

# Single View Refractive Index Tomography with Neural Fields

Brandon Zhao<sup>1\*</sup> Aviad Levis<sup>1</sup> Liam Connor<sup>2</sup> Pratul P. Srinivasan<sup>3</sup> Katherine L. Bouman<sup>1,2</sup>

<sup>1</sup>Department of Computing and Mathematical Sciences, California Institute of Technology

<sup>2</sup> Cahill Center for Astronomy and Astrophysics, California Institute of Technology

<sup>3</sup>Google Research

\*byzhao@caltech.edu

## Abstract

*Refractive Index Tomography is the inverse problem of reconstructing the continuously-varying 3D refractive index in a scene using 2D projected image measurements. Although a purely refractive field is not directly visible, it bends light rays as they travel through space, thus providing a signal for reconstruction. The effects of such fields appear in many scientific computer vision settings, ranging from refraction due to transparent cells in microscopy to the lensing of distant galaxies caused by dark matter in astrophysics. Reconstructing these fields is particularly difficult due to the complex nonlinear effects of the refractive field on observed images. Furthermore, while standard 3D reconstruction and tomography settings typically have access to observations of the scene from many viewpoints, many refractive index tomography problem settings only have access to images observed from a single viewpoint. We introduce a method that leverages prior knowledge of light sources scattered throughout the refractive medium to help disambiguate the single-view refractive index tomography problem. We differentially trace curved rays through a neural field representation of the refractive field, and optimize its parameters to best reproduce the observed image. We demonstrate the efficacy of our approach by reconstructing simulated refractive fields, analyze the effects of light source distribution on the recovered field, and test our method on a simulated dark matter mapping problem where we successfully recover the 3D refractive field caused by a realistic dark matter distribution.*

## 1. Introduction

When light passes from one medium to another, such as from air to water, it encounters a change in *refractive index* (RI). As a result, light rays are bent, or refracted, according to the magnitude of this change, causing effects like the shimmering (caustics) at the bottom of a pool of water. Media such as hot gas or translucent cell samples have RIs

that continuously vary in space, causing light to take paths curved by the gradient of the underlying RI field. Although the RI field is not directly visible, its effect on the scene's illumination provides clues to its structure.

Spatially varying RI fields are present in a diverse range of scientific applications. For example, certain fiber optic cables are designed to smoothly refract light away from the cable's outside padding [31]. In seismology, pressure waves from earthquakes are refracted continuously as they travel through ground structures of varying density [23]. In astrophysics, gravitational effects from massive dark matter structures lens the light from faraway galaxies resulting in observable effects in telescope images. When the degree of these gravitational lensing effects is relatively weak, they can be modeled as refraction [15].

In this work, we seek to tackle the problem of single-view RI tomography: 3D reconstruction of RI fields where image measurements are available only from a single viewpoint. This problem setting occurs when using telescope images or a single camera on scenes in the wild, and poses several challenges. First, image measurements are produced by rays traveling along continuously curving trajectories through the scene, where curvature is due to the structure of the underlying RI field. Unlike in linear tomography problems, where the projection geometry is independent of the medium, in RI tomography the ray paths are dictated by the unknown medium properties themselves. Thus the measurements are a non-linear function of the targeted refractive field. Second, limiting image measurements to a single viewing angle makes the problem severely ill-posed. That is, there are many possible refractive fields which could potentially produce the same image measurements. However, by leveraging the location of light sources distributed throughout the scene we can better constrain the recovery.

We propose a novel approach using neural fields with continuous ray tracing that leverages the physics of light propagation to recover an underlying refractive field. In particular, we represent the spatially varying RI field with




Figure 1. We formulate and present an approach for single-view refractive field reconstruction that explicitly models curved ray paths through a neural field. By leveraging spatial information and the emission profile of light sources located throughout the volume, we are able to localize refractive structures in 3D with just a single viewpoint. Here we demonstrate our approach on a simulation involving a realistic refractive field induced by simulated dark matter halos. Despite the large patches of space that are void of galaxy emitters, our method is still able to recover the rough structure of the 3D dark matter distribution from a single simulated telescope image of refracted galaxies. This reconstruction is also especially challenging because of the extremely weak degree of refraction, as seen in the relative scale of the intensity difference plot; the refractive effects of dark matter result in an intensity change of at most  $\pm 0.3\%$  of the maximum intensity of the original image. This experiment uses a realistic dark matter distribution derived from IllustrisTNG [27], but we leave the inclusion of realistic measurement noise to future work.

a coordinate-based neural network, which we optimize through a differentiable rendering procedure that traces paths of light through the refractive field to best reproduce measurements. Coordinate-based neural networks are advantageous for this problem because they provide a continuous, easily-differentiable representation of the refractive field as a function of space. We also find empirically that coordinate-based neural networks lead to better optimization behavior out-of-the-box than other representations such as linearly interpolated grids.

Our approach expands upon recent works in neural scene representation [21, 25]. While in these works ray paths through the scene are assumed to be known, our approach explicitly models and solves for unknown refracted ray paths. We first study the conditions under which single-view reconstruction is possible, then test our approach by analyzing its ability to recover weak refractive fields. Next, we analyze the sensitivity of our approach to the density of light sources in the scene. Finally, as shown in Fig. 1, we include a case study motivated by the application of weak lensing where our method recovers the refractive field generated by physically realistic simulations of dark matter.

## 2. Related Work

**Image-Based Refractive Field Recovery** The reconstruction of refractive objects and phenomena is an ongoing field of study, both in applications that make use of standard consumer cameras as well as scientific imaging applications. In Background-oriented Schlieren (BoS) imaging, images taken of a background with and without refractive effects are correlated to estimate an optical flow describing a projected measurement of the refractive field [4]. Combining these projected measurements from multiple viewpoints

through tomography gives an estimate of the 3D gradients, from which the 3D field can be estimated from a Poisson solver. A similar type of tomography can be done to passively model 3D turbulence statistics by combining time resolved deflection measurements from multiple viewpoints [1]. Optical flow estimates such as those in BoS can also be used to detect refractive “wiggles” in fluid flow that can then be used to track the projected motion of the refractive media; by observing these wiggles with stereo cameras the fluid’s depth can also be recovered [40]. Fluid refraction from water can also be estimated by submerged cameras viewing the sun as a guiding point [2].

While the aforementioned methods utilize the physics behind ray optics, they don’t explicitly model the continuously curved paths that rays take through these refractive media during 3D reconstruction. A method using the refractive curvature of traced rays to supervise an interpolated grid model has been shown to be effective for reconstructing and designing refractive media [37]. However, as with the previous methods, this method still relies on access to multiple viewpoints to reconstruct the refractive volume. In this work, we instead seek to tackle the challenge of accurately reconstructing 3D refractive fields from a single viewpoint.

**Applications** Recent works in optics have used differentiable rendering techniques to design refractive surfaces for various applications. The materials making up the lens typically have a constant RI, meaning light is refracted a discrete number of times at the interface between the lens and the surrounding air. Thus, these works solve for the optimal shape of the material rather than an RI field varying within the lens. Fast differentiable ray tracing techniques have been used for the design of a single lens to produce a

target image [22], as well as compound lenses for camera design [38]. These types of lens designs are aimed at reproducing a target image or favorable image effects. While we also aim to optimize a refractive field to match a set of observed image measurements, we are interested in recovering the properties of the refractive medium itself.

In biomedical imaging, Optical diffraction tomography (ODT) uses the intrinsic optical variation in a sample to reconstruct its RI field [11, 34]. Holographic measurements are taken from multiple illumination angles with coherent light modeled by wave optics. For example, in [12], the authors present a technique to recover RI by using intensity-only measurements from multiple viewpoints. DeCAF [24] uses a neural field representation to solve this inverse problem. Our problem setup differs from these biomedical applications in that we attempt to solve this problem without access to active illumination and coherent light sources. Additionally, we use ray optics rather than wave optics in our characterization of light transport.

We are motivated in particular by the application of gravitational lensing, where light coming from small elliptical sources (galaxies) curves on its path to Earth due to an invisible dark matter distribution. In the weak lensing regime, the small ray deflections caused by dark matter can be well-approximated as a refractive effect [15], allowing us to model the large-scale structure of the dark matter distribution as an RI field. Existing works use a linear tomography approach similar to that of BoS [4] by first computing a 2D projected map of mass density, then backprojecting into 3D to obtain a reconstruction. Because image measurements are limited to a single viewpoint, filtered backprojection-style methods [33, 39] are susceptible to heavy smearing along the line of sight, even in the absence of noise. In [19, 20] an approach is used that assumes the density field can be sparsely represented by a dictionary of dark matter halos. However, this assumption is only approximately true in practice; our method can instead reconstruct 3D RI fields induced by a general matter distribution.

**Coordinate-based Neural Representations** Recently coordinate-based neural networks have gained traction for applications in vision and graphics. These models typically parameterize a three dimensional field, such as a scene’s radiance, with the weights of a multi-layer perceptron (MLP). Coordinate-based neural representations have been used in graphics to model shapes with occupancy fields [10], as well as three-dimensional realistic scenes with neural radiance fields [25]. Both of these works rely on a fixed straight ray assumption to render scenes for training and inference. Coordinate-based MLPs have also been used in black hole imaging [21] to reconstruct emission fields with curved ray paths dictated by general relativity; however, the paths taken by rays of light are calculated then fixed beforehand according to a specific model of gravitational lens. In contrast, our

work seeks to optimize ray paths that are unknown and directly estimated during model training.

Eikonal ray-tracing with neural fields has been used for refractive tomography in the context of novel-view synthesis [6]. However, this work uses images of the scene from multiple viewpoints for supervision. Additionally, it focuses on recovering scenes which have a piecewise constant RI (e.g. a glass block with constant RI surrounded by air), where no light emission occurs within the refractive medium. While it achieves impressive results for rendering images from novel viewpoints, the reconstructed refractive indices of these volumes fail to match the ground truth. In contrast, our method leverages visible light sources from within the refractive medium to enable accurate reconstruction of a refractive volume from only a single viewpoint.

### 3. Forward: Image Formation Model

This section details the components of the image formation (forward) model. First, we introduce the equations governing the refraction and intensity measurements of individual rays. Next, we detail the numerical integration used to model image intensity measurements at a sensor.

#### 3.1. Refractive Ray Tracing

Purely refractive media are translucent and affect light rays by curving their paths rather than modifying their intensity. Light travels more slowly through more optically dense material, causing it to bend towards an increase in optical density. The RI  $\eta$  is a scalar used to quantify the optical density of a material relative to a vacuum; in natural settings,  $\eta \geq 1$  as light cannot exceed its speed in a vacuum. In this work, we are interested in volumes with a continuously varying refractive index. The effect of a continuously varying RI  $\eta$  on a light ray’s position  $x$  and direction  $v$  is described by Hamilton’s equations for ray tracing [16]:

$$\begin{cases} \frac{dx}{ds} = \frac{v(s)}{\eta(s)} \\ \frac{dv}{ds} = \nabla\eta(s) \end{cases} \quad (1)$$

Here the derivatives of the 3D ray position  $x$  and direction  $v$  are taken with respect to the ray path length  $s$ , and are described in terms of the RI  $\eta$  and its gradient  $\nabla\eta$ . The directional derivative of Eq. (2) can be understood as an attractive force towards local increases in RI.

Given an initial position and direction  $x_0, v_0 \in \mathbb{R}^3$ , equations (1) and (2) describe a first-order initial value problem whose solution  $\Gamma(s) = (x, v)$  is a curve describing the trajectory of a ray as a function of its path length  $s$ :

$$\Gamma(s) = (x_0, v_0) + \int_0^s (x'(\tilde{s}), v'(\tilde{s})) d\tilde{s} \quad (3)$$

An illustration of refractive ray tracing is given in Figure 2.

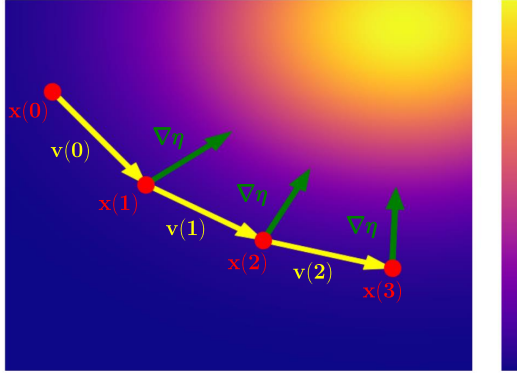


Figure 2. Illustration of ray tracing in a refractive field. Ray paths are calculated iteratively by solving the Hamiltonian equations (1), (2) parameterized by the refractive field  $\eta$ . Rays are curved locally towards areas of high RI in the direction of  $\nabla\eta$ .

### 3.2. Sensor and Light Emission Integration

The sensor is modeled as a single-wavelength pinhole camera positioned outside of the scene. The intensity image  $I$  of a given scene is represented as a discrete function indexed by pixel coordinates  $i, j$ . The specific intensity of a pixel  $I(i, j)$  is given by integration of light emission over a ray path solution  $\Gamma_{i,j}$ , to Eq. (3). The initial conditions  $\Gamma_{i,j}(0)$  are defined by the position and direction of a ray traveling backwards from the sensor through the pixel center.

Within the scene, light emission  $e(\mathbf{x})$  is represented as a continuous scalar function of a 3D coordinate  $\mathbf{x} = (x, y, z)$ , and is assumed to be known. This assumption is commonly used for refractive modeling problems in the forms of calibration images in BoS imaging [4] or galaxy catalogs in dark matter mapping [15]. Given the ray path  $\Gamma_{i,j}$  starting at the pixel in location  $(i, j)$ , we model the specific intensity  $I(i, j)$  by taking the line integral of emission coefficients  $e$  along the path  $\Gamma_{i,j}$ :

$$I(i, j) = \int_{\Gamma_{i,j}} e(\mathbf{x}(s)) ds \quad (4)$$

Eq. (4) corresponds to a simplified model of light transport in which the specific intensity is conserved along a ray’s path. In gravitational lens theory [32], the emission term in the integrand is decreased for sources in areas of high refractive index due to gravitational redshift. However, for gravitational RI fields with  $\eta$  close to 1, which we focus on in this work, this effect is negligible and we thus exclude it from our forward model [21]. In other applications of refractive tomography [6, 28], the refractive radiative transfer equation (RRTE) [3] may be more appropriate for the rendering forward model. Further details, including experiments with the RRTE, can be found in the supplement.

### 3.3. Numerical Integration

A ray path solution  $\Gamma = (x, v)$  to the Hamiltonian initial value problem given by equations (1) and (2) is fully determined by the initial position and direction  $\Gamma_{i,j}(0)$  of the ray traced through a pixel  $(i, j)$ , the refractive field  $\eta$ , and its gradient  $\nabla\eta$ . Thus, we rewrite the integral in Eq. (3) as:

$$\Gamma_{i,j} = \text{solveHam}(\Gamma_{i,j}(0), \eta, \nabla\eta). \quad (5)$$

Integrating the scene’s emission function  $e$  along the ray path  $\Gamma$  as in Eq. (4) allows us to calculate image intensities:

$$I(i, j) = \text{renderIm}(\Gamma_{i,j}, e). \quad (6)$$

In general, closed-form solutions for integrals (5) and (6) do not exist, so they are solved numerically. We use an adaptive step-size Dormand-Prince integrator [13] to solve Eqs. (5) and (6) in our experiments. Backpropagating through the integrator using autograd requires memory in the order of the number of solver steps. Instead, we follow previous works [6, 9, 37] in using the adjoint formulation [30] which has constant memory requirements for backpropagation.

## 4. Inverse: Estimating a Refractive Field

In this work, we focus on weakly refractive fields where the deflections of a single source are on the order of fractions of a pixel. Although this may seem small, we rely on the combined signal across an ensemble of lensed sources.

To solve the inverse problem of refractive tomography in this setting, we use an analysis-by-synthesis approach. We model the continuous refractive field with a coordinate-based MLP, which is optimized to fit a true observed image. In the following section we describe the neural representation and optimization scheme.

### 4.1. Continuous Refractive Field Representation

To form a candidate continuously-varying 3D refractive field, we use a coordinate-based MLP inspired by the NeRF [25] model. The MLP, parameterized by weights  $\theta$ , takes as input spatial coordinates  $\mathbf{x}$  and outputs a scalar index of refraction. To ensure that our reconstruction is physically plausible, we apply the Softplus function  $\sigma$  to restrict the output of our network to be greater than 1:

$$\hat{\eta}_\theta(\mathbf{x}) = \sigma(\text{MLP}(\gamma(\mathbf{x})); \theta). \quad (7)$$

Here,  $\gamma(\mathbf{x})$  is a positional encoding layer mapping an input coordinate to a set of Fourier basis functions [36] with exponentially-increasing frequencies:

$$\gamma(\mathbf{x}) = [\sin(\mathbf{x}), \cos(\mathbf{x}), \dots, \sin(2^{L-1}\mathbf{x}), \cos(2^{L-1}\mathbf{x})]^T. \quad (8)$$

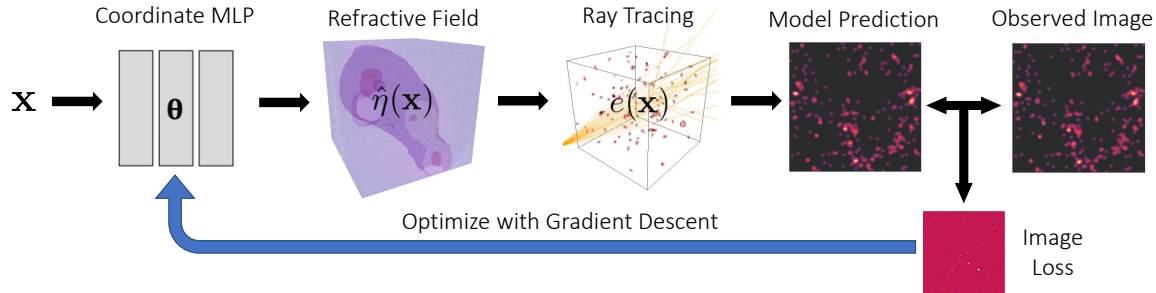


Figure 3. **Refractive Tomography Pipeline.** We model a 3D RI field  $\hat{\eta}(\mathbf{x})$  as a continuous function using a neural network parameterized by  $\theta$ . This refractive field, along with the known visible emission field  $e(x)$  from point-like light sources, is fed through a ray tracing simulator to produce a predicted 2D image measurement. An image loss is then taken between the model and true image measurements to solve for  $\theta$ . We differentially trace rays through this refractive field and accumulate the light emitted along these curved ray paths to produce a 2D image. This allows us to optimize MLP parameters  $\theta$  to minimize the loss between the rendered and true image measurements.

Adding this positional encoding allows the MLP to fit continuous fields with higher spatial frequencies as the degree  $L$  is increased [36]. We find empirically that a degree of  $L = 4$  is suitable for naturally smooth refractive fields.

Ray tracing through our model field requires integrating equation (5) during each training iteration. This requires evaluating the model  $\hat{\eta}_\theta$  and its spatial gradient  $\nabla \hat{\eta}_\theta$  at continuous spatial locations over the course of optimization, and for both of these functions to be differentiable with respect to the model parameters  $\theta$  for gradient descent. In other words, we need to be able to evaluate  $\nabla_\theta \nabla_x \hat{\eta}_\theta$  at arbitrary points throughout the modeled refractive field. We use the eLU activation function rather than a ReLU so that this gradient is smooth during optimization, and can be easily evaluated with existing autograd frameworks.

## 4.2. Optimization

Our neural field model defines a continuous refractive field  $\hat{\eta}_\theta$  and its gradient  $\nabla_x \hat{\eta}_\theta$  parameterized by its weights  $\theta$ . Thus, a rendered image  $\hat{I}$  obtained by ray-tracing through a neural field  $\hat{\eta}_\theta$  with Eqs. (5) and (6) is a function of  $\theta$ .

An overview of our pipeline is shown in Fig. 3. To train the neural network weights  $\theta$ , we solve a minimization problem whose objective is a mean-squared error between the observed image  $I(i, j)$ , and the rendered image  $\hat{I}(i, j; \theta)$ , plus a regularization term  $\mathcal{R}(\theta)$  to enforce a boundary condition of  $\eta = 1$  on the cube’s faces:

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{i,j} \|I(i, j) - \hat{I}(i, j; \theta)\|_2^2 + \lambda \mathcal{R}(\theta). \quad (9)$$

This loss is fully differentiable with respect to model parameters  $\theta$ , allowing us to use gradient descent to solve the optimization. The term  $\mathcal{R}(\theta)$  acts only as a constraint on the boundary conditions of the volume; unlike explicit regularizers, such as total variation, we rely on the implicit regu-

larization from the neural field for our reconstruction. More details on the  $\mathcal{R}(\theta)$  term can be found in the supplement.

## 4.3. Implementation Details

In our experiments, we use an MLP with 4 layers, where each layer is 256 units wide. The neural network is implemented in JAX [7] and trained with the Adam optimizer [17] with exponential learning rate decay from  $10^{-4}$  to  $5 \times 10^{-6}$  over 10K iterations. The weights of the neural network were randomly initialized with the He uniform variance scaling initializer [14]. Code implementation will be publicly available on the project page.

**Unmodeled Sources of Noise** Because our method leverages information about the location and shape of light sources to enable reconstruction from a single viewpoint, it is sensitive to noise in both our knowledge of light source locations and in image measurements. An analysis of our method’s sensitivity to these noise sources can be found in the supplement. We plan to investigate ways to overcome this sensitivity in future work.

## 5. Results

In this section we present results of our method on general refractive tomography problems. We generate ground truth refractive fields inside of a three-dimensional volume, where the RI takes a constant value of 1 outside the region of interest. Each ground truth refractive field is defined by evaluating a continuous function at regularly spaced grid points. The continuous field is then evaluated by linearly interpolating between these grid points. Image measurements are rendered using a pinhole camera with a  $64^2$  resolution. Throughout this section, we quantify the accuracy of our reconstructions by sampling the ground truth and reconstructed fields on the same uniformly spaced grid of sample points as the ground truth and report the PSNR. When calculating the PSNR, we first subtract 1 from the RI.

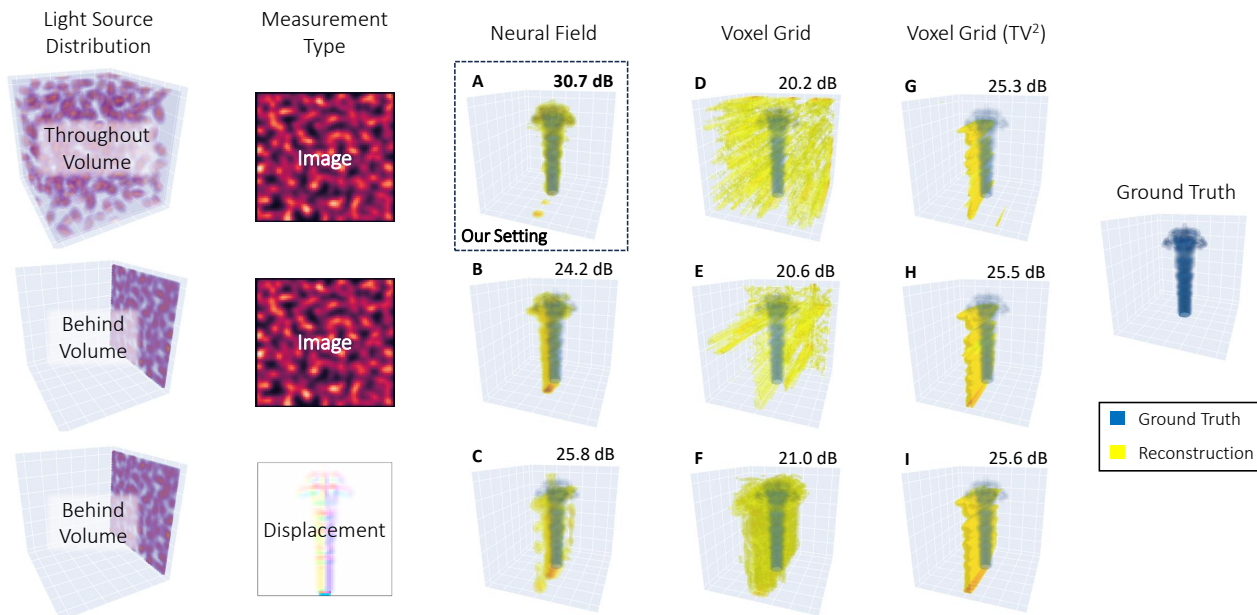


Figure 4. **Measurement and Reconstruction Conditions.** We compare our neural field reconstruction model to a linear voxel grid with and without  $TV^2$  regularization for gas flow reconstruction under three different measurement conditions. While single-view recovery is difficult with measurements of ray displacement, we find that image measurements of light sources spread throughout the refractive volume makes recovery possible with our method. In addition, a neural field’s implicit regularization results in better-constrained reconstructions than the linear voxel grid, even when an explicit regularizer is used. In each recovery we super-impose the ground-truth gas plume in blue over the reconstruction in yellow for comparison. We note that all reconstructions shown fit the measurements almost perfectly; differences in their structures highlight the ill-posedness of the problem.

### 5.1. Measurement and Reconstruction Conditions

In this section we compare refractive reconstructions under nine conditions: three sets of measurement conditions paired with three different reconstruction strategies. We show that single-view RI tomography is made possible by a combination of the measurement conditions of our problem setting and the implicit regularization given by our neural field model. Results are shown in Fig. 4.

*Measurement Conditions:* The first measurement conditions mirror the astronomical domain that motivates our work, in which image measurements are taken of light sources distributed throughout the refractive volume (Fig. 4 top row). Secondly, we compare to image measurements taken of light sources behind the volume (middle row). Third, we investigate using ray displacement measurements rather than image measurements; in this case, the displacements correspond to a plane of light sources behind the volume (bottom row). To simulate the displacement measurements, we ray trace through the ground truth refractive field to recover a perfect optical flow map at the same resolution as our sensor. Although this would be infeasible in practice, we use these measurements to represent the ideal limit of measurements from a single-view BoS setup [4].

*Reconstruction Conditions:* We also compare reconstructions from three models: our neural field representa-

tion, a linear voxel grid, and a linear voxel grid with  $TV^2$  regularization, with regularizer weight chosen to maximize PSNR with the ground truth volume. These correspond to the last three columns of Fig. 4, respectively. The goal of the regularized voxel grid is to show that our method can outperform even the highest-performing  $TV^2$  model; in practice, it would be infeasible to choose the regularizer weight in this way as we don’t have access to the ground truth.

**Data & Results** In Fig. 4 we demonstrate the effect of each pair of conditions in reconstructing the fuel injection dataset (a  $64^3$  RI grid) from SFB 382 of the German Research Council (DFG). Both a method using an unregularized voxel grid as well as BoS have been shown to reconstruct this volume given access to multiple viewpoints [4, 37]. We report the PSNR for each reconstruction, but note that it isn’t fully representative of reconstruction quality, which should also be evaluated visually.

*Effect of Measurement Conditions:* When optimizing our neural field model on image measurements of light sources located throughout the refractive volume, accurate 3D recovery of the gas flow is possible with minimal artifacts (A). However, in (B), where the light sources are instead located behind the volume, the neural field is able to recover the rough structure of the gas flow but not its absolute location. When using displacement measurements, none of the reconstructions (C, F, I) are able to correctly recover the shape




Figure 5. **Recovery of Smooth Refractive Fields.** Each row corresponds to a refractive tomography problem with a different random refractive field against a random distribution of 290 light sources. Although refractive recovery from a single viewpoint is highly ill-posed, our method can leverage source location information to detect and localize 3D refractive structure. However, as can be seen in the third example, due to the ill-posed and highly non-convex nature of this tomography problem, our model is still be susceptible to error. The yellow cones indicate sensor direction.

of the gas flow from a single viewpoint, even with a perfect optical flow map. While it has been shown that voxel grid optimization with displacement measurements can accurately reconstruct this volume [37], having access to only a single viewpoint causes the reconstruction to be smeared along the optical axis. An analysis of the voxel grid recovery with access to additional viewpoints can be found in the supplement. We conclude that using image measurements of light sources spread throughout the volume is a key ingredient enabling 3D reconstruction from a single viewpoint.

*Effect of Reconstruction Conditions:* An image-based loss on point-like light sources is highly nonconvex, and thus models such as an unregularized voxel grid can be prone to converging to bad minima. Although all reconstructions in Fig. 4 fit the measurements almost perfectly, each converges to a different refractive field. The unregularized grid reconstructions (D, E) optimize locally along ray paths to fit the image measurements, resulting in streak-like artifacts and a highly non-smooth recovery. Even when an explicit  $TV^2$  regularizer is used on the voxel grid, as in

recoveries (G) and (H), it still fails to recover the structure of the gas flow along the optical axis. The implicit regularization of our neural field model appears to favor compact reconstructions, which matches the gas flow volume well. In the following section, we show that this implicit regularization also allows us to recover more general smooth fields.

## 5.2. Recovery of Smooth Refractive Fields

We demonstrate our approach by recovering randomly generated smooth refractive fields in a dense field of light sources from a single viewpoint. We use Poisson Disk sampling [8] to generate an emission field of 290 light source locations. Light sources are amplitude scaled 3D Gaussians with independent random elliptical orientations. Ground truth refractive fields are generated from a random Gaussian Process with a smooth covariance kernel described in the supplement. These true fields were then scaled to an RI range of 1 to 1.003. To put this RI magnitude into context, the refractive effect of these fields causes a median deviation of 0.03 to 0.07 pixels on a  $64^2$  pixel image for rays shot parallel through the volume from the front to the back.

Fig. 5 shows that our method has the potential to accurately reconstruct complex, smooth fields. However, due to the ill-posedness and highly non-convex nature of this problem, our model is still susceptible to reconstruction artifacts, as seen in the last row. Note that no explicit regularization besides a boundary condition term was used in solving for these refractive fields. Thus, our results suggest that our neural network architecture combined with a low degree of positional encoding act as an implicit prior on the smoothness of our reconstruction, allowing it to effectively recover smooth refractive fields.

## 5.3. Sensitivity to Light Source Density

From a single viewpoint, refractive fields are detected when they lie between light sources and the sensor. Background refractive structures become difficult to detect with a sparse array of light sources since fewer light sources lie behind them. Because we supervise our model with curved ray paths throughout the scene, our model uses the deflections between adjacent light sources to localize the refractive field. Therefore, we expect that as available light sources in the scene become sparser, our model will struggle to detect and localize the structure of a refractive field.

We perform a refractive field reconstruction on a scene containing five elliptical refractive objects in light source emission fields of varying densities. We sample 250 light source locations uniformly, and create emission fields of decreasing density by taking random subsets of this initial sample. We find that reconstruction quality of our approach degrades with sparser light source fields, especially in the background portion of the field. However, even with only 50 light sources in the volume, our method is able to accu-

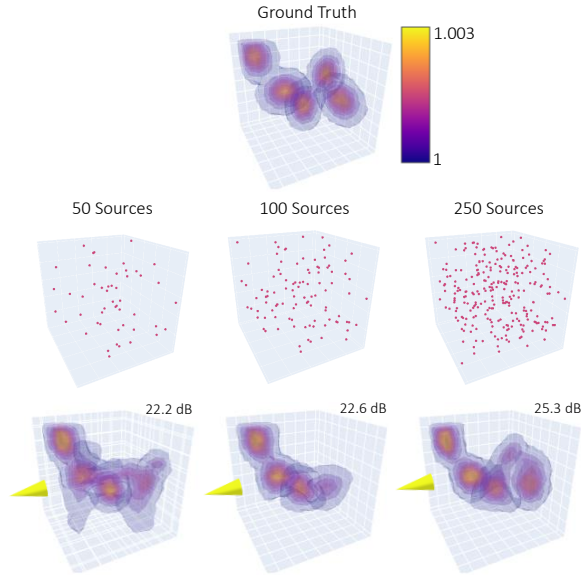


Figure 6. **Sensitivity to Light Source Density.** We analyze our approach’s reconstruction of a refractive field given light source fields with a range of densities. As the number of light sources decreases, it becomes harder to detect and localize refractive objects, especially towards the back of the volume. However, even with sparse light sources, our method is able to accurately resolve refractive objects in the foreground. The yellow cones indicate sensor direction.

rately resolve the structure of refractive objects in the foreground. Results are shown in Fig. 6.

## 6. Case Study: Dark Matter Mass Mapping

In this section we show an application to astronomical imaging of dark matter from simulations. In outer space, clumps of invisible mass called dark matter permeate the cosmic web. These clumps, called halos, are so massive that their gravitational pull warps the surrounding space, which lenses nearby passing rays of light. In the weak lensing regime, the effects of gravitational lensing are well-described by refraction from an “effective” RI field [26]. In fact, the degree of curvature in weak lensing is so small that its effect on galaxy images is often modeled by a rotation (rather than a translation) of the galaxy’s appearance on the sky [5]. By combining telescope images with photometric redshift measurements of the distances of galaxies from Earth, scientists are able to determine their approximate spatial locations. The inverse problem of dark matter mass mapping is to reconstruct an invisible distribution of dark matter given refracted images from telescope surveys and photometric measurements of galaxy distances [15].

### 6.1. Simulated Dark Matter Halo Recovery

We now use our method to recover a realistic dark matter refractive field. To generate this refractive field, we use the cosmological hydrodynamical simulation IllustrisTNG,

which includes both dark and baryonic matter [27, 29]. We use the catalog of dark matter halos produced by the TNG300-1 simulation, focusing on a  $20^3 \text{ Mpc}^3$  region containing 1747 massive halos. As galaxy placement is known to trace the dark matter distribution, we placed Gaussian emission functions emulating galaxies at the center of the 300 most massive halos. We aim to reconstruct a realistic dark matter distribution, but leave the inclusion of realistic measurement noise to future work.

RI tomography in the weak lensing regime requires sensors with sufficient resolution to detect the small shearing effects from gravitational lensing. In our experiments, we use an increased sensor resolution of  $512 \times 512$ ; this resolution in comparison with the scale of our scene is consistent with that of current and upcoming weak lensing surveys [18, 35]. We train our model using SGD, sampling batches of 4096 rays per iteration due to memory constraints.

Our reconstruction of the invisible dark matter field is shown in Fig. 1. Our method is able to accurately recover the structure of the dark matter field with slight damping on the largest peaks of halo centers. This recovery is especially challenging because of the distribution of galaxies throughout the scene; large patches of the image are dark as galaxies are not uniformly distributed throughout the volume. In addition, the amplitude of the refractive field generated by these halos is three orders of magnitude lower than in our previous results, resulting in a weak refractive signal; the median deflection caused by the refractive field is  $6e-3$  pixels. We find that despite these challenges, our method is able to recover a refractive field that correlates strongly with true simulated dark matter halos, and shows promise for application to 3D dark matter recovery in the future.

## 7. Conclusion

In this paper we present a novel method for the tomography of continuous refractive fields from single-view image measurements with neural fields. Our method succeeds at reconstructing simulated refractive fields with varying light source density. By using the known location of light emitters, our method reconstructs refractive volumes from a single viewpoint that rival those of multiview methods. Finally, our results on physically realistic simulations of dark matter show promise for the application of our method to real weak lensing survey data for dark matter mapping. Beyond dark matter fields, our method opens doors to reconstructing other refractive phenomena, underscoring the computer vision community’s potential to drive significant advances in scientific discovery with computational imaging.

**Acknowledgements** BZ was partially supported by the Gupta Fellowship. This work was supported by NSF award 2048237, Google, BP, and AWS Credits from the AI4Science Partnership. We also thank Arjun Teh, Adithya Pediredla and Andrew Chael for their helpful discussions.



## References

- [1] Marina Alterman, Yoav Y Schechner, Minh Vo, and Srinivasa G Narasimhan. Passive tomography of turbulence strength. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 47–60. Springer, 2014. [2](#)
- [2] Marina Alterman, Yohay Swirski, and Yoav Y Schechner. Stella maris: Stellar marine refractive imaging sensor. In *2014 IEEE International Conference on Computational Photography (ICCP)*, pages 1–10. IEEE, 2014. [2](#)
- [3] Marco Ament, Christoph Bergmann, and Daniel Weiskopf. Refractive radiative transfer equation. *ACM Transactions on Graphics (TOG)*, 33(2):1–22, 2014. [4](#)
- [4] 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. [2](#), [3](#), [4](#), [6](#)
- [5] Matthias Bartelmann and Peter Schneider. Weak gravitational lensing. *Physics Reports*, 340(4-5):291–472, 2001. [8](#)
- [6] Mojtaba Bemana, Karol Myszkowski, Jeppe Revall Frisvad, Hans-Peter Seidel, and Tobias Ritschel. Eikonal fields for refractive novel-view synthesis. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. [3](#), [4](#)
- [7] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. [5](#)
- [8] Robert Bridson. Fast poisson disk sampling in arbitrary dimensions. *SIGGRAPH sketches*, 10(1):1, 2007. [7](#)
- [9] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. [4](#)
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. [3](#)
- [11] Wonshik Choi, Christopher Fang-Yen, Kamran Badizadegan, Seungeun Oh, Niyom Lue, Ramachandra R Dasari, and Michael S Feld. Tomographic phase microscopy. *Nature methods*, 4(9):717–719, 2007. [3](#)
- [12] Shwetadwip Chowdhury, Michael Chen, Regina Eckert, David Ren, Fan Wu, Nicole Repina, and Laura Waller. High-resolution 3d refractive index microscopy of multiple-scattering samples from intensity images. *Optica*, 6(9):1211–1219, 2019. [3](#)
- [13] John R Dormand and Peter J Prince. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980. [4](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. [5](#)
- [15] Henk Hoekstra and Bhuvnesh Jain. Weak gravitational lensing and its cosmological applications. *Annual Review of Nuclear and Particle Science*, 58:99–123, 2008. [1](#), [3](#), [4](#), [8](#)
- [16] Ivo Ihrke, Gernot Ziegler, Art Tevs, Christian Theobalt, Marcus Magnor, and Hans-Peter Seidel. Eikonal rendering: Efficient light transport in refractive objects. *ACM Transactions on Graphics (TOG)*, 26(3):59–es, 2007. [3](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [18] Rene Laureijs, J Amiaux, S Arduini, J-L Augeres, J Brinckmann, R Cole, M Cropper, C Dabin, L Duvet, A Ealet, et al. Euclid definition study report. *arXiv preprint arXiv:1110.3193*, 2011. [8](#)
- [19] Adrienne Leonard, F-X Dupé, and J-L Starck. A compressed sensing approach to 3d weak lensing. *Astronomy & Astrophysics*, 539:A85, 2012. [3](#)
- [20] 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](#)
- [21] Aviad Levis, Pratul P Srinivasan, Andrew A Chael, Ren Ng, and Katherine L Bouman. Gravitationally lensed black hole emission tomography. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19841–19850, 2022. [2](#), [3](#), [4](#)
- [22] Zongling Li, Qingyu Hou, Zhipeng Wang, Fanjiao Tan, Jin Liu, and Wei Zhang. End-to-end learned single lens design using fast differentiable ray tracing. *Optics Letters*, 46(21):5453–5456, 2021. [3](#)
- [23] Fan-Chi Lin, Michael H Ritzwoller, and Roel Snieder. Eikonal tomography: surface wave tomography by phase front tracking across a regional broad-band seismic array. *Geophysical Journal International*, 177(3):1091–1110, 2009. [1](#)
- [24] Renhao Liu, Yu Sun, Jiabei Zhu, Lei Tian, and Ulugbek S Kamilov. Recovery of continuous 3d refractive index maps from discrete intensity-only measurements using neural fields. *Nature Machine Intelligence*, 4(9):781–791, 2022. [3](#)
- [25] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [2](#), [3](#), [4](#)
- [26] Ramesh Narayan and Matthias Bartelmann. Lectures on gravitational lensing. *arXiv preprint astro-ph/9606001*, 1996. [8](#)
- [27] Dylan Nelson, Annalisa Pillepich, Shy Genel, Mark Vogelsberger, Volker Springel, Paul Torrey, Vicente Rodriguez-Gomez, Debora Sijacki, Gregory F Snyder, Brendan Griffen, et al. The illustris simulation: Public data release. *Astronomy and Computing*, 13:12–37, 2015. [2](#), [8](#)
- [28] Adithya Pediredla, Yasin Karimi Chalmiani, Matteo Giuseppe Scopelliti, Maysamreza Chamanzar, Srinivasa Narasimhan, and Ioannis Gkioulekas. Path tracing estimators for refractive radiative transfer. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. [4](#)

- [29] Annalisa Pillepich, Volker Springel, Dylan Nelson, Shy Genel, Jill Naiman, Rüdiger Pakmor, Lars Hernquist, Paul Torrey, Mark Vogelsberger, Rainer Weinberger, et al. Simulating galaxy formation with the illustrisng model. *Monthly Notices of the Royal Astronomical Society*, 473(3):4077–4106, 2018. [8](#)
- [30] Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. CRC press, 1987. [4](#)
- [31] William A Reed, Man F Yan, and Mark J Schnitzer. Gradient-index fiber-optic microprobes for minimally invasive in vivo low-coherence interferometry. *Optics letters*, 27(20):1794–1796, 2002. [1](#)
- [32] Peter Schneider, Jürgen Ehlers, Emilio E Falco, et al. Gravitational lenses [electronic resource]. [4](#)
- [33] 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. [3](#)
- [34] Yongjin Sung, Wonshik Choi, Christopher Fang-Yen, Kamran Badizadegan, Ramachandra R Dasari, and Michael S Feld. Optical diffraction tomography for high resolution live cell imaging. *Optics express*, 17(1):266–277, 2009. [3](#)
- [35] Masahiro Takada. Subaru hyper supprime-cam project. In *AIP Conference Proceedings*, pages 120–127. American Institute of Physics, 2010. [8](#)
- [36] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. [4](#), [5](#)
- [37] Arjun Teh, Matthew O’Toole, and Ioannis Gkioulekas. Adjoint nonlinear ray tracing. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. [2](#), [4](#), [6](#), [7](#)
- [38] Ethan Tseng, Ali Mosleh, Fahim Mannan, Karl St-Arnaud, Avinash Sharma, Yifan Peng, Alexander Braun, Derek Nowrouzezahrai, Jean-Francois Lalonde, and Felix Heide. Differentiable compound optics and processing pipeline optimization for end-to-end camera design. *ACM Transactions on Graphics (TOG)*, 40(2):1–19, 2021. [3](#)
- [39] 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](#)
- [40] Tianfan Xue, Michael Rubinstein, Neal Wadhwa, Anat Levin, Fredo Durand, and William T Freeman. Refraction wiggles for measuring fluid depth and velocity from video. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part III 13*, pages 767–782. Springer, 2014. [2](#)