

Mapping Dark-Matter Clusters via Physics-Guided Diffusion Models

Diego Royo¹, Brandon Zhao², Adolfo Muñoz¹, Diego Gutierrez¹, and Katherine L. Bouman²

¹ Universidad de Zaragoza–I3A, Spain

² California Institute of Technology, USA

Abstract. Galaxy clusters are powerful probes of astrophysics and cosmology through gravitational lensing: their mass, dominated by 85% dark matter, distorts background light. Yet, mass reconstruction lacks the scalability and large-scale benchmarks to process the hundreds of thousands of clusters expected from forthcoming wide-field surveys. We introduce a fully automated method to reconstruct cluster surface mass density from photometry and gravitational lensing observables. Central to our approach is DARKCLUSTERS-15K, our new dataset of 15,000 mock cluster observations with paired mass and photometry maps, the largest to date, spanning multiple redshifts and simulation frameworks. We train a plug-and-play diffusion prior on DARKCLUSTERS-15K that learns the statistical relationship between mass and light, and draw approximate posterior samples constrained by weak- and strong-lensing observables, yielding principled reconstructions with well-calibrated empirical uncertainties. Our approach requires no expert tuning, runs in minutes rather than hours, achieves higher accuracy, and matches expert-tuned reconstructions of the MACS 1206 cluster. We release our method and DARKCLUSTERS-15K to support upcoming wide-field cosmological surveys.

Keywords: Gravitational lensing · Simulation-based priors · Inverse problem

1 Introduction

Studying galaxy clusters provides meaningful insights into the fundamental nature and formation of our Universe; their abundance and mass profiles place vital constraints on astrophysical and cosmological models. Because dark matter constitutes roughly 85% of their mass, accurately mapping their total mass distribution is essential to connect this invisible dark matter with observables such as galaxy luminosities, and to enable techniques like time-delay cosmography to measure the Universe’s expansion rate.

One of the most powerful ways to probe the mass of galaxy clusters is through *gravitational lensing*—the bending of light from background sources as it passes through the cluster’s curved space-time. We target massive clusters, exceeding $10^{13.5}$ times the mass of our Sun $M_{\odot}$. In these systems, the dense cores give rise to strong gravitational lensing that can produce multiple images of the same object, while the outer cluster regions produce weak lensing that only subtly distorts shapes. The rarest, most massive clusters of $10^{14-15} M_{\odot}$ provide abundant lensing constraints, hence most mass mapping studies focus on these systems [1, 14, 38, 40]. In contrast, clusters of $10^{13.5-14} M_{\odot}$ are ten to hundreds of times more common but yield far fewer lensing observables. Mass mapping analysis of these smaller systems is thus highly ill-posed, and often relies on

Table 1: Overview of galaxy cluster mass reconstruction methods. We differentiate between tools that require expert configuration (e.g., Lenstool [31], WSLAP+ [15]) and fully automatic approaches (MARS [10], the method from work by Napier et al. [40], and ours). Our free-form approach leverages a diffusion prior trained on DARKCLUSTERS-15K, enabling 512×512 free parameters with mass reconstructions within minutes, significantly faster than traditional methods. Unlike other methods that assume that Light Traces Mass (LTM) with a 1:1 ratio, we freely model photometry along with Strong Lensing (SL) and Weak Lensing (WL). * *Our reimplementation.*

Method name	Model type	# Param.	Inputs			Automation level	End-to-end time
			Photometry	SL	WL		
Lenstool [31]	Parametric	$\sim 10\text{--}100$	✓ (LTM 1:1)	✓	✓	Expert-guided	Hours [38]
WSLAP+ [15]	Hybrid	~ 1000	✓	✓	✓	Expert-guided	Hours
MARS [10]	Free-form	100×100	✗	✓	✓	Fully automatic	Hours [10]
Napier et al. [40]*	Parametric	$\sim 10\text{--}100$	✓ (LTM 1:1)	✗	✗	Fully automatic	Minutes
Ours	Free-form	512×512	✓	✓	✓	Fully automatic	Minutes

simplified parametric models which depend heavily on expert tuning [16], or on stacking many systems to recover only ensemble-average properties [23].

We present two main contributions. First, we release DARKCLUSTERS-15K, our dataset of 15,000 mock cluster observations derived from IllustrisTNG [45] and SIMBA [12] with paired surface mass density and photometry information. It is the largest of its kind, enabling rigorous benchmarking of mass mapping methods. Second, we present our *fully automated* method to estimate the non-parametric mass distribution of galaxy clusters. Since lensing information alone is insufficient to fully recover the underlying mass distribution, our method integrates a diffusion-based generative prior with gravitational lensing physics operators that ensure consistency with observations. Unlike other approaches that impose a rigid light-traces-mass (LTM) assumption, our prior captures more flexible and expressive dependencies by modeling the joint distribution of mass and light. By factorizing the prior and data likelihood, our approach establishes a physics-based principled framework. We formulate reconstruction as a reverse diffusion process based on posterior sampling, allowing us to quantify per-pixel uncertainties.

Compared to other approaches (Table 1), our method presents many advantages: it is non-parametric, can handle complex cluster structures where the mass-to-light ratio does not follow a strict 1:1 relationship, does not require cluster-specific expert tuning, and provides results in minutes instead of hours [9, 40]. Moreover, we obtain better quantitative and qualitative results on held-out DARKCLUSTERS-15K clusters from different redshifts and simulation frameworks. We also reconstruct the MACS 1206 cluster, qualitatively matching expert-level reconstructions.

Our method is timely: upcoming sky surveys with improved telescope optics are expected to discover hundreds of thousands of clusters above $10^{13.5} M_{\odot}$ [41, 49, 55]³, with higher quality lensing data [22]. Automated, scalable methods like ours will be essential to analyze these systems without expert intervention, enabling statistically robust tests of cosmology models. By relaxing rigid mass-to-light assumptions, our approach could even automate the detection of unobservable dark matter. We release our

³For instance, the Euclid space telescope recently discovered 26 million galaxies in just one week.

DARKCLUSTERS-15K dataset as a benchmark for future research, with code and further details in our project page⁴.

2 Related work

Galaxy cluster mass reconstruction. Galaxy cluster mass is typically constrained using both *strong* and *weak* gravitational lensing. Strong lensing effectively probes the dense inner regions of the cluster, which can be represented reasonably well with parametric halo models [7, 25, 30] like NFW [42] or PIEMD [28]. Weak lensing provides wider spatial coverage but suffers from severe noise. The seminal work of Kaiser and Squires [26] introduced a direct inversion method, which has since been improved via sparsity-based regularization [33, 34], inverse-variance filtering [51], or supervised deep learning [24, 47] to better handle the severe noise. A complementary approach relies on the light-traces-mass assumption [36], wherein parametric models scale mass according to observed luminosity [7, 28].

Modeling tools like Lenstool [31, 44] and WSLAP+ [15] incorporate additional observables like photometry, X-ray or the Sunyaev-Zel’dovich effect. While these tools achieve state-of-the-art performance [14, 40], both require expert guidance to define the number and type of parametric models for each cluster. We refer the reader to the comparison by Meneghetti et al. [38] for more details. Alternatively, MARS [9, 10] is a free-form method combining strong lensing with a simple maximum-entropy regularizer.

By replacing rigid parametric assumptions and simple regularizations with a diffusion prior trained on DARKCLUSTERS-15K, we capture an expressive mass-light joint distribution, constrained by weak and strong gravitational lensing, yielding non-parametric posterior estimates an order of magnitude faster than traditional optimizations.

Mass reconstruction using machine learning. With recent and upcoming surveys, both the quantity and quality of galaxy and galaxy-cluster observations are set to increase dramatically [41, 55]. This rapid growth has positioned machine learning (ML) as a powerful tool for automated lens analysis.

Early ML work focused primarily on strong-lensing systems produced by individual galaxies typically only a few arcseconds across, employing convolutional neural networks trained on simulations to estimate model parameters [20] and their uncertainties [35]. Subsequent efforts have combined parametric lens models with more flexible, free-form components using image-space [2], wavelet-based [18], or neural implicit [5] representations, which has also been used to reconstruct unlensed background sources [39]. Joint lens-and-source reconstruction has additionally been explored through variational inference [48] and diffusion models [3].

Conversely, ML for weak lensing operates over much larger angular scales, leveraging constraints from vast galaxy populations spanning several degrees⁵ but at much coarser resolution. In this regime, recent works utilize shear maps for diffusion-based posterior sampling [6] or neural implicit representations for 3D mass inference [57, 58].

⁴<https://graphics.unizar.es/projects/DarkMatterMapping/>

⁵1 degree is equivalent to 60 arcminutes and 3,600 arcseconds.

Bridging these extremes, our method targets galaxy clusters, which typically span a few arcminutes, while maintaining sub-arcsecond resolution. Furthermore, our work is the first to combine a data-driven score prior for galaxy clusters with weak and strong lensing and photometry measurements at this scale. Thanks to our DARKCLUSTERS-15K dataset, our model can learn fine structure in the joint mass/light distribution.

Cosmological simulations and datasets. Cosmological simulations track the structural evolution of the Universe over cosmic time, capturing the interplay between baryonic (visible) and dark matter components at different redshifts. Each varies in volume, resolution, and cosmology assumptions. Projects like TNG-Cluster [43] or The Three Hundred project [13] isolate and simulate hundreds of rare, extremely massive clusters with virial masses $\geq 10^{14} M_\odot$ at high resolution. Conversely, projects like MDPL2 [29], Magneticum Pathfinder [17] or BAHAMAS [37] simulate larger regions, which include smaller clusters, but at coarser resolutions.

Our DARKCLUSTERS-15K dataset is derived from the TNG300-1 dataset of the IllustrisTNG simulation suite [45], and the SIMBA simulations [12], which include many clusters in the $10^{13.5-14} M_\odot$ range, enabling us to extract 15,000 pairs of surface mass density and photometry maps from galaxy clusters at sub-arcsecond resolution. Compared to other benchmarks such as the two hand-crafted clusters from the Frontier Fields project [38], or the Camels Multifield Dataset [54] which targets larger volumes at coarser resolution, DARKCLUSTERS-15K enables large-scale benchmarking, being the largest dataset of its kind to date while also including our simulated gravitational lensing observations. Our work anticipates upcoming wide-field surveys that will discover many clusters in that mass range and their need for scalable reconstruction methods.

3 Background

3.1 Gravitational lensing

Here we outline the gravitational lensing formalism and refer the reader to Schneider et al. [50] for details. Figure 1 illustrates an example, with a mass concentration at a distance D_L and a light source at distance D_S from the observer. Given the vast distances involved, including the distance D_{LS} between the mass and the source, we model them as *lens* L and *source* S planes perpendicular to the optical axis. We will use θ and β for 2D *angular positions* of objects in L and S with respect to the optical axis, and $\xi \approx \theta D_L$ and $\eta \approx \beta D_S$ for planar coordinates.

Light from the source at η follows a curved path through the distorted space-time near the mass, thus we observe it at ξ instead. We simplify this continuous deflection into two straight rays bending by an angle $\tilde{\alpha}$ at the lens plane L , and establish from Figure 1:

$$\eta = \frac{D_S}{D_L} \xi - D_{LS} \tilde{\alpha}(\theta) \implies \beta = \theta - \alpha(\theta), \quad (1)$$

which simplifies to the expression on the right using angular positions and the *reduced* deflection angle $\alpha(\theta) \equiv \tilde{\alpha}(\theta) D_{LS}/D_S$, and which is used in all our other derivations.




Fig. 1: Sketch of a typical gravitational lens system. The light path of an object at angular position β on a source plane S is deflected by an angle $\tilde{\alpha}$ due to the presence of a galaxy cluster on the lens plane L , so the object appears at angular position θ instead.

The deflection $\alpha(\theta)$ depends on the lens's surface mass density, denoted as $\Sigma(\theta)$ and typically expressed via the dimensionless *convergence* $\kappa(\theta)$:

$$\kappa(\theta) = \frac{\Sigma(\theta)}{\Sigma_{cr}}, \quad \Sigma_{cr}(D_L, D_S) = \frac{c^2}{4\pi G} \frac{D_S}{D_L D_{LS}}, \quad (2)$$

where c and G are the speed of light and the gravitational constant, respectively, and $\Sigma_{cr}(D_L, D_S)$ is the critical surface mass density. Regions where $\Sigma \geq \Sigma_{cr}$ (i.e., $\kappa \geq 1$) can produce multiple images, defining the *strong* lensing regime. Conversely, the *weak* lensing regime measures small background distortions in outer regions where $\kappa \ll 1$.

Strong lensing regime. We use Equation 1, obtaining the reduced deflection angle $\alpha(\theta)$ from the convergence $\kappa(\theta)$ as a convolution over all angular positions θ' in the lens plane, computed efficiently in the Fourier domain:

$$\alpha(\theta) = \frac{1}{\pi} \int_{\mathbb{R}^2} d^2\theta' \kappa(\theta') \frac{\theta - \theta'}{|\theta - \theta'|^2}. \quad (3)$$

Weak lensing regime. Here, deflections caused by the lens plane are small enough to linearize $\alpha(\theta)$ in Equation 1 via its Jacobian matrix $\mathcal{A}$:

$$\beta \approx \mathcal{A}(\theta)\theta, \quad \mathcal{A}(\theta) = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix}, \quad (4)$$

where $\gamma(\theta) = \gamma_1(\theta) + i\gamma_2(\theta)$ is the complex-valued lens shear. Weak lensing assumes that all sources (i.e., background galaxies) have an elliptical morphology, thus Equation 4 maps their *intrinsic* ellipticity (in the source plane) to their *observed* ellipticity (in the lens plane) by altering their axis ratio and orientation. Because intrinsic shapes are unknown, deriving shear from observed ellipticities is highly noisy, with a signal-to-noise ratio (SNR) of ~ 0.01 - 0.1 . In our work, we compute shear from convergence as:

$$\gamma(\theta) = \frac{1}{\pi} \int_{\mathbb{R}^2} d^2\theta' \kappa(\theta') \mathcal{D}(\theta - \theta'), \quad \mathcal{D}(\theta) = \frac{\theta_2^2 - \theta_1^2 - 2i\theta_1\theta_2}{|\theta|^4}. \quad (5)$$

also using a convolution as in Equation 3, only now the kernel $\mathcal{D}$ uses the horizontal θ_1 and vertical θ_2 components of the angular position θ .

3.2 Diffusion models

Diffusion models provide a flexible framework for generative modeling by learning a score-based prior that captures the structure of complex data distributions. This generative process works by reversing the forward diffusion process which iteratively perturbs a data sample by adding noise [21]. Mathematically, the forward process $\{\mathbf{x}_t\}_{t=0}^T$ starts with an image $\mathbf{x}_0 \sim p_{\text{data}}$ from the data distribution p_{data} , and adds noise according to the following stochastic differential equation (SDE) [53]:

$$d\mathbf{x}_t = f(\mathbf{x}_t, t)dt + g(t)d\mathbf{w}, \quad (6)$$

where $\mathbf{w}$ represents n-dimensional Brownian motion. The drift and diffusion coefficients $f(\mathbf{x}_t, t)$ and $g(t)$ parameterize the noise schedule.

The reverse of Equation 6, which conversely transforms noisy images back to meaningful images in p_{data} , is another diffusion process given by the reverse-time SDE:

$$d\mathbf{x}_t = [f(\mathbf{x}_t, t) - g(t)^2 \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)] dt + g(t)d\tilde{\mathbf{w}}, \quad (7)$$

where $\tilde{\mathbf{w}}$ represents n-dimensional Brownian motion with reversed time, and dt is an infinitesimal negative timestep. Diffusion models learn the score function $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$, after which Equation 7 generates new samples in p_{data} starting from Gaussian noise. We later build on posterior sampling methods [11, 56] that augment this prior with a likelihood term that steers the reverse process towards solutions consistent with the data.

4 DARKCLUSTERS-15K dataset

The recent availability of large, high-fidelity cosmological simulations [43] has created a timely opportunity to build realistic datasets for cluster-scale lensing studies. Motivated by this development, we construct a new dataset that processes these large simulations to systematically isolate cluster environments and generate corresponding realistic lensing measurements. We additionally release the code used to create these simulations, enabling reproducibility and further expansion by the community.

4.1 Dataset generation

Our DARKCLUSTERS-15K dataset contains 15,000 projected galaxy cluster samples (split in 12,000 for training, 3,000 for test), derived from the IllustrisTNG [45] and SIMBA [12] cosmological simulations. It is the largest dataset of its kind. Each sample includes four 512×512 images: one surface mass density map $\Sigma(\theta)$ and three simulated Hubble Space Telescope (HST) images $\mathbf{P}_f(\theta)$ corresponding to wavelength filters $f \in \{125, 606, 814\}$, representing near-infrared (F125W), and optical (F606W, F814W) bands. All images are centered on the cluster and are provided in two field-of-view settings: 100×100 and 225×225 arcseconds. The total gravitationally-bound mass for each cluster in DARKCLUSTERS-15K is above $10^{13.5} M_{\odot}$, with most of them being below $10^{14} M_{\odot}$. See the Supplemental for the cluster distribution, units, and processing.




Fig. 2: Left: For the creation of our DARKCLUSTERS-15K dataset, we extract 600 simulated galaxy clusters from IllustrisTNG and SIMBA as 3D point clouds, and compute the surface mass density Σ for $n = 25$ different line of sight r_i projections. **Right:** For each projection, we then independently generate measurements for strong lensing (i.e., positions of multiple images for each source) and weak lensing (i.e., shear γ at the source positions).

Data extraction and train/test splits. We extract 600 distinct 3D cluster simulations as point clouds, 540 from the TNG300-1 subset of IllustrisTNG and 60 from SIMBA. To augment the data, we project each cluster along 25 different viewing directions r_i (Figure 2, left), which serve as lens planes L . Our training/validation set is formed exclusively from 480 TNG clusters (i.e., 12,000 projections) fixed at a distance D_L (Figure 1) corresponding to a redshift of $z_L = 0.5$. Our test set uses 120 clusters with two variations. First, we include 60 TNG clusters which are the same 20 clusters at three redshifts $z_L \in \{0.2, 0.5, 1.0\}$. This yields three TNG test sets of 500 projections each. Clusters at $z_L = 0.2$ are older, more massive clusters with complex merged structures, while clusters at $z_L = 1.0$ are younger and less evolved. Second, we use 60 clusters (1,500 projections) from SIMBA. Since SIMBA relies on different algorithms than TNG, this test set verifies our model’s performance across different simulation frameworks.

Generating strong and weak lensing observations. For every test dataset sample, we simulate both strong- and weak-lensing observations (Figure 2, right). These observations are used only at inference time and are not part of the diffusion-model training set. For strong lensing, we uniformly sample source angular positions β on the source plane S and solve Equation 1 with Equation 3 to identify configurations in which a single background galaxy is lensed into multiple, distinct image angular positions θ in the lens plane L . Depending on the cluster total mass, we simulate between 5 and 20 sources whose background galaxies are lensed into multiple distinct images, typical for clusters of $10^{13.5} M_\odot$ and above. For weak lensing, we compute the shear $\gamma(\theta)$ at each angular position θ using Equation 5. We adopt realistic source densities of 30 weakly lensed galaxies per arcmin² [33] and add Gaussian noise following prior work [22, 33] with variance $\sigma_w^2 \in \{0.03, 0.3\}$. The larger value corresponds to traditional weak lensing observations, where noise comes from the uncertainty in the galaxy’s intrinsic shape, while the smaller value corresponds to kinematic observations, which reduces this shape noise by including velocity information. Since γ is small, these correspond roughly to SNRs of 0.01 and 0.1, respectively. For both regimes, we distribute sources throughout redshifts z_S following a probability distribution $p(z_S) \propto \exp(-z_S/z_0) z_S^2 / (2z_0^3)$ [33]. We use $z_0 = 2/3$ which yields a median redshift of 1.75, and truncate at $z_S = 5$.

5 Method

Our method reconstructs surface mass density maps of individual galaxy clusters using a physics-guided diffusion framework. We train a diffusion prior on DARKCLUSTERS-15K, which provides realistic cluster mass and photometry distributions and enables the model to capture both their statistical structure and their mutual correlations (Section 5.1). To incorporate physical observations, we pair this prior with forward models for strong and weak gravitational lensing (Section 5.2), yielding a hybrid approach that automatically produces high-fidelity reconstructions within minutes while simultaneously quantifying uncertainty.

To perform inference, we adapt the Decoupled Annealing Posterior Sampling (DAPS) algorithm [56], as implemented in the InverseBench framework [59]. Our goal is to estimate the surface mass density Σ (Section 3.1), represented by the diffusion variable $\mathbf{x}_t$ at timestep t (Section 3.2). DAPS enables efficient posterior sampling by combining a data-driven prior with the likelihood given by physics-based lensing models, ultimately yielding mass maps that are both physically plausible and tightly constrained by the observed lensing data.

To generate these samples, DAPS alternates between two different updates. Consider the marginal update from time $t + \Delta t$ to time t , which will update the value of $\mathbf{x}_t$. In the first step, DAPS samples $\tilde{\mathbf{x}}_{0|\mathbf{y}}$:

$$\tilde{\mathbf{x}}_{0|\mathbf{y}} \sim p(\mathbf{x}_0 | \mathbf{x}_{t+\Delta t}, \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x}_0)p(\mathbf{x}_0 | \mathbf{x}_{t+\Delta t}), \quad (8)$$

from a conditional posterior that combines the likelihood term $p(\mathbf{y} | \mathbf{x}_0)$, which enforces consistency with the observed lensing measurements, with the prior term $p(\mathbf{x}_0 | \mathbf{x}_{t+\Delta t})$ derived from the diffusion model, which encourages $\mathbf{x}_0$ to follow the learned distribution of realistic clusters. In the second step, it samples:

$$\mathbf{x}_t \sim p(\mathbf{x}_t | \tilde{\mathbf{x}}_{0|\mathbf{y}}) = \mathcal{N}(\tilde{\mathbf{x}}_{0|\mathbf{y}}, \sigma_t^2 \mathbf{I}), \quad (9)$$

injecting the Gaussian noise with variance σ_t^2 . This two-step procedure is repeated for progressively smaller values of t until $\mathbf{x}_0$ approximates a sample from the target posterior distribution.

We now describe how observational constraints enter this posterior sampling. The weak- and strong-lensing observations are linked to the underlying mass distribution through well-defined physical forward models, so they naturally appear in the likelihood term $p(\mathbf{y} | \mathbf{x}_0)$. In contrast, the photometric data $\mathbf{P}_f$ does not admit a simple analytic relationship to the mass field. To incorporate this information, we augment the likelihood-driven constraints with a diffusion-model prior that is conditioned on photometry, allowing the model to learn the statistical relationship between light and mass directly from training data rather than imposing an explicit parametric form.

5.1 Photometry-conditioned diffusion prior

We train a DDPM++ network [53] to learn the diffusion prior used in DAPS, modeling realistic surface-mass-density distributions from DARKCLUSTERS-15K. The network learns to predict the (unconditional) score $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$ in Equation 7, at timestep t in

the generation. Importantly, although DAPS traditionally uses an unconditional diffusion prior, we condition the score network on the corresponding photometric images $\mathbf{P}_f$, which are typically highly correlated with the mass maps. This conditioning enables the model to learn $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{P}_f)$, injecting photometric information that cannot appear in the likelihood and guiding generation toward mass distributions consistent with the light. We train the DDPM++ network using the EDM loss from Karras et al. [27]. The network receives a four-channel image of mass and photometry where all channels are normalized to the range $[0, 1]$. For additional details about training parameters, see the Supplemental.

5.2 Lensing likelihood

In order to introduce weak and strong gravitational lensing observations into our generation pipeline, we include multiple physical *forward models* that we use to compute the likelihood $p(\mathbf{y} | \mathbf{x}_0)$ in Equation 8 for observations $\mathbf{y}$ given our mass density estimate $\mathbf{x}_0$ obtained using our diffusion prior. Specifically, we use the gradient of the log-likelihood $\nabla_{\mathbf{x}} \log p(\mathbf{y} | \mathbf{x}_0)$ to update our estimate. Rather than modeling the exact likelihood, we formulate this term as proportional to a sum of k independent loss terms $\mathcal{L}_k$ derived from our physical forward operators, providing a practical way to integrate gravitational lensing data into our pipeline. Each term is weighted by λ_k :

$$-\log p(\mathbf{y} | \mathbf{x}_0) \propto \sum_k \lambda_k \mathcal{L}_k, \quad (10)$$

In the rest of this section, we will use a hat, such as in $\hat{\mathbf{y}}_k$, to refer to our estimations, and no hat, such as in $\mathbf{y}_k$, to refer to the observed data.

Strong lensing forward model. In strong lensing, light from a background galaxy source on S at angular position β is deflected at multiple angular positions θ on the lens plane L . The true angular position β is unknown, but the lens equation (Equation 1) must hold. We compute our estimation $\hat{\beta}_i$ of the angular position of the i -th source on S :

$$\hat{\beta}_i = \frac{1}{m_i} \sum_{j=1}^{m_i} \hat{\beta}_{i,j}, \quad \text{where } \hat{\beta}_{i,j} = \theta_{i,j} - \hat{\alpha}(\theta_{i,j}), \quad (11)$$

as the average of individual estimations $\beta_{i,j}$ from all m_i images of the source. We compute $\hat{\alpha}$ through Equations 2 and 3 using the current estimation of the surface mass density $\hat{\Sigma}$ and the known source distance D_S . We formulate our *geometric* loss as the average squared error of all n sources:

$$\mathcal{L}_s^{\text{geo}} = \frac{1}{n} \sum_{i=1}^n \frac{1}{m_i} \sum_{j=1}^{m_i} \left| \hat{\beta}_{i,j} - \hat{\beta}_i \right|^2. \quad (12)$$

To prevent all images of all sources from collapsing to one unique point, we found it effective to incorporate photometry, comparing the observed image $\mathbf{P}_f(\theta)$ to our estimation $\hat{\mathbf{P}}_f(\theta; \hat{\beta})$ of the light at θ from a light source placed at $\hat{\beta}$. Concretely, we

model the light emission by this source using a parametric ellipsoidal (Sérsic) source [52] at angular position $\hat{\beta}_i$. This allows us to introduce a *photometric* loss:

$$\mathcal{L}_s^{\text{img}} = \frac{1}{n} \sum_{i=1}^n \frac{1}{m_i} \sum_{j=1}^{m_i} \left(\mathbf{P}_f(\boldsymbol{\theta}_{i,j}) - \hat{\mathbf{P}}_f(\boldsymbol{\theta}_{i,j}; \hat{\beta}_i) \right)^2. \quad (13)$$

Weak lensing forward model. We can compute our estimated shear $\hat{\gamma}$ through Equations 2 and 5. Since weak lensing happens in the $\kappa \ll 1$ regime, we approximate reduced shear $g = \gamma/(1 - \kappa) \approx \gamma$ in our computations. However, naive comparison between $\hat{\gamma}$ and γ for all angular positions $\boldsymbol{\theta}$ is not possible for two reasons: (i) these measurements come from sources at different distances D_S , thus from Equation 2 we would need to compute κ and later γ for each source, and (ii) measurement is only possible at sparse angular positions $\boldsymbol{\theta}$ with a source at S .

To solve these issues, we can compensate for the difference in distances D_S by scaling all observations to an arbitrary reference distance D_R by multiplying and dividing by the corresponding critical densities Σ_{cr} . Finally, we extrapolate sparse shear observations to the full field of view using a radial basis function (RBF) interpolator $\Phi(\gamma, \boldsymbol{\theta})$, yielding a dense shear map that covers the lens plane L from sparse values γ . Rather than an exact per-pixel noise model, this yields a smoothed data-fidelity term. This step mimics standard weak-lensing smoothing methods [26, 33], which suppress high-frequency noise while preserving the useful, large-scale part of the signal. It also enables computing data-fidelity gradients across the entire image rather than only at the observed points. We combine all these steps in our weak lensing loss:

$$\mathcal{L}_w = \frac{1}{\sigma_w^2} \sum_{\boldsymbol{\theta}} \left| \hat{\gamma}(\boldsymbol{\theta}) - \Phi \left(\gamma \frac{\Sigma_{\text{cr}}(D_L, D_S)}{\Sigma_{\text{cr}}(D_L, D_R)}, \boldsymbol{\theta} \right) \right|^2. \quad (14)$$

Shear measurements are very noisy; as per Section 4.1, we use values $\sigma_w^2 = 0.3$ and $\sigma_w^2 = 0.03$ depending on the observation regime.

6 Results and evaluation

Here we analyze the performance of our physically guided diffusion method for galaxy-cluster mass estimation. In Section 6.1 we show results of our posterior sampling and verify that our model’s predicted uncertainty is well calibrated. In Section 6.2, we leverage our principled framework to conduct an ablation study, evaluating our method’s performance using different combinations of inputs. Then, to provide fair and meaningful baselines, in Section 6.3 we compare against three automated approaches: Kaiser-Squires inversion [26], which relies solely on weak-lensing constraints; our implementation of the approach from Napier et al. [40], which operates directly on photometric images $\mathbf{P}_f$ with comparable runtime; and our UNet model baseline, which maps photometry $\mathbf{P}_f$ directly to surface mass density. In contrast, many of the methods listed in Table 1 either require substantial, cluster-specific expert tuning (e.g., Lenstool or WSLAP+) or lack publicly available implementations (e.g., MARS [10]). Finally, in Section 6.4 we apply our method to reconstruct the surface mass distribution of the real MACS 1206 cluster and compare it to an existing manual, expert-tuned reconstruction.




Fig. 3: Left: Our diffusion-based approach produces multiple plausible reconstructions; in each row we show one representative sample of the lens convergence κ , and the posterior mean and standard deviation over 20 samples. **Right:** Reliability diagram, showing close correlation between our model’s predicted uncertainty and the actual error.




Fig. 4: Ablation study on the different input components used by our method to estimate cluster masses. **Left:** Each row shows photometry $\mathbf{P}_f$ for one cluster, and mass reconstructions using only one input as stated. **Right:** Adding strong (SL) and weak lensing (WL) information (SNR 0.01) improves the photometry (P) results. We report PSNR of each image w.r.t. the ground truth (GT).

6.1 Posterior sampling and uncertainty calibration

Our model can sample a range of plausible mass density maps. In Figure 3, left, each row has one example for which we show one representative posterior sample, and the mean and standard deviation over 20 samples. The variance across these samples captures the range of plausible mass configurations consistent with the measurements and diffusion prior, with higher variance indicating regions where the data weakly constrain the solution. We verify that our predicted uncertainty closely tracks the actual reconstruction error. Figure 3, right, shows a reliability diagram using the procedure from work by Kuleshov et al. [32], which compares observed against expected confidence levels on a held-out test set. The agreement between the two indicates our marginal coverage is largely accurate, and suggests that our model’s uncertainty is well calibrated.

6.2 Ablation: effect of photometry and strong/weak lensing

In Figure 4, we examine how different inputs to our pipeline contribute to the final reconstruction. Referring to Equation 8, we disentangle the influence of the learned prior

$p(\mathbf{x}_t)$ —which captures the joint distribution of mass and photometry—from that of the likelihood terms $p(\mathbf{y} \mid \mathbf{x}_0)$, which arise from the physical forward models for strong and weak lensing. In the leftmost column, we show an RGB representation of the photometry $\mathbf{P}_f$ of the cluster as reference, while the remaining columns show reconstructions using a different single input. While photometric information $\mathbf{P}_f$ alone provides a useful initial estimate, it cannot completely account for all dark matter as it does not interact with light. Thus, strong and weak gravitational lensing observables are necessary to fully reconstruct the cluster’s surface mass density distribution. Strong lensing, in particular, constrains the surface mass density near the cluster core. Weak lensing contributes more subtly due to noise (with an SNR of 0.01-0.1) and sparsity (~ 420 observations is typical for a 225×225 field of view), but still provides valuable information about the outer regions of the cluster. As shown in Figure 4, right, relying solely on photometry can cause the estimation to overshoot. Adding strong and weak lensing constraints successfully corrects this issue, also improving the computed PSNR by more than five points.

6.3 Baseline comparisons

We compare against three baselines. First, we compare against Kaiser-Squires inversion [26], which relies on weak-lensing signal alone. This approach performs poorly as these constraints are very sparse and noisy. For reference, we also show our strong/weak-lensing-only reconstructions in Section 6.2. Second, we implement the approach from Napier et al. [40], which employs the publicly available SE++ tool [4] to estimate the surface mass density of observed galaxy clusters. Their work proceeds in two steps. First, they use SE++ to extract the positions, shapes and orientations of elliptical light sources from a photometric image $\mathbf{P}_f$ (we use the wavelength filter $f = 125$ for this). These light sources correspond to the galaxy members of the cluster. Second, they estimate the cluster’s mass distribution by summing multiple parametric mass components, each modeled with a PIEMD profile [28] and assigned to each of the extracted galaxy members. The mass assigned to each parametric model has a 1:1 ratio with respect to light, following the relations discussed in Limousin et al. [36]. Finally, we also include our UNet baseline, which directly maps photometry $\mathbf{P}_f$ to surface mass density. Note that neither Napier et al. nor UNet account for gravitational lensing observables, so they can only infer the invisible dark matter indirectly from luminous sources. In contrast, our approach integrates gravitational lensing observables to reconstruct the complete mass distribution.

Quantitative results. Table 2 displays quantitative results comparing our method with Kaiser-Squires [26], Napier et al. [40], and our UNet baseline. We report the median PSNR, Structure Similarity Index (SSIM), Pearson Correlation Coefficient (PCC), and the deflection angle α , convergence κ , and shear γ RMSE at held-out points (i.e., points without photometry or lensing signals). We do so across 100 clusters by comparing each reconstructed mass map to its ground truth in DARKCLUSTERS-15K. Our method outperforms all baselines on practically all metrics, both in our TNG test set (which uses the same simulation framework as the train set), and our SIMBA test set (which does not), demonstrating generalization towards an entirely unseen simulation environment.

Table 2: Evaluation results for galaxy clusters at $z_L = 0.5$ in the TNG and SIMBA test sets of DARKCLUSTERS-15K. We report the Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Pearson Correlation Coefficient (PCC), and the κ , α and γ RMSE at held-out points (i.e., points without any photometry or lensing signals).

Metric	KS [26]		Napier et al. [40]		UNet		Ours	
	TNG	SIMBA	TNG	SIMBA	TNG	SIMBA	TNG	SIMBA
PSNR $\uparrow$	20.9	21.1	26.0	22.1	31.1	29.6	33.3	30.9
SSIM $\uparrow$	0.443	0.437	0.703	0.565	0.894	0.855	0.927	0.884
PCC $\uparrow$	0.189	0.197	0.832	0.616	0.924	0.863	0.975	0.930
Held-out α RMSE $\downarrow$	3.697	3.919	1.475	3.295	0.411	1.597	0.530	1.485
Held-out κ RMSE $\downarrow$	0.0552	0.0521	0.0173	0.0169	0.0067	0.0073	0.0058	0.0067
Held-out γ RMSE $\downarrow$	0.0493	0.0437	0.0188	0.0217	0.0094	0.0174	0.0084	0.0151

Qualitative results. Figure 5 displays five galaxy clusters together with their reconstructed convergence maps κ . For each cluster, our method generates multiple reconstructions that integrate the learned prior with the lensing constraints. We report one representative sample, the posterior mean over 20 samples, and the standard deviation, which characterizes the uncertainty in the reconstruction.

We evaluate generalization at different redshifts and simulation frameworks. For this, we show reconstructions on our SIMBA test set, and for the IllustrisTNG test set we also show different distances D_L of the lens plane, represented by their redshifts $z_L \in \{0.2, 0.5, 1.0\}$. Clusters at $z_L = 0.2$ are older and have had more time to develop complex structures, while clusters at $z_L = 1.0$ are younger and less evolved. Even though our method was trained only on clusters at $z_L = 0.5$ from IllustrisTNG, it generalizes well to clusters at different redshifts and simulated using different techniques.

The approach of Napier et al. [40] shows two main limitations. First, it overly smooths the mass distributions as SE++ uses one simple parametric model for each detected source, ignoring per-pixel differences in intensity. In contrast, our physically-guided diffusion pipeline captures these finer variations. Second, because the mass of each source is constrained to follow a 1:1 relation with light intensity, small deviations from this assumption can lead to significant overestimation (e.g., spurious features in the fourth and fifth rows) or underestimation (e.g., the outer regions of the first row) of the mass. Also, compared to our UNet baseline, our diffusion approach leverages gravitational lensing data to physically constrain the unseen dark matter, and produces uncertainty estimates to robustly quantify the model’s confidence.

6.4 Mass reconstruction of MACS 1206

As a proof of concept, we further evaluate our physically guided diffusion pipeline outside of DARKCLUSTERS-15K on the MACS 1206 galaxy cluster. This cluster, located at redshift $z_L \approx 0.44$, contains an unusually large number of multiply imaged background galaxies, which has allowed other works to study the mass distribution of this cluster in detail. As detailed in the Supplemental, we use HST photometry $\mathbf{P}_f$ in the same three bands $f \in \{125, 606, 814\}$ used to train our model, together with the strong lensing constraints from Caminha et al. [8].




Fig. 5: Comparison of our method with Kaiser-Squires (KS) [26], Napier et al. [40] and our UNet baselines. We show lens convergence κ for five test samples of DARKCLUSTERS-15K at different redshifts z_L and simulations. Napier et al. produce an overly smooth mass prediction, and yield inaccurate estimates in certain regions due to its strict 1:1 mapping of mass and photometry. Also, contrary to both Napier et al. and our UNet baselines, our method includes gravitational lensing observables and generates samples from a posterior distribution with calibrated uncertainty estimation; we show one sample of this posterior along with the mean and std. dev. computed for 20 generated samples.

In Figure 6, we compare our inferred surface mass density of MACS 1206 with the reconstruction using Lenstool from Caminha et al. [8]. We show the resulting dimensionless surface mass density κ together with its isolines, and overlay the contours from Caminha et al. which trace the diffuse cluster halo, a smoother variant that does not account for light sources in the lens plane. Although the two reconstructions are derived under different assumptions, this side-by-side comparison highlights agreement in overall halo morphology and orientation, while small-scale peaks near individual cluster members naturally differ. We also compare the critical curve at $z_S \approx 1.44$ (the curve marking the boundary of the strong-lensing region where multiple images appear), and find similar morphology. Our strong-lensing multiple-image Root Mean Square scatter, which measures the offset between the predicted and observed multiple-image positions, is 0.56 arcseconds. This is comparable to the 0.44 arcseconds of the expert-tuned reconstruction of Caminha et al. [8]. Finally, our estimate of the total projected mass enclosed within 100 kpc is $(7.61 \pm 0.04) \times 10^{13} M_\odot$ across 20 generated samples. Compared to the estimate from Caminha et al., which is $(7.25 \pm 0.02) \times 10^{13} M_\odot$, our enclosed mass error is just $5.03\% \pm 0.49\%$. As a reference, independent expert-tuned




Fig. 6: Our mass estimate for the MACS 1206 cluster, compared with the one shown in work by Caminha et al. [8]. We show our dimensionless surface mass density κ estimate along with its isolines and critical curve at $z_S \approx 1.44$; the reference work shows the diffuse cluster halo which does not account for light sources in the lens plane and thus is naturally smoother.

models of the same cluster routinely differ by $\sim 10\%$ [38]. Notably, our reconstruction requires no expert tuning, and matches expert-derived results in just a few minutes.

7 Conclusion

We have introduced the first fully automatic framework for reconstructing galaxy cluster surface mass densities from both weak and strong lensing data. Central to this effort is DARKCLUSTERS-15K, our new collection of 15,000 cluster observations derived from IllustrisTNG and SIMBA, and rendered with realistic lensing observables, which enables training, benchmarking, and future community development. By coupling a diffusion prior that learns the joint statistics of mass and light with physically grounded forward models, our approach delivers accurate, high-resolution mass reconstructions in minutes without parametric assumptions or manual tuning. Together, our diffusion-based method and the DARKCLUSTERS-15K benchmark establish a foundation for scalable, reproducible, and fully automated cluster mass inference, strengthening our tools for probing the Universe.

Acknowledgements

We are grateful to Liam Connor and Samuel McCarty for their helpful discussions, and to Edurne Bernal-Berdún for her assistance with the figures.

References

1. AbdelSalam, H.M., Saha, P., Williams, L.L.: Nonparametric reconstruction of abell 2218 from combined weak and strong lensing. *The Astronomical Journal* **116**(4), 1541 (1998)

2. Adam, A., Perreault-Levasseur, L., Hezaveh, Y.: Pixelated reconstruction of gravitational lenses using recurrent inference machines (2022), <https://arxiv.org/abs/2207.01073>
3. Barco, G.M., Legin, R., Stone, C., Hezaveh, Y., Perreault-Levasseur, L.: Blind strong gravitational lensing inversion: Joint inference of source and lens mass with score-based models. arXiv preprint arXiv:2511.04792 (2025)
4. Bertin, E., Arnouts, S.: SExtractor: Software for source extraction. *Astronomy and astrophysics supplement series* **117**(2), 393–404 (1996)
5. Biggio, L., Vernardos, G., Galan, A., Peel, A., Courbin, F.: Modeling lens potentials with continuous neural fields in galaxy-scale strong lenses. *Astronomy & Astrophysics* **675**, A125 (2023)
6. Boruah, S.S., Jacob, M., Jain, B.: Diffusion-based mass map reconstruction from weak lensing data. *Physical Review D* **111**(8), 083542 (2025)
7. Broadhurst, T., Benítez, N., Coe, D., Sharon, K., Zekser, K., White, R., Ford, H., Bouwens, R., Blakeslee, J., Clampin, M., et al.: Strong-lensing analysis of a1689 from deep advanced camera images. *The Astrophysical Journal* **621**(1), 53 (2005)
8. Caminha, G., Grillo, C., Rosati, P., Meneghetti, M., Mercurio, A., Ettori, S., Balestra, I., Biviano, A., Umetsu, K., Vanzella, E., et al.: Mass distribution in the core of macs j1206-robust modeling from an exceptionally large sample of central multiple images. *Astronomy & Astrophysics* **607**, A93 (2017)
9. Cha, S., HyeonHan, K., Scofield, Z.P., Joo, H., Jee, M.J.: Precision mars mass reconstruction of a2744: Synergizing the largest strong-lensing and densest weak-lensing data sets from jwst. *The Astrophysical Journal* **961**(2), 186 (2024)
10. Cha, S., Jee, M.J.: Mars: a new maximum-entropy-regularized strong lensing mass reconstruction method. *The Astrophysical Journal* **931**(2), 127 (2022)
11. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. *The Eleventh International Conference on Learning Representations* (2023)
12. Davé, R., Anglés-Alcázar, D., Narayanan, D., Li, Q., Rafieferantsoa, M.H., Appleby, S.: Simba: Cosmological simulations with black hole growth and feedback. *Monthly Notices of the Royal Astronomical Society* **486**(2), 2827–2849 (2019)
13. De Luca, F., De Petris, M., Yepes, G., Cui, W., Knebe, A., Rasia, E.: The three hundred project: dynamical state of galaxy clusters and morphology from multiwavelength synthetic maps. *Monthly Notices of the Royal Astronomical Society* **504**(4), 5383–5400 (2021)
14. Diego, J., Congedo, G., Gavazzi, R., Schrabback, T., Atek, H., Jain, B., Weaver, J., Kang, Y., Hartley, W., Mahler, G., et al.: Euclid: Early release observations. a combined strong and weak lensing solution for abell 2390 beyond its virial radius. arXiv preprint arXiv:2507.08545 (2025)
15. Diego, J.M., Tegmark, M., Protopapas, P., Sandvik, H.: Combined reconstruction of weak and strong lensing data with wslap. *Monthly Notices of the Royal Astronomical Society* **375**(3), 958–970 (2007)
16. Ding, H., Shu, Y., Chen, Y., Li, N., Cao, X., Nightingale, J., He, Q., Wang, L., Li, H.: Pixel-level modeling of group-scale strong lens cassowary 19. *Research in Astronomy and Astrophysics* **25**(6), 065013 (May 2025). <https://doi.org/10.1088/1674-4527/add0a0>, <http://dx.doi.org/10.1088/1674-4527/add0a0>
17. Dolag, K., Remus, R.S., Valenzuela, L.M., Kimmig, L.C., Seidel, B., Fortune, S., Stoiber, J., Ivleva, A., Hoffmann, T., Biffi, V., et al.: Encyclopedia magneticum: Scaling relations from cosmic dawn to present day. arXiv preprint arXiv:2504.01061 (2025)
18. Galan, A., Vernardos, G., Peel, A., Courbin, F., Starck, J.L.: Using wavelets to capture deviations from smoothness in galaxy-scale strong lenses. *Astronomy & Astrophysics* **668**, A155 (2022)

19. Gonzalez, A.H., Zabludoff, A.I., Zaritsky, D.: Intracluster light in nearby galaxy clusters: relationship to the halos of brightest cluster galaxies. *The Astrophysical Journal* **618**(1), 195 (2005)
20. Hezaveh, Y.D., Levasseur, L.P., Marshall, P.J.: Fast automated analysis of strong gravitational lenses with convolutional neural networks. *Nature* **548**(7669), 555–557 (2017)
21. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
22. Huff, E., Krause, E., Eifler, T., Everett, S., Huang, Y.H., Xu, J., et al.: Kinematic lensing inference—ii. cluster lensing with $o(1)$ galaxies. *Monthly Notices of the Royal Astronomical Society* **540**(4), 2877–2893 (2025)
23. Ingoglia, L., Covone, G., Sereno, M., Giocoli, C., Bardelli, S., Bellagamba, F., Castignani, G., Farrens, S., Hildebrandt, H., Joudaki, S., Jullo, E., Lanzieri, D., Lesci, G.F., Marulli, F., Maturi, M., Moscardini, L., Nanni, L., Puddu, E., Radovich, M., Roncarelli, M., Sapio, F., Schimd, C.: Amico galaxy clusters in kids-dr3: measurement of the halo bias and power spectrum normalization from a stacked weak lensing analysis. *Monthly Notices of the Royal Astronomical Society* **511**(1), 1484–1501 (Jan 2022). <https://doi.org/10.1093/mnras/stac046>, <http://dx.doi.org/10.1093/mnras/stac046>
24. Jeffrey, N., Lanusse, F., Lahav, O., Starck, J.L.: Deep learning dark matter map reconstructions from des sv weak lensing data. *Monthly Notices of the Royal Astronomical Society* **492**(4), 5023–5029 (2020)
25. Jullo, E., Kneib, J.P., Limousin, M., Eliasdottir, A., Marshall, P., Verdugo, T.: A bayesian approach to strong lensing modelling of galaxy clusters. *New Journal of Physics* **9**(12), 447 (2007)
26. Kaiser, N., Squires, G.: 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)
27. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems* **35**, 26565–26577 (2022)
28. Kassiola, A., Kovner, I.: Elliptic mass distributions versus elliptic potentials in gravitational lenses. *Astrophysical Journal* v. 417, p. 450 **417**, 450 (1993)
29. Klypin, A., Yepes, G., Gottlöber, S., Prada, F., Hess, S.: Multidark simulations: the story of dark matter halo concentrations and density profiles. *Monthly Notices of the Royal Astronomical Society* **457**(4), 4340–4359 (2016)
30. Kneib, J.P., Ellis, R.S., Smail, I., Couch, W., Sharples, R.: Hubble space telescope observations of the lensing cluster abell 2218. *The Astrophysical Journal* **471**(2), 643–656 (1996)
31. Kneib, J.P., Bonnet, H., Golse, G., Sand, D., Jullo, E., Marshall, P.: Lenstool: a gravitational lensing software for modeling mass distribution of galaxies and clusters (strong and weak regime). *Astrophysics Source Code Library* pp. ascl-1102 (2011)
32. Kuleshov, V., Fenner, N., Ermon, S.: Accurate uncertainties for deep learning using calibrated regression. In: *International conference on machine learning*. pp. 2796–2804. PMLR (2018)
33. Lanusse, F., Starck, J.L., Leonard, A., Pires, S.: High resolution weak lensing mass mapping combining shear and flexion. *Astronomy & Astrophysics* **591**, A2 (2016)
34. Leonard, A., Lanusse, F., Starck, J.L.: Glimpse: accurate 3d weak lensing reconstructions using sparsity. *Monthly Notices of the Royal Astronomical Society* **440**(2), 1281–1294 (2014)
35. Levasseur, L.P., Hezaveh, Y.D., Wechsler, R.H.: Uncertainties in parameters estimated with neural networks: Application to strong gravitational lensing. *The Astrophysical Journal Letters* **850**(1), L7 (2017)
36. Limousin, M., Kneib, J.P., Natarajan, P.: Constraining the mass distribution of galaxies using galaxy-galaxy lensing in clusters and in the field. *Monthly Notices of the Royal Astronomical Society* **356**(1), 309–322 (2005)

37. McCarthy, I.G., Schaye, J., Bird, S., Le Brun, A.M.C.: The bahamas project: calibrated hydrodynamical simulations for large-scale structure cosmology. *Monthly Notices of the Royal Astronomical Society* **465**(3), 2936–2965 (2017). <https://doi.org/10.1093/mnras/stw2792>
38. Meneghetti, M., Natarajan, P., Coe, D., Contini, E., De Lucia, G., Giocoli, C., Acebron, A., Borgani, S., Bradac, M., Diego, J.M., et al.: The frontier fields lens modelling comparison project. *Monthly Notices of the Royal Astronomical Society* **472**(3), 3177–3216 (2017)
39. Mishra-Sharma, S., Yang, G.: Strong lensing source reconstruction using continuous neural fields (2022), <https://arxiv.org/abs/2206.14820>
40. Napier, K., Sharon, K., Dahle, H., Bayliss, M., Gladders, M.D., Mahler, G., Rigby, J.R., Florian, M.: Hubble constant measurement from three large-separation quasars strongly lensed by galaxy clusters. *The Astrophysical Journal* **959**(2), 134 (2023)
41. Natarajan, P., Williams, L., Bradač, M., Grillo, C., Ghosh, A., Sharon, K., Wagner, J.: Strong lensing by galaxy clusters. *Space Science Reviews* **220**(2), 19 (2024)
42. Navarro, J.F., Frenk, C.S., White, S.D.: A universal density profile from hierarchical clustering. *The Astrophysical Journal* **490**(2), 493–508 (1997)
43. Nelson, D., Pillepich, A., Ayromlou, M., Lee, W., Lehle, K., Rohr, E., Truong, N.: Introducing the tng-cluster simulation: Overview and the physical properties of the gaseous intracluster medium. *Astronomy & Astrophysics* **686**, A157 (2024)
44. Niemiec, A., Jauzac, M., Jullo, E., Limousin, M., Sharon, K., Kneib, J.P., Natarajan, P., Richard, J.: hybrid-lenstool: a self-consistent algorithm to model galaxy clusters with strong- and weak-lensing simultaneously. *Monthly Notices of the Royal Astronomical Society* **493**(3), 3331–3340 (2020)
45. Pillepich, A., Nelson, D., Hernquist, L., Springel, V., Pakmor, R., Torrey, P., Weinberger, R., Genel, S., Naiman, J.P., Marinacci, F., Vogelsberger, M.: First results from the illustrious simulations: the stellar mass content of groups and clusters of galaxies. *Monthly Notices of the Royal Astronomical Society* **475**(1), 648–675 (2018). <https://doi.org/10.1093/mnras/stx3112>, <http://dx.doi.org/10.1093/mnras/stx3112>
46. Postman, M., Coe, D., Benitez, N., Bradley, L., Broadhurst, T., Donahue, M., Ford, H., Graur, O., Graves, G., Jovel, S., et al.: The cluster lensing and supernova survey with hubble: an overview. *The Astrophysical Journal Supplement Series* **199**(2), 25 (2012)
47. Remy, B., Lanusse, F., Jeffrey, N., Liu, J., Starck, J.L., Osato, K., Schrabback, T.: Probabilistic mass-mapping with neural score estimation. *Astronomy & Astrophysics* **672**, A51 (2023)
48. Rüstig, J., Guardiani, M., Roth, J., Frank, P., Enßlin, T.: Introducing lenscharm-a charming bayesian strong lensing reconstruction framework. *Astronomy & Astrophysics* **682**, A146 (2024)
49. Sartoris, B., Biviano, A., Fedeli, C., Bartlett, J., Borgani, S., Costanzi, M., Giocoli, C., Moscardini, L., Weller, J., Ascaso, B., et al.: Next generation cosmology: constraints from the euclid galaxy cluster survey. *Monthly Notices of the Royal Astronomical Society* **459**(2), 1764–1780 (2016)
50. Schneider, P.: Introduction to gravitational lensing and cosmology. In: *Gravitational Lensing: Strong, Weak and Micro*, pp. 1–89. Springer (2006)
51. Seljak, U.: Weak lensing reconstruction and power spectrum estimation: minimum variance methods. *The Astrophysical Journal* **506**(1), 64 (1998)
52. Sérsic, J.: Influence of the atmospheric and instrumental dispersion on the brightness distribution in a galaxy. *Boletín de la Asociación Argentina de Astronomía La Plata Argentina* **6**, 41–43 (1963)
53. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations (2021), <https://arxiv.org/abs/2011.13456>

54. Villaescusa-Navarro, F., Genel, S., Angles-Alcazar, D., Thiele, L., Dave, R., Narayanan, D., Nicola, A., Li, Y., Villanueva-Domingo, P., Wandelt, B., et al.: The camels multifield data set: Learning the universe’s fundamental parameters with artificial intelligence. *The Astrophysical Journal Supplement Series* **259**(2), 61 (2022)
55. Wen, Z., Han, J.: A catalog of 1.58 million clusters of galaxies identified from the desi legacy imaging surveys. *The Astrophysical Journal Supplement Series* **272**(2), 39 (2024)
56. Zhang, B., Chu, W., Berner, J., Meng, C., Anandkumar, A., Song, Y.: Improving diffusion inverse problem solving with decoupled noise annealing. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 20895–20905 (2025)
57. Zhao, B., Levis, A., Connor, L., Srinivasan, P.P., Bouman, K.: Revealing the 3d cosmic web through gravitationally constrained neural fields. In: *The Thirteenth International Conference on Learning Representations* (2025)
58. Zhao, B., Levis, A., Connor, L., Srinivasan, P.P., Bouman, K.L.: Single view refractive index tomography with neural fields. In: *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 25358–25367. IEEE Computer Society (2024)
59. Zheng, H., Chu, W., Zhang, B., Wu, Z., Wang, A., Feng, B., Zou, C., Sun, Y., Kovachki, N.B., Ross, Z.E., Bouman, K., Yue, Y.: Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. In: *The Thirteenth International Conference on Learning Representations* (2025), <https://openreview.net/forum?id=U3PBITXNG6>

A Construction of DARKCLUSTERS-15K

To construct our DARKCLUSTERS-15K dataset, we mostly used the TNG300-1 subset of the IllustrisTNG cosmological simulation [45]. This dataset spans a cubic volume with a side length of $205,000 \text{ ckpc}/h$ that includes both baryonic matter (visible matter: stars, black holes and gas) and dark matter. TNG300-1 contains thousands of galaxy clusters resolved at high mass and spatial resolution. Additionally, we used the SIMBA simulations [12], which are similar in terms of contents, volume, and resolution, but rely on different algorithms than IllustrisTNG. Our SIMBA test set allows us to evaluate generalization of our method under different simulation settings.

The surface mass density maps $\Sigma(\theta)$ include both baryonic (visible) and dark matter, and are measured in $(10^{10} M_{\odot}/h)/(\text{ckpc}/h)^2$, i.e., in terms of 10^{10} times the mass of the Sun $M_{\odot}$ per comoving kiloparsec squared, all scaled by the Hubble parameter h . We also release maps where both components have been separated to facilitate analysis. The photometry information $\mathbf{P}_f(\theta)$ is measured in janskys (Jy). Since all DARKCLUSTERS-15K images are 512×512 pixels, each pixel corresponds to an area of $1.56 (\text{ckpc}/h)^2$ for images with a field of view of 100×100 arcseconds, and $7.87 (\text{ckpc}/h)^2$ in the case of 225×225 arcseconds, for clusters at $z_L = 0.5$. These correspond to image resolutions of 0.2 and 0.44 arcseconds respectively. For reference, dark matter distributions are substantially smoother, and existing mass reconstruction methods typically operate at resolutions of 0.5 to 1 arcsecond [10, 14].

Selection of galaxy clusters. TNG300-1 provides a friends-of-friends catalog of individual galaxy clusters within the simulation. For each identified cluster, we measure its total gravitationally-bound mass (the *virial* mass) and the corresponding enclosing radius (the virial radius). We use a similar strategy for clusters from SIMBA, using their provided catalogs to filter by total cluster masses. To construct our dataset, we select 600 galaxy clusters with virial masses of at least $10^{13.5}$ times the mass of our Sun $M_{\odot}$. Figure 7 shows the distribution of selected galaxy clusters by their virial masses and radii. As noted in our introduction, our analysis focuses on clusters in the range $10^{13.5-14} M_{\odot}$, which form the majority of our dataset.

A.1 Parsing data from IllustrisTNG

Galaxy clusters contain volumetric information on mass and emitted light, represented as 3D particle clouds. Out of the 600 clusters in DARKCLUSTERS-15K, 540 come from IllustrisTNG. In our work, we process this information to obtain 25 samples per cluster, where each sample consists of four images: one surface mass density map $\Sigma(\theta)$, and three simulated HST images $\mathbf{P}_f$ (Section 4.1). Below, we describe the processing steps for each one.

Surface mass density maps $\Sigma(\theta)$. For baryonic matter, we use the `Coordinates` and `Masses` fields of gas, star and black hole particles. For dark matter, we use the `Coordinates` and `SubfindDensity` fields together with the `MassTable` attribute. This allows us to provide visible and dark matter components in DARKCLUSTERS-15K, although in our work we use only their combined mass.




Fig. 7: Distribution of the galaxy clusters selected from IllustrisTNG and SIMBA to construct DARKCLUSTERS-15K. We show their virial masses (i.e., their total gravitationally-bound masses) and their virial radii (i.e., the radius that encloses that mass).

We generate 25 different 2D projections from each 3D particle cloud (Figure 2). These correspond to 25 different viewing directions that are uniformly distributed over the surface of a hemisphere enclosing the cluster, each defined by a unit line-of-sight direction $\mathbf{r}_i$. In practice, each corresponds to a rotation of the particle cloud and a projection along $\mathbf{r}_i$. Then, to compute the mass at each pixel θ in $\Sigma(\theta)$, we find the 300 nearest particles and perform kernel density estimation with a box filter.

Photometry maps $\mathbf{P}_f$. We use `Coordinates` and `GFM_StellarPhotometrics` from star particles for each cluster. This last field contains stellar magnitudes in eight bands (U, B, V, K, g, r, i, z) [45], which we convert to flux densities f_ν in Jy. As with the surface mass density maps, we generate the same projections by rotating the 3D particle cloud and projecting along each line-of-sight direction $\mathbf{r}_i$. Unlike the mass maps, we do not perform kernel density estimation, and instead we just bin each particle to its nearest pixel. To simulate HST observations, this binning accounts for dust attenuation and redshift, and then it uses the public response curves of the HST F125W, F606W and F814W wavelength filters for the bands $f \in \{125, 606, 814\}$ respectively. This produces initial photometry maps, which we refine through post-processing as described below.

First, we add background light using publicly available mosaics from other HST captures.

Second, we add intracluster light (ICL), which corresponds to light from stars that are not bound to the galaxy cluster. Following Gonzalez et al. [19], we model ICL using two Sérsic source parametric models [52] fit to the brightest cluster galaxies. We use SE++ [4] to do the fitting, then we create the Sérsic models and add the resulting light to the image.

Finally, we simulate several strongly lensed light sources. Following the process described in Section 4.1 and Figure 2, we use a Sérsic parametric model for each light on the source plane S , and simulate the corresponding observed image on the lens plane L , which we add to the photometry.

A.2 Parsing data from SIMBA

Our pipeline for the SIMBA simulations [12] follows a very similar approach. We extract 60 galaxy clusters and also generate 25 projections per cluster, resulting in a test set of 1,500 samples. Just as with IllustrisTNG, each sample consists of one surface mass density map $\Sigma(\theta)$ and three HST images $\mathbf{P}_f$ in identical units. The main differences from the IllustrisTNG setup are:

Surface mass density maps $\Sigma(\theta)$. While we also use the `Coordinates` and `Masses` for gas, stars and black holes, we now use the `Masses` field for dark matter particles as well.

Photometry maps $\mathbf{P}_f$. Since SIMBA does not directly provide stellar magnitudes, we derive them from the `Masses`, `StellarFormationTime` and `Metallicity` fields of the star particles. We map these properties using stellar emission models to determine their base brightness, which we then scale by the particle’s total mass to calculate the final magnitudes.

B Parsing data from MACS 1206

To show that our method works on real observations, we applied it to the galaxy cluster MACS 1206. We chose this cluster since it is not affected by external convergence or shear (i.e., gravitational lensing deflection in the space before the cluster itself). Our estimation uses photometry $\mathbf{P}_f$ and strong lensing data, which we obtain as follows:

Photometry maps $\mathbf{P}_f$. We use the publicly available CLASH drizzled mosaics of MACS 1206 [46], taken in the same three filters as our synthetic data (F125W, F606W and F814W, i.e., $f \in \{125, 606, 814\}$). We align the three mosaics and then crop and resample to 512×512 pixels, so that their field of view and resolution match those of DARKCLUSTERS-15K. Finally, we convert the mosaics from their native units (electron counts per second) to flux densities f_ν in janskys (Jy), using the photometric keywords in the FITS headers.

Strong lensing constraints. We extract strong lensing constraints from work by Caminha et al. [8]. From their catalog, we filter 55 multiple images that fall within the region covered by our photometry; these belong to 15 distinct background sources, with redshifts ranging from $z_S \approx 1.0$ to $z_S \approx 3.8$.

C Training and inference details

For training and inference, we normalize the surface mass density maps Σ to the $[0, 1]$ range using a logarithmic transform followed by shift and scale operations. We apply the same normalization, independently, to the photometry maps $\mathbf{P}_f$.

Training follows default configuration provided by InverseBench [59]. Our 137M parameter model trains in around 10 hours on a single RTX 4090 GPU, where we run all

of our experiments. All our experiments use a DDPM++ backbone. We also evaluated our method using a Diffusion Transformer backbone, which achieved comparable performance while requiring higher computational cost. This suggests we are not significantly bottlenecked by prior model choice in our experiments.

During inference, the loss weights λ_k for our strong and weak lensing forward operators in Equation 10 are set as follows. For strong lensing, the geometric and photometric losses use $\lambda_s^{\text{geo}} \approx 0.6 \cdot 10^{-2}$ and $\lambda_s^{\text{img}} \approx 0.6 \cdot 10^{-3}$ respectively, which puts both losses on roughly the same scale. For weak lensing, we use $\lambda_w \approx 0.05$. Additionally, for the DAPS algorithm, we set the parameters $\tau \approx 2.6 \cdot 10^{-3}$ and learning rate of $2.8 \cdot 10^{-6}$ over 100 steps. The rest of the parameters use the default values from InverseBench [59]. All these values were chosen through parameter sweeps on a small validation subset of DARKCLUSTERS-15K excluded from both training and final evaluation.

D Light-traces-mass in DARKCLUSTERS-15K

Here we measure the extent to which the light-traces-mass assumption holds in our DARKCLUSTERS-15K dataset. There are several previous works [31, 40] that assume a 1:1 relation of mass to light, while our work incorporates a diffusion-based prior that can learn complex relations. For this, we compute the Pearson Correlation Coefficient (PCC) between the mass and photometry maps at the $f = 125$ band, which we also use for our implementation of the method from Napier et al. [40] in Section 6.3. We evaluate this metric on the three IllustrisTNG test subsets of DARKCLUSTERS-15K, corresponding to galaxy cluster redshifts $z_L \in \{0.2, 0.5, 1.0\}$, obtaining values of 0.58, 0.61 and 0.65 respectively. These results indicate a moderate correlation. One of the main differences comes from the presence of massive clumps of dark matter which do not interact with light. There are also noticeable discrepancies primarily caused by background light, intracluster light (Section A.1) and the absence of the scaling relations from Limousin et al. [36] (Section 6.3).

Comparing these PCC values to those reported in Table 2, both the method of Napier et al. [40], our UNet baseline, and, especially, our method achieve substantially higher correlations. As shown in Figure 5, our approach learns how to effectively compensate for the mismatches mentioned before without assuming an explicit 1:1 mass-light relation, resulting in improved overall performance.