

Generative Diffusion Priors for 3D Mapping of the Dark Universe

Brandon Zhao^{1*} Diana Scognamiglio² Olivier Doré^{2,3} Katherine L. Bouman^{1,3}

¹Department of Computing and Mathematical Sciences, California Institute of Technology

²Jet Propulsion Laboratory, California Institute of Technology

³Cahill Center for Astronomy and Astrophysics, California Institute of Technology

*byzhao@caltech.edu

Abstract

Reconstructing the three-dimensional distribution of dark matter from weak-lensing observations is a central but highly ill-posed inverse problem in cosmology. Unlike standard 3D reconstruction with multiple viewpoints, we observe the universe from a single line of sight, through noisy shape distortions of galaxies with uncertain distances, so meaningful recovery of the 3D matter field requires strong prior assumptions. Existing methods either produce point estimates with handcrafted priors or use neural ensembles for approximate Bayesian uncertainty, and struggle to capture the non-Gaussian, filamentary structure of the cosmic web. With the advent of new high-resolution cosmological simulations, we now have an alternative source of prior knowledge that captures the nonlinear statistics of structure formation with far greater fidelity than analytic prescriptions. We leverage these simulations to build a new dataset Conicus3D, which enables us to learn a data-driven diffusion-model prior capturing the full 3D distribution of dark matter structure across cosmic time. Building on recent plug-and-play approaches, we modify a diffusion-based posterior sampling scheme to the 3D weak-lensing setting, combining the learned prior with a differentiable physical forward model. On realistic simulations targeting a modern weak lensing survey, our approach yields substantially improved 2D and 3D reconstruction accuracy over baseline methods. Moreover, it produces posterior samples whose statistics closely track the underlying simulations, while remaining robust to moderate shifts in cosmology.

1. Introduction

Reconstructing the three-dimensional distribution of dark matter is a central challenge in modern cosmology. Because dark matter drives structure formation, recovering its 3D distribution is essential for testing gravity, constraining dark energy, and interpreting galaxy surveys to model the dynamics of the universe at large scales.

Dark matter does not emit or absorb light, so its struc-

ture must be inferred indirectly through its gravitational effects, most notably through the weak lensing (WL) of background galaxies. Recovering the underlying mass distribution from such indirect observations is an ill-posed inverse problem, and traditional reconstruction methods typically rely on smoothness-based priors which tend to overly damp high-frequency features.

Recent advances in cosmological simulations now offer high-resolution, large-volume realizations of structure formation across cosmic time [27, 29, 30, 44]. These datasets capture the full statistical complexity of the evolving dark matter field, enabling machine-learning approaches that learn directly from simulated universes. Modern generative models, in particular, provide powerful tools for representing these high-dimensional distributions as priors in inverse problems [48].

Building on these developments, we introduce a new dataset, Conicus3D, and diffusion-based generative prior for 3D cosmological mass mapping. Our dataset, derived from state-of-the-art N-body simulations, consists of dark matter light cones paired with 2D lens-plane projections that preserve the relevant physics for lensing-based inference. Using this dataset, we train a score-based diffusion model that learns the 3D statistical distribution of cosmic structure and serves as a prior within a posterior sampler to generate physically realistic mass reconstructions.

We demonstrate our method on the inverse problem of WL mass mapping, performing posterior sampling for tomographic reconstruction within simulations modeled after the James Webb Space Telescope (JWST) Cosmos-Web survey [9]. Our approach achieves substantially higher spatial fidelity than previous methods [39, 46], produces posterior samples whose statistics match true simulations, and generalizes robustly to cosmologies beyond the training domain. Our key contributions are:

- Conicus3D: A first-of-its-kind 3D dark matter light-cone dataset, derived from state-of-the-art N-body simulations [15, 27] and tailored for WL mass mapping, made




Figure 1. Lightcone structure and redshift-conditional diffusion model. Left: The large-scale matter distribution is represented as a 3D lightcone, built from a sequence of lens planes slicing the dark matter density field at increasing comoving distance (redshift). Right: Each lens plane is modeled with a redshift-conditional 2D diffusion prior that learns the statistics of simulated overdensity maps at that redshift. During generation, we start from Gaussian noise on each plane, append one-hot redshift encodings, and denoise with a U-Net score model to obtain a coherent 3D lightcone. This redshift-conditioned factorization enables efficient generation of realistic volumetric dark matter fields consistent with lightcone geometry.

publicly available for training and benchmarking by the computer vision and astrophysics communities;

- A diffusion-based posterior sampling framework that combines data-driven priors with WL observables to enable robust, uncertainty-aware cosmological inference.

All data, simulated measurements, and code will be publicly released with documentation, establishing a reproducible testbed that invites broad participation in advancing 3D dark matter mass mapping.

2. Background

In the standard Λ CDM cosmological model, the universe began hot and nearly uniform, seeded by tiny quantum fluctuations in the primordial density field. As the universe expanded and cooled, gravity amplified those fluctuations into the cosmic web of dark matter halos, filaments, and voids that now host galaxies and gas. A central goal of cosmology is to determine how matter is arranged and how it evolves: the statistics of the initial conditions, the growth of structure over time, and the astrophysical processes that modulate it. Observationally, we infer this history indirectly, by measuring how light is deflected as it travels through the universe.

The focus of this work is a probe known as weak gravitational lensing, which measures the tiny, coherent “shearing” of galaxy shapes as their light passes through uneven distributions of matter. For a single galaxy, the induced distortion is subtle and masked by the galaxy’s unknown intrinsic shape, but averaging over many galaxies recovers a shear field that statistically traces the underlying dark matter.

In cosmology, distance is often expressed in terms of the redshift z , which measures how much cosmic expansion has stretched a galaxy’s light. Larger redshift corresponds to *more distant* galaxies that we observe at *earlier* cosmic times. Throughout this work, we treat redshift as a coordinate along the line of sight. Moreover, galaxy distances can be estimated from their colors via photometric redshifts; these redshift estimates provide a noisy but informative 3D signal along the line of sight that can be used to constrain the radial structure of the matter field.

Because light travels at a finite speed, increasing redshift also means looking deeper into the past. As a result, cosmo-

logical probes such as WL measure integrated signals along our past light cone, offering a layered view of the evolving matter distribution (see Fig. 1). 3D reconstruction therefore reveals not only where matter structures exist, but when they emerge in the history of the Universe.

3. Related Work

Weak Lensing Mass Mapping Weak Lensing (WL) mass mapping is the inverse problem of reconstructing the distribution of total (mostly dark) matter along our past light cone from its gravitational-lensing imprints on galaxy images. In WL, deflections are small and can be approximated as slight rotations and axis-ratio distortions of intrinsically elliptical galaxy shapes. The induced shape change depends on a weighted line-of-sight projection of the matter overdensity, known as the convergence κ , which traces the integrated mass between the source galaxy and the observer. Because the cosmic matter field is dominated by dark matter, this task is referred to as *dark matter mass mapping*.

Kaiser and Squires [20] first showed that the 2D shear field γ describing coherent shape distortions as a function of sky position, can be analytically inverted to yield a projected 2D convergence map of the intervening mass. In practice, however, we only observe noisy estimates of γ inferred from galaxy ellipticities, contaminated by each galaxy’s unknown intrinsic shape (“shape noise”). Previous work proposes smoothing the observed galaxy shape field before inversion [20], while regularized approaches such as a Gaussian prior [37] and sparsity [25] have also been proposed. More recent developments include deep-learning-based reconstructions employing U-Net convolutional architectures [19] and diffusion models [34], which learn nonlinear priors directly from simulated data

While significant progress has been made in the development and implementation of 2D mass mapping algorithms, the problem of 3D mass mapping — deprojecting the 2D mass map into a 3D matter field — is extremely challenging and remains a largely unsolved problem. While observed galaxy redshifts (distances from the observer) can be estimated to give a 3D signal, the combi-

nation of intrinsic shape noise, large uncertainties in redshift estimation, and a single viewpoint limitation make the 3D mass mapping problem highly ill-posed. Existing approaches from the astrophysics community involve the use of strong regularizers such as a Gaussian prior [39], sparsity constraints [26], singular vector truncation [41] or multiresolution wavelet methods [28] to obtain highly over-smoothed point estimates. Recent machine learning-based works [46, 47] have explored neural-field representations and ensemble models that provide approximate uncertainty quantification but do not correspond to exact Bayesian posterior sampling. Our approach targets the exact posterior with respect to the chosen simulation-derived prior. Alternatively, rigorous Bayesian forward modeling approaches such as the BORG framework [32, 33] infer the 3D density field by sampling initial conditions and evolving them through gravity solvers, ensuring physical consistency but at significant computational expense. In this work, we introduce a diffusion-based framework that leverages outputs from state-of-the-art N-body simulations as physically motivated, data-driven priors. This approach produces high-resolution posterior samples of the 3D matter field consistent with both WL observations and structure-formation physics. While we can use the mean of this posterior as a point estimate, being able to also sample the posterior enables crucial uncertainty quantification as well as statistical estimation for downstream cosmological analyses. We are inspired by previous works [5, 6, 34] which use a similar data-driven prior for 2D mass mapping; however, we present a fully 3D dataset from a more advanced N-Body simulation, and estimate the posterior directly in 3D, while retaining the ability to also output 2D mass maps.

Diffusion Models for Inverse Problems Diffusion models have recently achieved state-of-the-art performance across a wide range of scientific inverse problems [48]. A straightforward yet common approach for general scientific inverse problems treats a diffusion model as a posterior estimator, incorporating the measurements as conditioning through concatenation to the model input. However, this ignores the physics underlying the actual measurement model and thus can produce physically implausible solutions. From a Bayesian perspective, if the measurement process is defined by a known likelihood $p(y|x)$, the diffusion model acts as a score estimator for the data-driven prior $p(x)$, approximating its log-gradient $\nabla_x \log p(x)$. To generate an unconditional prior sample x_0 , a diffusion model smoothly removes noise from an unstructured initial noise sample x_T by solving a stochastic differential equation. Many methods have been proposed for combining an unconditional diffusion model prior with a known measurement likelihood to sample from the resulting posterior. Early work such as DPS [11] uses guidance-based techniques to “nudge” samples along the likelihood gradient

during training; however, the resulting samples are not theoretically guaranteed to correspond to a proper posterior. To address this limitation, a family of *Plug-and-Play* [42] diffusion samplers has been developed, which incorporate the likelihood score term directly into the reverse process in a principled way. In this work, we propose a modified version of the DAPS (Decoupled Annealing Posterior Sampling, [48]) framework, which unifies the learned diffusion prior with a known forward model to perform physically consistent posterior sampling. This enables us to reconstruct matter fields that fulfill observational constraints and preserve the statistics learned from cosmological simulations.

4. Conicus3D - 3D Lightcone Dataset

We introduce Conicus3D, a large-scale 3D lightcone dataset designed for training and evaluating dark matter mass-mapping algorithms on state-of-the-art, realistic N-body simulations of the universe. By providing a comprehensive and well-documented resource derived from high-fidelity cosmological simulations, our goal is to make the challenge of 3D WL mass reconstruction more accessible to both the astrophysics and computer-vision communities. Upon release, Conicus3D will contain:

- 20,000 in-distribution dark matter lightcones, generated from a fiducial cosmology corresponding to the baseline parameters of the AbacusSummit suite [15, 27].
- 800 out-of-distribution lightcones, produced under alternative cosmological parameters that yield different statistical properties, enabling controlled evaluation of generalization performance across cosmological shifts.
- Mock galaxy shape catalogs, containing realistic galaxy positions, redshifts, and projection weights, along with open-source code for sampling additional shape catalogs to simulate observational noise and selection effects.

Dataset Construction Conicus3D is derived from the AbacusSummit simulations [15, 27], which evolve over 300 billion particles under varying cosmological models. Each N-body simulation can be viewed as a function of (1) the cosmological parameters, which determine the statistical properties of the initial matter field and the cosmic expansion history, and (2) a random phase field, which sets the initial perturbations of the simulation. Evolving these initial conditions forward generates realistic projections of the dark matter field across redshift.

To create each lightcone volume, we process particle counts from the AbacusSummit base simulation, projecting the simulated particles directly onto the observer’s lightcone geometry. We compute density contrasts (overdensities) for a sequence of redshift shells extending to $z \approx 2$, and bin the results into 20 lens planes equally spaced in comoving distance. Finally, we crop each plane to the COSMOS survey footprint so that the resulting lightcone volumes match the geometry and sky coverage of real observations.




Figure 2. Weak Lensing Measurements. (Left) As light from distant galaxies propagates through the universe, it is deflected by intervening matter, producing small differences between the intrinsic (unlensed) shape e_{int} and the observed shape e_{obs} . These subtle distortions provide a powerful probe of the 3D dark matter distribution. (Right) To characterize cosmic shear, galaxy shapes are often represented in the complex ellipticity plane, whose components encode the axis ratio and orientation of an ellipse. In typical WL surveys, shape noise is more than two orders of magnitude larger than the shear signal, so the lensing effect, described by the shear γ and approximately additive in ellipticity space, must be extracted statistically by combining many measurements from a dense field of galaxies to recover the coherent imprint of the underlying matter distribution. Figure adapted from [35, 46].

Broader Applications While our dataset is primarily designed for 3D dark matter mass mapping, its structure makes it naturally extensible to a wide range of inverse problems on the lightcone. Many astrophysical observables—such as Fast Radio Bursts (FRBs) [31], CMB secondary anisotropies [1, 8], and 21-cm intensity maps [7]—can be interpreted as line-of-sight projections of the underlying matter field across cosmic time. The same lightcone geometry, redshift binning, and comoving-distance sampling used in this work provide a unified framework for simulating and reconstructing these observables.

5. Methods

5.1. Weak Lensing Measurements

As light from distant galaxies travels toward us, its trajectory is continuously deflected by the intervening (dark-dominated) matter distribution. In the WL regime, these deflections are small and manifest as slight distortions in the observed galaxy shapes. To quantify these small changes, approximately elliptical galaxy shapes are described using a complex ellipticity parameter e . Given an approximately elliptical galaxy with axis ratio r and orientation angle ϕ , we can define the magnitude and phase of its ellipticity e as in [35]:

$$|e| = \frac{1-r}{1+r} \quad \angle e = 2\phi. \quad (1)$$

WL measurements relate a galaxy’s observed ellipticity e_{obs} to its intrinsic ellipticity e_{int} , which describes what its observed shape would be in the absence of lensing. In the regime of WL, the lensing effects of intervening matter can be approximately described as a shear γ , which is additive

in the ellipticity domain ([35]):

$$e_{\text{obs}} - e_{\text{int}} = \gamma. \quad (2)$$

An overview of shear measurements and the complex ellipticity can be found in Fig. 2. In practice, although the observed shape e_{obs} can be measured directly from an image, a galaxy’s intrinsic shape e_{int} is much more difficult to estimate. In traditional WL surveys, e_{int} is completely unknown [35]. The standard assumption is that the true shape of galaxies are randomly distributed without a preferred overall direction, so e_{int} is modeled as an additive Gaussian “shape noise”. To mitigate the effects of the shape noise, WL surveys combine measurements from many galaxies into a galaxy catalog containing the sky position, redshift (distance), and observed shapes of millions of galaxies, constituting the measurements for 3D mass mapping.

The shear γ observed in a survey image arises from the cumulative lensing occurring along the line of sight between us and the galaxy. More precisely, the shear is a projection followed by a 2D convolution of the *overdensity* $\delta = (\rho - \bar{\rho})/\bar{\rho}$ intersected by a ray of light as it travels towards earth. Here ρ is the continuous matter density field and $\bar{\rho}$ is the mean density of the universe at the time of the intersection. In this section we introduce only the specific equations governing the forward model; a more detailed derivation can be found in the supplement.

Projection To compute the shear from the matter overdensity field δ , we first introduce the convergence κ [35], which for a source galaxy with sky position θ and comoving distance¹ w is:

$$\kappa(\theta, w) = \mathbf{Q}\delta = \frac{3H_0^2\Omega_m}{2c^2} \int_0^w dw' \frac{w'(w-w')}{w} \frac{\delta(w'\theta, w')}{a(w')}, \quad (3)$$

where H_0 and Ω_m are cosmological parameters and c is the speed of light. The scale parameter $a(w)$ represents the known expansion of the universe that occurs as light travels towards us, and δ is the matter overdensity. In the case of standard WL surveys, w has a degree of uncertainty, as it is inferred from photometric estimates of the galaxy redshift. In this case, we take the convergence to be the expected value over the estimated distribution of the galaxy comoving distance w , thereby incorporating photometric redshift uncertainty directly into the forward model.

Convolution γ is obtained by convolving the convergence κ with a complex kernel $\mathcal{D}(\theta) = -1/(\theta^*)^2$ ([35]):

$$\gamma(\theta, w) = \mathbf{P}\kappa = \frac{1}{\pi} \int_{\mathbb{C}^2} d^2\theta' \mathcal{D}(\theta - \theta') \kappa(\theta', w). \quad (4)$$

Here, the angular position $\theta = \theta_1 + i\theta_2$ is expressed in complex coordinates, and the asterisk $*$ represents complex

¹Comoving distance refers to a distance measure that factors out the expansion of the universe.

conjugation. Notably, the forward operators in Eqns. 3 and 4 are both linear with respect to the overdensity δ , so obtaining the overdensity from the shear measurements is a linear inverse problem. Combining Eqns. 2, 3 and 4 gives the full formulation with the effects of the additive shape noise ε :

$$e_{\text{obs}} = \mathbf{P}\mathbf{Q}\delta + \varepsilon \quad \varepsilon \sim \mathcal{N}(0, \sigma_{\text{shape}}). \quad (5)$$

This assumption of additive Gaussian noise implies a Gaussian likelihood for the observed ellipticities across the galaxy catalog. Reconstructing $\kappa = \mathbf{Q}\delta$ rather than the full 3D overdensity is known as 2D mass mapping, and is much more well-constrained. A majority of the uncertainty in 3D mass mapping comes from the inversion of the projection operator \mathbf{Q} . Although some resolution in the radial direction can be inferred using the known distances to (i.e., redshifts of) galaxies, uncertainty in these redshift estimations coupled with shape noise and single-view geometry make the backprojection severely ill-posed.

5.2. Lightcone Diffusion Model

Diffusion models are powerful generative models which generate data by learning to reverse a diffusion noising process applied to the data distribution. In this work, we derive our data-driven prior by estimating the score of the data distribution $p(\delta)$ of possible lightcones δ . As is common in WL analyses, we represent each lightcone δ as a set of discrete lensplanes $\{\delta^{(z)}\}_{z=1}^M$, and take the distributions of lensplanes at different redshifts to be independent, i.e. $p(\delta) = \prod_{z=1}^M p(\delta^{(z)})$. Because the lensplanes cover such large distance scales, this is a reasonable assumption to make [39]; this can also be seen empirically in the simulated dataset (see Fig. 5). We then train a conditional diffusion model using Conicus3D to estimate the prior distributions $p(\delta^{(z)}|z)$ for each of M redshift values, which are fixed beforehand. Our diffusion model uses a standard DDPM UNet architecture operating on 128×128 lens-plane images. In this way, we can use a traditional image denoising model to generate 3D lightcones, by denoising a batch of images with a list of one-hot vectors corresponding to different redshifts; see Fig. 1 for an illustration.

5.3. Decoupled Annealing Posterior Sampling

Given our data-driven lightcone prior $p(\delta)$, and the physics-based likelihood $p(\gamma|\delta)$ from the forward model above, we wish to sample from the resulting posterior distribution

$$p(\delta|\gamma) \propto p(\gamma|\delta)p(\delta). \quad (6)$$

To do so, we propose a modified version of the Decoupled Annealing Posterior Sampling (DAPS) algorithm, which has been shown to be successful on a wide range of scientific inverse problems in the InverseBench framework [45].

DAPS samples a target “clean” posterior $p(\delta_0|\gamma)$ by sequentially generating samples from noise annealed posteriors $p(\delta_t|\gamma)$ from $t = T$ to 0. To generate samples δ_t from $\delta_{t+\Delta t}$, DAPS alternates between (1) sampling $\delta_0|\gamma \sim p(\delta_0|\delta_{t+\Delta t}, \gamma)$, and (2) sampling $\delta_t \sim \mathcal{N}(\delta_0|\gamma, \sigma_{t_2}^2 \mathbf{I})$. This process is repeated for decreasing t until δ_0 are sampled from the target posterior [45].

The original DAPS algorithm samples $p(\delta_0|\delta_t, \gamma) \propto p(\delta_0|\delta_t)p(\gamma|\delta_0)$ with score-based methods such as Langevin dynamics, approximating the intractable $p(\delta_0|\delta_t)$ with a Gaussian distribution with diagonal covariance Σ . While this approximation enables efficient posterior sampling, it does not encode any information about possible spatial correlations that may exist in δ . Motivated by the cosmological principles of translation invariance and isotropy, we propose a physically informed covariance which is instead diagonalized in Fourier Space according to the matter power spectrum (PS) $P(k)$:

$$\Sigma = F^{-1} \text{diag}(P_k) F. \quad (7)$$

In our experiments, we estimate $P(k)$ empirically by averaging over the training data. This Fourier-diagonal covariance structure arises in any inverse problem over approximately stationary fields, including adaptive optics [24], geophysical inversion [18], and CMB map-making [13], suggesting that our spectral sampling strategy could improve diffusion-based posterior sampling in these settings; see Supplement Sec. 15 for further discussion. A detailed explanation and pseudocode of our specific implementation can also be found in the supplement.

While the posterior mean provides a strong map-level reconstruction, most WL science evaluates non-linear functionals of the field posterior. Our sampler targets the full posterior, combining a data-driven diffusion prior with the physical likelihood so that each draw preserves small-scale power and realistic radial structure. This ensures that statistics computed from the ensemble—not only from the mean—are unbiased and well calibrated, which is crucial when moving from simulations to real data.

6. Experiments

6.1. Mock Weak Lensing Survey

We evaluate our method on simulated WL data modeled after the JWST COSMOS-Web survey [9], which targets one of the most extensively studied extragalactic fields, large enough to capture a representative portion of the cosmic large-scale structure. While the shear catalogs for JWST COSMOS-Web are not publicly available, we consider mock surveys that reproduce the footprint, source density, and shape-noise characteristics of the real COSMOS-Web observations [36]. We assume a footprint of 0.54 deg^2 with a mean redshift catalog source density of $\approx 261 \text{ galaxies arcmin}^{-2}$, an intrinsic shape-noise dispersion of $\sigma_e \approx 0.25$ [36], and Gaussian photo- z dispersion $\sigma_z = 0.11(1+z)$.




Figure 3. Two-dimensional mass reconstruction results from simulated WL data. For three representative lightcone volumes (rows), we show the projected convergence (κ) maps for the ground truth, our method, the Neural Ensemble baseline, and the Wiener filter reconstruction (columns); for each method, we use the posterior mean as a point estimate. Our reconstructions recover sharp cluster peaks and filamentary features that closely match the true κ maps, while the baselines exhibit oversmoothing and reduced contrast in high-density regions. The third row highlights performance under a cosmology mismatch; even in this out-of-distribution setting, our method maintains high-fidelity 2D mass reconstructions that better preserve the locations and amplitudes of prominent structures.

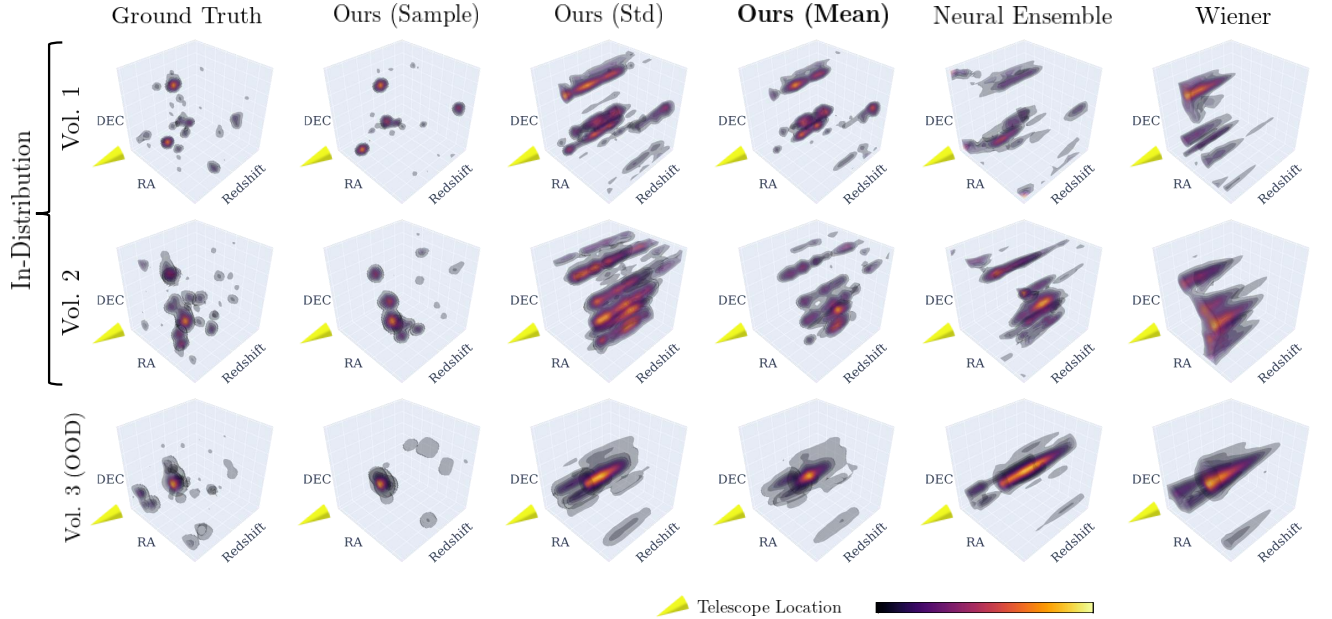


Figure 4. Three-dimensional mass reconstruction results from simulated WL data. Each row corresponds to a different simulated lightcone volume, and columns compare the ground-truth overdensity field to reconstructions from our diffusion-based posterior sampler, the Neural Ensemble baseline [46], and the Wiener filter method [39]; for each method, we use the posterior mean as a point estimate. Our reconstructions better recover the spatial morphology of clusters and filaments and place structures at the correct redshift slices, whereas the baselines tend to smear overdensities along the line of sight or miss small-scale features. The third row illustrates the challenging case of a misspecified cosmology with massless neutrinos, where our method still tracks the true 3D structure more faithfully despite a mismatch between the prior cosmology and the ground-truth simulation.




Figure 5. Sample quality. Our diffusion-based posterior sampler produces 3D mass-field realizations whose statistics match those of realistic dark matter simulations. Left: Angular PS C_ℓ^κ of the convergence κ for individual samples. Our samples retain the correct high- ℓ (small-scale) power, whereas the Neural Ensemble [46] and Wiener filter [39] baselines over-smooth small-scale modes when the noise power exceeds the signal power and the reconstruction becomes prior-dominated. Center: Radial PS of the overdensity δ across lens planes. Our samples are nearly flat in this radial spectrum, indicating decorrelation between different redshift planes and avoiding the spurious line-of-sight correlations seen in the neural ensemble outputs. Because many downstream WL statistics are nonlinear functionals of each sample (e.g., peak counts, voids), matching these sample-level power spectra is essential for unbiased cosmological inference and calibrated uncertainties. Right: Our diffusion prior is trained under a fiducial cosmology, which implies an expected angular PS for κ (“Prior”). We plot five power spectra: the ground truth and reconstruction for an in-distribution cosmology (In-Dist GT / In-Dist Recon), the ground truth and reconstruction for an out-of-distribution cosmology with massless neutrinos (OOD GT / OOD Recon), and the prior from our diffusion model. When the true cosmology matches the training cosmology, prior samples and reconstructions reproduce the same spectrum. In the out-of-distribution case, the true power is significantly higher, but our posterior samples still match the ground-truth spectrum due to likelihood guidance, showing that the method adapts to the observed universe even when the prior cosmology is misspecified.

We compare against two publicly available baseline methods for 3D mass mapping: the analytical Wiener filter reconstruction [39] and the neural ensemble estimator [46].

We first present results for recovering the projected 2D convergence field κ (a weighted integral of the overdensity δ along the line of sight) from a simulated light cone. Figure 3 and Table 1 show the reconstructed κ maps and quantitative comparisons. For each method, we compute the posterior mean projected overdensity across 128 samples and evaluate agreement with the ground-truth map using the mean-subtracted Pearson cross-correlation coefficient, which is appropriate since the WL forward model is invariant under constant additive shifts in κ . Our method achieves significantly higher correlation with the true convergence field than both the Wiener filter and neural ensemble baselines, demonstrating the ability of the diffusion-based posterior sampling framework to recover fine-scale structure while maintaining physically consistent large-scale modes. In Fig. 4 we reconstruct the 3D overdensity field δ for the same simulated light cones up to redshift $z = 1$. We again compute the posterior mean across 128 samples as a point estimate of the overdensity. Because of single-view limitations, and large amounts of noise, it is only possible to reconstruct the overdensity field with a limited radial resolution. Uncertainty in the radial position of structures results in point estimates with a blur along the line of sight. For each method, we additionally blur the ground truth and reconstructed volumes with a Gaussian filter of length $\sigma = 4$ lensplanes in the radial z direction (see supplement for detailed analysis of z blur lengths). We then report the cross-correlation (av-

Vol.	Method	$\rho_{\text{blur}}^{3\text{D}}(\uparrow)$	$\rho^{3\text{D}}(\uparrow)$	$\rho^{2\text{D}}(\uparrow)$
1	Ours	0.83	0.23	0.87
	Neur. [46]	0.79	0.21	0.86
	Wien. [39]	0.71	0.21	0.77
2	Ours	0.83	0.27	0.88
	Neur. [46]	0.80	0.21	0.87
	Wien. [39]	0.72	0.23	0.83
3	Ours	0.92	0.18	0.98
	Neur. [46]	0.86	0.09	0.96
	Wien. [39]	0.84	0.13	0.92

Table 1. Quantitative comparison of reconstruction quality. We summarize 3D and 2D mass-mapping performance for three simulated JWST-scale lightcone volumes. We report the Pearson cross-correlation coefficients averaged over lensplanes between reconstructed and ground-truth maps: $\rho_{\text{blur}}^{3\text{D}}$ where both lightcones are first blurred with a Gaussian kernel of width $\sigma = 4$ lensplanes in the line-of-sight direction, $\rho^{3\text{D}}$ for the full-resolution 3D fields, and $\rho^{2\text{D}}$ for the projected 2D convergence maps (higher is better in all cases). Our diffusion-based method consistently achieves higher correlations than both the Neural Ensemble and Wiener filter baselines across all three volumes in 2D and 3D.

eraged across lensplanes) of each reconstruction with the ground truth, blurred and unblurred. Our method outperforms the 3D reconstruction baselines in each volume.

6.2. Cosmology Mismatch

Because our method is supervised, a key question is whether or not we can generalize well under a reasonable error in our prior assumptions. Thus, we repeat the experiment described in Sec. 6.1 but for a ground truth lightcone generated under different cosmological assumptions.

In particular, we consider a cosmology with massless neutrinos, which results in significantly higher amplitudes in the PS. Recovery results are shown in row 3 of Figs. 3, 4. Our method generalizes well despite a mismatch in the statistics of the prior and ground truth volumes. In this case, the OOD cosmology produces higher power and thus stronger lensing signal, allowing the likelihood to partially overcome the mismatched prior.

In addition, we compare the PS of both the fiducial (assumed) cosmology and the out-of-distribution cosmology, the PS for both of our reconstructions, and the PS for unconditional PS in Fig. 5. Although our prior samples match the fiducial cosmology of the training set, our method is able to generate posterior samples with angular PS matching the ground truth due to likelihood guidance.

6.3. Sample Quality

Beyond point estimation, our method’s principal advantage is the ability to *sample* from a realistic posterior whose statistics are consistent with the true dark matter distribution. We assess sample realism using the angular PS of the recovered 2D convergence maps (Fig. 5); matching both amplitude and scale-dependence indicates correct two-point statistics.

Although all methods have access to the same theoretical PS, smoothness regularization in the Neural Ensemble and Wiener filter mean estimators inevitably damps small-scale structure. This suppression wipes out high-frequency power and produces convergence maps that look overly smoothed (Fig. 5). In contrast, our diffusion-based posterior generates high-resolution samples whose angular power spectra match the ground truth across all scales. This fidelity is essential for extracting reliable cosmological information from WL data. Many key probes, such as peak and void counts [22], Minkowski functionals [23], bispectra [12], halo and filament finders [21, 40], cross-correlation with galaxies, clusters [10, 43], or CMB lensing [17], depend directly on the detailed small-scale and radial structure in each possible 3D realization of the matter field. If the samples are oversmoothed or exhibit spurious correlations across redshift planes, the resulting cross-correlation amplitudes, cosmological parameter constraints, and covariance estimates will be systematically biased. Our samples robustly reproduce the correct angular power and remain decorrelated along the line of sight, ensuring that downstream non-Gaussian statistics remain unbiased and that ensemble-based uncertainty propagation is physically meaningful.

6.4. Uncertainty Calibration

Another important aspect of our method is the ability to estimate uncertainty from the statistics of samples recovered from our posterior sampler. Results are shown in Fig. 6. We show a very clear correlation between sample standard deviation and error with the ground truth, with a near perfect




Figure 6. Uncertainty calibration. For each voxel, we compute the standard deviation of the posterior samples (“sample std”) and the mean absolute error (MAE) between the posterior mean and the ground truth. We then bin voxels by sample std and plot the average MAE per bin, along with the fraction of voxels in each bin. Both our method ($r = 0.92$) and the neural ensemble estimator ($r = 0.90$) [46] show a strong correlation between predicted uncertainty and actual error, but only our samples correspond to a well-defined Bayesian posterior distribution.

correspondence for a majority of the lightcone voxels. Additionally, although both our method and the neural ensemble method offer “samples” for uncertainty quantification, only our method can be interpreted as a proper Bayesian posterior. Together with the sample-level power-spectrum agreement in Fig. 5, this behavior indicates that posterior-predictive intervals for non-linear summaries (e.g., peak counts, cross-correlations in COSMOS-Web) should be correctly calibrated on real data.

7. Conclusion

In this work, we introduce a first-of-its-kind 3D dark matter lightcone dataset tailored for WL mass mapping. Building on this resource, we propose a diffusion-based posterior sampler that reconstructs the 3D matter distribution from noisy 2D shear observations. Our approach significantly outperforms Neural Ensemble and Wiener filter baselines in 3D/2D correlation and power-spectrum fidelity on simulations designed to mimic JWST observations, and remains effective even under cosmology mismatch, indicating robustness to realistic modeling uncertainties. Important directions for future work include analyzing the sensitivity of reconstructions to the choice of simulation-derived prior, including different cosmological parameter settings and simulation codes. By releasing our dataset, code, and evaluation protocol, we aim to create a reproducible testbed that opens 3D cosmological mass mapping to a far broader research community. Our work brings humanity closer to seeing the dark universe in three dimensions, yielding reconstructions that not only reveal its structure with unmatched fidelity, but also unlock richer and more powerful downstream analyses.

8. Acknowledgements

Part of this work was done at Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004) and funded through the President’s and Director’s Research and Development Fund (PDRDF). This work was also supported by the Stanback Innovation Fund and NSF Career award 2048237.

References

- [1] Nabila Aghanim, Subhabrata Majumdar, and Joseph Silk. Secondary anisotropies of the cmb. *Reports on Progress in Physics*, 71(6):066902, 2008. 4
- [2] Bradley Atcheson, Ivo Ihrke, Wolfgang Heidrich, Art Tevs, Derek Bradley, Marcus Magnor, and Hans-Peter Seidel. Time-resolved 3d capture of non-stationary gas flows. *ACM transactions on graphics (TOG)*, 27(5):1–9, 2008. 4
- [3] Matthias Bartelmann and Peter Schneider. Weak gravitational lensing. *Physics Reports*, 340(4-5):291–472, 2001. 3, 4
- [4] Max Born and Emil Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013. 4
- [5] Supranta S Boruah, Michael Jacob, and Bhuvnesh Jain. Diffusion-based mass map reconstruction from weak lensing data. *Physical Review D*, 111(8):083542, 2025. 3
- [6] Supranta S Boruah, Michael Jacob, Bhuvnesh Jain, Riya Maiya, and Raghav Venkataramanan. High-resolution weak lensing mass mapping from des-y3 data using diffusion-based prior. *arXiv preprint arXiv:2511.14667*, 2025. 3
- [7] Philip Bull, Pedro G Ferreira, Prina Patel, and Mario G Santos. Late-time cosmology with 21 cm intensity mapping experiments. *The Astrophysical Journal*, 803(1):21, 2015. 4
- [8] John E Carlstrom, Gilbert P Holder, and Erik D Reese. Cosmology with the sunyaev-zel’dovich effect. *Annual Review of Astronomy and Astrophysics*, 40(1):643–680, 2002. 4
- [9] Caitlin M Casey, Jeyhan S Kartaltepe, Nicole E Drakos, Maximilien Franco, Santosh Harish, Louise Paquereau, Olivier Ilbert, Caitlin Rose, Isabella G Cox, James W Nightingale, et al. Cosmos-web: an overview of the jwst cosmic origins survey. *The Astrophysical Journal*, 954(1):31, 2023. 1, 5
- [10] Chihway Chang, Vinu Vikram, Bhuvnesh Jain, David Bacon, A Amara, MR Becker, G Bernstein, C Bonnett, S Bridle, D Brout, et al. Wide-field lensing mass maps from dark energy survey science verification data. *Physical review letters*, 115(5):051301, 2015. 8
- [11] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022. 3
- [12] Asantha Cooray and Wayne Hu. Weak gravitational lensing bispectrum. *The Astrophysical Journal*, 548(1):7, 2001. 8
- [13] Franz Elsner and Benjamin D Wandelt. Efficient Wiener filtering without preconditioning. *Astronomy & Astrophysics*, 549:A111, 2013. 5, 7
- [14] Lawrence C Evans. *Partial differential equations*. American mathematical society, 2022. 5
- [15] Lehman H Garrison, Daniel J Eisenstein, Douglas Ferrer, Nina A Maksimova, and Philip A Pinto. The abacus cosmological n-body code. *Monthly Notices of the Royal Astronomical Society*, 508(1):575–596, 2021. 1, 3, 5
- [16] Boryana Hadzhiyska, Sihan Yuan, Chris Blake, Daniel J Eisenstein, J Aguilar, Steven Ahlen, David Brooks, Todd Claybaugh, Axel de la Macorra, Peter Doel, et al. Synthetic light-cone catalogues of modern redshift and weak lensing surveys with abacussummit. *Monthly Notices of the Royal Astronomical Society*, 525(3):4367–4387, 2023. 5, 6
- [17] Nick Hand, Alexie Leauthaud, Sudeep Das, Blake D Sherwin, Graeme E Addison, J Richard Bond, Erminia Calabrese, Aldee Charbonnier, Mark J Devlin, Joanna Dunkley, et al. First measurement of the cross-correlation of cmb lensing and galaxy lensing. *Physical Review D*, 91(6):062001, 2015. 8
- [18] Thomas Mejer Hansen, Andre G Journal, Albert Tarantola, and Klaus Mosegaard. Linear inverse Gaussian theory and geostatistics. *Geophysics*, 71(6):R101–R111, 2006. 5, 7
- [19] Niall Jeffrey, François Lanusse, Ofer Lahav, and Jean-Luc Starck. Deep learning dark matter map reconstructions from des sv weak lensing data. *Monthly Notices of the Royal Astronomical Society*, 492(4):5023–5029, 2020. 2
- [20] Nick Kaiser and Gordon Squires. Mapping the dark matter with weak gravitational lensing. *Astrophysical Journal, Part 1 (ISSN 0004-637X)*, vol. 404, no. 2, p. 441–450., 404:441–450, 1993. 2, 5
- [21] Alexander Knebe, Steffen R Knollmann, Stuart I Muldrew, Frazer R Pearce, Miguel Angel Aragon-Calvo, Yago Ascasibar, Peter S Behroozi, Daniel Ceverino, Stephane Colombi, Juerg Diemand, et al. Haloes gone mad: the halo-finder comparison project. *Monthly Notices of the Royal Astronomical Society*, 415(3):2293–2318, 2011. 8
- [22] Jan M Kratochvil, Zoltán Haiman, and Morgan May. Probing cosmology with weak lensing peak counts. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 81(4):043519, 2010. 8
- [23] Jan M Kratochvil, Eugene A Lim, Sheng Wang, Zoltán Haiman, Morgan May, and Kevin Huffenberger. Probing cosmology with weak lensing minkowski functionals. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 85(10):103513, 2012. 8
- [24] Richard G Lane, Andreas Glindemann, and John C Dainty. Simulation of a Kolmogorov phase screen. *Waves in Random Media*, 2(3):209–224, 1992. 5, 7
- [25] Francois Lanusse, J-L Starck, Adrienne Leonard, and Sandrine Pires. High resolution weak lensing mass mapping combining shear and flexion. *Astronomy & Astrophysics*, 591:A2, 2016. 2
- [26] Adrienne Leonard, François Lanusse, and Jean-Luc Starck. Glimpse: accurate 3d weak lensing reconstructions using sparsity. *Monthly Notices of the Royal Astronomical Society*, 440(2):1281–1294, 2014. 3, 5
- [27] Nina A Maksimova, Lehman H Garrison, Daniel J Eisenstein, Boryana Hadzhiyska, Sownak Bose, and Thomas P

- Satterthwaite. Abacussummit: a massive set of high-accuracy, high-resolution n-body simulations. *Monthly Notices of the Royal Astronomical Society*, 508(3):4017–4037, 2021. 1, 3, 5
- [28] Richard Massey, Jason Rhodes, Richard Ellis, Nick Scoville, Alexie Leauthaud, Alexis Finoguenov, Peter Capak, David Bacon, Hervé Aussel, Jean-Paul Kneib, et al. Dark matter maps reveal cosmic scaffolding. *Nature*, 445(7125):286–290, 2007. 3
- [29] Dylan Nelson, Volker Springel, Annalisa Pillepich, Vicente Rodriguez-Gomez, Paul Torrey, Shy Genel, Mark Vogelsberger, Ruediger Pakmor, Federico Marinacci, Rainer Weinberger, et al. The illustrious simulations: public data release. *Computational Astrophysics and Cosmology*, 6(1):2, 2019. 1
- [30] Nathanaël Perraudin, Ankit Srivastava, Aurelien Lucchi, Tomasz Kacprzak, Thomas Hofmann, and Alexandre Réfrégier. Cosmological n-body simulations: a challenge for scalable generative models. *Computational Astrophysics and Cosmology*, 6(1):5, 2019. 1
- [31] Emily Petroff, JWT Hessels, and DR Lorimer. Fast radio bursts. *The Astronomy and Astrophysics Review*, 27(1):4, 2019. 4
- [32] Natalia Porqueres, Alan Heavens, Daniel Mortlock, and Guilhem Lavaux. Bayesian forward modelling of cosmic shear data. *Monthly Notices of the Royal Astronomical Society*, 502(2):3035–3044, 2021. 3
- [33] Natalia Porqueres, Alan Heavens, Daniel Mortlock, and Guilhem Lavaux. Lifting weak lensing degeneracies with a field-based likelihood. *Monthly Notices of the Royal Astronomical Society*, 509(3):3194–3202, 2022. 3
- [34] Benjamin Remy, Francois Lanusse, Niall Jeffrey, Jia Liu, J-L Starck, Ken Osato, and Tim Schrabback. Probabilistic mass-mapping with neural score estimation. *Astronomy & Astrophysics*, 672:A51, 2023. 2, 3
- [35] Peter Schneider. Weak gravitational lensing. In *Gravitational lensing: strong, weak and micro*, pages 269–451. Springer, 2006. 4, 3
- [36] Diana Scognamiglio et al. The highest resolution map of (dark) matter. *accepted for Nature Astronomy*, 2025. 5
- [37] Uroš Seljak. Weak lensing reconstruction and power spectrum estimation: minimum variance methods. *The Astrophysical Journal*, 506(1):64, 1998. 2
- [38] Marko Shuntov, Hollis B Akins, Louise Paquereau, Caitlin M Casey, Olivier Ilbert, Rafael C Arango-Toro, Henry Joy McCracken, Maximilien Franco, Santosh Harish, Jeyhan S Kartaltepe, et al. Cosmos2025: The cosmos-web galaxy catalog of photometry, morphology, redshifts, and physical parameters from jwst, hst, and ground-based imaging. *arXiv preprint arXiv:2506.03243*, 2025. 5
- [39] Patrick Simon, AN Taylor, and Jan Hartlap. Unfolding the matter distribution using three-dimensional weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 399(1):48–68, 2009. 1, 3, 5, 6, 7
- [40] Thierry Sousbie. The persistent cosmic web and its filamentary structure–i. theory and implementation. *Monthly notices of the royal astronomical society*, 414(1):350–383, 2011. 8
- [41] JT VanderPlas, AJ Connolly, Bhuvnesh Jain, and Mike Jarvis. Three-dimensional reconstruction of the density field: An svd approach to weak-lensing tomography. *The Astrophysical Journal*, 727(2):118, 2011. 3
- [42] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE global conference on signal and information processing*, pages 945–948. IEEE, 2013. 3
- [43] Vinu Vikram, Chihway Chang, Bhuvnesh Jain, D Bacon, Adam Amara, Matthew R Becker, G Bernstein, Christopher Bonnett, Sarah Bridle, Dillon Brout, et al. Wide-field lensing mass maps from dark energy survey science verification data: Methodology and detailed analysis. *Physical Review D*, 92(2):022006, 2015. 8
- [44] Francisco Villaescusa-Navarro, ChangHoon Hahn, Elena Massara, Arka Banerjee, Ana Maria Delgado, Doogesh Kodi Ramanah, Tom Charnock, Elena Giusarma, Yin Li, Erwan Allys, et al. The quijote simulations. *The Astrophysical Journal Supplement Series*, 250(1):2, 2020. 1
- [45] Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 20895–20905, 2025. 5, 7
- [46] Brandon Zhao, Aviad Levis, Liam Connor, Pratul P Srinivasan, and Katherine Bouman. Revealing the 3d cosmic web through gravitationally constrained neural fields. In *The Thirteenth International Conference on Learning Representations*. 1, 3, 4, 6, 7, 8
- [47] Brandon Zhao, Aviad Levis, Liam Connor, Pratul P Srinivasan, and Katherine L Bouman. Single view refractive index tomography with neural fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25358–25367, 2024. 3
- [48] Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy T Feng, Caifeng Zou, Yu Sun, Nikola Kovachki, Zachary E Ross, et al. Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. *arXiv preprint arXiv:2503.11043*, 2025. 1, 3

Generative Diffusion Priors for 3D Mapping of the Dark Universe

Supplementary Material

9. Visualization of Sample Quality

A central goal of our diffusion-based posterior sampler is not only to produce a high-fidelity *posterior mean* as a point estimate of the mass distribution, but also to generate *realistic posterior samples* that reflect the variability and statistical structure of the underlying dark matter field. In this section, we provide qualitative visualizations of 2D and 3D samples (Figs. 7 and 8) that complement the quantitative sample-quality diagnostics in the main text.

In the two-dimensional case, both our method and the neural ensemble produce posterior means that broadly recover the ground-truth convergence field κ . However, the individual samples reveal clear differences in sample quality. Samples from the neural ensemble contain spurious small-scale artifacts that are not present in the true κ maps and that only disappear after averaging over many samples. In contrast, samples drawn from our diffusion-based posterior are relatively stable from realization to realization, without significant sample-wise artifacts. This indicates that our learned prior and sampling procedure preserve physically plausible structure at the level of individual convergence maps, not just in the mean.

The three-dimensional visualizations highlight an additional limitation of neural-ensemble-based uncertainty estimates. Neural ensemble samples exhibit spurious correlations along the line of sight: overdense regions are stretched into radially elongated features that persist across neighboring redshift slices, inconsistent with the statistics of the ground-truth overdensity field δ and with the radial power-spectrum behavior discussed in the main text. Our diffusion-based posterior samples, by contrast, maintain realistic slice-to-slice variability and remain relatively decorrelated across lens planes. The posterior mean naturally exhibits modest radial blurring, reflecting the intrinsic limits of 3D weak-lensing reconstruction, but individual samples themselves do not suffer from spurious line-of-sight coherence. These properties are crucial for downstream non-Gaussian analyses that operate directly on posterior samples, ensuring that ensemble-based statistics remain physically meaningful and well calibrated.

10. Additional Recovery Results

To further validate the robustness of our method, we present recovery results for three additional simulated lightcones (Volumes 4, 5 and 6), expanding upon the qualitative and quantitative comparisons shown in the main paper. Similar to the reconstruction results in the main text, volumes 4 and 5 correspond to the in-distribution cosmology of

Vol.	Method	$\rho_{\text{blur}}^{3D}(\uparrow)$	$\rho^{3D}(\uparrow)$	$\rho^{2D}(\uparrow)$
4	Ours	0.88	0.25	0.93
	Neur. [46]	0.85	0.24	0.92
	Wien. [39]	0.84	0.24	0.90
5	Ours	0.85	0.25	0.91
	Neur. [46]	0.81	0.18	0.87
	Wien. [39]	0.80	0.22	0.86
6	Ours	0.95	0.24	0.88
	Neur. [46]	0.91	0.15	0.82
	Wien. [39]	0.86	0.23	0.77

Table 2. Additional quantitative comparison of recovery performance. We report 3D and 2D reconstruction quality for three additional simulated JWST-scale lightcone volumes: two in-distribution (Volumes 4 and 5) and one out-of-distribution with massless neutrinos (Volume 6). As in the main paper, we compute (1) the blurred 3D correlation ρ_{3D}^{blur} after applying a Gaussian filter of width $\sigma = 4$ lensplanes along the line of sight, (2) the full-resolution 3D correlation ρ_{3D} , and (3) the 2D convergence correlation ρ_{2D} . Across all metrics and both volumes, our diffusion-based posterior sampler achieves the highest agreement with the ground truth, outperforming the Neural Ensemble and Wiener filter baselines.

the training set, while volume 6 corresponds to an out-of-distribution cosmology with massless neutrinos. Figure 9 and Figure 10 illustrate the 2D convergence κ reconstruction and the full 3D overdensity δ recovery, respectively. Consistent with the three volumes shown in the main text, our diffusion-based posterior sampler reconstructs high-contrast cluster peaks and filamentary structure that closely match the ground truth, while the Neural Ensemble baseline tends to smooth overdensities and the Wiener filter over-regularizes fine features. These qualitative trends persist across all test volumes—including those with complex multi-scale structure—and indicate that our sampler reliably maintains small-scale power and accurate radial localization.

Table 2 provides quantitative cross-correlation metrics for these additional volumes. As before, we report (1) the blurred 3D correlation using $\sigma = 4$ lensplanes, (2) the full-resolution 3D correlation, and (3) the projected 2D convergence correlation. Across all metrics, our method again achieves the highest agreement with the ground truth.




Figure 7. 2D Convergence κ Posterior Samples. Left: Ground-truth convergence field for a representative lightcone. Top row: Posterior mean and individual posterior samples from our diffusion-based method. Bottom row: Posterior mean and samples from the neural ensemble baseline. While both methods produce reasonable posterior means, the neural ensemble samples contain spurious small-scale artifacts that are washed out only when averaging. In contrast, our diffusion-based posterior produces visually coherent, artifact-free samples consistent with the statistics of the true κ field.

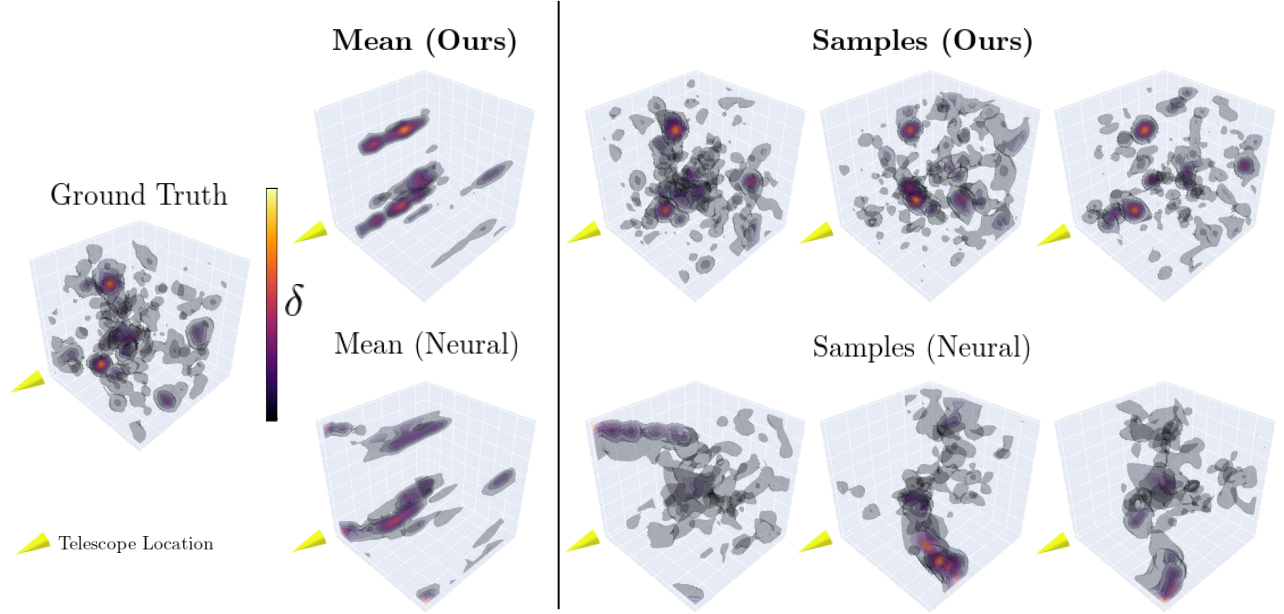


Figure 8. 3D Overdensity δ Posterior Samples. Left: Ground-truth 3D overdensity for a representative volume. Top row: Posterior mean and individual samples from our diffusion-based sampler. Bottom row: Corresponding mean and samples from the neural ensemble. Neural-ensemble samples display artificial correlations along the line of sight, producing radially elongated structures not present in the true simulation. Our samples maintain realistic slice-to-slice variability and avoid spurious line-of-sight correlations; only the posterior mean exhibits expected radial blurring. These properties underscore the validity of our method’s 3D posterior samples for downstream non-Gaussian cosmological analyses.




Figure 9. Additional 2D Recovery Results. For three additional simulated lightcones (Volumes 4, 5 and 6), we compare the reconstructed convergence κ maps from our method, the Neural Ensemble baseline, and the Wiener filter against the ground truth. Consistent with the results in the main text, our method preserves high-contrast peaks and sharp filamentary features, whereas the baselines exhibit over-smoothing and reduced small-scale power. These trends persist across diverse structures and demonstrate the robustness of our 2D recovery performance.

11. Detailed Resolution Analysis

In the main paper, we evaluated 3D reconstruction fidelity using a single Gaussian blur scale ($\sigma = 4$ lensplanes) as a summary metric for radial resolution. Figure 11 shows the blurred 3D cross-correlation $\rho_{3D}^{\text{blur}}(\sigma)$ between each reconstruction and the ground-truth volume as a function of σ , where both volumes are convolved with the same Gaussian kernel and larger σ corresponds to more aggressive smoothing along the line of sight. For all five test volumes, our method outperforms the baselines at every blur scale, demonstrating that the improvements provided by our diffusion prior are not tied to any particular choice of radial smoothing.

12. Weak Gravitational Lensing Formalism

12.1. Deflection and Shape Distortion

A more explicit derivation of Eq. (3)–Eq. (5) is presented below. A more explicit derivation of Eq. (3)–Eq. (5) is presented below. Further details can be found in [3, 35].

Consider a background galaxy (“source”) whose image is being weakly lensed by an extended 3D mass distribution (“lens”). The source is observed at angular position θ on the sky, whereas in the absence of lensing it would appear at β . We therefore define the scaled deflection angle α through the lens equation:

$$\beta = \theta - \alpha. \quad (8)$$

For an extended 3D mass distribution, the deflection an-

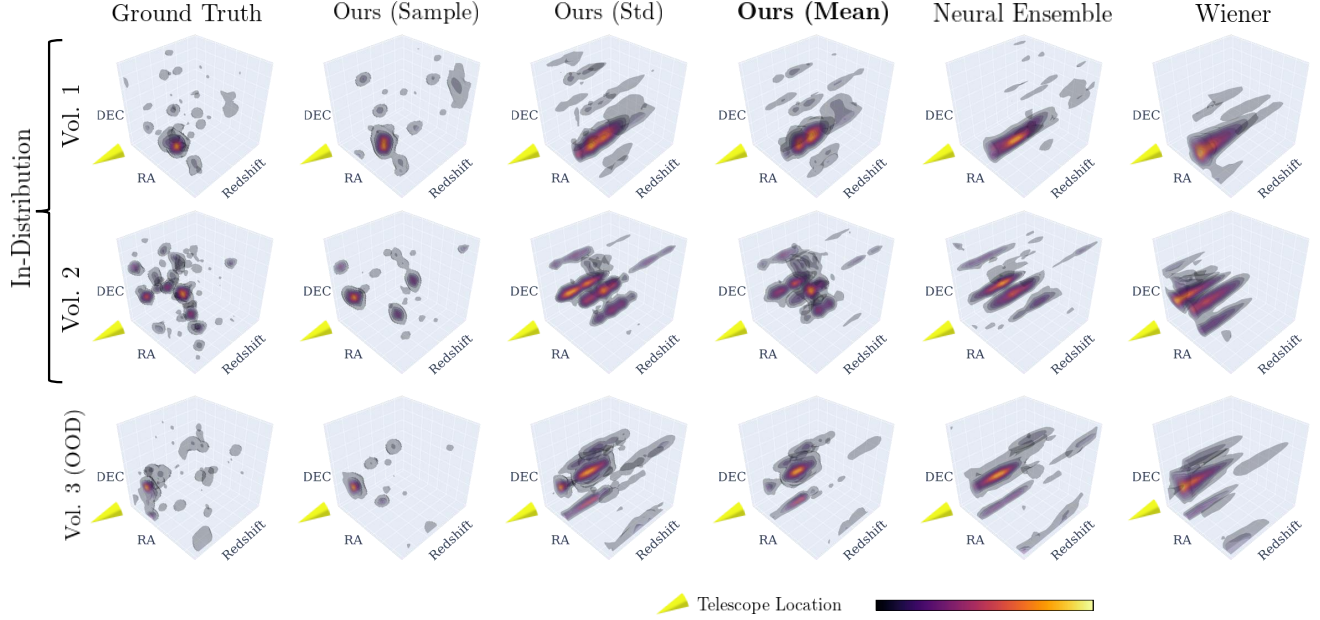


Figure 10. Additional 3D Recovery Results. We show full 3D overdensity δ reconstructions for Volumes 4, 5 and 6, comparing our diffusion-based posterior sampler with the Neural Ensemble and Wiener filter baselines. Our method more accurately localizes structures to the correct redshift slices and recovers finer-scale morphology along the line of sight. The baselines tend to smear overdense regions or miss small-scale features entirely. Together with Fig. 9, these results confirm that the improvements observed in the main paper hold across all evaluated lightcones.

gle depends on the gravitational potential Φ encountered by the light ray along its path from source to observer. Under the Born approximation [3], this becomes the line-of-sight integral:

$$\alpha(\theta, w) = \frac{2}{c^2} \int_0^w dw' \frac{w - w'}{w} \nabla_{\perp} \Phi(w' \theta, w'). \quad (9)$$

This expression is closely analogous to the Eikonal ray equation from optics, which describes the deflection of a ray by a refractive-index field [2, 4]. For convenience, we pull the transverse gradient ∇_{\perp} outside the integral and define the *effective lensing potential* $\psi(\theta, w)$, which represents the weighted integral of Φ along the line of sight:

$$\psi(\theta, w) = \frac{2}{c^2} \int_0^w dw' \frac{w - w'}{w \cdot w'} \Phi(w' \theta, w'), \quad (10)$$

$$\alpha = \nabla_{\theta} \psi \quad (11)$$

While gravitational lensing deflects the apparent position of light rays, our primary interest lies in how it distorts the ellipticity e of a galaxy image (Eq. (1)). In the weak-lensing regime, a first-order expansion shows that the shear components γ (Eq. (2)) are given by the partial derivatives of the scaled deflection angle [3]:

$$\gamma_1 \approx \frac{1}{2}(\alpha_{11} - \alpha_{22}) \quad (12)$$

$$\gamma_2 \approx \alpha_{12} = \alpha_{21} \quad (13)$$

Here, we adopt the notation $\alpha_{ij} = \partial \alpha_i / \partial \theta_j$, and $\alpha_{12} = \alpha_{21}$ follows from the fact that α is a gradient field. Consequently, the shear can also be expressed in terms of the effective lensing potential:

$$\gamma_1 = \frac{1}{2} \left(\frac{\partial^2 \psi}{\partial \theta_1^2} - \frac{\partial^2 \psi}{\partial \theta_2^2} \right) \quad (14)$$

$$\gamma_2 = \frac{\partial^2 \psi}{\partial \theta_1 \partial \theta_2} \quad (15)$$

Intuitively, instead of a point-like ray, one may picture a small ellipse representing a galaxy image. The gradient (first-order derivatives) of Φ determines the deflection α of the ellipse's center, while the second-order derivatives of Φ describe the shear γ that distorts its shape. To connect this potential to the underlying matter field, quantified by the density overdensity δ , we now introduce the convergence

κ :

$$\kappa = \frac{1}{2} \Delta_{\theta} \psi \quad (16)$$

$$= \frac{1}{c^2} \Delta_{\theta} \int_0^w dw' \frac{w - w'}{w \cdot w'} \Phi(w' \theta, w') \quad (17)$$

$$= \frac{1}{c^2} \int_0^w dw' \frac{w'(w - w')}{w} \Delta \Phi(w' \theta, w') \quad (18)$$

In Eq. (18) we assume that the Φ has negligible derivative in the radial direction. Next, applying the Poisson equation for gravitational potential in comoving coordinates,

$$\Delta \Phi = \frac{3H_0^2 \Omega_m \delta}{2a} \quad (19)$$

which leads directly to the weighted projection of the matter overdensity given in Eq. (3). To relate the convergence κ to the shear γ , we first invert the Poisson equation given in Eq. (16) using the method of Green's functions [14, 20]:

$$\tilde{\psi}(\theta) = \frac{1}{\pi} \int_{\mathbb{R}^2} d^2 \theta' \ln |\theta - \theta'| \kappa(\theta') \quad (20)$$

Finally, plugging $\tilde{\psi}$ into Eq. (14) and Eq. (15) yields Eq. (4).

12.2. Photometric Redshift Errors

When a survey relies on photometric redshifts, the projection defined in equation Eq. (3) must be marginalized over the uncertainty in the source's true redshift. For a galaxy with a redshift probability distribution $p(z)$, the expected convergence is therefore obtained by averaging $\kappa(\theta, w)$ over this distribution

$$\kappa(\theta) = \mathbb{E}_{z \sim p(z), z=z(w)} [\kappa(\theta, w)] \quad (21)$$

Conceptually, this replaces the upper limit in the line-of-sight projection with an expectation over all comoving distances consistent with the photometric redshift estimate. This leads to the standard expression in which the convergence is written as an integral over comoving distance

$$\kappa(\theta) = \frac{3H_0^2 \Omega_m}{2c^2} \int_0^w dw \frac{\delta(w\theta, w)}{a(w)} g(w) w \quad (22)$$

Here, $g(w)$ denotes the lensing efficiency, which incorporates the probability that the galaxy lies behind a given lens plane. It is defined by

$$g(w) = \int_w^{w_k} dw' \frac{w' - w}{w'} \left(p(z) \frac{dz}{dw} \right)_{z=z(w')}. \quad (23)$$

Here w_k is some reasonable upper limit on the source comoving distance. In practice, the efficiency function $g(w)$

can be precomputed for each galaxy at any set of lens-plane positions. To evaluate the required derivatives, such as dz/dw and the mapping between redshift and comoving distance, we use the cosmology package `jax-cosmo`, which provides differentiable cosmological functions.

Throughout this work, we model the photometric-redshift distribution $p(z)$ as a Gaussian centered on the measured redshift,

$$p(z) \sim \mathcal{N}(z, \sigma_z), \sigma_z = 0.11(1 + z) \quad (24)$$

with the scatter σ_z estimated from the JWST COSMOS-Web shear catalog [36, 38]. This Gaussian approximation is standard and has been adopted in several previous weak-lensing studies [26, 39].

13. Conicus3D - Dataset Construction

Conicus3D is a dataset of 3D lightcones, binned into lensplanes, derived from the AbacusSummit simulations [15, 27]. We use the 25 base boxes with the AbacusSummit fiducial cosmology (AbacusSummit_base_c000_ph{000-024}), as well as one base box from the cosmology corresponding to massless neutrinos (AbacusSummit_base_c009_ph000). Each box contains two lightcones spanning approximately $28 \times 28 \text{ deg}^2$, from which we crop smaller patches to train the diffusion model on a chosen survey footprint.

To construct the overdensity lensplanes, we follow the methodology in [16] to compute the overdensity maps on spherical shells. Given a particle shell map with particle count N_{ij} at pixel index ij , we compute the overdensity as:

$$\delta(ij) = \frac{\rho(ij)}{\bar{\rho}} - 1 \quad (25)$$

where

$$\bar{\rho} = (N_{\text{part}}/L_{\text{box}}^3) \Delta \Omega \chi_j^2 d\chi_j \quad (26)$$

is the mean density. Here $N_{\text{part}} = 6912^3$ is the number of particles in a simulation box, $L_{\text{box}} = 2 \text{ Gpc}/h$ is the simulation box length, $\Delta \Omega$ is the solid-angle pixel area, and χ_j and $d\chi_j$ are the shell center and width, respectively. The density per pixel is therefore

$$\rho(ij) = \frac{N_{ij}}{dV_j} = \frac{N_{ij}}{\Delta \Omega \chi_j^2 d\chi_j}. \quad (27)$$

After computing an overdensity map on each shell, we combine the shells into 20 lensplanes using a weighted average. These planes span comoving distances from approximately 290 to 3580 Mpc/h. Each lensplane is assigned an effective scale factor and comoving distance; values for a subset of these planes are provided in Table 3.

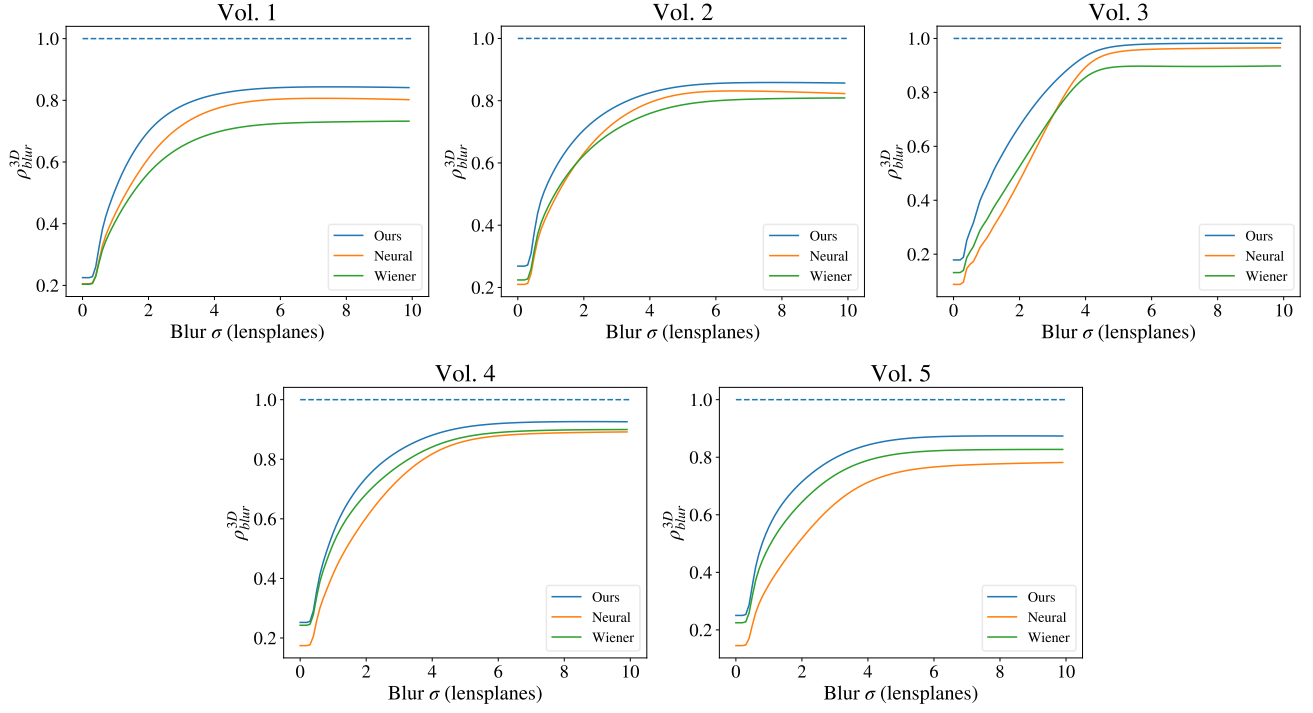


Figure 11. Resolution Analysis across radial blur scales. We evaluate the blurred 3D cross-correlation $\rho_{3D}^{\text{blur}}(\sigma)$ between each reconstruction and the ground truth as a function of the Gaussian blur width σ (in lensplanes). Larger σ corresponds to more aggressive smoothing along the line of sight. Across all five simulated volumes, our method achieves higher correlation than both baselines for every blur scale, demonstrating improved depth resolution at both fine and coarse radial smoothing regimes. The consistently superior performance indicates that our diffusion-based prior enhances radial structure recovery independently of the chosen blur width.

Plane Index	Redshift z	Scale Factor a	Center (Mpc/h)	Width (Mpc/h)
0	0.13	0.88	372.25	164.5
2	0.25	0.80	701.25	164.5
6	0.52	0.66	1359.25	164.5
12	1.04	0.49	2346.25	164.5
19	1.92	0.34	3497.75	164.5

Table 3. Geometry of selected Conicus3D lensplanes. For reference, we list the redshift, scale factor, comoving distance, and physical width for several representative lensplanes used in the Conicus3D dataset. These planes are derived from overdensity shells computed following the methodology of [16] and subsequently binned into 20 evenly spaced lensplanes in comoving distance. The planes span the full radial range of the lightcone (approximately 290–3580 Mpc/h, corresponding to $z \approx 2$), and their geometry is used both in constructing the dataset and in defining the line-of-sight distances relevant for weak-lensing projections. This subset illustrates the non-uniform evolution of redshift and scale factor across the lightcone and provides context for the depth-dependent analysis presented in the supplement.

Algorithm 1 Modified Decoupled Annealing Posterior Sampling

Require: Observation γ , diffusion prior $p(\delta)$, weak-lensing forward model, annealing schedule $\{\sigma_T, \dots, \sigma_0\}$, power spectrum $P(k)$

```
1: Initialize  $\delta_T \sim \mathcal{N}(0, \sigma_T^2 \mathbf{I})$ 
2: for  $t = T, T-1, \dots, 0$  do
3:   Prior projection via diffusion model:
4:    $\hat{\delta}_0 \leftarrow \text{REVERSEDIFFUSION}(\delta_t, \sigma_t)$ 
5:   Likelihood-guided refinement (Langevin dynamics):
6:   Define prior covariance  $\Sigma$  using Fourier diagonal  $P(k)$ 
7:   Sample  $\delta_{0|\gamma}$  with Langevin updates targeting  $\log p(\gamma | \delta) + \log \mathcal{N}(\delta; \hat{\delta}_0, \Sigma)$ 
8:   if  $t > 0$  then
9:     Draw  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ 
10:     $\delta_{t-1} \leftarrow \delta_{0|\gamma} + \sigma_{t-1} \epsilon$ 
11:   end if
12: end for
13: return samples  $\{\delta_0\}$  approximating  $p(\delta | \gamma)$ 
```

14. Decoupled Annealing Posterior Sampling

To draw samples from the posterior distribution $p(\delta | \gamma)$, we use a modified version of Decoupled Annealing Posterior Sampling (DAPS) [45]. DAPS constructs a sequence of noise-annealed posteriors $p(\delta_t | \gamma)$ indexed by a decreasing noise scale σ_t . Sampling proceeds from a high-noise initial distribution toward the target posterior, alternating between (1) projecting the current sample toward the learned prior using the diffusion model, and (2) refining it using likelihood guidance from the weak-lensing forward model.

At each stage, the diffusion model provides a clean estimate $\hat{\delta}_0$ consistent with the data-driven prior. This estimate is then refined using Langevin dynamics that incorporates both the WL likelihood and a Gaussian prior term. Unlike the standard DAPS formulation, we replace the diagonal covariance of this Gaussian prior with a Fourier-space covariance derived from the estimated matter power spectrum. This modification enforces statistical isotropy and translation invariance and helps ensure that radial and angular correlations of the overdensity field are preserved throughout sampling.

After refinement, the sample is re-noised according to the next noise level σ_{t-1} , enabling a coherent annealing schedule. Iterating this process yields approximate samples from $p(\delta | \gamma)$, which accurately capture both small-scale structure and the full posterior variability needed for downstream cosmological analyses.

15. Broader Applicability of the Spectral Covariance Formulation

The prior covariance in our formulation (Eq. 7) is diagonalized in Fourier space by the power spectrum, a structure that arises naturally whenever the field of interest is ap-

proximately translation-invariant. This property is common across scientific inverse problems.

In adaptive optics, the Kolmogorov power spectrum characterizes atmospheric wavefront phase covariance, and phase screens are standardly generated by spectral filtering in Fourier space [24]. In geophysical inversion, subsurface properties such as seismic velocity and porosity are modeled as stationary fields with spectrally defined covariance for posterior sampling in tomography and reservoir characterization [18]. In CMB analysis, the angular power spectrum C_ℓ defines a harmonic-diagonal signal covariance for Wiener filtering and constrained realizations of the microwave sky [13].

In each case, translation invariance permits efficient spectral representation of the prior. We believe our spectral sampling strategy could be adapted to improve diffusion-based posterior sampling in these and similar inverse problems, as properly modeling the covariance structure allows the sampler to more faithfully represent the learned prior distribution during inference.