

Inference of Black Hole Fluid-Dynamics from Sparse Interferometric Measurements

Aviad Levis^{1*}, Daeyoung Lee², Joel A. Tropp¹, Charles F. Gammie², Katherine L. Bouman¹

¹California Institute of Technology, ²University of Illinois

*alevis@caltech.edu

Abstract

We develop an approach to recover the underlying properties of fluid-dynamical processes from sparse measurements. We are motivated by the task of imaging the stochastically evolving environment surrounding black holes, and demonstrate how flow parameters can be estimated from sparse interferometric measurements used in radio astronomical imaging. To model the stochastic flow we use spatio-temporal Gaussian Random Fields (GRFs). The high dimensionality of the underlying source video makes direct representation via a GRF's full covariance matrix intractable. In contrast, stochastic partial differential equations are able to capture correlations at multiple scales by specifying only local interaction coefficients. Our approach estimates the coefficients of a space-time diffusion equation that dictates the stationary statistics of the dynamical process. We analyze our approach on realistic simulations of black hole evolution and demonstrate its advantage over state-of-the-art dynamic black hole imaging techniques.

1. Introduction

The first-ever black hole image was produced by the *Event Horizon Telescope* (EHT) collaboration in 2019 [35]. This image, of the black hole in the center of the M87 galaxy was computationally constructed from radio measurements collected over an entire night [2]. The massive size of the M87 black hole gives rise to evolution timescales greater than acquisition length, thus, the image could be reconstructed assuming a static source [8]. In contrast, the Milky Way's galactic center hosts a smaller [18] more dynamic [15, 23] black hole: *Sgr A**.

While black holes are not directly observable by radio telescopes, it is their imprint on the fluid-dynamic backdrop that can be observed. In this work, our goal is to recover the fluid-dynamics of quickly evolving phenomena, such as *Sgr A**, from sparse remote sensing data. Recovering these dynamic properties could reveal information about the spin and mass of *Sgr A** and pose a stronger test [11] to General Relativity than M87 (due to the tighter mass con-

straints). While motivated by the evolving plasma surrounding a black hole, our approach is also applicable to other inference problems with fluid-dynamics at their core.

Inference of dynamics from sparse measurements can be placed on a spectrum of approaches ranging from model-free, with many degrees of freedom (DOFs), to fully physics-based, with only a handful of DOFs. On the model-free end of the spectrum are approaches that first recover an entire video sequence from the sparse observations [7]. In theory, video estimation is extremely flexible and could capture any evolution. In practice, to estimate many DOFs (every pixel in the video) strong inter-frame regularization is imposed and much of the dynamic information is lost.

On the other end of the spectrum, are approaches that seek to infer parameters of high fidelity physics-based models [36]. While advantageously characterized by a small number of parameters, these models lack the flexibility for novel scientific discovery. For instance, it is crucial to avoid a model that assumes the presence of a black hole obeying our current numerical models of physics, as this could skew the results and interpretation. Furthermore, in many scenarios physics-based models are complex and non-linear, resulting in computationally intensive simulations that are intractable for use in an inverse problem. For instance, high-fidelity models used for black hole require thousands of core hours for a single simulation¹.

We instead seek to find a middle ground between these two ends of the spectrum, using a *physics-motivated* surrogate model that is able to capture key features of the fluid-dynamical process. In particular, we model the spatio-temporal flow surrounding a black hole as an output of a stochastic partial differential equation (SPDE). This model enables the efficient drawing of video samples with spatio-temporal correlation at multiple scales (Fig. 1), and was previously shown to capture the spatio-temporal variability of astrophysical disks [24]. The inference task we propose is to recover the SPDE coefficients, which characterize pixel interactions. These coefficients are interpretable and can be related back to the local correlations of physics-based simulations [19].

¹General relativistic magneto-hydrodynamics (GRMHD) [31]




Figure 1. The anisotropic spatio-temporal SPDE generative forward model. The PDE transforms uncorrelated random noise samples, $\mathcal{W}(\mathbf{x})$, into correlated Gaussian Markov Random Fields (GRFs), $\rho(\mathbf{x})$. Different GRF samples have different pixel values but the same statistics. The inference problem we formulate is to infer the PDE parameters Θ from either a single video sample or interferometric measurements of a source exhibiting similar dynamic properties. We are not interested in recovering exact pixel values of an evolving source, which is dependent on the unknown random noise. Rather we wish to recover the source’s underlying stationary statistics.

2. Relevant Prior Work

Prior work on imaging dynamic black holes has focused on the recovery of an entire video from sparse interferometric measurements [7, 22, 5]. Although these video reconstruction methods do not estimate flow properties directly, subsequent analysis (e.g., optical flow) could be used for this purpose. However, in practice, estimating an entire video of unknown pixel values from very sparse data requires strong temporal regularization that suppresses dynamic information within the recovered video. In contrast, our work seeks to directly recover the flow model *statistics* from measurements. As we demonstrate in Sec. 5, model-fitting statistical flow parameters outperforms recovery via video reconstruction when measurements are very sparse.

To model complex fluid flows surrounding a black hole, we rely on fitting random spatio-temporal fields to observed data. In the field of computer vision, there has been significant work dedicated to the analysis of dynamic video textures. This includes detection [4], classification [32, 17], estimation [25], and generation [12]. More recently, generative neural nets [38, 10] have shown great promise in capturing statistics of complex dynamical videos. Nevertheless, statistical properties captured by network weights lack scientific interpretability. In contrast, we wish to recover dynamic-parameters that relate to the underlying physics and could be used in subsequent analysis. Closest in spirit to our work are [39, 33], which also solve a related inverse problem, recovering fluid properties from sparse data. Nevertheless, these works recover fluid properties at each time frame, whereas our approach recovers motion statistics.

In the following subsections we introduce the generation and modeling of stochastic flows as Gaussian Random Fields (Section 2.1) that are efficiently captured by a Stochastic Partial Differential Equation (Section 2.2).

2.1. Gaussian Random Fields (GRFs)

Gaussian Random Fields are commonly used to represent random spatial fields [34]. GRFs capture second order statistics with a mean and continuous covariance kernel (function). On a discrete grid (e.g. pixels), a continuous GRF can be approximated as a multivariate Gaussian. For small grids the full covariance matrix can be stored in memory and used to draw GRF samples.

Using a covariance matrix to model GRFs is common for spatial data [34]. However, in this work we model spatio-temporal fields, for which the covariance matrix can easily reach impractical memory and computation requirements².

2.2. Stochastic Partial Differential Equations

To efficiently model spatio-temporal random fields (i.e., GRFs) we employ an approach that generates a random field as a solution to a stochastic partial differential equation (SPDE) [37, 26, 16]. Thus, instead of requiring full rank covariance matrices to specify and sample a field, only local interactions need to be specified. In particular, in Sec. 2.2.1 we describe the *anisotropic* spatio-temporal diffusion equation that is able to capture the (in-homogeneous and non-isotropic) variability of astrophysical accretion disks [24]. Anisotropic (directional) correlations are introduced through a diffusion tensor, as detailed in Sec. 2.2.2.

2.2.1 Anisotropic Spatio-Temporal Diffusion

Let $\mathbf{x}=(x_0, x_1, x_2)$ and $\nabla=(\frac{\partial}{\partial x_0}, \frac{\partial}{\partial x_1}, \frac{\partial}{\partial x_2})^\top$ denote a space-time grid point and derivative. Subscript 0 denotes the temporal axis whereas 1 and 2 denote spatial axes. For notational compactness we denote only the spatial coor-

²Drawing samples from an $N \times N$ covariance matrix typically involves computing its Cholskey decomposition which scales as $\mathcal{O}(N^3)$

ordinates as $\mathbf{x}=(x_1, x_2)$ (note that $\mathbf{x} \neq x$). The anisotropic spatio-temporal diffusion equation is given by [24]

$$(1 - \nabla \cdot \mathbf{\Lambda}(\mathbf{x}) \nabla) \rho(\mathbf{x}) = \gamma \det(\mathbf{\Lambda}(\mathbf{x}))^{1/4} \mathcal{W}(\mathbf{x}). \quad (1)$$

Here $\mathcal{W}$ is an input emission source modeled by a Gaussian white noise process [37], $\rho(\mathbf{x})$ is the spatio-temporal output GRF, and $\mathbf{\Lambda}(\mathbf{x})$ is a spatially dependent diffusion tensor that encodes local interactions. The heterogeneity of $\mathbf{\Lambda}(\mathbf{x})$ across the image produces GRFs with spatially varying temporal variance. The normalization factor $\det(\mathbf{\Lambda}(\mathbf{x}))^{1/4}$ on the right hand side of Eq. (1) ensures constant (flat) temporal variance across image pixels. The scaling factor γ determines the magnitude of variations.

In practice, Eq. (1) is solved by discretization in space and time; the output GRF, ρ , is obtained as a solution to a linear set of equations:

$$\mathbf{D}\rho = \mathbf{B}\mathbf{w}, \quad (2)$$

$$\rho = (\rho(\mathbf{x}^1), \dots, \rho(\mathbf{x}^N))^T, \quad \mathbf{w} = (\mathcal{W}(\mathbf{x}^1), \dots, \mathcal{W}(\mathbf{x}^N))^T,$$

where $\mathbf{D}$ is a sparse matrix that approximates the differential operators via finite differences,

$$\mathbf{D} \simeq (1 - \nabla \cdot \mathbf{\Lambda}(\mathbf{x}) \nabla), \quad (3)$$

and $\mathbf{B}$ is a diagonal normalization matrix,

$$\mathbf{B} = \text{diag} \left[\gamma \det(\mathbf{\Lambda}(\mathbf{x}^1))^{1/4}, \dots, \gamma \det(\mathbf{\Lambda}(\mathbf{x}^N))^{1/4} \right]. \quad (4)$$

The SPDE formulation reduces the problem of sampling a correlated Gaussian distribution, with a large dense covariance matrix, to that of sampling *i.i.d.* Gaussian variables, $\mathbf{w}$, and solving a sparse linear set of equations. The latter can be done very efficiently³.

2.2.2 Diffusion Tensor

The correlations in the output GRF are completely determined by the diffusion tensor $\mathbf{\Lambda}(\mathbf{x})$. In contrast to a scalar coefficient, the tensor captures non-isotropic local interactions. The diffusion tensor is a symmetric 3×3 matrix with positive coefficients specifying correlations in each of the axes and their combinations. Locally $\mathbf{\Lambda}$ can be described by three directional axes

$$\mathbf{\Lambda} = \lambda_0^2 \xi_0 \xi_0^T + \lambda_1^2 \xi_1 \xi_1^T + \lambda_2^2 \xi_2 \xi_2^T, \quad (5)$$

where $\lambda_0, \lambda_1, \lambda_2$ denote correlation time and lengths along axes specified by ξ_0, ξ_1 , and ξ_2 , respectively⁴. The correlation time λ_0 describes a characteristic timescale in which features persist in the flow and the correlation lengths λ_1, λ_2 describe the spatial direction and extent of the features.

³Efficient sparse numerical algorithms have memory requirements which scale as $\mathcal{O}(N)$ and computations as $\mathcal{O}(N^{3/2})$

⁴Here the spatial coordinates dependence was omitted for brevity




Figure 2. The bottom row shows four sequential frames of a clockwise (CW) rotating GRF, dictated by $\mathbf{v}$ (blue in the top figure). The spiral spatial correlations are dictated by ξ_1 (green in the top figure). The correlation lengths $\lambda_0, \lambda_1, \lambda_2$ are azimuthally symmetric and are a function of radius only (top left).

We define a parameterization of the correlation axis as

$$\text{Temporal axis: } \xi_0 = (1, v_x, v_y)^T, \quad (6)$$

$$\text{Major spatial axis: } \xi_1 = (0, \cos \phi, \sin \phi)^T, \quad (7)$$

$$\text{Minor spatial axis: } \xi_2 = (0, -\sin \phi, \cos \phi)^T. \quad (8)$$

Here $\mathbf{v} = (v_x, v_y)$ is a velocity field that dictates the motion of the GRF and ϕ defines the spatial direction of correlation at each position $\mathbf{x}$ (Fig. 2). Note that the normalization factor defined in Eq. (1) is given by $\det(\mathbf{\Lambda}) = \lambda_0^2 \lambda_1^2 \lambda_2^2$.

3. Forward Video Generation Model

In Sec. 2.2 we described a computationally tractable method for sampling spatio-temporal GRFs, $\rho(\mathbf{x})$, via an SPDE. While these GRFs capture flow dynamics, they are not suitable for modeling static image features, such as the photon ring in a black hole shadow [30]. Thus, we define a video formation model where the GRF is a multiplicative perturbation to a static “envelope” image, I . The formation of a video sample, $V(\mathbf{x})$, is modeled by

$$V(\mathbf{x}) = I(\mathbf{x}) e^{\rho(\mathbf{x})}. \quad (9)$$

Here, negative and positive values of ρ describe a fractional deviation from the envelope’s local intensity.

3.1. Astrophysical Accretion Disks as GRFs

Equations (5)–(8) define a diffusion tensor field where the correlation length and direction can, in principle, vary arbitrarily in space. Nevertheless, for the purpose of black hole imaging, it is constructive to examine the special case of astrophysical accretion disks. Accretion disks are structures formed by diffuse material orbiting around a massive central body. To capture the structural variability of a face-on accretion disk we define a simple, azimuthally symmetric, parameterization of the evolution. Derived from Keplerian orbital motion [24], we define the correlation lengths $\lambda_0, \lambda_1, \lambda_2$ (Eq. 5) and velocity magnitude (Eq. 6) as a func-

tion of radius r from the origin (Fig. 2). We define a constant ratio $\lambda_2 = \epsilon \lambda_1$ for $\epsilon \in [0, 1]$ that dictates the directionality of spatial correlations ($\epsilon = 1.0$ yields isotropic correlations). In our experimental results (Sec. 5) we fix $\epsilon = 0.1$.

Moreover, we parameterize the directions of flow velocity, $\mathbf{v}$, and spatial correlation, ϕ , with two angles φ_0 and φ_1 , respectfully. These angles dictate the flow direction⁵ and spiral *opening angle* (Fig. 2). For the accretion disk model, the inverse problem we formulate seeks to recover the unknown angles (φ_0, φ_1) . Recovering (φ_0, φ_1) reveals disk rotation that could shed light on important questions in fundamental physics [3, 19]. Refer to the supplemental material [1] for mathematical definitions of this parameterization.

4. Inverse Problem Formulation

Let Θ denote the unknown dynamic parameters (e.g., $\Theta = \{\varphi_0, \varphi_1\}$) of $\rho(\mathbf{x})$. We seek to estimate Θ from measurements $\mathbf{y}$ (yet to be defined). Recall that Θ only dictates the *statistics* of evolution seen in $\rho(\mathbf{x})$, not the individual pixel values. In fact, as is shown in Figure 1, multiple ρ 's sampled from the same PDE parameters, Θ , result in movies with very different pixel values over time. Therefore, in order to naively solve for parameters, Θ , that lead to movies that directly match ρ (e.g., via an ℓ_p -norm), we would also have to recover the unknown random noise source, $\mathbf{w}$. Unfortunately, jointly estimating Θ and $\mathbf{w}$ is impractical, as $\mathbf{w}$ is *i.i.d.* random noise and as high-dimensional as the video itself. Thus we seek a metric $\mathcal{D}$, that is invariant to $\mathbf{w}$ when solving for the statistical parameters of ρ 's evolution, Θ :

$$\mathbf{I}^*, \Theta^* = \arg \min_{\mathbf{I}, \Theta} \mathcal{D}(\mathbf{I}, \Theta | \mathbf{y}). \quad (10)$$

Our key insight to designing $\mathcal{D}$ is recognizing that the underlying dynamic video, ρ , will be best captured by a low-dimensional subspace describing videos drawn from the true parameters, Θ_{true} . The challenge is to compute the low-dimensional subspace for every possible Θ , and then identify which subspace best captures ρ .

Since we model ρ as a (zero-mean) GRF (Sec. 2.1), its statistics are fully captured by a covariance matrix parametrized by Θ . If we had access to this covariance matrix we could simply compute an eigenvalue decomposition to reduce the dimensionality – selecting the top eigenvectors, which explain most of the variability seen in videos drawn with parameters Θ . However, recall that a spatio-temporal covariance matrix describing the video ρ would be intractably large to store in memory; thus, we cannot use standard eigenvalue decompositions to recover the low-dimensional subspace. An alternative, but inefficient, ap-

⁵Direction with respect to counter-clockwise (CCW) rotation. For example, $\varphi_0 = \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ describe counter-clockwise (CCW) flow, radial outflow, clockwise (CW) flow, and radial inflow

proach to reduce the dimensionality is by sampling multiple GRFs with parameters Θ (via the SPDE in Sec. 2.2) and applying Principle Component Analysis (PCA) [29]. For a large number of GRF samples, PCA vectors should coincide with the eigenvectors of the covariance. However, this PCA approach is either very computationally intensive or inaccurate (see further analysis in the supplemental material [1]). Alternatively, in Section 4.1 we describe a tractable approach for dimensionality reduction without explicit access to the covariance matrix.

4.1. Dimensionality Reduction

In this section we describe a matrix-free approach to reduce the dimensionality of an SPDE system parameterized by Θ . It is convenient to *symbolically* re-write Eq. (2) as

$$\rho = \mathbf{A}\mathbf{w}, \quad (11)$$

where $\mathbf{A} = \mathbf{D}^{-1}\mathbf{B}$. The dependency on Θ is omitted throughout for clarity (i.e. $\mathbf{A} \leftrightarrow \mathbf{A}_\Theta$). We seek to find a low-rank approximation for the forward operator

$$\mathbf{A} \simeq \hat{\mathbf{A}} = \mathbf{U}\Sigma\mathbf{Q}^\top, \quad (12)$$

where $\mathbf{U}$ and $\mathbf{Q}^\top$ are of size $N \times K$, Σ is a $K \times K$ diagonal matrix, and $K \ll N$. The challenge is that $\mathbf{A}$ is a large, dense matrix that we do not have access to; thus, we must rely on methods to recover a low-dimensional decomposition without having access to $\mathbf{A}$.

We employ randomized matrix-free computations [28] to solve for $\mathbf{U}$ and Σ . A matrix-free approach is advantageous for both high-dimensional spaces in which matrix operations (i.e. direct SVD or QR) are too costly, as well as for numerical simulations in which an input-output relation is established without an explicit matrix representation.

Randomized subspace iteration (RSI)⁶ [27, 21] is a robust approach that builds a low-rank approximation of the forward operator, $\mathbf{A}$, through sequential applications to a set of K random (typically *i.i.d.* Gaussian) vectors. The *block-size* K is (roughly) the desired dimension of the computed subspace. While the “application” of $\mathbf{A}$ is symbolically written as an input-output relation (Eq. 11), in practice it means solving a linear set of equations (Eq. 2). In the context of the SPDE model, we iteratively compute an output GRF by replacing the input source with the previous output.

Algorithm 1 outlines RSI where the key outputs are

$$\mathbf{U} = (\mathbf{u}_0, \dots, \mathbf{u}_{K-1}), \quad \Sigma = \text{diag}(\sigma_0, \dots, \sigma_{K-1}). \quad (13)$$

Here $\mathbf{u}_k, \sigma_k$ are the top singular vectors and values of $\mathbf{A}$. The orthonormal vectors $\mathbf{u}_k$ are aligned with the axes of most variability, where σ_k is the variance along each axis.

⁶RSI typically computes a good approximation of the leading singular subspaces of $\mathbf{A}$ after a modest number of iterations. For our purposes, RSI is simple and highly effective. In cases where very precise estimates of the leading singular values and vectors are needed, more sophisticated algorithms could be considered (e.g. randomized block Krylov methods [27]).




Figure 3. The measurement $\mathbf{y} = \rho$ was generated with $\Theta_{\text{true}} = (0^\circ, 70^\circ)$. Here $\varphi_0 = 0^\circ$ defines CCW rotation and $\varphi_1 = 70^\circ$ defines the opening angle of the spiral. The top and bottom rows show projections of $\mathbf{y}$ onto the ($K=60$) modes generated with $\Theta_1 = (10^\circ, 71^\circ)$ and $\Theta_2 = (-28^\circ, -17^\circ)$, respectively. Since Θ_1 is closer to Θ_{true} , the projection $\mathcal{P}(\mathbf{y}, \Theta_1)$ retains more of the measurement flow features.

Algorithm 1 Random Subspace Iteration

```

1:  $\mathbf{b}_0 \leftarrow \text{qr\_economic}(\text{randn}(N, K))$   $\triangleright$  Draw  $N \times K$ 
   orthonormal block
2: for  $t=1, 2, 3, \dots, q$  do
3:    $\mathbf{b}_t \leftarrow \mathbf{A}\mathbf{b}_{t-1}$   $\triangleright$  Solve Eq. 2 for each vector in  $\mathbf{b}_{t-1}$ 
4:    $\mathbf{b}_t \leftarrow \text{qr\_economic}(\mathbf{b}_t)$   $\triangleright$  Orthogonalize
5: end for
6:  $\mathbf{U}, \Sigma, \mathbf{Q} \leftarrow \text{svd}(\mathbf{A}\mathbf{b}_q)$ 

```

Once $\mathbf{U}$ and Σ have been identified, the low-dimensional random field representation, $\hat{\rho}$, is given by

$$\hat{\rho} = \hat{\mathbf{A}}\mathbf{w} = \mathbf{U}\Sigma\mathbf{Q}^\top \mathbf{w} = \mathbf{U}\Sigma\mathbf{c} = \sum_{k=0}^{K-1} c_k \sigma_k \mathbf{u}_k, \quad (14)$$

where

$$c_k = \mathbf{q}_k^\top \mathbf{w} \sim \mathcal{N}(0, 1) \quad (15)$$

Note that the series given by the *de-correlated* coefficients c_k is in fact a truncated Karhunen–Loève expansion [6].

4.2. Measurement Models

Denote a measurement operation by $\mathcal{M}$, which takes as input a video source, ρ or $V(\mathbf{x})$, and outputs measurements $\mathbf{y}$. In the following sections we explore three measurement models with an increasing level of complexity and realism.

4.2.1 Direct GRF Measurement

We first analyze the inverse problem in a setting where the static envelope has no effect on the measurements. In particular, we define a simplistic measurement process with direct access to GRF pixel values:

$$\mathbf{y} = \mathcal{M}[\rho(\mathbf{x})] + \mathbf{n} = \rho + \mathbf{n}. \quad (16)$$

Here ρ is the discretized GRF and $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \Sigma_y)$ is measurement noise. Since the measurements do not depend on $\mathbf{I}$, Eq. (10) can be written as

$$\Theta^* = \arg \min_{\Theta} \mathcal{D}_\rho(\Theta|\mathbf{y}). \quad (17)$$

Using $\hat{\rho}$ (Eq. 14) as the forward model yields

$$\mathcal{D}_\rho(\Theta|\mathbf{y}) = \min_{\mathbf{c}} \|\mathbf{y} - \mathbf{U}_\Theta \Sigma_\Theta \mathbf{c}\|_{\Sigma_y}^2 + \|\mathbf{c}\|^2. \quad (18)$$

Note that Eq. (18) is equivalent to the *maximum a posteriori* (MAP) estimation of $\mathbf{c}$, where the prior comes from Eq. (15). Based on Eq. (18) we define a projection operator

$$\mathcal{P}(\mathbf{y}, \Theta) = \mathbf{U}_\Theta \Sigma_\Theta \mathbf{c}^*, \quad (19)$$

$$\mathbf{c}^* = \arg \min_{\mathbf{c}} \|\mathbf{y} - \mathbf{U}_\Theta \Sigma_\Theta \mathbf{c}\|_{\Sigma_y}^2 + \|\mathbf{c}\|^2.$$

For a fixed Θ , this is a projection of the measurements onto a subspace spanned by the top modes of the SPDE system. Although gradients can be derived, in practice, we evaluate $\mathcal{D}_\rho$ for each proposed parameter Θ (which includes solving for $\mathbf{U}_\Theta$ and Σ_Θ), and identify the Θ^* that minimizes $\mathcal{D}_\rho$.

For the accretion disk model (Sec. 3.1) we define the unknown parameters as $\Theta = \{\varphi_0, \varphi_1\}$ and assume measurement noise $\Sigma_y = \mathbb{I}$. Figure 3 shows the top modes, and projections for two different parameter settings Θ_1, Θ_2 . When Θ is closer to Θ_{true} the projection retains more flow features. We analyze the two-dimensional (2D) data-fit manifold defined by Eq. (18) in Figure 4. This figure shows the manifold for four cases where measurements are simulated with different Θ_{true} . Note that $\mathcal{D}_\rho$ is periodic with respect to φ_0 and φ_1 with 2π and π periods, respectively. Additional loss manifolds are shown in the Supp. material.

4.2.2 Direct Video Measurement

We now analyze the inverse problem in a setting with direct access to the modulated envelope pixel values. In this model (defined in Eq. 9), measurements are affected by both static and dynamic parameters:

$$\tilde{\mathbf{y}} = \mathcal{M}[V(\mathbf{x})] + \tilde{\mathbf{n}} = \mathcal{M}\left[I(\mathbf{x}) e^{\rho(\mathbf{x})}\right] + \tilde{\mathbf{n}} = \mathbf{V} + \tilde{\mathbf{n}}. \quad (20)$$

Here $\mathbf{V}$ is the discretized video. Taking into account both static and dynamic parameters, Eq. (10) is given by

$$\mathbf{I}^*, \Theta^* = \arg \min_{\mathbf{I}, \Theta} \mathcal{D}_V(\mathbf{I}, \Theta|\tilde{\mathbf{y}}), \quad (21)$$




Figure 4. A 2D data-fit manifold over possible $\Theta = (\varphi_0, \varphi_1)$ for four different $\mathbf{y} = \boldsymbol{\rho}$ that correspond to different Θ_{true} . The global minima track the ground-truth. Note that due to periodicity the manifolds wrap around in both the horizontal and vertical axes.

To estimate the dynamic parameters we define a non-linear pre-processing operation. Let $\log(\cdot)$ denote an element-wise logarithm, then

$$\log(\tilde{\mathbf{y}}) \simeq \log(\mathbf{I}) + \boldsymbol{\rho} + \mathbf{n}. \quad (22)$$

Here the noise is approximated as additive Gaussian $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Since the temporal mean over the GRF and measurement noise converges to zero,

$$\mathbb{E}_{x_0}[\log(\tilde{\mathbf{y}})] \simeq \log(\mathbf{I}). \quad (23)$$

Thus, by subtracting the temporal mean from Eq. (22), we can isolate the dynamic video component

$$\mathbf{y} = \log(\tilde{\mathbf{y}}) - \mathbb{E}_{x_0}[\log(\tilde{\mathbf{y}})] \simeq \boldsymbol{\rho} + \mathbf{n}. \quad (24)$$

Overall the measurement model of Eq. (24) is approximately equivalent to Eq. (16) and identical expressions to Eqs. (17)–(18) can be derived to solve for Θ^* .

4.2.3 Interferometric Measurements

In radio astronomy, interferometry is an imaging approach that can achieve high angular resolution. This is done by correlating signals captured at multiple telescopes simultaneously observing a radio source. The EHT array is an extreme example of Very Long Baseline Interferometry (VLBI), where telescopes across the globe are synchronized to image the emission around a single black hole.

Interferometric measurements, known as *visibilities* denote by ν , are related to the source video via a sparse set spatial frequencies, k_1, k_2 . We decompose $\mathcal{M}$ into a dense Fourier transform and a sparse sampling operator:

$$\mathbf{y} = \mathcal{M}[V(\mathbf{x})] + \mathbf{n} = \mathcal{S} \circ \mathcal{F}[I(\mathbf{x}) e^{\rho(\mathbf{x})}] + \mathbf{n}. \quad (25)$$

Here $\mathcal{F}$ is the 2D spatial Fourier transform as a function of




Figure 5. EHT synthetic measurement locations of a ~ 12 hour observation period (color-coded by time) observing the black hole Sgr A*. The EHT2017 array is comprised of the eight radio telescopes that were observing during the 2017 campaign. The EHT++ array contains an additional 13 prospected telescopes.

time x_0 . The operator $\mathcal{S}$ samples the 2D Fourier transform at specific times, and spatial frequencies. For VLBI arrays, sampled frequencies are dictated by the *projected baselines* of the observing telescopes. Thus, as the Earth rotates with respect to the source, EHT measurements carve out elliptical trajectories through the Fourier plane (Fig. 5). The challenge of imaging a dynamic source is that the source is evolving as measurements are being acquired.

Similarly to Eq. (21), estimation takes into account both the unknown static image and dynamic parameters. The data-fit for visibility measurements is given by

$$\mathcal{D}_\nu(\mathbf{I}, \Theta | \mathbf{y}) = \min_{\mathbf{c}} \|\mathbf{y} - \mathcal{S} \circ \mathcal{F}[I(\mathbf{x}) e^{\hat{\rho}(\mathbf{x})}]\|_{\Sigma_y}^2 + \|\mathbf{c}\|^2 \quad (26)$$

In practice, we solve Eq. (21) using block-coordinate descent, alternating between two minimization problems:

$$\mathbf{I}^* = \arg \min_{\mathbf{I}} \mathcal{D}_\nu(\mathbf{I}, \Theta^* | \mathbf{y}) + \mathcal{R}(\mathbf{I}), \quad (27)$$

$$\Theta^* = \arg \min_{\Theta} \mathcal{D}_\nu(\mathbf{I}^*, \Theta | \mathbf{y}). \quad (28)$$

We discuss optimization of $\mathbf{I}^*$ in detail in the supplemental material [1]. To efficiently optimize Eq. (28) for recovery of Θ^* , we linearize the measurement model:

$$\mathcal{S} \circ \mathcal{F}[I(\mathbf{x}) e^{\hat{\rho}(\mathbf{x})}] \simeq \mathcal{S} \circ \mathcal{F}[I(\mathbf{x})] + \mathcal{S} \circ \mathcal{F}[I(\mathbf{x}) \hat{\rho}(\mathbf{x})]. \quad (29)$$

This linearization enables approximation of Eq. (26) as a linear convex minimization (see Supp. for more details)

$$\hat{\mathcal{D}}_\nu(\mathbf{y}, \mathbf{I}^*, \Theta) = \min_{\mathbf{c}} \|\tilde{\mathbf{y}} - \mathbf{Z}_{\Theta, \mathbf{I}^*} \Sigma_{\Theta} \mathbf{c}\|_{\Sigma_y}^2 + \|\mathbf{c}\|^2 \quad (30)$$

for $\tilde{\mathbf{y}} = \mathbf{y} - \mathcal{S} \circ \mathcal{F}[I^*(\mathbf{x})],$

where $\tilde{\mathbf{y}}$ can be intuitively interpreted as the *current estimate* (since it depends on $\mathbf{I}^*$) of the dynamic portion of the visibility measurements. Here the k 'th column of $\mathbf{Z}_{\Theta, \mathbf{I}^*}$ is the sampled Fourier transform of mode u_k : $\mathcal{S} \circ \mathcal{F}[I(\mathbf{x}) u_k(\mathbf{x})]$. Note that $\mathbf{Z}_{\Theta, \mathbf{I}^*}$ depends on both the static envelope and dynamic parameters.




Figure 6. Joint estimation of $\mathbf{I}^*$ and $\Theta^* = (\varphi_0^*, \varphi_1^*)$ with four different sets of measurements, simulated with the same envelope $\mathbf{I}^{\text{true}}$. Each row shows recovery from video measurements (Sec. 4.2.2) with different dynamic parameters $\Theta_1^{\text{true}}, \Theta_2^{\text{true}}$. The second column shows a single frame of the true GRFs with CCW rotation [top] and radial inflow [bottom]. For each EHT array, recovery is shown by $\mathbf{I}^*$ and $\mathcal{D}$ where Θ^* is given by the global minima. Note that due to periodicity the manifolds wrap in both the horizontal and vertical axes. Our method is able to identify flow parameters close to truth and demonstrates robustness to errors in the envelope reconstruction.

5. Numerical Simulations

We analyze the performance of our proposed method on recovering the static envelope $\mathbf{I}$ and dynamic parameters $\Theta = (\varphi_0, \varphi_1)$ of simulated astronomical accretion disks. In Section 5.1 we demonstrate recovery when the underlying video obeys the assumed GRF forward model. In Section 5.2 we show that our method is able to generalize to realistic complex black hole simulations that are *not* derived from our assumed GRF fluid-flow model.

Synthetic Data Generation For each specified source video, we generate synthetic visibility measurements corresponding to 12 hours of observation. Visibilities are sampled from two telescope arrays: EHT2017 and EHT++. The spatial frequency coverage sampled by the two arrays is shown in Fig. 5. EHT2017 contains only existing telescopes that were used during the 2017 campaign. EHT++ is a future projection that has 13 additional radio telescopes. Thus, recovery with EHT2017 data is more challenging due to severe measurement sparsity. VLBI measurements are synthesized using the `eht-imaging` [9] library, with a field-of-view (FOV) of $160\mu\text{as}$ and no additional noise.

Stochastic Parameter Representation For stochastic parameter estimation (Eq. 28) we grid $\Theta = (\varphi_0, \varphi_1)$ with 20×20 grid points. For each angle-pair we compute $K=60$ modes of the SPDE using RSI with $q=10$ iterations (see Alg. 1). The values of K and q were chosen empirically. During each iteration of RSI we solve Eq. (1) on a discrete spatio-temporal grid of size $N_t \times N_x \times N_y = 64^3$, using the `HYPRE` [13] computing library.

Envelope Estimation For the static envelope reconstruction (Eq. 27) we use Maximum Entropy Method

(MEM) [20] regularization as $\mathcal{R}(\mathbf{I})$. The MEM regularization uses a prior image of 2D Gaussian with a standard deviation of $50\mu\text{as}$ [35] (see Supp. material for more details).

5.1. SPDE Sampled Videos

To validate our approach, with access to ground-truth parameters Θ , we generate videos that mimic black hole emission using the SPDE forward model (Sec. 3). The true envelope $\mathbf{I}^{\text{true}}$ contains a ring structure [30] with a radius of $\sim 30\mu\text{as}$ that resembles a black hole shadow (Fig. 6 left).

Figure 6 shows estimation results obtained from EHT2017 and EHT++ measurements for two different evolution parameters $(\Theta_1^{\text{true}}, \Theta_2^{\text{true}})$ with the same envelope, $\mathbf{I}^{\text{true}}$. For both EHT arrays, we recover an envelope with a ring shape. The 2D loss manifolds (Eq. 28) are also shown for the two parameter settings. The true parameters Θ^{true} are indicated by a white triangle on the plots; the global minimum, Θ^* , obtained by our proposed method is indicated by the red circle. We find that Θ^* is able to roughly track Θ^{true} , even with very sparse EHT2017 measurements. In particular, note that the envelope recovery with EHT2017 is degraded compared to EHT++; nevertheless, we find that the dynamic parameter estimation is somewhat robust to these errors.

We compare our results to a state-of-the-art VLBI imaging approach, *StarWarp*s (SW) [7], which was developed to handle time-varying interferometric datasets. The output of SW is an estimated full length video, V^{SW} , from which we extract the mean OpticalFlow (OF) field [14] to approximate the velocity field. Figure 7 highlights the advantage of having a model for the dynamic evolution in the setting of very sparse measurements made by EHT2017. In particular, for the EHT2017 array, V^{SW} 's low-fidelity prevents




Figure 7. Comparison of the proposed (model-based) approach with StarWarps (SW) [7] for EHT2017 measurements. The top row shows the true velocity field (captured in φ_0) [left], the estimated velocity field [center], and the mean OpticalFlow (OF) field recovered from V^{SW} [right]. The bottom row shows a frame from the true movie sequence [right], a synthetic frame generated with the recovered parameters and envelope $V^* = I^* e^{\rho^*}$ [center], and a frame from V^{SW} [left]. Green ticks represent the true and recovered spatial correlations (captured in φ_1). Our model fitting approach is better able to recover the stochastic evolution of the flow field from sparse data (see animations in the Supp. clip).

accurate recovery of dynamic properties using subsequent video analysis. Additional comparisons are provided in the supplemental material [1].

5.2. GRMHD Black Hole Videos

We test our approach in a more realistic setting, using physics-based video simulations of black hole accretion. These movies are generated with General Relativistic Magneto Hydrodynamics (GRMHD) [31]. Each simulation takes around 50K core hours to compute.

While we do not have access to ground-truth parameters for GRMHD movies, we can compare the recovery obtained from VLBI measurements to the recovery obtained directly from the GRMHD video, as formulated in Sec. 4.2.2. Recall that we expect the pre-processing step of Eq. (24) to yield a video with similar statistics to the sampled GRFs. Moreover, the 2D loss manifold recovered directly from the pre-processed video is an upper bound of what we can expect to achieve from sparse EHT measurements. Figure 8 shows the recovered envelope and loss manifolds for video data, as well as EHT++ and EHT2017 measurements. In all three cases we are able to estimate the clockwise (CW) rotational field. The loss manifolds for the EHT based recoveries resemble the loss from the video data, with global minima (red dot) around $\varphi_1 = -60^\circ$. Note that the local minima at $\varphi_0 \simeq 0.0$ correspond to an opposite rotation (CCW) with the same $\varphi_1 = -60^\circ$. While in this work we consider the entire 2D parameter space, these local minima could be regularized or suppressed as an unlikely (non-physical) flow.




Figure 8. Recovery of the envelope and dynamic parameters of a realistic black hole simulation from video data [top], EHT++ measurements [center], and EHT2017 measurements [bottom]. Left column: a frame from the GRMHD video and its preprocessed dynamic component (Eq. 24). Right column: a frame from a GRF sampled with the estimated parameters Θ^* . The correlations in $\rho(\Theta^*)$ resemble those in the preprocessed GRMHD.

6. Discussion and Conclusion

Motivated by imaging black hole evolution, we formulate an approach to infer fluid-dynamics from interferometric measurements. The sparsity of the measurements makes direct imaging challenging and highly ill-posed. In contrast, our approach relies on a stochastic surrogate model that captures key features of the flow. We analyze our approach through simulations and demonstrate its advantage over state-of-the-art interferometric imaging.

The primary limitation of our approach is the assumption of stationarity; however, in Sec. 5.2 we demonstrate recovery on non-stationary GRMHD flows. In the future, we plan to extend the flow parameterization (Sec. 3.1) to capture black holes observed at high inclination angles, and infer parameters from noisy measurements [9].

Finally, we believe key insights from this work are valuable for other scientific domains with complex fluid-dynamics at their core. Potential applications include, for example, remote sensing of atmospheric patterns (storms, clouds, dust) from geo-stationary satellites.

7. Acknowledgements

The authors would like to thank George Wong for his help with GRMHD simulations. AL is supported by the Zuckerman and Viterbi postdoctoral fellowships. This work was supported by NSF award 1935980: “Next Generation Event Horizon Telescope Design,” and Beyond Limits, and NSF awards 1743747, 1716327, and 2034306, XSEDE allocation TG-AST170024, and TACC Frontera LSCP AST20023. JAT was supported by ONR BRC Award N00014-18-1-2363 and NSF FRG Award 1952735.

References

- [1] Inference of black hole fluid-dynamics from sparse interferometric measurements: project page. http://imaging.cms.caltech.edu/stochastic_inference/.
- [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. ii. array and instrumentation. *The Astrophysical Journal Letters*, 875(1):L2, 2019.
- [3] 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. V. physical origin of the asymmetric ring. *The Astrophysical Journal Letters*, 875(1):L5, 2019.
- [4] Tomer Amiaz, Sándor Fazekas, Dmitry Chetverikov, and Nahum Kiryati. Detecting regions of dynamic texture. In *International conference on scale space and variational methods in computer vision*, pages 848–859. Springer, 2007.
- [5] Philipp Arras, Philipp Frank, Philipp Haim, Jakob Knollmüller, Reimar Leike, Martin Reinecke, and Torsten Enßlin. M87* in space, time, and frequency. *arXiv preprint arXiv:2002.05218*, 2020.
- [6] Wolfgang Betz, Iason Papaioannou, and Daniel Straub. Numerical methods for the discretization of random fields by means of the karhunen–loève expansion. *Computer Methods in Applied Mechanics and Engineering*, 271:109–129, 2014.
- [7] 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*, 4(4):512–527, 2018.
- [8] Katherine L Bouman, Michael D Johnson, Daniel Zoran, Vincent L Fish, Sheperd S Doeleman, and William T Freeman. Computational imaging for vlbi image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 913–922, 2016.
- [9] 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*, 857(1):23, 2018.
- [10] Mengyu Chu, You Xie, Jonas Mayer, Laura Leal-Taixé, and Nils Thuerey. Learning temporal coherence via self-supervision for gan-based video generation. *ACM Transactions on Graphics (TOG)*, 39(4):75–1, 2020.
- [11] Sheperd S Doeleman, Jonathan Weintraub, Alan EE Rogers, Richard Plambeck, Robert Freund, Remo PJ Tilanus, Per Friberg, Lucy M Ziurys, James M Moran, Brian Corey, et al. Event-horizon-scale structure in the supermassive black hole candidate at the galactic centre. *Nature*, 455(7209):78–80, 2008.
- [12] Gianfranco Doretto, Alessandro Chiuso, Ying Nian Wu, and Stefano Soatto. Dynamic textures. *International Journal of Computer Vision*, 51(2):91–109, 2003.
- [13] Robert D Falgout and Ulrike Meier Yang. hypre: A library of high performance preconditioners. In *International Conference on Computational Science*, pages 632–641. Springer, 2002.
- [14] Gunnar Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer, 2003.
- [15] Vincent L Fish, Sheperd S Doeleman, Christopher Beaudoin, Ray Blundell, David E Bolin, Geoffrey C Bower, Richard Chamberlin, Robert Freund, Per Friberg, Mark A Gurwell, et al. 1.3 mm wavelength vlbi of sagittarius a*: Detection of time-variable emission on event horizon scales. *The Astrophysical Journal Letters*, 727(2):L36, 2011.
- [16] Geir-Arne Fuglstad, Finn Lindgren, Daniel Simpson, and Håvard Rue. Exploring a new class of non-stationary spatial gaussian random fields with varying local anisotropy. *Statistica Sinica*, pages 115–133, 2015.
- [17] Koki Fujita and Shree K Nayar. Recognition of dynamic textures using impulse responses of state variables. In *Proc. Third International Workshop on Texture Analysis and Synthesis (Texture 2003)*, pages 31–36, 2003.
- [18] Andrea M Ghez, S Salim, NN Weinberg, JR Lu, Tea Do, JK Dunn, K Matthews, MR Morris, S Yelda, EE Becklin, et al. Measuring distance and properties of the milky way’s central supermassive black hole with stellar orbits. *The Astrophysical Journal*, 689(2):1044, 2008.
- [19] Xiaoyue Guan, Charles F Gammie, Jacob B Simon, and Bryan M Johnson. Locality of mhd turbulence in isothermal disks. *The Astrophysical Journal*, 694(2):1010, 2009.
- [20] Stephen F Gull and John Skilling. Maximum entropy method in image processing. In *IEEE Proceedings F (Communications, Radar and Signal Processing)*, volume 131, pages 646–659. IET, 1984.
- [21] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [22] 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.
- [23] Michael D Johnson, Vincent L Fish, Sheperd S Doeleman, Daniel P Marrone, Richard L Plambeck, John FC Wardle, Kazunori Akiyama, Keiichi Asada, Christopher Beaudoin, Lindy Blackburn, et al. Resolved magnetic-field structure and variability near the event horizon of sagittarius a. *Science*, 350(6265):1242–1245, 2015.
- [24] Daeyoung Lee and Charles F Gammie. Disks as inhomogeneous, anisotropic gaussian random fields. *The Astrophysical Journal*, 906(1):39, 2021.
- [25] Dahua Lin, Eric Grimson, and John Fisher. Modeling and estimating persistent motion with geometric flows. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2010.
- [26] Finn Lindgren, Håvard Rue, and Johan Lindström. An explicit link between gaussian fields and gaussian markov ran-

dom fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011.

- [27] Gunnar Martinsson, Adrianna Gillman, Edo Liberty, Nathan Halko, Vladimir Rokhlin, Sijia Hao, Yoel Shkolnisky, Patrick Young, Joel Tropp, Mark Tygert, et al. Randomized methods for computing the singular value decomposition (svd) of very large matrices. *Works. on Alg. for Modern Mass. Data Sets, Palo Alto*, 2010.
- [28] Per-Gunnar Martinsson and Joel A Tropp. Randomized numerical linear algebra: Foundations and algorithms. *Acta Numerica*, 29:403–572, 2020.
- [29] Lia Medeiros, Tod R Lauer, Dimitrios Psaltis, and Feryal Özel. Principal component analysis as a tool for characterizing black hole images and variability. *The Astrophysical Journal*, 864(1):7, 2018.
- [30] Ramesh Narayan, Michael D Johnson, and Charles F Gammie. The shadow of a spherically accreting black hole. *The Astrophysical Journal Letters*, 885(2):L33, 2019.
- [31] Oliver Porth, Koushik Chatterjee, Ramesh Narayan, Charles F Gammie, Yosuke Mizuno, Peter Anninos, John G Baker, Matteo Bugli, Chi-kwan Chan, Jordy Davelaar, et al. The event horizon general relativistic magnetohydrodynamic code comparison project. *The Astrophysical Journal Supplement Series*, 243(2):26, 2019.
- [32] Avinash Ravichandran, Rizwan Chaudhry, and Rene Vidal. Categorizing dynamic textures using a bag of dynamical systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(2):342–353, 2012.
- [33] Roi Ronen, Yoav Y Schechner, and Eshkol Eytan. 4D Cloud Scattering tomography. In *Proc. IEEE International Conference on Computer Vision*, 2021.
- [34] Havard Rue and Leonhard Held. *Gaussian Markov random fields: theory and applications*. CRC press, 2005.
- [35] 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.
- [36] 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.
- [37] Peter Whittle. On stationary processes in the plane. *Biometrika*, pages 434–449, 1954.
- [38] Jianwen Xie, Song-Chun Zhu, and Ying Nian Wu. Synthesizing dynamic patterns by spatial-temporal generative convnet. In *Proceedings of the ieee conference on computer vision and pattern recognition*, pages 7093–7101, 2017.
- [39] Guangming Zang, Ramzi Idoughi, Congli Wang, Anthony Bennett, Jianguo Du, Scott Skeen, William L Roberts, Peter Wonka, and Wolfgang Heidrich. Tomofluid: Reconstructing dynamic fluid from sparse view videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1870–1879, 2020.