal covariance and N is the noise covariance). If we were to Wiener filter the QE result, we would get the same behavior, i.e., the QE reconstructed power spectrum would fall at large ℓ . To estimate the unbiased power spectrum from the Wiener-filtered field, one would include a power spectrum transfer matrix, as is commonly done in power spectrum estimation. Finally, a somewhat surprising result is that the improvement of MAP over QE is larger at $z = 0.7$ than at $z = 1$, even though the QE performs better at $z = 1$. However, it is difficult to predict the improvement factor analytically because it depends on several redshift-dependent factors, in particular, the number density n per sky angle (which first grows and then falls toward higher z), as well as the amplitude of the cross power spectrum $P_{g\tau}$ (which grows

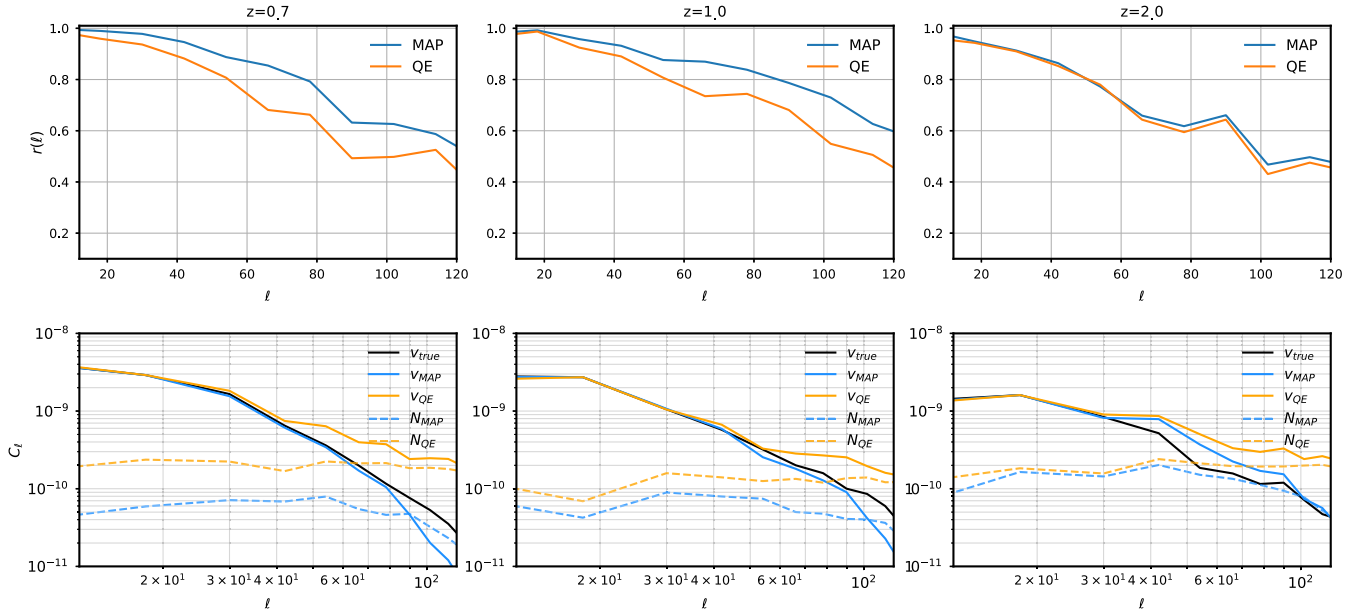


FIG. 4. Velocity reconstruction based on δ_m and θ^{CMB} . Cross-correlation coefficient r_ℓ (upper row) and noise power spectra (lower row) of QE and MLE as a function of multipole number ℓ for three different redshifts. CMB noise and galaxy density are, respectively, $N_{1,1}^{\text{pCMB}}$ ($N_T = 1 \mu\text{K arcmin}$ and $\theta_{\text{FWHM}} = 1 \text{ arcmin}$, without foregrounds) and n_{LSST}^g (LSST gold sample). We include data in 60 redshift bins ($0.5 < z < 3$) with angular resolution $\ell_{\text{max}} = 6000$ in the reconstruction, which is the resolution limit with our GPU memory.

toward lower z) and perhaps the non-Gaussianity of the small-scale fields (which grows toward lower z). As a final note, we cannot exclude the fact that our likelihood is not fully converged on small scales, despite our best effort to achieve convergence (see Sec. III F). In particular, in the high redshift bin $z = 2$, at high ℓ , the MAP power should be below the true velocity field power, due to the inverse variance weighting. However, on large scales, which are those most interesting for cosmology, convergence appears to be robust.

In all plots of the reconstruction power spectrum, we have adjusted an overall (scale-independent) amplitude parameter to match the theory power. This is equivalent to fixing the “velocity bias” described in [14], which also has to be marginalized over in the case of the QE. The velocity bias is physically associated with the unknown normalization of the electron density in the real Universe. In simulations, it could be fixed by carefully tracking normalizations and the masking bias, but we have not done this here. We also estimated the small-scale power spectrum that appears in the prior term from simulations. In the future, we aim to fit a parametrization of the small-scale power spectrum jointly with the field parameters to take into account baryonic uncertainty.

We now explore the influence of different experimental configurations. As a futuristic case, we consider an experiment with 10 times the number density of Rubin Observatory. While no experiment with such a low shot noise is planned in the near term, this configuration allows us to probe how the signal-to-noise scales with galaxy number density. As described in Sec. IV B we also consider three different CMB noise configurations, CMB-S4 (ILC), “1,1 noise,” and CMB-HD. As before, we include data in 60 redshift bins up to $\ell_{\max} = 6000$ in the reconstruction, which is the resolution limit with our GPU memory (three GPUs with each 16 GB). The absolute SNR is thus lower than could be achieved with a larger ℓ_{\max} (especially for a CMB noise better than CMB-S4), but our main concern here is to demonstrate the improvement over the QE, for which our GPU setup suffices. For these configurations, we provide signal-to-noise improvement factors in Table I, ranging in values from 1 to 4. Unfortunately, with CMB-S4 ILC noise curves we do not find an improvement. Note, however, that this result could be different with halos, a more complicated halo-based kSZ model, and halo mass weighting in the galaxy field, which we aim to explore in the future with a more high-resolution simulation. With the present simulation, the question of whether an improvement with CMB-S4 is possible cannot be definitively answered. Comparing the 1,1 noise case and the CMB-HD noise, we only find a substantial improvement in the $z = 0.7$ bin, indicating that the galaxy shot noise is the limiting factor at higher redshifts for this extremely low noise configuration. In some bins, we find a slightly better result with the 1,1 noise than with the HD noise, which may

TABLE I. Improvement of the likelihood over the QE for futuristic noise levels (reconstruction based on δ_m with Poisson shot noise). $N_{\text{QE}}^v/N_{\text{MAP}}^v$, averaged over $\ell = 10$ –60 at three different redshifts for various shot noise and CMB noises configurations. ukSZ means that we include an uncorrelated kSZ component of the same magnitude ($\mathcal{D}_{\ell=2000}^{\text{ukSZ}} = \mathcal{D}_{\ell=2000}^{\text{kSZ}} \approx 1 \text{ } [\mu\text{K}]^2$) as the simulation kSZ to approximate kSZ contributions from outside the redshift range of the simulation. We include data in 60 redshift bins ($0.5 < z < 3$) with angular resolution $\ell_{\max} = 6000$ in the reconstruction. For CMB-S4 ILC noise we do not find an improvement.

	$z = 0.7$	$z = 1$	$z = 2$
$N_{1,1}^{\text{pCMB}}$ and n_g^{LSST}	3.2	1.7	1.3
$N_{1,1}^{\text{pCMB}}$ and n_g^{LSST} and ukSZ	2.2	1.6	1.1
$N_{1,1}^{\text{pCMB}}$ and $10 \times n_g^{\text{LSST}}$	3.4	2.6	1.6
$N_{1,1}^{\text{pCMB}}$ and $10 \times n_g^{\text{LSST}}$ and ukSZ	2.8	2.2	1.5
$N_{\text{HD}}^{\text{pCMB}}$ and $10 \times n_g^{\text{LSST}}$	4.0	2.7	1.4
$N_{\text{HD}}^{\text{pCMB}}$ and $10 \times n_g^{\text{LSST}}$ and ukSZ	2.9	2.0	1.4

be due to small fluctuations in the convergence of our optimization procedure.

Next, we explore how additional kSZ power influences our results. In Fig. 4, the CMB map contained the primary CMB, the simulation kSZ, as well as the instrument noise. We now include an additional uncorrelated kSZ component as described in Sec. IV B, by adding it to the map and modifying the covariance of the CMB prior. This situation is more realistic, as our simulation generates kSZ only at the redshift of 0.5–3, but there is an additional kSZ signal coming from $z > 3$ including the reionization epoch. Since the level of the reionization kSZ is not well known, we set it to be equal in power to the modeled kSZ signal ($\mathcal{D}_{\ell=2000}^{\text{ukSZ}} = \mathcal{D}_{\ell=2000}^{\text{kSZ}} \approx 1 \text{ } [\mu\text{K}]^2$). We summarize the average improvement of MAP reconstruction with respect to QE, $N_{\text{QE}}^v/N_{\text{MAP}}^v$ both with and without adding unresolved kSZ signal in Table I. We find that the extra kSZ reduces the improvement factor, but significant improvement remains.

As an interesting consistency check, we now investigate whether our likelihood has succeeded in reconstructing the “kSZ-corrupted” primary CMB. The likelihood approach naturally provides a kSZ-reduced temperature anisotropy field. In Fig. 5 (left) we demonstrate this showing the cross-correlation coefficients $r(\ell)$ for $\theta^{\text{pCMB}} - \hat{\theta}^{\text{MAP}}$ and $\theta^{\text{pCMB}} - \theta^{\text{obs}}$ for two MAPs obtained for two different noise configurations. We see that, indeed, the reconstructed primary CMB cross-correlates better with the true primary CMB than the observed CMB. In Fig. 5 (right) we show the reduced power in the reconstructed primary CMB map. Despite the futuristic experimental configuration, de-kSZing is not very efficient in our setup. This is because we only reconstruct velocities on large scales, while all velocity scales contribute to the total kSZ signal. For




FIG. 5. Left: “de-kSZing” of primary CMB. Cross-correlation coefficient $r(\ell)$ of $\hat{\theta}^{\text{MAP}}$ or θ^{obs} with θ^{pCMB} for two noise configurations: $N_{\text{HD}}^{\text{pCMB}}, 10 \times n_{\text{LSST}}^g$ and $N_{1,1}^{\text{pCMB}}, n_{\text{LSST}}^g$. As expected, the reconstructed primary CMB cross-correlates better with the true primary CMB than the observed CMB does. The de-kSZing is not large here because we do not include an estimate of small-scale velocities from small-scale galaxies (see main text). Right: power spectra of reconstructed and true kSZ signal (red and purple) and pCMB (green and orange) for the $N_{\text{HD}}^{\text{pCMB}}, 10 \times n_{\text{LSST}}^g$ noise configuration (correspondent to blue curves on the left).

example, in [35], galaxies were used to make a template for velocities on small scales. It is possible to modify the likelihood to reconstruct the velocities jointly from the kSZ and the galaxy catalog, which would substantially improve de-kSZing. However, we leave this to future work as it involves nonlinear modeling.

D. Results for observed halos δ_h

In this section, we perform the same MAP analysis, but now using halos (rather than matter plus Poisson noise) from the AGORA simulations. The posterior is now given by Eq. (44) and the likelihood is given by Eq. (43). A difference to the previous section is that we now observe the halo/galaxy field δ_g , but we model the kSZ from $\delta_m = \delta_e$. As discussed in Sec. III D 1 we need to make a template for δ_e given observed δ_g . In this section, we use the simplest possible template, given by Eq. (42), with small-scale power spectra P_{gg} and P_{ge} estimated from the simulations.

Because the Agora halo density is lower compared to the LSST “gold sample” (Fig. 2) by an order of magnitude, we expect less signal-to-noise and MAP improvement than in the previous section. The cross-correlation $r(\ell)$ and noise power spectra for reconstruction with halos and $N^{\text{pCMB}} = N_{1,1}^{\text{pCMB}}$ are shown on Fig. 6 (left), finding about a factor of 2 improvement. Interestingly, we find improvements even though the galaxy density is not as large as for LSST. The same plots but for higher CMB noise, $N^{\text{pCMB}} = N_{\text{S4}}^{\text{pCMB}}$, are depicted in Fig. 6 (right). We see that, at CMB-S4 noise, the MAP estimator performs equivalent to the QE in terms of signal-to-noise.

V. A MACHINE-LEARNING-BASED ESTIMATOR FOR THE ELECTRON DENSITY

To estimate the velocity field, with both the QE or the likelihood, one needs an estimate or statistical connection

between δ_e and δ_g , as we discussed in Sec. III D. In the previous section, we used the linear template estimate, Eq. (42), which is also the assumption usually made for QE reconstruction. However, in principle, this estimator can be an arbitrary nonlinear function of δ_g . In this section, we explore a neural-net-based estimator which we train on the Agora simulations. In a more complete exploration of this idea, one would train on hydro simulations to reconstruct a true baryon distribution, while on Agora we can only reconstruct $\delta_m \approx \delta_e$. We defer an exploration of this idea on CAMELS [52] or IllustrisTNG [53] to future work.

Our goal is to train a neural network to increase the cross-correlation between the electron template derived from the galaxy field and the true electron density field. Let us first show how this cross-correlation affects the reconstruction noise of the QE, where we have an analytic expression for the noise. The QE velocity reconstruction noise is sensitive to the cross-correlation between the true electron field and its tracer as follows:

$$(N_\ell^v)^{-1} \propto \int d\ell \frac{(P_\alpha^{g\tau}(\ell))^2}{\tilde{P}^{\theta\theta}(\ell)\tilde{P}^{gg}(\ell)} = \int d\ell \frac{f_\tau^2 r_{eg}^2(\ell) P^{ee}}{\tilde{P}^{\theta\theta}(\ell)},$$

$$r_{eg}^2(\ell) = \frac{(P^{eg})^2}{\tilde{P}^{gg} P^{ee}}, \quad (55)$$

where we had δ_g as a tracer of δ_e . Hence, we can reduce noise by having a tracer with better r_{eg}^2 . In the region of high ℓ , primary CMB is suppressed. If CMB noise is small, the cross-correlation coefficient of the true electron density field with its tracer defines noise completely. Indeed, assuming that at high ℓ , $\tilde{P}^{\theta\theta}(\ell) \sim \text{const}$, we know that $P^{ee}(\ell) \sim \frac{\ell^2}{\chi_*^2} P^{ee}(|\mathbf{k}| = \frac{\ell}{\chi_*}) \sim \text{const} \times \ell^{-1}$, so that $(N_\ell^v)^{-1} \propto \int r^2(\ell) d\ell$. Hence, in the case of perfect reconstruction, the noise is only limited by resolution.

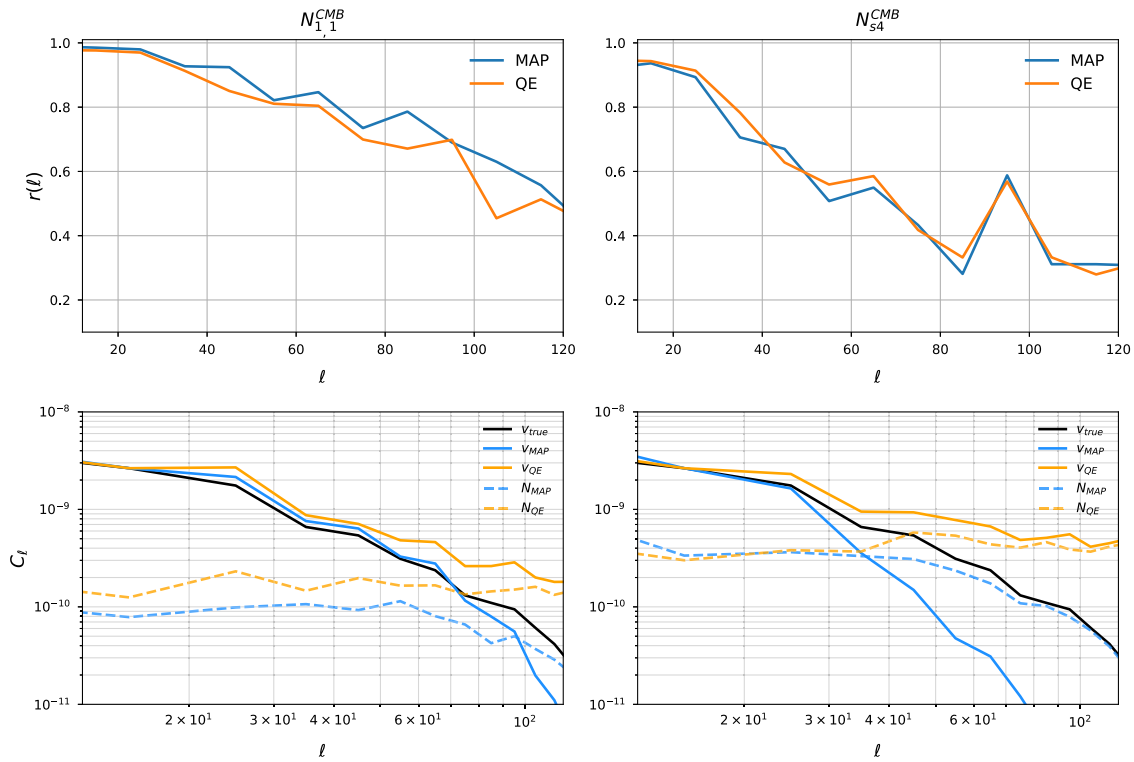


FIG. 6. Velocity reconstruction based on δ_h (Agora simulation halos) and θ^{CMB} . Cross-correlation coefficient r_ℓ (upper row) and noise power spectra (lower row) of QE and MLE as a function of multipole number ℓ for two different CMB noise levels: $N_{1,1}^{\text{pCMB}}$ (left) and N_{S4}^{pCMB} (right) at $z = 0.8$. In the reconstruction, we include data in 60 redshift bins ($0.5 < z < 3$) with angular resolution $\ell_{\text{max}} = 6000$ in the reconstruction. For $N_{1,1}^{\text{pCMB}}$, we see an improvement by a factor of 2, while for N_{S4}^{pCMB} (including ILC) we see no improvement.

We train an eight-block residual network (ResNet) [54] with 3×3 convolutional layers on pairs of (δ_e, δ_g) to minimize the standard rms loss for the three cases of matter tracer used in the preceding sections. In our simulation $\delta_e = \delta_m$ and the matter tracer δ_g here is either Poisson noise corrupted δ_m with noise $N = 1/n_{g,\text{LSST}}$ and $N = 0.1/n_{g,\text{LSST}}$ or the Agora halo field. The network converges fast (the whole training takes a couple of minutes on RTX-A4000). We implement it with JAX and Haiku—a neural network library, compatible with JAX. This allows us to effortlessly include this piece in likelihood optimization for velocity reconstruction (although the neural network template is equally applicable to the QE).

As can be seen in Fig. 7, the neural network (NN) estimator for electron density gives an improvement in cross-correlation, especially for higher multipoles, where the kSZ signal is observable. We do not see any improvement on the lower-density Agora halo catalogs. Instead of a ResNet, it would also be interesting to train a more physical low-parameter model such as Lagrangian deep learning [55].

Next, we investigate how the improved electron template affects velocity reconstruction. Here we consider a more realistic case of high CMB noise—Gaussian white noise with $N_T = 5 \mu\text{K arcmin}$ and $\theta_{\text{FWHM}} = 1.5 \text{ arcmin}$, which

is closer to the target noise of Simons Observatory. We also take halo shot noise at the LSST level, so this forecast is oriented toward an LSST \times SO analysis. We use the QE velocity estimator, but at these noise levels, we expect the likelihood to perform likewise, as was seen in the previous section. Figure 8 shows approximately a factor of 1.25 improvement in reconstruction noise in the low ℓ region resulting from our NN-based estimator of electron density applied to the QE.

As we saw, in our setup the possible improvements due to the neural network are not very large. However, this conclusion could change with more realistic hydrodynamic simulations where the baryon density is being simulated, as well as by considering the properties (such as mass) of halos and galaxies. We plan to explore this question in more detail in future work on hydrodynamic simulations. One may also ask whether the neural network technique is sensitive to baryonic uncertainty. Here we point out that an analytic template, as used in the QE, is also not correct in nature. As shown in [14], on large scales, such baryonic uncertainty can be included by marginalizing over a velocity bias on large scales. This is still true in the neural network case. All we require is that the neural network increases the cross-correlation coefficient in reality over the analytic template. Further, we did not take into account




FIG. 7. Cross-correlation $1 - r_{ee}^2(\ell)$ of the true electron field with its estimate from the analytic template, Eq. (42), compared to the neural network estimate for two halo density levels: $n_{g,\text{LSST}}$ (left) and $10 \times n_{g,\text{LSST}}$ (right). We find a significant improvement in the cross-correlation due to the neural network, which turns out to be larger with higher shot noise (lower galaxy density). Note, however, that in our simulations $\delta_e = \delta_m$ and results may change with realistic hydrodynamic simulations and galaxy properties.

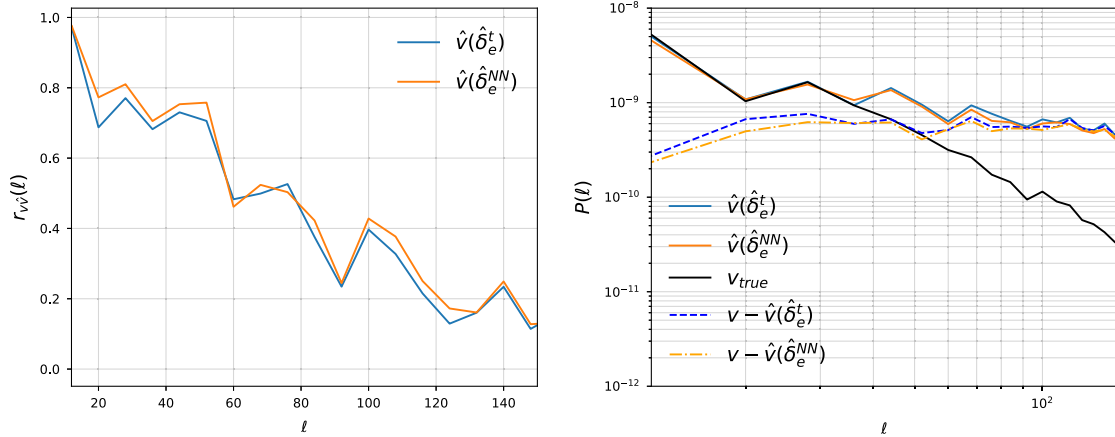


FIG. 8. Left: cross-correlation coefficient $r_{vv}(\ell)$ of true and QE-reconstructed velocity field based on the standard and the neural network estimate of the electron density. Right: corresponding power spectra and residuals. The experimental parameters are similar to an LSST \times SO analysis. We find a small 25% improvement when the QE uses the neural network template, rather than the analytic template [Eq. (42)].

CMB foreground residuals, and a detailed study of our method should take them into account.

VI. CONCLUSION

In this paper, we developed an autodifferentiable likelihood pipeline for the cross-correlation of the CMB and a large-scale structure survey due to the kSZ effect. As is the case for CMB lensing, which has been developed in more detail in the literature because of its larger signal-to-noise for current experiments, a likelihood pipeline is more statistically sensitive than the quadratic estimator at low experimental noises. Beyond the increased statistical sensitivity, a likelihood approach provides a general framework to fit a complete model to the data, which can include

both physical and experimental parameters. Likelihood-based kSZ velocity reconstruction is computationally tractable because the relevant scales of the large-scale velocity field and the small-scale matter field can be separated to a good approximation, so nonlinear differentiable forward modeling of structure formation is not necessary. We show that our approach is computationally tractable for a realistic survey size and provide expected improvement factors over the QE for different idealized experimental configurations. We further developed a machine learning approach to the estimation of the electron density given observed galaxies, which may be able to improve the signal-to-noise for Simons Observatory, although a more detailed study using hydrodynamic simulations is required.

There are two main directions of future work which we will explore. The first direction is to make our forecasts more realistic. Based on the results in the present work, we do not find significant gains in signal-to-noise at the resolution of CMB-S4 and Rubin Observatory (LSST). However, the available Agora simulations do not fully satisfy the required simulation constraints for this forecast, and to probe LSST density we made the simplified assumption that we observe the matter field of Agora with Poisson shot noise. In a simulation with realistic galaxies, with sufficient density, one could take into account galaxy properties, in particular, tracers of the halo mass. Halo masses significantly reduce the shot noise in the galaxy field [56] and could improve our forecast. Fortunately, more realistic high-resolution simulations are starting to become available. A more elaborate forecast would also take into account photo- z errors beyond the simple redshift binning we have performed here. A future study should also take into account foregrounds that correlate with large-scale structure and can affect both the QE results and the likelihood. In particular, the tSZ and CIB correlate with the galaxy field and their remnants (after foreground cleaning) could bias the reconstruction. Our improvement factors can be sensitive to these effects, however, this is a complicated topic that is beyond the scope of the present work. The second main goal for future work is the inclusion of cosmological and astrophysical parameters in the likelihood as outlined in Sec. III E. Jointly fitting such parameters together with fields is, in principle, possible, but achieving convergence in the fit can be difficult. In addition, to obtain unbiased parameter measurements and error bars, one would need to integrate out the field level variables, perhaps with a Gaussian approximation around their MAP. Recent related work in the case of CMB lensing includes [27,31,45]. We will explore these questions in future work.

To obtain results quoted in this work, we extensively used JAX [57] for construction of the autodifferentiable posterior, as well as tools from DeepMind JAX Ecosystem [58], such as Optax for the optimization via gradient descent and Haiku [59] for the construction of a NN-based estimator of electron density. We also benefited from PIXELL and PYTHON implementation of HEALPIX⁵ [60] HEALPY for data preprocessing, along with NUMPY, SCIPY, and MATPLOTLIB [61–63] for various calculations and plotting, and CAMB [64] for theory computations.

ACKNOWLEDGMENTS

We thank Matthew Johnson for comments on the manuscript and the members of the Simons Observatory SZ working group for comments on this work. M. M. acknowledges support from DOE Grant No. DE-SC0022342.

⁵<http://healpix.sourceforge.net>.

Support for this research was provided by the University of Wisconsin–Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation.

APPENDIX A: RELATIONS BETWEEN ANGULAR POWER SPECTRUM C_ℓ , FLAT-SKY APPROXIMATION $P(l)$, AND POWER SPECTRUM OF THE 3D FIELD $P(k)$

We review the known correspondence between multipole and flat-sky coordinates, apply them to our case of binned flat-sky coordinates, and calculate the radial velocity power spectrum.

1. Equivalence of C_ℓ and $P(l)$

Given some field on a 2D sphere $A(\hat{n})$, we define its spherical harmonics transform coefficients as

$$a_{lm} = \int d\Omega Y_{lm}^*(\theta, \phi) A(\hat{n}), \quad (\text{A1})$$

and if the field is rotationally symmetric, its angular power spectrum is defined, as usual

$$\langle a_{l_1 m_1} a_{l_2 m_2}^* \rangle = C_{l_1} \delta_{l_1 l_2} \delta_{m_1 m_2}. \quad (\text{A2})$$

If we consider a small enough solid angle, we can approximate that part of a sphere as flat. Without loss of generality, we can choose \hat{z} as radial direction. As in [65], we may write then

$$\begin{aligned} \hat{n}_\Omega &= \{\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta\} \xrightarrow{\theta \rightarrow 0} \hat{z} + \alpha \\ &= \{\theta \cos \phi, \theta \sin \phi, 1\}, \end{aligned} \quad (\text{A3})$$

so that

$$a_{lm} \approx \int d^2\alpha Y_{lm}^*(\theta, \phi) A(\hat{n}). \quad (\text{A4})$$

Now we introduce continuous flat-sky fields

$$a(l) = \int e^{-il\alpha} A(\alpha) d^2\alpha, \quad (\text{A5})$$

with power spectrum $P(l)$ defined as

$$\langle a(l_1) a(l_2)^* \rangle = (2\pi)^2 \delta^{(2)}(l_1 - l_2) P(l_1). \quad (\text{A6})$$

We are going to show that, defined this way, $P(l)$ is an exact continuous analog of C_l . Approximation from [65] tells us that

$$Y_{lm} \xrightarrow{\theta \rightarrow 0} \sqrt{\frac{l}{2\pi}} J(l\theta) e^{im\phi_\alpha}. \quad (\text{A7})$$

Using this formula and Jacobi-Anger expansion, $e^{ial} = e^{il\theta \cos(\phi_\alpha - \phi_l)} = \sum_n (i)^n J_n(l\theta) e^{in(\phi_\alpha - \phi_l)}$, we can write

$$a(l) = \sum_n (-i)^n \int A(\alpha) d^2\alpha J_n(l\theta) e^{-in(\phi_\alpha - \phi_l)}$$

$$a_{lm} = \int A(\alpha) d^2\alpha J_n(l\theta) \sqrt{\frac{l}{2\pi}} e^{-im\phi_\alpha}. \quad (\text{A8})$$

Comparing two formulas and employing an orthogonality relation $\int \frac{d\phi}{2\pi} e^{i\phi(m-n)} = \delta_{mn}$, we can finally get

$$a(l) = \sqrt{\frac{2\pi}{l}} \sum_m (-i)^m a_{lm} e^{im\phi_l},$$

$$a_{lm} = \sqrt{\frac{l}{2\pi}} i^m \int \frac{d\phi_l}{2\pi} a(l) e^{-im\phi_l}. \quad (\text{A9})$$

Then,

$$\langle a_{l_1 m_1} a_{l_2 m_2}^* \rangle = \sqrt{\frac{l_1}{2\pi}} \sqrt{\frac{l_2}{2\pi}} \int \frac{d\phi_{l_1}}{2\pi} \int \frac{d\phi_{l_2}}{2\pi}$$

$$\times \langle a(l_1) a^*(l_2) \rangle e^{im_1 \phi_{l_1}} e^{-im_2 \phi_{l_2}} \quad (\text{A10})$$

$$= \sqrt{\frac{l_1}{2\pi}} \sqrt{\frac{l_2}{2\pi}} \int d\phi_{l_1} P(l_1) e^{i\phi_{l_1}(m_1 - m_2)} \frac{\delta(l_1 - l_2)}{l_1} \quad (\text{A11})$$

$$= P(l_1) \delta_{l_1, l_2} \delta_{m_1 m_2}. \quad (\text{A12})$$

The other way round, the result may be obtained if we formally consider high l , so that $\sum_{m=-l}^l e^{im\phi_l} \approx \sum_{m=-\infty}^{\infty} e^{im\phi_l} = 2\pi \delta(\phi_l)$,

$$\langle a(l_1) a^*(l_2) \rangle = \sum_{m_1} \sum_{m_2} (-i)^{m_1} i^{m_2} \sqrt{\frac{2\pi}{l_1}} \sqrt{\frac{2\pi}{l_2}} \langle a_{l_1 m_1} a_{l_2 m_2}^* \rangle$$

$$(\text{A13})$$

$$= \sum_{m_1} \frac{2\pi}{l_1} \delta_{l_1, l_2} e^{im_1(\phi_{l_1} - \phi_{l_2})} C_{l_1} \quad (\text{A14})$$

$$\approx (2\pi)^2 C_{l_1} \frac{\delta(l_1 - l_2)}{l_1} \delta(\phi_{l_1} - \phi_{l_2})$$

$$= (2\pi)^2 C_{l_1} \delta(l_1 - l_2). \quad (\text{A15})$$

In these derivations, we made associations of discrete and continuous δ functions, mainly $\delta_{l_1, l_2} = \delta(l_1 - l_2)$, and considered the expression of δ function in polar coordinates: $\delta(l) = \frac{\delta(l)}{l} \delta(\phi_l)$.

2. Relation between C_ℓ and $P(k)$

We define redshift binned fields via integrals with a window function,

$$f_{lm}^i = \int dz W_{z^i}(z) \int d\Omega_x Y_{lm}^*(\hat{\mathbf{x}}) f(\mathbf{x}, z). \quad (\text{A16})$$

In our case, the window function is a top-hat filter: $W_{z^i}(z) = \frac{1}{\Delta z} (\theta(z + (z^i + \frac{\Delta z}{2})) - \theta(z - (z^i - \frac{\Delta z}{2})))$. Here $\theta(z)$ is a Heaviside step function. Then for the densitylike field δ^m , it is easy to get

$$C_l^{\delta\delta}(z^i, z^j) = \frac{2}{\pi} \int dz dz' W_{z^i}(z) W_{z^j}(z')$$

$$\times \int dk k^2 P_{mm}(k) j_l(k\chi(z)) j_l(k\chi(z')), \quad (\text{A17})$$

where we have used the Rayleigh expansion: $e^{ik\mathbf{x}} = 4\pi \sum_{lm} i^l j_l(k\chi) Y_{lm}(\hat{\mathbf{k}}) Y_{lm}^*(\hat{\mathbf{x}})$. For velocity $\mathbf{v}(\mathbf{x})$, we have a continuity equation that connects it to density mode on linear scales $\mathbf{v}(\mathbf{k}) = \frac{if aH}{k} \delta^m(\mathbf{k}) \hat{\mathbf{k}}$. To compute radial velocity angular power spectrum, we define

$$v_{lm}^r(z^*) = -i \int dz W_{z^*}(z) \int d\Omega_x Y_{lm}^*(\hat{\mathbf{x}}) (\mathbf{v}(\mathbf{x}) \cdot \hat{\mathbf{x}}) \quad (\text{A18})$$

$$= \int dz W_{z^*}(z) \int d\Omega_x Y_{lm}^*(\hat{\mathbf{x}})$$

$$\times \int \frac{d^3 k}{(2\pi)^3} \frac{f a H}{k} \delta^m(\mathbf{k}) (\hat{\mathbf{k}} \cdot \hat{\mathbf{x}}) e^{ik\mathbf{x}}. \quad (\text{A19})$$

We can notice that $(\hat{\mathbf{k}} \cdot \hat{\mathbf{x}}) e^{ik\mathbf{x}} = \frac{\partial}{\partial(k\chi)} e^{ik\mathbf{x}} = \frac{\partial}{\partial(k\chi)} e^{ik\chi \hat{\mathbf{k}} \cdot \hat{\mathbf{x}}}$, so that the radial velocity angular correlation function is

$$C_l^{vv}(z^i, z^j) = \frac{2}{\pi} \int dz dz' W_{z^i}(z) W_{z^j}(z')$$

$$\times \int dk f(z) a(z) H(z) (f(z') a(z') H(z'))$$

$$\times P_{mm}(k) j_l'(k\chi(z)) j_l'(k\chi(z')). \quad (\text{A20})$$

Here $j_l(x)$ and $j_l'(x)$ are the spherical Bessel function of the first kind and its derivative correspondingly.

APPENDIX B: ERROR BARS FROM THE HESSIAN WITH THE VELOCITY-ONLY kSZ LIKELIHOOD

As we discussed in Sec. III G, autodifferentiation makes it easy to evaluate the Hessian which can be used to obtain error bars under the approximation that the posterior is Gaussian around the MAP. For our full posterior, including the small-scale fields, we cannot evaluate the full Hessian. We will explore methods to integrate out the small-scale

fields in future work. Here instead we consider the simpler likelihood of Sec. III A [Eq. (23)] where the only degrees of freedom are the velocity modes $v(\mathbf{k})$. We further simplify the problem to a single redshift bin and make all fields Gaussian and periodic (no mask). We consider the case where all kSZ signals comes from one bin with known electron density so that the optical depth field τ is fully known. While in this case, in principle, there is an analytic MAP estimator, as discussed in Sec. III B, we instead use optimization to find the MAP as in our main analysis.

To obtain error bars after finding the MAP, we evaluate the Hessian. For a complex random field with symmetry $v(\mathbf{k}) = v(-\mathbf{k})^\dagger$, we have

$$\text{cov}(v(\mathbf{k})v(\mathbf{k}')) = \frac{1}{2} \left(\frac{-\partial^2 \ln \mathcal{P}}{\partial v(\mathbf{k}) \partial v(\mathbf{k}')^\dagger} \right)^{-1} \Big|_{v=\hat{v}}. \quad (\text{B1})$$

We show the results of the MAP reconstruction in Fig. 9. It depicts power spectra of MAP estimator \hat{v} for two different levels of generated kSZ signal, power spectra of residuals $\epsilon = v_{\text{true}} - \hat{v}$, and analytic errors σ_{hess} . Analytic errors were evaluated via inversion of Hessian at MAP according to Eq. (B1). Because reconstructed modes are (nearly) independent, Hessians can be computed in narrow




FIG. 9. Power spectra of MAP-estimated velocities $\hat{v}_{1,2}$ and residuals with the true fields $\epsilon_{1,2}$ for two levels of generated kSZ signal along with estimated errors from the Hessian. The two errors match very well, i.e., the Hessian errors are identical to the true residual errors.

bins of ℓ , which significantly simplifies the evaluation and inversion process. Depicted values correspond to an average over the ℓ bins. We see that the error bar from the Hessian matches the true residual error very well.

-
- [1] R. A. Sunyaev and Ya. B. Zeldovich, The observations of relic radiation as a test of the nature of x-ray radiation from the clusters of galaxies, *Comments Astrophys. Space Phys.* **4**, 173 (1972).
 - [2] Nick Hand *et al.*, Evidence of galaxy cluster motions with the kinematic Sunyaev-Zel'dovich effect, *Phys. Rev. Lett.* **109**, 041101 (2012).
 - [3] Simon DeDeo, David N. Spergel, and Hy Trac, The kinetic Sunyaev-Zel'dovich effect as a dark energy probe, *arXiv: astro-ph/0511060*.
 - [4] Carlos Hernandez-Monteagudo, Licia Verde, Raul Jimenez, and David N. Spergel, Correlation properties of the kinematic Sunyaev-Zel'dovich effect and implications for dark energy, *Astrophys. J.* **643**, 598 (2006).
 - [5] Suman Bhattacharya and Arthur Kosowsky, Dark energy constraints from galaxy cluster peculiar velocities, *Phys. Rev. D* **77**, 083004 (2008).
 - [6] Eva-Maria Mueller, Francesco de Bernardis, Rachel Bean, and Michael D. Niemack, Constraints on massive neutrinos from the pairwise kinematic Sunyaev-Zel'dovich effect, *Phys. Rev. D* **92**, 063501 (2015).
 - [7] Joseph Kuruvilla and Cristiano Porciani, The n-point streaming model: How velocities shape correlation functions in redshift space, *J. Cosmol. Astropart. Phys.* **07** (2020) 043.
 - [8] Daisy S. Y. Mak, Elena Pierpaoli, Fabian Schmidt, and Nicolo' Macellari, Constraints on modified gravity from Sunyaev-Zeldovich cluster surveys, *Phys. Rev. D* **85**, 123513 (2012).
 - [9] Zhen Pan and Matthew C. Johnson, Forecasted constraints on modified gravity from Sunyaev-Zel'dovich tomography, *Phys. Rev. D* **100**, 083522 (2019).
 - [10] Neha Anil Kumar, Selim C. Hotinli, and Marc Kamionkowski, Uncorrelated compensated isocurvature perturbations from kinetic Sunyaev-Zeldovich tomography, *Phys. Rev. D* **107**, 043504 (2023).
 - [11] Moritz Münchmeyer, Mathew S. Madhavacheril, Simone Ferraro, Matthew C. Johnson, and Kendrick M. Smith, Constraining local non-Gaussianities with kinetic Sunyaev-Zel'dovich tomography, *Phys. Rev. D* **100**, 083508 (2019).
 - [12] Neha Anil Kumar, Gabriela Sato-Polito, Marc Kamionkowski, and Selim C. Hotinli, Primordial trispectrum from kinetic Sunyaev-Zel'dovich tomography, *Phys. Rev. D* **106**, 063533 (2022).
 - [13] Anne-Sylvie Deutsch, Emanuela Dimastrogiovanni, Matthew C. Johnson, Moritz Münchmeyer, and Alexandra Terrana, Reconstruction of the remote dipole and quadrupole fields from the kinetic Sunyaev Zel'dovich and polarized Sunyaev Zel'dovich effects, *Phys. Rev. D* **98**, 123501 (2018).

- [14] Kendrick M. Smith *et al.*, KSZ tomography and the bispectrum, [arXiv:1810.13423](#).
- [15] Utkarsh Giri and Kendrick M. Smith, Exploring KSZ velocity reconstruction with N-body simulations and the halo model, *J. Cosmol. Astropart. Phys.* **09** (2022) 028.
- [16] Juan Cayuso, Richard Bloch, Selim C. Hotinli, Matthew C. Johnson, and Fiona McCarthy, Velocity reconstruction with the cosmic microwave background and galaxy surveys, *J. Cosmol. Astropart. Phys.* **02** (2023) 051.
- [17] Juan I. Cayuso, Matthew C. Johnson, and James B. Mertens, Simulated reconstruction of the remote dipole field using the kinetic Sunyaev Zel'dovich effect, *Phys. Rev. D* **98**, 063502 (2018).
- [18] Utkarsh Giri and Kendrick M. Smith, Exploring KSZ velocity reconstruction with N-body simulations and the halo model, *J. Cosmol. Astropart. Phys.* **09** (2022) 028.
- [19] Alexandra Terrana, Mary-Jean Harris, and Matthew C. Johnson, Analyzing the cosmic variance limit of remote dipole measurements of the cosmic microwave background using the large-scale kinetic Sunyaev Zel'dovich effect, *J. Cosmol. Astropart. Phys.* **02** (2017) 040.
- [20] Pengjie Zhang, The dark flow induced small-scale kinetic Sunyaev-Zel'dovich effect, *Mon. Not. R. Astron. Soc.* **407**, L36 (2010).
- [21] Pengjie Zhang and Matthew C. Johnson, Testing eternal inflation with the kinetic Sunyaev Zel'dovich effect, *J. Cosmol. Astropart. Phys.* **06** (2015) 046.
- [22] Wayne Hu and Takemi Okamoto, Mass reconstruction with CMB polarization, *Astrophys. J.* **574**, 566 (2002).
- [23] Christopher M. Hirata and Uros Seljak, Analyzing weak lensing of the cosmic microwave background using the likelihood function, *Phys. Rev. D* **67**, 043001 (2003).
- [24] Julien Carron and Antony Lewis, Maximum a posteriori CMB lensing reconstruction, *Phys. Rev. D* **96**, 063510 (2017).
- [25] Marius Millea, Ethan Anderes, and Benjamin D. Wandelt, Bayesian delensing of CMB temperature and polarization, *Phys. Rev. D* **100**, 023509 (2019).
- [26] Marius Millea, Ethan Anderes, and Benjamin D. Wandelt, Sampling-based inference of the primordial CMB and gravitational lensing, *Phys. Rev. D* **102**, 123542 (2020).
- [27] M. Millea *et al.*, Optimal cosmic microwave background lensing reconstruction and parameter estimation with SPTpol data, *Astrophys. J.* **922**, 259 (2021).
- [28] Louis Legrand and Julien Carron, Lensing power spectrum of the cosmic microwave background with deep polarization experiments, *Phys. Rev. D* **105**, 123519 (2022).
- [29] Dagoberto Contreras, Fiona McCarthy, and Matthew C. Johnson, Maximum likelihood kinetic Sunyaev-Zel'dovich velocity reconstruction, *Phys. Rev. D* **107**, 023521 (2023).
- [30] Nhat-Minh Nguyen, Jens Jasche, Guilhem Lavaux, and Fabian Schmidt, Taking measurements of the kinematic Sunyaev-Zel'dovich effect forward: Including uncertainties from velocity reconstruction with forward modeling, *J. Cosmol. Astropart. Phys.* **12** (2020) 011.
- [31] Federico Bianchini and Marius Millea, Inference of gravitational lensing and patchy reionization with future CMB data, *Phys. Rev. D* **107**, 043521 (2023).
- [32] Yin Li *et al.*, pmwd: A differentiable cosmological particle-mesh N-body library, [arXiv:2211.09958](#).
- [33] Vanessa Böhm *et al.*, MADLENS, a PYTHON package for fast and differentiable non-Gaussian lensing simulations, *Astron. Comput.* **36**, 100490 (2021).
- [34] Denise Lanzieri, François Lanusse, Chirag Modi, Benjamin Horowitz, Joachim Harnois-Déraps, and Jean-Luc Starck, Forecasting the power of higher order weak lensing statistics with automatically differentiable simulations, *Astron. Astrophys.* **679**, A61 (2023).
- [35] Simon Foreman, Selim C. Hotinli, Mathew S. Madhavacheril, Alexander van Engelen, and Christina D. Kreisch, "De-kSZing" the cosmic microwave background with surveys of large-scale structure, *Phys. Rev. D* **107**, 083502 (2023).
- [36] J. Colin Hill, Simone Ferraro, Nick Battaglia, Jia Liu, and David N. Spergel, Kinematic Sunyaev-Zel'dovich effect with projected fields: A novel probe of the baryon distribution with Planck, WMAP, and WISE Data, *Phys. Rev. Lett.* **117**, 051301 (2016).
- [37] Fuyu Dong, Yu Yu, Jun Zhang, Xiaohu Yang, and Pengjie Zhang, Measuring the integrated Sachs-Wolfe effect from the low-density regions of the Universe, *Mon. Not. R. Astron. Soc.* **500**, 3838 (2020).
- [38] Selim C. Hotinli, Joel Meyers, Neal Dalal, Andrew H. Jaffe, Matthew C. Johnson, James B. Mertens, Moritz Münchmeyer, Kendrick M. Smith, and Alexander van Engelen, Transverse velocities with the moving lens effect, *Phys. Rev. Lett.* **123**, 061301 (2019).
- [39] Adam Rouhiainen and Moritz Münchmeyer, De-noising non-Gaussian fields in cosmology with normalizing flows, *36th Conference on Neural Information Processing Systems* (2022), [arXiv:2211.15161](#).
- [40] Adam Rouhiainen, Utkarsh Giri, and Moritz Münchmeyer, Normalizing flows for random fields in cosmology, [arXiv:2105.12024](#).
- [41] Uros Seljak, Grigor Aslanyan, Yu Feng, and Chirag Modi, Towards optimal extraction of cosmological information from nonlinear data, *J. Cosmol. Astropart. Phys.* **12** (2017) 009.
- [42] Jens Jasche and Benjamin D. Wandelt, Bayesian physical reconstruction of initial conditions from large-scale structure surveys, *Mon. Not. R. Astron. Soc.* **432**, 894 (2013).
- [43] Adrian E. Bayer, Chirag Modi, and Simone Ferraro, Joint velocity and density reconstruction of the Universe with nonlinear differentiable forward modeling, *J. Cosmol. Astropart. Phys.* **06** (2023) 046.
- [44] David Alonso, Javier Sanchez, and Anže Slosar, A unified pseudo- C_ℓ framework, *Mon. Not. R. Astron. Soc.* **484**, 4127 (2019).
- [45] Marius Millea and Uros Seljak, Marginal unbiased score expansion and application to CMB lensing, *Phys. Rev. D* **105**, 103531 (2022).
- [46] Yuuki Omori, Agora: Multi-component simulation for cross-survey science, [arXiv:2212.07420](#).
- [47] Anatoly Klypin, Gustavo Yepes, Stefan Gottlöber, Francisco Prada, and Steffen Heß, MultiDark simulations: The story of dark matter halo concentrations and density profiles, *Mon. Not. R. Astron. Soc.* **457**, 4340 (2016).
- [48] LSST Science Collaboration *et al.*, LSST Science Book, Version 2.0, [arXiv:0912.0201](#).

- [49] Hongbo Cai *et al.*, Bias to cosmic microwave background lensing reconstruction from the kinematic Sunyaev-Zel'dovich effect at reionization, *Phys. Rev. D* **105**, 043516 (2022).
- [50] Srinivasan Raghunathan, CMB-S4: Dark Radiation Anisotropy Flowdown Team (DRAFT) tool, GitHub repository, <https://github.com/sriniraghunathan/DRAFT>.
- [51] Dongwon Han and Neelima Sehgal, Mitigating foreground bias to the CMB lensing power spectrum for a CMB-HD survey, *Phys. Rev. D* **105**, 083516 (2022).
- [52] Francisco Villaescusa-Navarro *et al.*, The CAMELS project: Cosmology and astrophysics with machine learning simulations, *Astrophys. J.* **915**, 71 (2021).
- [53] Dylan Nelson *et al.*, The IllustrisTNG simulations: Public data release, [arXiv:1812.05609](https://arxiv.org/abs/1812.05609).
- [54] Kaiming He *et al.*, Deep residual learning for image recognition, *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, New York, 2016), [arXiv:1512.03385](https://arxiv.org/abs/1512.03385).
- [55] Biwei Dai and Uroš Seljak, Learning effective physical laws for generating cosmological hydrodynamics with Lagrangian deep learning, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2020324118 (2021).
- [56] Nico Hamaus, Uroš Seljak, Vincent Desjacques, Robert E. Smith, and Tobias Baldauf, Minimizing the stochasticity of halos in large-scale structure surveys, *Phys. Rev. D* **82**, 043515 (2010).
- [57] James Bradbury *et al.*, JAX: Composable transformations of Python + NumPy programs, Version 0.3.13, 2018.
- [58] Igor Babuschkin *et al.*, The DeepMind JAX Ecosystem, 2020.
- [59] Tom Hennigan *et al.*, Haiku: Sonnet for JAX, Version 0.0.9, 2020.
- [60] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPIX—A framework for high resolution discretization, and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [61] Charles R. Harris *et al.*, Array programming with NUMPY, *Nature (London)* **585**, 357 (2020).
- [62] Pauli Virtanen *et al.*, SCIPY1.0: Fundamental algorithms for scientific computing in PYTHON, *Nat. Methods* **17**, 261 (2020).
- [63] J. D. Hunter, MATPLOTLIB: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [64] Antony Lewis and Anthony Challinor, CAMB: Code for anisotropies in the microwave background, Astrophysics Source Code Library, record ascl:1102.026, 2011.
- [65] Kanan K. Datta, T. Roy Choudhury, and Somnath Bharadwaj, The multifrequency angular power spectrum of the epoch of reionization 21-cm signal, *Mon. Not. R. Astron. Soc.* **378**, 119 (2007).