

Gravitationally Lensed Black Hole Emission Tomography

Aviad Levis^{1*}, Pratul P. Srinivasan^{2*}, Andrew A. Chael^{3†}, Ren Ng⁴, Katherine L. Bouman¹

¹California Institute of Technology ²Google Research ³Princeton ⁴UC Berkeley

Abstract

Measurements from the Event Horizon Telescope enabled the visualization of light emission around a black hole for the first time. So far, these measurements have been used to recover a 2D image under the assumption that the emission field is static over the period of acquisition. In this work, we propose *BH-NeRF*, a novel tomography approach that leverages gravitational lensing to recover the continuous 3D emission field near a black hole. Compared to other 3D reconstruction or tomography settings, this task poses two significant challenges: first, rays near black holes follow curved paths dictated by general relativity, and second, we only observe measurements from a single viewpoint. Our method captures the unknown emission field using a continuous volumetric function parameterized by a coordinate-based neural network, and uses knowledge of Keplerian orbital dynamics to establish correspondence between 3D points over time. Together, these enable *BH-NeRF* to recover accurate 3D emission fields, even in challenging situations with sparse measurements and uncertain orbital dynamics. This work takes the first steps in showing how future measurements from the Event Horizon Telescope could be used to recover evolving 3D emission around the supermassive black hole in our Galactic center.

1. Introduction

The Event Horizon Telescope (EHT), which recently captured the first image of a black hole, opens the door to a new era in black hole research [42]. We are now able to image the direct environment around a black hole and resolve image features on the scale of an event horizon. This unprecedented resolution is achieved by linking radio telescopes across the globe to synchronously observe an astronomical source [1, 5, 11], a technique known as Very Large Baseline Interferometry [43]. Although a black hole itself is not visible, its imprint or “shadow” on the emission from the surrounding hot gas is what the EHT aims to image.

The two candidate black holes within reach of the EHT are M87* and Sagittarius A* (SgrA*). The first black hole

* Authors contributed equally to this work.

† NASA Hubble Fellowship Program, Einstein Fellow.




Figure 1. We formulate and present an approach for a novel tomography: recovering the emission distribution of flares in orbit around a black hole using observations captured from a single viewpoint over time. This is achieved by leveraging expected physics around a black hole: gravitational lensing and emission flow modeled by Keplerian orbital dynamics. In this example, observations are simulated telescope measurements that place a sparse set of constraints on the Fourier components of the 2D images over time (top row); The image formation model accounts for ray trajectories that are curved due to gravitational lensing. The orientation of the observed projection plane is indicated in green. Our method, *BH-NeRF*, recovers a representation of the underlying 3D emission from the sparse telescope measurements.

image of M87* [42] was constructed from measurements collected over an entire night, with the underlying assumption that its emission is static throughout acquisition. In contrast, SgrA*, the black hole in our Galactic center, is far more dynamic and displays strong flares of emission up to several times a day [19]. A potential explanation for these flares are compact bright emission spots [8] in orbit around the black hole, referred to as *hot-spots*.

In this work, we formulate and propose an approach for a novel tomography problem to reconstruct these dynamic events (Fig. 1). In contrast to previous works [4, 23, 27], which focus on recovering 2D image plane dynamics, our

work recovers the underlying dynamic 3D emission distribution (i.e., 4D spatiotemporal volume). However, reconstructing volumetric emission in this setting faces two salient challenges. The first challenge is that light rays which reach the EHT telescopes do not follow straight lines that regularly sample space. Instead, the rays follow curved paths due to the strong gravitational field around black holes. This effect, predicted by general relativity (GR), is known as *gravitational lensing* since massive objects, such as black holes, act as lenses and bend light. The second challenge is the EHT only observes from a single fixed viewpoint, as opposed to the multi-view observations acquired in other tomography settings.

In this paper, we introduce a novel tomography that leverages expected physics around a black hole to recover the time-varying volumetric emission from a single observational direction. In particular, we present a promising solution, BH-NeRF, that represents the 3D emission as a continuous function using a coordinate-based neural network and makes use of our knowledge of Keplerian orbit dynamics to relate 3D emission points across time.

Our approach builds upon recent advances in neural representations of radiance fields [34] and generalizes them to account for curved ray trajectories. Our experiments demonstrate promising first steps towards recovering the 3D dynamics around a black hole, a high-priority task for astronomy research over the upcoming decade [36].

2. Related Work

Tomographic 3D Reconstruction Tomography aims to recover densities of a medium from lower dimensional projection measurements. In medical imaging, where X-ray computed tomography (CT) [25] is extensively used for diagnostics, 2D projections from different orientations are used to recover a 3D volume. These CT approaches are a well established technology, partly due to the fact that simple photon-medium interactions result in straight-rays through the medium [20].

In many scientific fields, complex interactions result in non-straight light paths (e.g. scattering, refraction, lensing) and non-linear measurement models, rendering tomography a largely open problem for research. In Earth Science, atmospheric tomography using scattered light gives rise to broken-ray paths and non-linear measurements [28–30]. In Seismology, imaging the interior of the Earth involves resolving refraction and propagation along curved trajectories [18]. For optical imaging, a similar effect is caused by a medium with variable refractive index [3, 46]. In Astrophysics, curved light paths caused by gravitational lensing around high density regions in space has been used to constrain the 3D structure of dark matter [22]. Analogously, the strong gravitational field near a black hole causes light paths to curve (refer to Fig. 3).

In the context of black hole imaging, the work of Tiede *et al.* [44] showed recovery of a handful of parameters related to orbiting hot-spots using Monte-Carlo sampling. While limited to a parameterized model, their work showed promise in recovering 3D information from 2D projections. Our work is inherently different from [44] as it is not limited to a simple parameteric form. Rather, we aim to recover a *continuous* unknown 3D volume.

Coordinate-based Neural Representations A recent trend in computer vision and graphics is to replace traditional discrete representations of geometry, such as triangle meshes and voxel grids, with continuous “coordinate-based” neural representations [13, 33, 37]. This approach parameterizes continuous 3D fields with a multilayer perceptron (MLP) that maps from an input 3D coordinate to a representation of scene content at that location [34].

The use of neural representations for rendering unobserved views of scenes was popularized by Neural Radiance Fields (NeRF) [34], which outputs the volume density and view-dependent radiance at any continuous 3D point in the scene. Our work is related to extensions of NeRF to dynamic scenes [17, 32, 38, 39, 45]. In dynamic settings, a key challenge is relating 3D points across time so that the scene can be represented by deformations of a canonical 3D representation. Previous methods use priors on deformation sparsity and rigidity [38, 45], pre-trained monocular depth estimation [32], or explicit models of human faces and bodies [17, 39] to relate 3D points across time. In our work, we leverage knowledge of Keplerian orbit dynamics to model the dynamic evolution around a black hole.

We take inspiration from the success of coordinate-based neural representations in other scientific imaging domains, such as cryo-electron microscopy [47] and X-ray tomography [41], and use it to represent the evolving 3D emission around a black hole as described in the following sections.

3. Image & Measurement Formation Model

In this section, we discuss the three components of our “forward model”, which models measurements observed by the EHT as a function of the emission surrounding a black hole: 1) the emission dynamics that describe the evolution over time, 2) the gravitationally lensed ray tracing that describes how 3D emission is integrated along curved rays to form a 2D image, and 3) the measurement model that describes how EHT measurements relate to these images.

3.1. Emission Dynamics

We represent emission, $e(t, \mathbf{x})$, as a continuous function of time t and 3D coordinate $\mathbf{x} = (x, y, z)$. Emission around a black hole can be the result of either inflowing plasma from an accretion disk or outflowing plasma from a jet. Here, we assume the emission originates in the equatorial




Figure 2. An illustration of a Keplerian velocity field about a rotation axis ξ , where the velocity at 3D locations is depicted by arrows whose length signifies the velocity’s magnitude. The velocity magnitude decreases with distance, as described by Eq (1). Our approach does not assume knowledge of the axis ξ which is jointly estimated with the underlying 3D emission (Sec. 4.2).

plane, where the temporal dynamics are modeled by a Keplerian angular velocity $\omega(r)$ that decreases with distance r from the black hole center [9]. In particular, the angular frequency of an orbit at radius r around a black hole with mass M and zero angular momentum is given by:

$$\omega(r) = \frac{1}{2\pi\sqrt{r^3/GM}} \propto r^{-3/2}. \quad (1)$$

where G is the gravitational constant. The extension of (1) for cases where the black hole angular momentum (i.e., spin) is nonzero is given in the supplemental material [31].

The continuous change in angular velocity with radial distance causes a *shearing* of structures (Fig. 2). The shearing operation can be represented by a coordinate transformation of a canonical emission at time $t = 0$:

$$e(t, \mathbf{x}) = e_0(\mathbf{R}_{\xi, \phi} \mathbf{x}), \quad (2)$$

where $\mathbf{R}_{\xi, \phi}$ is a rotation matrix of angle ϕ about the axis ξ ¹. The position and time dependent angle ϕ is given by

$$\phi(t, r) = t\omega(r) \propto tr^{-3/2}. \quad (3)$$

3.2. Ray Tracing with Gravitational Lensing

The radiance incident any pixel $p_n(t)$ at time t is given by “backward” integration of 3D emission over a ray path, Γ_n , originating at the image plane. For $N \times N$ pixels the discretized image plane is defined by the vector

$$\mathbf{I}(t) = [p_1(t), \dots, p_{N^2}(t)]. \quad (4)$$

The curved ray path in General Relativity (GR) can be described as a 4D (one time and three spatial dimensions) parametric function of the distance along the ray, s . If we

¹e.g. as given by the *Rodrigues’ Rotation Formula*




Figure 3. 2D Illustration of an emission hot-spot orbiting around a black hole and the resulting 1D image plane projection. As time progresses, $t_0 \rightarrow t_1$, the spot is sheared due to the inner radii moving faster than the outer radii (Sec. 3.1). The color-coded measurements illustrate the projected image at two different times. Selected rays are traced to highlight interesting image-features caused by the gravitationally curved trajectories.

know the path $\Gamma_n = (t(s), \mathbf{x}(s))$ that ends at the n ’th pixel, we can compute its intensity by integrating over s :

$$p_n(t) = \int_{\Gamma_n} e(t, \mathbf{x}) ds \approx \sum_{\mathbf{x}_i \in \Gamma_n} e(t, \mathbf{x}_i) \Delta s_i. \quad (5)$$

Here Δs_i is the integration weight associated with the i -th sample along the path Γ_n and can be thought of as the distance along the ray’s path taken in step i . In contrast to standard ray tracing in Euclidean space, the black hole’s gravitational field causes light ray trajectories to curve, and these curved ray paths are computed by integrating a differential equation [15, 21]; we provide explicit equations in the supplementary material [31]. Note that since in this work we assume ray paths are fixed according to the mass and spin of the black hole, pre-computed trajectories are used throughout the tomographic optimization.

In Eq. (5), we make three simplifying assumptions. First, we assume that the attenuation (due to absorption and scattering) of light by the emitting gas is negligible. Attenuation is important for general astrophysical images but can be neglected for EHT images of black holes Sgr A* and M87 [16]. Second, although gas is moving at relativistic velocities, causing emission to change as the light ray propagates at a finite speed, we do not account for this [12, 14, 35]. Finally, we do not account for relativistic redshift, which decreases the emission when material is deep in the gravitational field or moving away from the observer.

Figure 3 illustrates how an emission hot-spot behind the black hole is “lensed” onto the image plane as a consequence of the bending of light rays. In the figure, selected rays are illustrated to showcase interesting image features.




Figure 4. EHT spatial frequency sampling pattern as a function of time. Due to the Earth rotation, EHT telescopes probe different frequencies at different times. The EHT coverage (left) includes existing telescopes whereas the ngEHT [40] also includes candidate telescope sites, resulting in a denser coverage (right). The observation time was chosen to capture ~ 1 full orbit around SgrA* for emission at distances used in our simulations (Sec. 5).

A key image feature predicted by GR is bright “photon rings” produced by light rays that traverse full orbits around the black hole before reaching the sensor [24].

3.3. Event Horizon Telescope Measurements

The Event Horizon Telescope (EHT) records simultaneous radio signals from the black hole source at telescopes located around the globe. The correlation of the recorded radio signals from telescope pairs results in a measurement that probes a single spatial frequency of the underlying 2D image [43]. The spatial frequency measured is proportional to the projected baselines between the telescopes: short baselines probe low spatial frequencies and long baselines probe high spatial frequencies. These measurements, known as *complex visibilities*, can be represented as:

$$\mathbf{y}(t) = \mathbf{F}_t \mathbf{I}(t) + \boldsymbol{\varepsilon}. \quad (6)$$

Here, $\mathbf{F}_t$ is the 2D discrete-time Fourier transform (DTFT) matrix containing frequency components sampled by the EHT. The measurement *thermal noise*, $\boldsymbol{\varepsilon} = (\epsilon_1, \dots, \epsilon_K)^\top$, is Gaussian distributed according to $\epsilon_k \sim \mathcal{N}(0, \sigma_k^2)$, where σ_k is related to the sensitivity of the telescope pair [43]. In reality, additional noise sources exist due to the atmosphere and instrumental error [2], but in this work we limit ourselves to the fundamental telescope sensitivity.

Figure 4 shows the EHT frequency sampling coverage as a function of time during the night of April 7, 2017. The EHT array includes the existing telescopes which were used during the 2017 campaign. The next-generation EHT (ngEHT) [40] array includes additional candidate telescope sites that have been proposed for future EHT observations. These telescope sites, chosen according to favorable atmospheric transmittance and geographic location [40], yield denser frequency coverage.

4. BH-NeRF: Black Hole Neural Radiance Fields

In this section we discuss how the measurement model (detailed in Sec. 3) is used to supervise a coordinate-based neural network reconstruction of the 3D emission around a black-hole. Figure 5 gives a block diagram overview.

4.1. Representing Emission as an Neural Network

We take inspiration from NeRF [34], parameterizing the 3D volume at time $t = 0$ with the weights, $\boldsymbol{\theta}$, of an MLP neural network. This MLP takes continuously-valued coordinates $\mathbf{x}$ as input, and outputs the corresponding emission:

$$e_0(\mathbf{x}) = \text{MLP}_{\boldsymbol{\theta}}(\gamma(\mathbf{x})) \quad (7)$$

where $\gamma(\mathbf{x})$ is a positional encoding of the input coordinates that projects each input coordinate onto a set of sinusoids with exponentially-increasing frequencies:

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

This positional encoding controls the underlying interpolation kernel used by the MLP, where the parameter L determines the bandwidth of the interpolation kernel [41].

In our experiments, we use an MLP with 4 layers, where each layer is 128 units wide and uses ReLU activations. We use a maximum positional encoding degree of $L = 3$. The low degree of L is suitable for volumetric emission fields which are naturally smooth (see Suppl. [31]). Moreover, EHT measurements observe a limited range of frequencies, thus, it is important to use a positional encoding degree that does not introduce spurious high frequencies that cannot be supervised by the measurements.

4.2. Rotation Axis Estimation

In the previous section we described how we model the 3D emission as a continuous function using a neural network. To recover the unknown emission from measurements, our approach relies on Keplerian motion to model the orbit and evolution over time (Sec. 3.1). We consider a case where the angular velocity is well modeled by Eq. (3), however, the rotation axis $\boldsymbol{\xi}$ (Eq. 2) is unknown.

For SgrA*, past studies have attempted to constrain the rotation axis [7]. However, these constraints have large uncertainty due to measurement and modeling error. Therefore, our approach estimates an unknown axis jointly with the 3D emission (Fig. 5). The estimated values from other observations may be able to serve as a good prior or initialization for the optimization problem.

Since the forward model is fully differentiable, gradients can be back-propagated to directly optimize the unknown rotation axis, represented as a normalized 3-vector, $\boldsymbol{\xi}$. As highlighted in Fig. 6, we identify a persistent local minimum that occurs at the symmetric axis, equal to the neg-




Figure 5. Block diagram of our proposed tomography approach. We model an initial 3D emission e_0 as a continuous function using an MLP (parameterized by θ). The input to the network are 3D coordinates $\mathbf{x}$ (transformed by a positional encoding—see Sec. 4.1). The orbit dynamics, parameterized by a rotation axis ξ , dictate how the initial emission evolves over time $e_0(\mathbf{x}) \rightarrow e(t, \mathbf{x})$ (Sec. 3.1). BH-NeRF solves for θ and ξ jointly using a physically motivated loss that accounts for material orbit and gravitational lensing effects. Telescope measurements are used to constrain the optimization over θ and ξ .




Figure 6. Symmetry induces a local minimum at the opposite rotation direction. The top row shows the ground truth emission of a single hot-spot on the equatorial plane directly behind the black hole. Top row images are rendered with a rotation axis $\xi_{\text{true}} = (0, 0, 1)$. When initializing ξ_0 with a random vector in the southern hemisphere, the converged solution finds a local minimum. Simply initializing the rotation axis to the opposite sign vector (in the northern hemisphere) identifies the global minimum and the hot-spot emission is recovered accurately.

ative of the true rotation axis. This can be intuitively explained: even without temporal aliasing, two opposing orbits share a common rotation plane. Thus, the projected emission results in a common “line” formed on the image plane, that results in a lower loss than a perturbed rotation axes. This local minimum exists in all presented experiments. To avoid this local minimum, we initialize the rotation axis with a random unit vector drawn from the unit sphere and its symmetric counterpart from the opposing

hemisphere. We present the resulting emission with the lowest converged loss.

4.3. Optimization

We now detail how all the the components described in Secs. 3, 4 are integrated in the full optimization problem. The optimization is solved by minimizing a loss function over network parameters θ and rotation axis ξ :

$$\mathcal{L}(\theta, \xi) = \sum_t \|\mathbf{y}(t) - \mathbf{F}_t \mathbf{I}_{\theta, \xi}(t)\|_{\Sigma}^2, \quad (9)$$

Here, t denotes discrete observational time frames and $\|\mathbf{v}\|_{\Sigma}^2 \equiv \mathbf{v}^\top \Sigma^{-1} \mathbf{v}$. The thermal noise covariance matrix, $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_K^2)$, is determined by each telescope pair antenna parameters [40]. Using Eqs. (2), (5), (7), each image pixel p_n (elements of $\mathbf{I}$) can be expanded in terms of the unknown parameters:

$$p_n(\theta, \xi) = \sum_{\mathbf{x}_i \in \Gamma_n} \text{MLP}_{\theta}[\gamma(\mathbf{R}_{\xi} \mathbf{x}_i)] \Delta s_i. \quad (10)$$

5. Experiments

In this section we describe the simulation setup and show experimental results. These synthetic experiments assume that we are observing SgrA*, the Galactic center black hole that has a mass, M , of four million solar masses and zero angular momentum. These intrinsic parameters are necessary to set the magnitude of $\omega(r)$, as described in Eq. (1). Emitting hot-spots are randomly placed roughly at a radius of $1.16 \times r_{\text{ms}}$, resulting in an orbital period of ~ 40 minutes. Each hot-spot is generated as a 3D Gaussian with an isotropic standard deviation of $0.4 \times GM/c^2$.

Ground truth emissions were generated with 128 temporal frames and 64^3 3D voxel grid resolution. The network was implemented in JAX [6] and optimized using an




Figure 7. Comparison of BH-NeRF with two alternative representations. We display the recovered emission for three experimental setups, where the ground truth emission increases in complexity from a single hot-spot (left column) to four hotspots (middle column) to eight hotspots (four flaring up before $t = 0$). In all approaches, we recover emissions from image-domain measurements (projection plane indicated in green) and visualize the emission recovered at time $t = 0$. We find that BH-NeRF performs better visually and quantitatively. For quantitative metrics we show the PSNR value comparing recovered emissions the ground truth. While BH-NeRF significantly outperforms the 4D MLP, we note that this more generalized representation could prove useful in environments with unknown dynamics.

ADAM optimizer [26] with a polynomial learning rate transitioning from $1e^{-4} \rightarrow 1e^{-6}$ over 5K iterations. EHT measurements were generated using `eht-imaging` [11] with realistic thermal noise. Recovery run times were ~ 15 minutes on two NVIDIA Titan RTX GPUs. Code implementation is available at the project page [31].

Exploring Alternative Emission Representations Our proposed approach, BH-NeRF, described in Sec. 4 represents the 4D emission through time as a 3D volume parameterized with a coordinate-based continuous neural network and a parametric flow field. In the experiments presented in Fig. 7 we explore two alternative representations

and compare them to BH-NeRF.

The first approach, which we term *3D Grid*, relies on the Keplerian orbital motion model (described in Sec. 3.1), however, instead of a continuous neural representation, it uses a discrete voxel grid. The second approach, *4D MLP*, ignores the orbital motion model and recovers a 4D continuous emission field, represented as a MLP similar to BH-NeRF (i.e., taking as input a 4D spatiotemporal coordinate rather than a 3D spatial coordinate).

In Fig. 7 we compare emissions recovered with full image-plane measurements and the true rotation axis (not estimated). Results are shown for three different ground truth emission patterns with increasing complexity (from left to right). In the simplest case (left column), a single emission hot-spot flares at $t = 0$ in front of the black hole, thus, it is directly visible. The top row of Fig. 7 shows the observations at the initial and final frames: $t = \{0, t_1\}$. The middle column shows a reconstruction of four hot-spots which flare at $t = 0$, where two are directly visible and two are “hidden” behind the black-hole and are gravitationally lensed. The last experiment (right column) shows eight hot-spots where four flared up before $t = 0$ and four at $t = 0$.

In our experiments, we find that BH-NeRF slightly outperforms the 3D voxel grid with minimal tuning, potentially due to its continuous and adaptive allocation of resolution. In contrast, the 4D representation performs significantly worse. This is expected, as it makes no use of any assumption on the emission dynamics. However, we highlight that a 4D representation could be particularly useful in situations where the true dynamics are significantly different from our assumption of Keplerian flow.

Increasing Measurement Sparsity In the following simulations, we demonstrate recoveries for three emission patterns orbiting about different axes with known and unknown (estimated) rotation axis. The initial emission in the first experiment (top two rows of Fig. 8) has three hot-spots with two “hidden” behind the black hole and one directly visible in-front of the black hole. The image plane, where rays are traced from, is highlighted in green. For each experiment we show the recovered volume (with PSNR under each reconstruction) and the recovered axis (the dot product of the recovered and true axis specified under the reconstruction). Furthermore, results are shown at $t=10$ min, with the estimated axis, $\hat{\xi}$, used to advance time.

A key aspect of BH-NeRF is the ability to recover arbitrary emission distributions which enables novel scientific discovery. The third experiment in Fig. 8 demonstrates this ability by recovering 3D MNIST [10] digits (‘0’–‘4’) orbiting at a random orientation. Despite the fact that the ‘2’ initially appears behind the black hole, BH-NeRF is able to recover its features (further analysis given in the Suppl. [31]).

While in practice we only have access to sparse visibility measurements of the image plane (Sec. 3.3), the image-




Figure 8. Simulation results for three experiments with different rotation axes and emission patterns. The top two experiments show recoveries of Gaussian emission hotspots at two different times. The third experiment (bottom) shows recovery of 3D digits which illustrates the flexibility of our approach which models arbitrary 3D emission. Each experiment shows: the image-plane projection, ground-truth and recoveries. The orbital period of both is ~ 40 minutes. Green highlights the image-plane where rays are traced from. Recoveries are shown for both unknown (estimated) and known rotation axis. At $t = 0$ some of the emission is directly visible and some is “hidden” behind the black-hole and is lensed onto the image-plane. Quantitative PSNR values for each recovered emission compared to the ground truth are given below each recovery. For the rotation axis the dot product with the ground truth is indicated below each recovery.

based recoveries are useful as a form of upper bound (analogous to tiling the Earth with telescopes) on EHT-based results. As expected, increasing the measurements with additional ngEHT telescopes improves the recovery visually and quantitatively over using current EHT telescopes only.

Velocity Model Mismatch The Keplerian motion model is a good approximation, however, it does not account for

complex dynamics (e.g. due to turbulence or outflowing motion in a magnetically-driven jet [16]). To analyze the robustness of our approach to a mismatch in the velocity model we perturb it with correlated Gaussian noise parameterized by magnitude m and correlation length ℓ . For each parameter set, (m, ℓ) , five velocity profiles were sampled to generate data. Sample reconstructions are shown from the




Figure 9. Recovery performance with model mismatch. We perturb the radial velocity model with correlated Gaussian noise. 1) The top panel shows samples of noisy velocity profiles with different noise magnitudes (m) and correlation lengths (ℓ). Red highlights the emission orbit region. 2) The mean PSNR is shown in the top right. 3) The bottom panel shows the volumetric ground-truth, no velocity noise reconstruction, and reconstructions from the highlighted parameter sets. Results are obtained from synthetic ngEHT measurements.

best and worst PSNR parameter sets, highlighted in blue and red respectively. While reconstruction degrades with noise magnitude, our approach is still able to capture the emission structure and estimate the rotation axis with a reasonable level of model mismatch.

6. Limitations

While our model is able to recover 3D emission in challenging scenarios with sparse observations, we rely on a few assumptions to do so. First, we assume knowledge of the black hole’s spin and mass in order to pre-compute the curved trajectories of rays and velocity profile. Although reasonable ranges of values for these variables can be obtained from other observations, estimating them independently could provide tighter constraints on these scientifically impactful parameters. However, estimating these parameters requires back-propagating through the gravitational lensing ray-tracing equations (see supplemental material [31]). We therefore leave this challenging extension to future work. Second, we assume a Keplerian dynamics model where the angular velocity is determined by the distance from the black-hole center. Nevertheless, in Sec. 5 we show that our method is robust to some uncertainty in the velocity profile. Third, in this work we do not consider additional flares appearing during the observation window. However, as we expect to be able to identify and isolate the time that flares appear by looking for changes in the source’s integrated flux (total brightness), we do not expect this to be a significant limitation. Fourth, while our experiments include realistic thermal noise, we do not include instrumental errors or particularly challenging atmospheric

noise, which would need to be studied before applying our model to real EHT or, in the future, ngEHT measurements.

7. Conclusion

This work takes the first steps in showing how ground-based observations from the Event Horizon Telescope could be used to recover evolving 3D emission around the supermassive black hole in our Galactic center. Our contributions are two-fold. First, we formulate the novel tomography problem of recovering dynamic 3D emission orbiting a black hole. Second, we generalize neural radiance fields to work in scenarios with curved ray trajectories and orbital dynamics induced by the gravitational field around a black hole. This enables the recovery of dynamic 3D emission from sparse measurements of a single viewpoint over time.

Our simulations and analysis show great promise for tomographic reconstruction of the dynamic environment around a black hole, a direction that can open the door to new insights into our dynamic universe. Moreover, we believe that our work can pave the way for other scientific tomography applications, such as refraction and scattering, which rely on complex ray-tracing at their core.

Acknowledgements

AL is supported by the Zuckerman and Viterbi postdoctoral fellowships. AC is supported by Hubble Fellowship grant HST-HF2-51431.001-A awarded by the STScI, operated by the AURA Astronomy for NASA (NAS5-26555). This work was supported by NSF awards 1935980, 2034306, 2048237 and Beyond Limits.

References

- [1] Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First m87 event horizon telescope results. ii. array and instrumentation. *The Astrophysical Journal Letters*, 875(1):L2, 2019. 1
- [2] Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First m87 event horizon telescope results. iii. data processing and calibration. *The Astrophysical Journal Letters*, 2019. 4
- [3] 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
- [4] Katherine L. Bouman, Michael D. Johnson, Adrian V. Dalca, Andrew A. Chael, Freek Roelofs, Sheperd S. Doeleman, and William T. Freeman. Reconstructing video of time-varying sources from radio interferometric measurements. *IEEE Transactions on Computational Imaging*, 2018. 1
- [5] Katherine L. Bouman, Michael D. Johnson, Daniel Zoran, Vincent L. Fish, Sheperd S. Doeleman, and William T. Freeman. Computational imaging for VLBI image reconstruction. *CVPR*, 2016. 1
- [6] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Nectra, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. 5
- [7] Avery E. Broderick, Vincent L. Fish, Michael D. Johnson, Katherine Rosenfeld, Carlos Wang, Sheperd S. Doeleman, Kazunori Akiyama, Tim Johannsen, and Alan L. Roy. Modeling seven years of event horizon telescope observations with radiatively inefficient accretion flow models. *The Astrophysical Journal*, 2016. 4
- [8] Avery E. Broderick and Abraham Loeb. Imaging bright-spots in the accretion flow near the black hole horizon of sgr a. *Monthly Notices of the Royal Astronomical Society*, 2005. 1
- [9] Avery E. Broderick and Abraham Loeb. Imaging optically-thin hotspots near the black hole horizon of sgr a* at radio and near-infrared wavelengths. *Monthly Notices of the Royal Astronomical Society*, 367(3):905–916, 2006. 3
- [10] David de la Iglesia Castro. 3D MNIST: A 3D version of the MNIST database of handwritten digits. <https://www.kaggle.com/daavoo/3d-mnist>. 6
- [11] Andrew A. Chael, Michael D. Johnson, Katherine L. Bouman, Lindy L. Blackburn, Kazunori Akiyama, and Ramesh Narayan. Interferometric imaging directly with closure phases and closure amplitudes. *The Astrophysical Journal*, 2018. 1, 6
- [12] Chi-kwan Chan, Dimitrios Psaltis, and Feryal Özel. GRay: A Massively Parallel GPU-based Code for Ray Tracing in Relativistic Spacetimes. *The Astrophysical Journal*, 2013. 3
- [13] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *CVPR*, 2019. 2
- [14] Jason Dexter. A public code for general relativistic, polarised radiative transfer around spinning black holes. *MNRAS*, 2016. 3
- [15] Jason Dexter and Eric Agol. A Fast New Public Code for Computing Photon Orbits in a Kerr Spacetime. *The Astrophysical Journal*, 2009. 3
- [16] Event Horizon Telescope Collaboration et al. First M87 Event Horizon Telescope Results. V. Physical Origin of the Asymmetric Ring. *The Astrophysical Journal Letters*, 2019. 3, 7
- [17] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. *CVPR*, 2021. 2
- [18] Angela F. Gao, Jorge C. Castillo, Yisong Yue, Zachary E. Ross, and Katherine L. Bouman. DeepGEM: Generalized expectation-maximization for blind inversion. *NeurIPS*, 2021. 2
- [19] R. Genzel, R. Schödel, Thomas Ott, A. Eckart, Tal Alexander, F. Lacombe, D. Rouan, and B. Aschenbach. Near-infrared flares from accreting gas around the supermassive black hole at the galactic centre. *Nature*, 2003. 1
- [20] Richard Gordon, Robert Bender, and Gabor T. Herman. Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography. *Journal of Theoretical Biology*, 1970. 2
- [21] Samuel E. Gralla and Alexandru Lupsasca. Null geodesics of the Kerr exterior. *Physical Review D*, 2020. 3
- [22] Henk Hoekstra and Bhuvnesh Jain. Weak gravitational lensing and its cosmological applications. *Annual Review of Nuclear and Particle Science*, 2008. 2
- [23] Michael D. Johnson, Katherine L. Bouman, Lindy Blackburn, Andrew A. Chael, Julian Rosen, Hotaka Shiokawa, Freek Roelofs, Kazunori Akiyama, Vincent L. Fish, and Sheperd S. Doeleman. Dynamical imaging with interferometry. *The Astrophysical Journal*, 850(2):172, 2017. 1
- [24] Michael D. Johnson, Alexandru Lupsasca, Andrew Strominger, George N. Wong, Shahar Hadar, Daniel Kapec, Ramesh Narayan, Andrew Chael, Charles F. Gammie, Peter Galison, et al. Universal interferometric signatures of a black hole’s photon ring. *Science advances*, 2020. 4
- [25] A.C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988. 2
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2014. 6
- [27] Aviad Levis, Daeyoung Lee, Joel A. Tropp, Charles F. Gammie, and Katherine L. Bouman. Inference of black hole fluid-dynamics from sparse interferometric measurements. *CVPR*, 2021. 1
- [28] Aviad Levis, Yoav Y. Schechner, Amit Aides, and Anthony B. Davis. Airborne three-dimensional cloud tomography. *ICCV*, 2015. 2
- [29] Aviad Levis, Yoav Y. Schechner, and Anthony B. Davis. Multiple-scattering microphysics tomography. *CVPR*, 2017. 2
- [30] Aviad Levis, Yoav Y. Schechner, Anthony B. Davis, and Jesse Loveridge. Multi-view polarimetric scattering cloud tomography and retrieval of droplet size. *Remote Sensing*, 12(17):2831, 2020. 2

- [31] Aviad Levis, Pratul P. Srinivasan, Andrew A. Chael, Ren Ng, and Katherine L. Bouman. Gravitationally lensed black hole emission tomography: project page. <http://imaging.cms.caltech.edu/bhnerf/>. 3, 4, 6, 8
- [32] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *CVPR*, 2021. 2
- [33] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: learning 3D reconstruction in function space. *CVPR*, 2019. 2
- [34] 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. *ECCV*, 2020. 2, 4
- [35] M. Mościbrodzka and C. F. Gammie. IPOLE - semi-analytic scheme for relativistic polarized radiative transport. *Monthly Notices of the Royal Astronomical Society*, 2018. 3
- [36] Engineering National Academies of Sciences and Medicine. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. The National Academies Press, 2021. 2
- [37] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. *CVPR*, 2019. 2
- [38] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. 2
- [39] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. *ICCV*, 2021. 2
- [40] Alexander W. Raymond, Daniel Palumbo, Scott N. Paine, Lindy Blackburn, Rodrigo Córdova Rosado, Sheperd S. Doeleman, Joseph R. Farah, Michael D. Johnson, Freek Roelofs, Remo P.J. Tilanus, et al. Evaluation of new sub-millimeter vlbi sites for the event horizon telescope. *The Astrophysical Journal Supplement Series*, 2021. 4, 5
- [41] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020. 2, 4
- [42] The Event Horizon Telescope Collaboration. First M87 Event Horizon Telescope Results. IV. Imaging the central supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L4, 2019. 1
- [43] A. Richard Thompson. Fundamentals of radio interferometry. *Synthesis Imaging in Radio Astronomy II*, 1999. 1, 4
- [44] Paul Tiede, Hung-Yi Pu, Avery E. Broderick, Roman Gold, Mansour Karami, and Jorge A. Preciado-López. Spacetime tomography using the event horizon telescope. *The Astrophysical Journal*, 892(2):132, 2020. 2
- [45] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhoefer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. *ICCV*, 2021. 2
- [46] Jinhui Xiong and Wolfgang Heidrich. In-the-wild single camera 3D reconstruction through moving water surfaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12558–12567, 2021. 2
- [47] Ellen D. Zhong, Tristan Bepler, Bonnie Berger, and Joseph H. Davis. CryoDRGN: reconstruction of heterogeneous cryo-em structures using neural networks. *Nature Methods*, 18(2):176–185, 2021. 2