

Single View Refractive Index Tomography with Neural Fields: Supplement

Brandon Zhao^{1*} Aviad Levis¹ Liam Connor² Pratul P. Srinivasan³ Katherine L. Bouman^{1,2}

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

² Cahill Center for Astronomy and Astrophysics, California Institute of Technology

³Google Research

*byzhao@caltech.edu

Contents

1. Robustness to Strong Refraction	2
2. Sensitivity to Noise	2
2.1. Image-Level Noise	2
2.2. Uncertainty in Light Source Orientation	3
3. Multi-View Recovery with Voxel Grids	4
4. Boundary Regularization	5
5. Generating Random Refractive Fields	5
6. Refractive Radiative Transfer Equation	6
7. Recovery with Light Source Grid	8

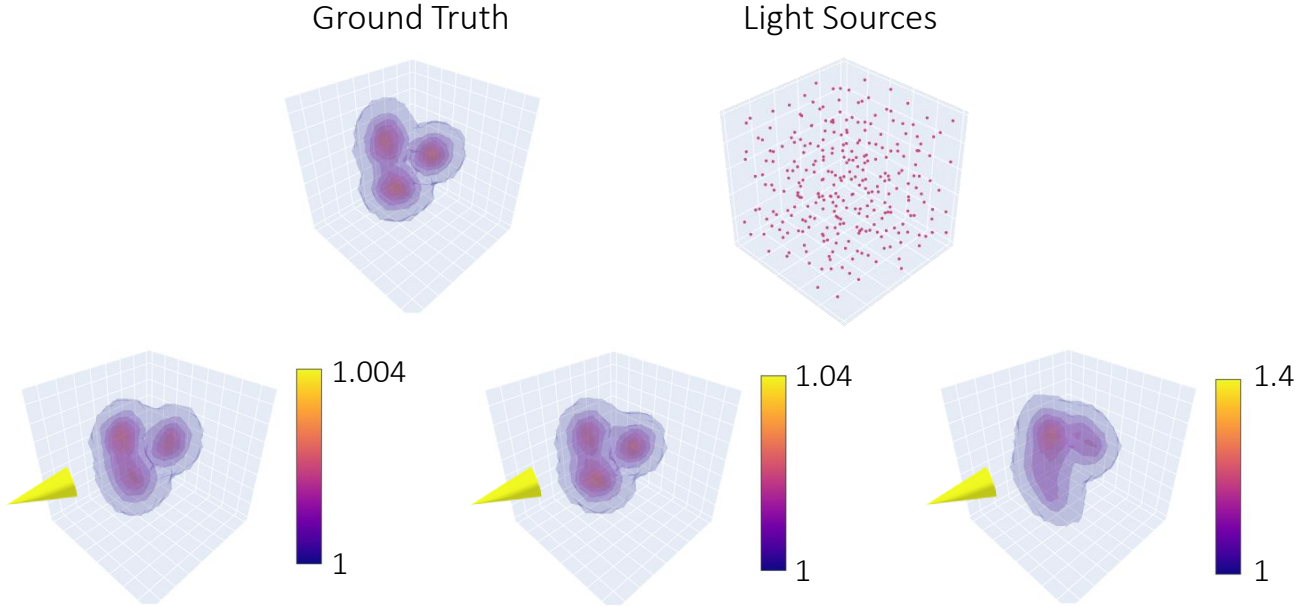


Figure 1. **Robustness to Strong Refraction:** Our approach explicitly models the continuous nature of refraction, and therefore curving ray paths, within an optimization framework. Using a continuous model that allows for curving ray paths enables robust reconstructions of refractive fields with a wide range of magnitude. We highlight how our approach is able to reconstruct the underlying refractive field across three orders of magnitude. As the magnitude of refraction increases the problem becomes more ill-posed due to the potential for large deviations in ray paths, leading to the observed reconstruction degradation in the high magnitude case. The yellow cones indicate sensor direction.

1. Robustness to Strong Refraction

Typically, refraction tomography (e.g. Schlieren tomography [2]) recoveries rely on estimation of the total deflection of light rays entering and exiting a refractive volume. The underlying assumptions are that light is refracted once and that ray paths are approximately linear. Although these assumptions are valid for refractive fields with a refractive index (RI) concentrated at a small volume (e.g. transparent gas flows) they become inaccurate in large volumes with a high RI. In contrast, our method explicitly models the continuous curvature of ray paths throughout the volume. By doing so our approach is able to reconstruct strongly refractive fields where ray paths are continuously curving and do not adhere to the approximately linear model or single refraction. We illustrate the strength of our approach in its ability to recover a range of refractive fields with increasing magnitudes. Figure 1 shows that our method is able to accurately recover the volume in all cases with minimal artifacts. However, it is important to note that as the magnitude of refraction increases the problem becomes more ill-posed due to the increased potential for crossing ray paths; this leads to some degradation of results when the magnitude of refraction is large.

2. Sensitivity to Noise

Our reconstruction method uses image measurements and knowledge of the locations and shapes of light sources to reconstruct the 3D refractive medium. In this section, we examine the effects of different noise sources. The first is additive pixel noise that affects image intensity measurements. The second is noise in our estimate of light source properties (position/orientation). We analyze the effects these noise sources would have on the reconstruction with a range of noise levels.

2.1. Image-Level Noise

To simulate sensor noise in our image measurements, we add white Gaussian noise to each pixel measurement. Reconstruction is performed on a RI field consisting of five elliptical Gaussian “blobs”. Light sources are modeled as randomly oriented elliptical Gaussians uniformly distributed throughout the volume. Each light source emits a maximum of ~ 1 unit of radiance

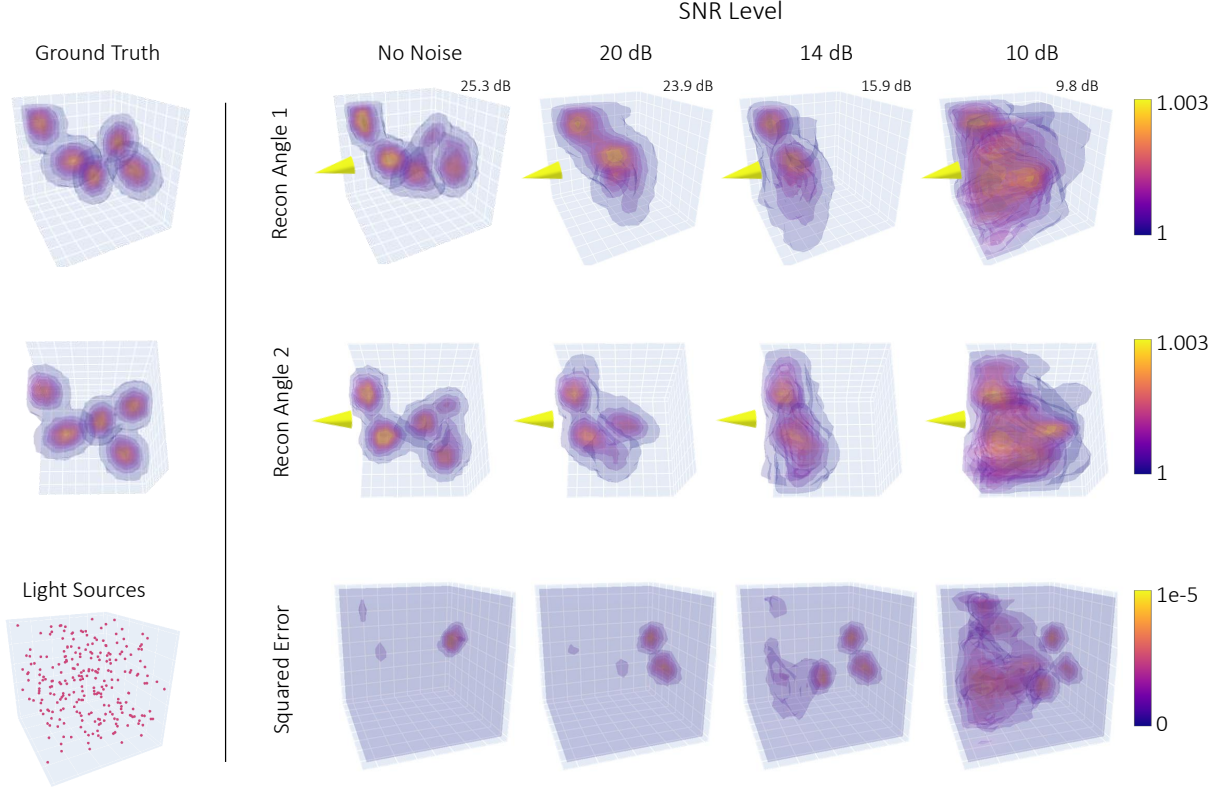


Figure 2. **Image-Level Noise:** In these simulations, we highlight the effect of image measurement noise on the reconstruction of a weakly refractive field. The left column shows the ground-truth refractive field together with the distribution of light sources. Each reconstruction (each column) highlights the effects of a different measurement SNR level on the reconstructed volume. The volumetric error plot highlights the non-uniform effects of image noise resulting in recovery error increasing with distance from the sensor. Note that no explicit spatial regularization is being used in reconstruction. The yellow cones indicate sensor direction.

resulting in an image signal-to-noise ratio (SNR) of approximately: $\text{SNR} = 10 \log(\sigma_{\text{noise}})$. To prevent over-fitting to the measurement noise, we rely on early stopping. We stop optimization when the data fit term in the optimization loss reaches the expected noise mean square error (MSE). The underlying assumption of using this criterion is that we have a reasonable model for the additive sensor noise.

In Figure 2 we show results for three SNR levels: 20, 14, 10 [dB] which are typical for current observational instruments [5, 6]. We highlight the impact of image noise on the reconstructions from two viewing angles (top rows), as well as the error volume comparing the recovery to the ground truth (bottom row). At higher SNR our approach is able to reconstruct several refractive Gaussian blobs with minor blurring artifacts. Lower SNR levels impact the recovered structures resulting in reconstruction artifacts. In this work we show a general reconstruction method as a first step that does not rely on explicit regularization. We expect that future applications incorporating specific informative priors could help to mitigate these noise artifacts. Furthermore, although many galaxies (light sources) lie within the $\text{SNR} \approx 10$ regime, there are many brighter or closer sources in a higher SNR regime that could be used for a reliable albeit partial reconstruction of the universe.

2.2. Uncertainty in Light Source Orientation

Our approach relies on knowing the shapes and positions of the light sources. In the astronomical setting, we have access to an estimated distance through spectral observations (further sources appear redder due to an effect called cosmological red-shift). Furthermore, weak refraction causes small distortions in the image plane that are best approximated as rotations of an elliptical source (rather than translations). Therefore, in this section, we analyze the dominant noise source which comes from the uncertainty in the source orientation.

To analyze this effect we reconstruct a refractive field (see Sec. 2.1) with randomly perturbed light source orientations.

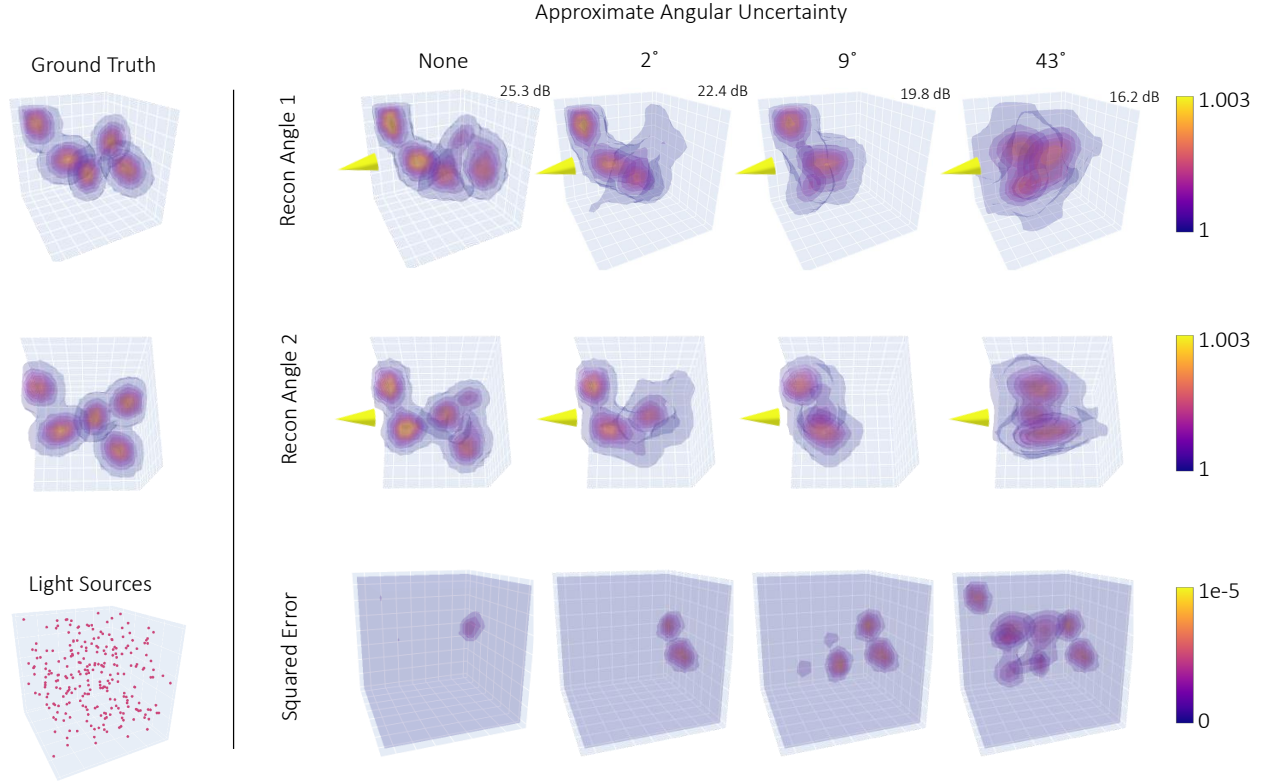


Figure 3. **Sensitivity to Noise in Light Source Orientation:** In these simulations, we highlight the effect of uncertainty in the orientation of each light source on the reconstruction. The left column shows the ground-truth refractive field together with the distribution of light sources. Each reconstruction (each column) corresponds to a different level of uncertainty in the light source orientation. Due to symmetry, a 90° rotation is the maximum orientation difference. The volumetric error plot highlights the non-uniform effects of source orientation noise resulting in increasing recovery error with distance from the sensor. Note that no explicit spatial regularization is being used in reconstruction. The yellow cones indicate sensor direction.

Each light source is randomly rotated in the image plane (about the depth axis) at an angle sampled from a mean-zero Gaussian distribution. The noisy orientations are subsequently used to generate observations for reconstruction. Analogously to the early stopping criteria described in Sec. 2.1 we prevent over-fitting to noise by estimating the image noise induced by the noisy orientation. We estimate the expected rotation MSE by comparing 10 random sampled orientations to an unrotated image. This empirical MSE estimate is used as a stopping criterion for the optimization.

In figure 3 we show results for three levels of uncertainty in the source orientation corresponding to an uncertainty of $\sim 2^\circ$, $\sim 9^\circ$, and $\sim 43^\circ$, where the uncertainty is given by a range of ± 3 standard deviations of the orientation noise distribution. Fig. 3 highlights that our approach is able to reconstruct the underlying structure well given good estimates of light source orientations. As expected, the reconstruction quality degrades with a degree of uncertainty. Note that again no explicit regularization is being used in these reconstructions.

3. Multi-View Recovery with Voxel Grids

In Sec. 5.1 of the main paper we highlight that single-view RI tomography is made possible by a combination of specific measurement conditions and our neural field reconstruction strategy. We show single-view reconstruction results on a synthetic fuel injection volume; both a method using an unregularized voxel grid as well as BoS techniques have been shown to reconstruct this volume given access to 32 viewpoints [8] and displacement measurements. In this section we reproduce multi-view reconstruction results for voxel grid optimization on this volume, varying the number of viewpoints from single to multi-view. In particular, as the number of viewpoints is decreased, the voxel grid reconstruction suffers from projection

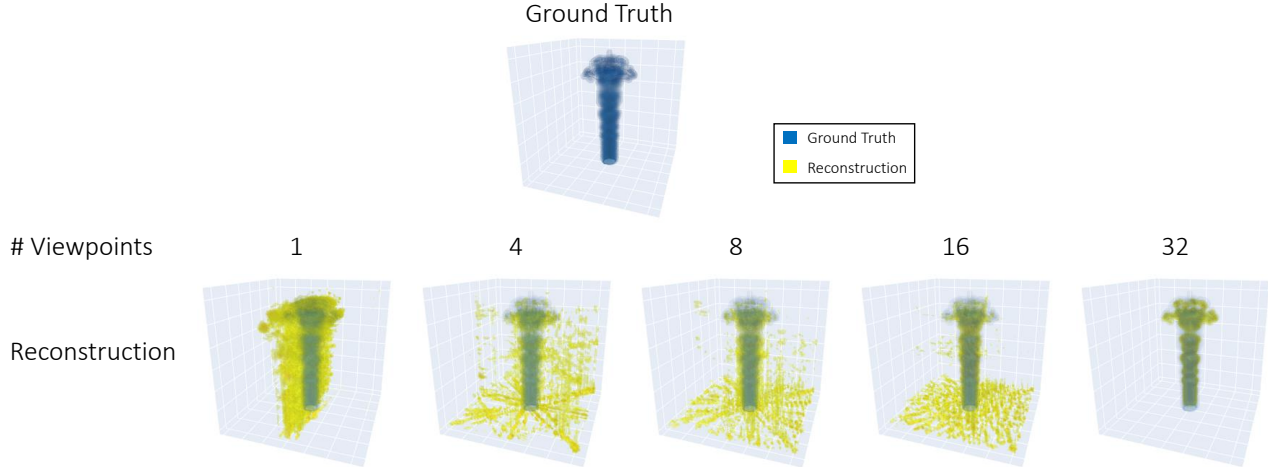


Figure 4. **Multiview Recovery with Voxel Grids:** In these simulations, we analyze the reconstruction performance of an unregularized voxel grid on a fuel injection RI volume with a varying number of viewpoints. With fewer than 32 viewpoints, reconstructions suffer from projection artifacts due to the underconstrained nature of the recovery. In particular, for a single viewpoint, the reconstruction is smeared along the optical axis. In each recovery we super-impose the ground-truth gas plume in blue over the reconstruction in yellow for comparison. All reconstructions shown fit the measurements almost perfectly; differences in the structures are due to the ill-posedness of the problem.

artifacts as the problem is underconstrained. Only with a large number of viewpoints can the unregularized voxel grid reconstruct the volume accurately with few artifacts, which motivates the use of an alternate strategy for single-view recovery. Results are shown in Fig. 4.

4. Boundary Regularization

In our simulations, we reconstruct refractive volumes assuming a constant value of 1 outside of their boundary. In this way, we avoid modeling refraction outside of the volume of interest. In our reconstructions, we regularize the predicted field boundary to have a value of 1 by minimizing the following loss function with respect to network parameters θ :

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{i,j} \|I(i,j) - \hat{I}(i,j,\theta)\|_2^2 + \sum_{k=1}^n (\hat{\eta}_{\theta}(\mathbf{x}_k) - 1)^2. \quad (1)$$

The second term in Eq. (1) enforces the model to output 1 at points \mathbf{x}_k chosen from a uniform grid on the volume boundary.

5. Generating Random Refractive Fields

In Section 4.1, we show reconstruction results on smooth, randomly generated refractive fields. We use a Gaussian random process to generate fields as follows: Let G be a mean-zero Gaussian process with the Matern kernel, with length scale 0.25 and smoothness parameter $\nu = 1.5$. We sample G at an evenly spaced $16 \times 16 \times 16$ grid of points within the unit cube, then pass these outputs through a soft-plus function to give non-negative values. Finally, we scale the soft-plus outputs, add 1, then pad these points at the cube boundaries with a constant value of 1. We then use trilinear interpolation to evaluate the refractive field and its spatial gradient for ray tracing. The refractive field η is given as:

$$\eta(\mathbf{x}) = \text{interp}(\sigma(G - 5) * \tau + 1, \mathbf{x}) \quad (2)$$

where σ is the soft-plus function output, G is the sampled Gaussian Process output, and τ is the scaling factor set to control the magnitude of the RI field. For Section 4.1 in the paper, we set τ to be 1/50, 1/50, and 1/90 for the first, second and third rows respectively.




Figure 5. **Recovery with the RRTE:** In many applications the RRTE forward model [1] more accurately describes how pixel intensities in the image plane should be calculated. Our method also works well in these cases. Yellow cones indicate sensor direction.

6. Refractive Radiative Transfer Equation

In our work we calculate the intensity of a given image pixel by first solving for refracted ray paths with backwards ray-tracing then integrating emission coefficients along the traced ray path according to the Radiative Transfer Equation (RTE) [3]. This equation (Eq. 4 in the paper) relies on the fact that specific intensity I , otherwise known as radiance, is conserved along a ray bundle traveling in free space. However, for many applications of refractive tomography, this is not exactly true; instead, a quantity called basic radiance, equal to $\frac{I}{\eta^2}$, is conserved along a ray path, according to the Refractive Radiative Transfer Equation (RRTE) [1]. Our method also works for this forward model; recovery results for the experiment given in paper Figure 4 but with the RRTE forward model are shown in Fig. 5.

In the regime of gravitational lensing, we approximate the effects of general relativity on rays of light as refraction. We make the simplifying assumptions that we are imaging a stationary refractive field, and that the effects of redshift on emitted intensities are negligible; similar assumptions have been made in previous works [4]. Under these assumptions, the regular




Figure 6. **Smooth Field Recovery with Light Source Grid:** In this experiment we analyze our method’s performance imaging a uniform grid of light sources rather than a random distribution. Recovery performance is similar to experiments with a random light source distribution.

RTE is sufficient to accurately calculate image pixel intensities [7, 9] ; thus, experiments in the main paper are performed with the RTE forward model.

7. Recovery with Light Source Grid

In this experiment, we analyze our method’s performance on recovering smooth random fields as generated in Fig. 5 of the main paper, but with an evenly spaced grid of light sources rather than a random distribution. Results are shown in Fig. 6. Recovery performance is similar to situations with random light source distributions; however, due to the grid structure of the light source distribution, background light sources overlap with those in the foreground in the image, making reconstruction of deeper parts of the field difficult.

References

- [1] Marco Ament, Christoph Bergmann, and Daniel Weiskopf. Refractive radiative transfer equation. *ACM Transactions on Graphics (TOG)*, 33(2):1–22, 2014. 6
- [2] Bradley Atcheson, Ivo Ihrke, Wolfgang Heidrich, Art Tevs, Derek Bradley, Marcus Magnor, and Hans-Peter Seidel. Time-resolved 3d capture of non-stationary gas flows. *ACM transactions on graphics (TOG)*, 27(5):1–9, 2008. 2
- [3] Subrahmanyan Chandrasekhar. *Radiative transfer*. Courier Corporation, 2013. 6
- [4] 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. 6
- [5] Rachel Mandelbaum, Hironao Miyatake, Takashi Hamana, Masamune Oguri, Melanie Simet, Robert Armstrong, James Bosch, Ryoma Murata, François Lanusse, Alexie Leauthaud, et al. The first-year shear catalog of the subaru hyper supprime-cam subaru strategic program survey. *Publications of the Astronomical Society of Japan*, 70(SP1):S25, 2018. 3
- [6] Masamune Oguri, Satoshi Miyazaki, Chiaki Hikage, Rachel Mandelbaum, Yousuke Utsumi, Hironao Miyatake, Masahiro Takada, Robert Armstrong, James Bosch, Yutaka Komiyama, et al. Two-and three-dimensional wide-field weak lensing mass maps from the hyper supprime-cam subaru strategic program s16a data. *Publications of the Astronomical Society of Japan*, 70(SP1):S26, 2018. 3
- [7] George B Rybicki and Alan P Lightman. *Radiative processes in astrophysics*. John Wiley & Sons, 1991. 7
- [8] Arjun Teh, Matthew O’Toole, and Ioannis Gkioulekas. Adjoint nonlinear ray tracing. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 4
- [9] Frederic H Vincent, Thibaut Paumard, Ericourgoulhon, and Guy Perrin. Gyoto: a new general relativistic ray-tracing code. *Classical and Quantum Gravity*, 28(22):225011, 2011. 7