NeuralDMD: Interpretable Neural Representation of Dynamics from Sparse and Noisy Measurements

Ali SaraerToosi[1,2], Renbo Tu[1,2], Esther Y. H. Lin[1,2], Kamyar Azizzadehnesheli[3], and Aviad Levis[1,2]

[1] University of Toronto, Toronto, ON, Canada, {asaraert, tutubo, lin, alevis}@cs.toronto.edu
[2] Vector Institute, Toronto, ON, Canada
[3] NVIDIA Corporation, USA, kamyar@nvidia.com

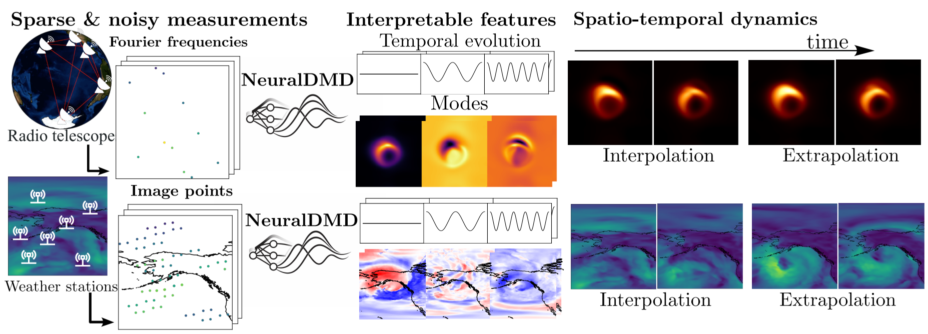

Fig. 1: **NeuralDMD** reconstructs spatio-temporal dynamics from sparse observations in both indirect and direct sensing regimes, including sparse Fourier (top) and image-domain (bottom) measurements. It models the field as a low-rank linear combination of spatial modes with simple spectral time evolution, yielding an interpretable decomposition that supports forecasting. We demonstrate this on black hole imaging (top) and atmospheric data assimilation (bottom).

Abstract. Many challenges in scientific imaging involve solving ill-posed inverse problems, where the goal is to recover spatio-temporal fields from indirect, noisy, and highly sparse measurements — often without access to ground truth data or reliable simulators. To address this challenging scenario, we present NeuralDMD, an interpretable, untrained (per-instance) reconstruction framework that combines neural implicit representations with Dynamic Mode Decomposition (DMD) to reconstruct continuous spatio-temporal dynamics directly from measurements. NeuralDMD parameterizes DMD modes as continuous neural fields, and imposes a low-rank linear dynamics prior with spectral time evolution to enforce temporal continuity. This formulation enables both forecasting under sparsity, and yields interpretable modes and spectra. We find that NeuralDMD outperforms baselines on a wide variety of tasks: from weather data assimilation from sparse station observations to interferometric (Fourier domain) observations of Sagittarius A*, the black hole at the center of our galaxy. Moreover, NeuralDMD remains stable when extrapolating into the future. While this framework is most naturally suited to linear dynamics, we show that it can be applied to nonlinear regimes, though with extrapolation performance that degrades with increasing nonlinearity. Together, these results show that NeuralDMD enables interpretable reconstruction and forecasting of spatio-temporal dynamics directly from sparse and indirect measurements without relying on numerical simulators or training data.

1 Introduction

A central challenge in scientific imaging is to recover continuous spatio-temporal dynamics from sparse and noisy measurements. Such measurements may be direct image-domain samples or indirect observations such as sparse Fourier measurements. To further compound the challenge, scientific imaging settings are data-constrained or lack known governing equations, making it difficult to generate simulation data. In this work, we propose NeuralDMD, a per-instance reconstruction method that supports both direct and indirect sensing regimes. We demonstrate it on weather data assimilation (image domain) and black hole imaging from the Event Horizon Telescope (EHT) interferometric observations (Fourier domain).

There have been efforts to address this problem of spatio-temporal reconstruction under sparse and noisy observations without simulation training. Neural fields (coordinate-based Multilayered Perceptrons or MLPs) can model high-dimensional signals of real world physical phenomena [41, 80] but require dense data, lack interpretability, and are sensitive to overfitting under sparse data conditions [52]. To mitigate issues, physics informed neural networks (PINNs) learn explicit analytical equations that can model fluids [50], but assume that the equations are known, making them biased and not suitable when physics is unknown [35]. Operator learning approaches are also capable of modeling physical phenomena but suffer from the same simulation-to-real domain gap, making them heavily biased on prior assumptions. Purely data-driven models (no simulation training) such as dynamic mode decomposition (DMD; [36]) preserve interpretability but are extremely difficult to use under high noise and for sparse data. To the best of our knowledge, no existing approach is simultaneously simulation-free, supports indirect operators, and provides interpretable forecasting prior to extreme sparsity.

We propose NeuralDMD, a model-free approach that enables spatio-temporal reconstruction from sparse, noisy, and indirect measurements. NeuralDMD fits a compact dynamical model directly to observations through the measurement operator. It enforces temporal coherence via a low-rank modal decomposition with explicit form of time evolution, and parameterizes modes as neural fields to enable grid-free evaluation. This yields an interpretable set of modes and spectra and supports forecasting beyond the observation window.

In summary, our contributions include:

- We propose a new representation for modeling dynamic physical processes: NeuralDMD.[4]
- We show how to optimize this representation on physical data from different imaging scenarios without requiring simulation training.
- We thoroughly evaluate the method and show state of the art results.

2 Related Work

We review prior approaches according to their data requirements; emphasizing whether they rely on simulation-trained priors, known governing equations, or densely observed ground-truth states.

Simulation-driven modeling. This class of modeling includes operator learning methods [3,15,33,34,44,75], principal component analysis (PCA; [45]), video predictor models [27], including recent approaches trained on numerical simulations to predict black hole dynamics [73], and some variations of INRs such as CORAL [64] that are designed for operator learning. These approaches, although very expressive for predicting in-distribution dynamics, are not designed for scientific discovery as they are heavily biased by the numerical simulations that were used for training and struggle to express out-of-distribution features. Moreover, most video predictor models rely on the availability of dense frame sequences and are incapable of reconstructing dynamics from sparse observations.

For NeuralDMD, we explicitly avoid training on simulations and directly fit the sparse and noisy measurements to enable scientific discovery. We only compare NeuralDMD to one simulation-trained model [64] to demonstrate the failure point of such models when the observed physical dynamics differ from what was seen in simulations.

Physics-informed modeling This class of models includes physics informed neural networks (PINNs) and some variations of implicit neural representations (INRs) that take advantage of prior physical knowledge of the underlying system. Existing latent-state approaches mainly (i) estimate parameters of a known model [14,24,29,74,82], or (ii) embed predefined governing equations into a network [26,74,82], which ties dynamics to assumed physics and limits performance under sparse, noisy, or poorly constrained conditions. Similar ideas appear in black hole imaging. BHNeRF [39] and PI-DEF [20] recover three-dimensional structure from observations, with PI-DEF additionally recovering the velocity field. In contrast, Levis et al. [38] recover image-plane dynamics through a modal decomposition, but assume a known PDE. NeuralDMD instead learns image-plane dynamics directly from observations without assuming a predefined dynamical model.

[4] Code publicly available here: https://github.com/pi-vision/NeuralDMD.git.

We do not compare NeuralDMD to this class of models, as they suffer from the same limitations as simulation-driven modeling: they either require simulation data for training or assume access to the PDEs that control the dynamics, a limiting assumption when the goal is scientific discovery rather than parameter estimation under a fixed physical model.

Untrained modeling. In this paper, we consider the set of problems within this category. Most INR and DMD variations, including NeuralDMD, are fully data driven, i.e., they rely only on the available data with little to no reliance on numerical simulations. Many scientific problems fall within this category, where it is necessary to minimize physical assumptions for scientific discovery (we discuss some scientific applications in Sec. 3).

Neural fields have been shown to have immense expressive power [42, 81]. They have been used to reconstruct 3-dimensional videos from a dense or sparse set of camera views [25, 47, 56]. However, typically it is assumed that we have complete image-plane measurements, i.e., no missing pixels and minimal measurement noise. In principle, neural fields can be used as a model-free approach to learn dynamics from sparse measurements. However, in practice, direct neural representation of the dynamics is used only when the physics is already known [26, 74, 82]. This is due to the large degree of freedom in a multi-layered perceptron (MLP), where even a small model would have millions of parameters, quickly leading to overfitting. As such, neural fields such as NeRF and dynamic variants [7, 21, 23, 40, 47, 56, 57, 79, 84], including sparse-view NeRFs [25], still require dense, high-quality observations and typically address only simple physical systems.

Modal decomposition methods offer a complementary perspective by distilling complex dynamics into the evolution of a few characteristic modes, i.e. a small number of degrees of freedom. Dynamic Mode Decomposition (DMD) learns a continuous-time linear operator directly from time-ordered snapshots, constructing a matrix A whose eigen-decomposition yields spatial modes associated with dominant coherent structures [61, 72]. As a result, DMD is model-free and data-driven [31, 37], and its low-rank structure makes it well-suited for reduced-order modeling and forecasting [46]. Related ideas also arise in dynamical-texture models [11,12,18,53], extending texture analysis to stochastic spatio-temporal motion such as smoke, water, and turbulence. However, standard DMD and its variations require access to dense video frames rather than sparse measurements and are sensitive to initialization. Furthermore, because they operate on grids as opposed to continuous fields, their memory and computational costs scale with the number of pixels, making them both memory-intensive and ill-suited for sparse and noisy observations.

3 Background and Notation

We review linear dynamical systems and Dynamic Mode Decomposition, which form the foundation of our approach.

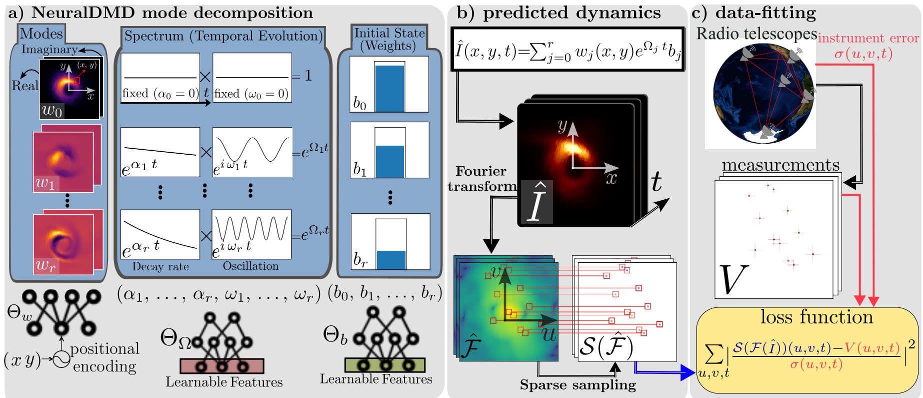

Fig. 2: (a) Block diagram of the NeuralDMD architecture. The model comprises three MLPs defining the spatio-temporal dynamics: Θ_w, mapping positionally encoded spatial coordinates (x, y) to mode values $w_j(x, y)$; Θ_Ω, producing complex spectral components (decay rate α, frequency ω) from a learnable latent vector; and Θ_b, mapping a learnable latent vector to the initial state b_j for temporal evolution. **(b), (c)** The learned modes and coefficients are combined to produce the predicted dynamics $\hat{I}(x, y, t)$. For the black hole setting, predictions are Fourier transformed ($\hat{F}$) and sampled at measured baseline coordinates ($\mathcal{S}(\hat{\mathcal{F}})$). The loss is computed between the observed visibilities V and the corresponding model prediction. For image-domain experiments (weather data assimilation) the loss is instead computed directly on the sparse measurements.

Linear dynamical systems. We consider an autonomous linear dynamical system $I_{t+1} = AI_t$, where A governs the temporal evolution of the system state. We assume A to be constant in time (see [83] for non-constant A) and no external inputs act on the system, so the dynamics are fully described by the state evolution [54, 55]. Such systems admit a modal decomposition in which the spatio-temporal field can be expressed as a low-rank superposition of spatial modes evolving in time,

$$I(x, y, t) = \sum_{j=1}^{r} w_j(x, y)\, a_j(t), \quad a_j(t) = b_j e^{\Omega_j t} \tag{1}$$

where $w_j(x, y)$ are spatial modes, $a_j(t)$ capture the temporal evolution, b_j are coefficients determined by the initial state of the system, and $\Omega_j = \alpha_j + i\omega_j$ are complex spectral coefficients whose real and imaginary parts encode growth/decay rates and oscillation frequencies. Having established the modal form of a linear dynamical system, we now turn to Dynamic Mode Decomposition (DMD), which provides a practical framework for estimating the parameters $\{w_j, \Omega_j\}$ from observation.

Dynamic Mode Decomposition (DMD). DMD estimates the modal decomposition of a dynamical system directly from time-ordered observations. Classical DMD models the system dynamics as a linear mapping between consecutive states or frames,

$$[I_1, I_2, \ldots, I_M] = A\,[I_0, I_1, \ldots, I_{M-1}] = A\,X, \tag{2}$$

where X is a data matrix whose columns are vectorized snapshots of the state at successive time steps. DMD estimates spatial modes w_j and spectral coefficients Ω_j from the eigendecomposition of a low-rank approximation of A, computed using singular value decomposition (SVD) [16, 36, 62]. While computationally efficient [58], classical DMD assumes direct access to densely observed system states. In many scientific imaging problems, however, the underlying state is never observed directly and must be reconstructed from indirect measurements.

Optimized DMD (optDMD). To address these limitations, optDMD [1, 59] reformulate modal extraction as a nonlinear least-squares problem,

$$\min_{\mathbf{w}_j, \Omega_j} \sum_{t \in \{t_1, \ldots, t_M\}} \left\| X_t - \sum_{j=1}^{r} \mathbf{w}_j e^{\Omega_j t} b_j \right\|_2^2, \tag{3}$$

where X_t denotes the observed system state at time t. The initial-state coefficients are computed as $b_j = \mathbf{w}_j^{\dagger} X_0$, where $\mathbf{w}_j^{\dagger}$ is the complex conjugate transpose of $\mathbf{w}_j$. This formulation improves robustness to noise and directly estimates a continuous-time spectrum Ω_j. However, optDMD still assumes densely observed states defined on a fixed spatial grid and direct access to X_0, limiting its applicability to sparse or indirect measurements.

4 NeuralDMD

We introduce NeuralDMD, which combines neural representations with the optDMD formulation by parameterizing the spatial modes $w_j(x, y)$ as neural fields and fitting the model directly to sparse measurements, without requiring densely observed states on a fixed grid, and inferring the initial state from all measurements rather than only those at the initial time.

4.1 NeuralDMD Representation

Building on the modal representation in Eq. (1), NeuralDMD extends optDMD by parameterizing the spatial modes as *continuous* neural fields [13, 47, 49, 85] with trainable weights $\Theta = [\Theta_w, \Theta_\Omega, \Theta_b]$ and learning the modal parameters directly from sparse measurements. Θ_w, Θ_Ω, and Θ_b parameterize the networks for the spatial modes $w_j(x, y)$, spectral coefficients Ω_j, and initial coefficients b_j, respectively (Fig. 2). In particular, we represent each spatial mode as a coordinate multilayer perceptron (MLP) network:

$$w_j(x, y) = w_j(x, y; \Theta_w), \tag{4}$$

allowing evaluation at arbitrary spatial coordinates.

To support stable optimization, we treat the spectral parameters and initial coefficients as learnable quantities and parameterize them with lightweight networks:

$$\Omega_j = \Omega_j(\Theta_\Omega), \qquad b_j = b_j(\Theta_b). \tag{5}$$

Crucially, in contrast to optDMD, this formulation eliminates the need for direct access to the initial state X_0. Instead b_j is inferred jointly from all measurements, since its evolution under the dynamical model must remain consistent with the entire observation sequence. The resulting prediction $\hat{I}(x, y, t)$ follows Eq. (1) with the specified neural parameterizations; the next subsections define how Θ is optimized from sparse (i) direct image-domain observations and (ii) indirect Fourier-domain measurements.

4.2 Reconstruction from Sparse Image-Domain Observations

The problem of reconstruction from direct image domain sparse observations is present in weather data assimilation [9]. To reconstruct a field from sparse spatio-temporal measurements at coordinates $\mathcal{S} = \{(x_k, y_k, t_k)\}_{k=1}^{N}$, with corresponding measurements $\{I(x_k, y_k, t_k)\}$, we fit the parameters, $\Theta^\star$, by minimizing the residual over observed samples,

$$\Theta^\star = \arg\min_{\Theta} \sum_{(x,y,t)\in\mathcal{S}} \left\| I(x, y, t) - \hat{I}(x, y, t) \right\|_2^2, \tag{6}$$

where $\hat{I}(x, y, t)$ is given by Eq. (1) with the neural parameterizations:

$$\hat{I}(x, y, t) = \sum_{j=1}^{r} w_j(x, y; \Theta_w) b_j(\Theta_b) e^{\Omega_j(\Theta_\Omega)t} \tag{7}$$

To ensure stable dynamics, we constrain the real part of the spectrum to be non-positive, i.e., $\alpha_j = \Re(\Omega_j) \le 0$ (see Supp. S1). This enforces physically meaningful decay and oscillatory behavior while preventing exponential growth. After optimization, forecasting can be done by evaluating $\hat{I}(x, y, t)$ at future times t. In Sec. 5, we explore the accuracy of extrapolation beyond the observed times.

4.3 Reconstruction from Indirect Fourier Measurements

In some imaging systems, observations are indirect and obtained as sparse Fourier samples or projections, such as MRI [67], Fourier ptychography [68], coded diffraction imaging [2], computational microscopy [8], and cryo-EM imaging [86]. In this work, we consider interferometry with the Event Horizon Telescope (EHT), a global interferometer used to image black holes (see e.g. [19]). Each telescope pair measures signal correlations, known as visibilities, which correspond to samples of the underlying source's Fourier transform. Mathematically, the visibility at spatial frequency (u, v) and time t is given by

$$V(u, v, t) = \mathcal{F}[I_t](u, v), \tag{8}$$

where I_t is the unknown image frame to be reconstructed and $\mathcal{F}$ denotes the 2D Fourier transform.

Observations are available only at a sparse set of Fourier coordinates: $\mathcal{S}_t = \{(u_k, v_k, t_k)\}_{k=1}^{N}$, which vary over time due to Earth-rotation [71]. We fit NeuralDMD by matching predicted and observed visibilities,

$$\Theta^\star = \arg\min_{\Theta} \sum_{(u,v,t)\in\mathcal{S}_t} \left\| V(u, v, t) - \mathcal{F}[\hat{I}_t](u, v) \right\|_2^2, \tag{9}$$

More generally, the Fourier transform in Eq. (9) can be replaced by an arbitrary differentiable measurement operator, extending NeuralDMD beyond Fourier imaging to nonlinear observation models and more general inverse problems.

For EHT interferometry, we fit models in the visibility domain using a χ^2 data term, defined as:

$$\chi^2(\mathbf{\Theta}) = \frac{1}{N} \sum_{i=1}^{N} \frac{|V_i - \hat{V}_i(\mathbf{\Theta})|^2}{\sigma_i^2}, \tag{10}$$

where N is the total number of observations and $\hat{V}_i$ is the model visibility. In our experiments, fitted models typically achieve $\chi^2 \approx 1$, indicating agreement with the measurements within their uncertainties (see Supp. S5.4 for details).

5 Experiments

5.1 Experimental Details

Experimental Setup For all experiments, we fit NeuralDMD by minimizing the appropriate measurement residual (Sec. 4) Here, an *epoch* denotes a fixed number of gradient steps. We uniformly batch over observed coordinate sets (image-domain $\mathcal{S}$ or Fourier-domain $\mathcal{S}_t$), using spatial batch sizes of 64–256 and selecting a per-epoch sampling fraction between 0.1 and 1.0 solely to control memory footprint; results are insensitive to this choice (see Supp). The number of modes r is chosen based on the sparsity (typically $r \in [5, 75]$). Ablations on sparsity versus mode count provided in the Supp.

Implementation Details We use a fixed architecture for NeuralDMD in all experiments. The mode network Θ_w comprises four MLP blocks, each consisting of two linear layers with 256 hidden units and SiLU activations. The spectrum and initial-coefficient networks, Θ_Ω and Θ_b, share a lightweight architecture: a learnable 16D latent input followed by a two-layer MLP with 64 hidden units

(SiLU) and a final linear layer producing per-mode outputs; Θ_Ω uses a sigmoid output mapping, while Θ_b uses a softplus on amplitudes and no activation on phases. Spatial inputs use NeRF-style sinusoidal positional encoding [47] with $L \in [0, 10]$ frequencies (powers of two), with smaller L used under higher sparsity. We use Adam optimizer with learning rate 10^{-4} and a reduce-on-plateau scheduler for 10^4–10^6 epochs depending on the experiment. All parameters are randomly initialized except for black hole imaging. There, we use structured initializations (disk pretraining for INR baselines and Zernicke initialization for NeuralDMD), with full details in the Supp. S5.

Baselines For image-domain sparse observations (weather), we compare NeuralDMD against a naive data-assimilation baseline: 3D-Var [4, 43], and two spatio-temporal coordinate networks that regress $I(x, y, t)$ from (x, y, t): a vanilla MLP and SIREN [66]. For Fourier-domain sparse observation experiments (black hole imaging), we compare against (i) a spatio-temporal INR baseline fit directly to visibilities, (ii) optDMD, (iii) StarWarps [6], which is a gold-standard dynamical imaging method used in the EHT community [19], and (iv) resolve [32], a more recent Bayesian image reconstruction method.

Ablations and Evaluation Metrics We ablate the INR backbone (SIREN [66], MFN [65], Fourier-feature MLP [69]) and the parameterization of (Ω, b), direct optimization versus networks Θ_Ω, Θ_b. For fairness, all INR baselines use the same MLP backbone as NeuralDMD 's mode network Θ_w, with time as an additional input and comparable positional encoding. All methods are fit under the same observation operator and sampling pattern. We report RMSE, PSNR, SSIM, and LPIPS for all reconstructions.

5.2 Image-Domain Sparse Observations

Weather Data Assimilation. We evaluate NeuralDMD on atmospheric data assimilation, reconstructing a complete spatio-temporal wind field $I(x, y, t)$ from image-domain observations at sparse station locations. We use hourly ERA5 wind-speed magnitude at 10m above the surface over North America from April 1–7, 2025 (168 frames; [77]), and sample observations at ~ 5000 real station coordinates per frame derived from WMO OSCAR/WIGOS [77,78]; with some stations dropping in and out on different days as their operational status changes. We fit models on the assimilation window and evaluate both reconstruction within the window and nowcasting beyond it (April 8, 2025). We compare against a 3D-Var baseline and spatio-temporal INR baselines (vanilla MLP and SIREN [66]) fit to the same sparse samples.

We report image metrics computed against dense ERA5 ground truth (mean $\pm$ std across frames). As shown in Fig. 3 and Tab. 1, NeuralDMD achieves the lowest error across all metrics and maintains higher-quality reconstructions under sparsity, while also providing stable short-term forecasts beyond the observation window (additional sparsity studies are provided in the Supp.).

Table 1: Quantitative evaluation: Weather data assimilation (ERA5). Quantitative comparison (mean±std over frames) for dense (~ 5000) and sparse (~ 500) observations per frame. We report results at both sparsity levels to highlight performance in a highly undersampled, operationally realistic (~ 500) and in a denser regime (~ 5000), where 3D-Var is applicable. NeuralDMD achieves the best overall performance at both sparsity levels.

Method	Samples per frame	RMSE↓	PSNR↑	SSIM↑	LPIPS↓
NeuralDMD (ours)	~ 500	**2.3±0.1**	**19.1±0.8**	**0.53±0.03**	**0.59±0.02**
	~ 5000	**1.5±0.1**	**24.7±0.5**	**0.70±0.01**	**0.32±0.01**
3D-Var	~ 500	-	-	-	-
	~ 5000	3.6±0.8	18±4	0.48±0.09	0.5±0.1
Vanilla MLP	~ 500	2.7±0.2	17±1	0.49±0.04	0.65±0.03
	~ 5000	2.3±0.2	21.0±0.6	0.56±0.01	0.6±0.2
SIREN	~ 500	3.0±0.3	17±1	0.50±0.04	0.63±0.02
	~ 5000	2.3±0.3	21±1	0.63±0.02	0.52±0.02

Table 2: In distribution (ID) and out of distribution (OOD) quantitative performance of CORAL vs. NeuralDMD. Sparsity rates of 10% and 40% of the total image pixels are used for fitting for ID and OOD cases respectively.

	CORAL ID	NeuralDMD ID	CORAL OOD	NeuralDMD OOD
RMSE ↓	0.02±0.01	**(3.6±0.8)e-3**	0.05±0.02	**0.014±0.004**
PSNR ↑	39±1	**56±1.3**	30±0.3	**43±1**
SSIM ↑	0.89±0.0	**0.9994±0.0002**	0.69±0.05	**0.997±0.001**
LPIPS ↓	0.012±0.001	**0.0002±0.0001**	0.020±0.001	**0.002±0.001**

Wave Equation: Simulation-Trained vs. Untrained Modeling. Operator-learning models leverage large numerical simulation datasets to learn reusable solution operators, but their performance can degrade under distribution shift. Here we compare this paradigm against per-instance fitting with NeuralDMD in a sparse-observation inverse problem.

We generate 2D wave-equation sequences and compare NeuralDMD against CORAL [64], a representative operator-learning baseline trained on 5000 64-frame videos of low-frequency waves of resolution 96×96. At test time, we evaluate both an in-distribution (ID) regime within the training frequency range and an out-of-distribution (OOD) regime containing higher spatial frequencies, while keeping the sparse sampling pattern fixed across methods. Table 2 shows quantitative results (qualitative results in Supp.), illustrating that CORAL generalizes well in-distribution (PSNR $= 39$) but can degrade under frequency shift, whereas NeuralDMD remains reliable because it is fit directly to the measurements of each sequence (details in Supp. S4.2). This highlights a key advantage of per-instance fitting, particularly relevant in scientific discovery settings, where the true dynamics may be only partially understood and representative training simulations are often unavailable.

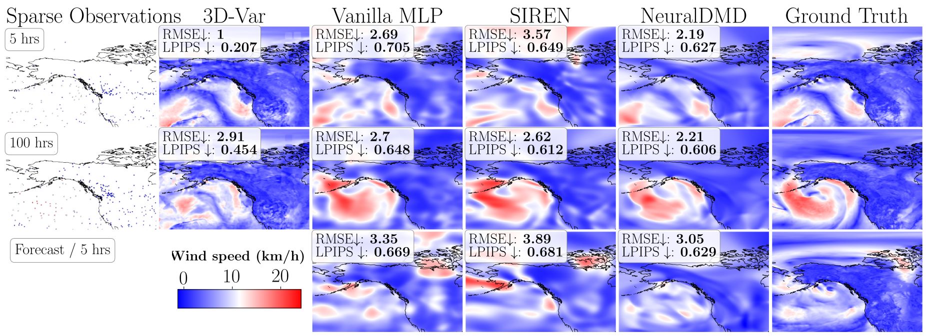

Fig. 3: Qualitative evaluation: Weather data assimilation (ERA5). Rows show t=5 hrs, t=100 hrs, and extrapolation, with $\sim$ 500 samples per-frame. 3D-Var is accurate near initialization but degrades over time and cannot extrapolate. NeuralDMD better preserves storm structure and forecasts more reliably than INR baselines, which lack explicit dynamics.

5.3 Fourier-Domain Sparse Observations

We evaluate NeuralDMD on black hole video reconstruction from sparse interferometric measurements. We simulate EHT observations using the array configuration and observing schedule from the 2017 Sagittarius A* campaign, and report reconstruction and extrapolation performance against established baselines.

Interferometric Observation Model Simulations. We use physically realistic General Relativistic Magnetohydrodynamics (GRMHD) simulations of black hole accretion flows [17,22,48] for evaluation. Corresponding ray-traced image sequences ($I_t(x, y)$) are used as ground truth and passed through an interstellar scattering screen to mimic the Galactic-center observing conditions [28, 30].

To simulate sparse visibilities, we use `eht-imaging` [10] together with `ngehtsim` [51] which model Earth-rotation synthesis and sparse Fourier sampling of the EHT 2017 array configuration for Sgr A* (`EHT2017`; [70]). This yields a set of complex visibilities $\{V_i\}$ at time-varying spatial frequencies $\{(u_i, v_i, t_i)\}$ with associated uncertainty estimates $\{\sigma_i\}$. Full simulation details, uv-coverage patterns, and additional configurations (e.g., edge-on views and denser arrays) are provided in Supp. S5.

Reconstruction Results and Comparisons. Figure 4 compares NeuralDMD against StarWarps [6], resolve [32], optDMD, and a spatio-temporal INR baseline (vanilla MLP). Under the extremely sparse `EHT2017` sampling, NeuralDMD yields lower reconstruction error and sharper, more temporally consistent reconstructions throughout the data-fitting window. This suggests that representing the dynamics through a small number of modes provides a strong inductive bias for reconstructing dynamics. Additional baseline sweeps and robustness studies under varying noise, sparsity, and corruptions are provided in Supp. S5.7. To

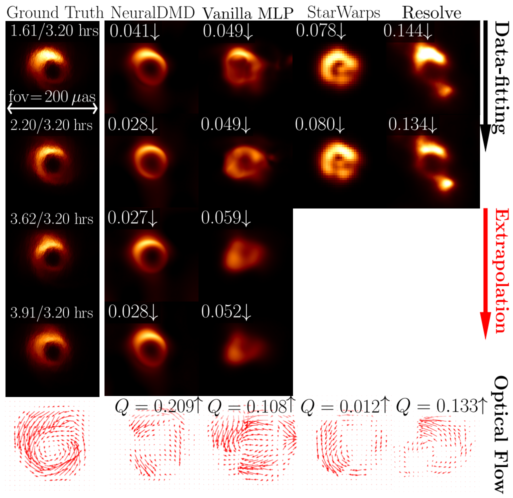

(a) GRMHD reconstruction from sparse interferometric data.

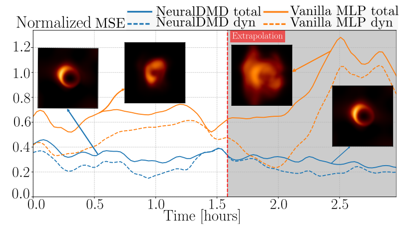

(b) Reconstruction error over time.

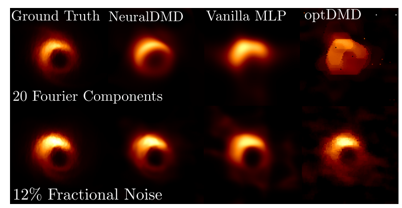

(c) Robustness to noise and sparsity.

Fig. 4: Spatio-temporal reconstruction from sparse EHT 2017 visibilities. (a) Reconstructions of a GRMHD sequence from sparse observations during fitting and extrapolation (forecasting). Values denote image-space MSE (lower is better). NeuralDMD produces lower error and better temporal consistency than `StarWarps`, resolve, and the spatio-temporal INR (vanilla MLP); The bottom row shows the average optical flow, with Q denoting cosine similarity to the ground truth flow (higher is better). **(b)** Per-frame normalized image-space MSE. Gray region denotes the extrapolation window, where NeuralDMD remains stable in contrast to the vanilla MLP diverges. **(c)** Robustness under increased visibility noise and reduced Fourier sampling. Fractional noise is relative to the data amplitude. A 100×100 Fourier grid is used; 20 Fourier components correspond 0.2% sparsity.

ensure a fair comparison, we use random initialization with the same seed values for NeuralDMD and the vanilla MLP, and use the same architecture for all baselines (see Supp. for details). In real-world setting, each station observes through a different, time-varying atmospheric column and instrumental chain, introducing unknown complex-valued gain corruptions (affecting both amplitude and phase) in the measured signal. To handle these realistic station-based effects, we marginalize over the complex gains during fitting (details in Supp. S5 and S15.2).

Forecasting Beyond Observations. After fitting the model to the observations, we forecast future frames by evaluating the learned dynamics at times beyond the last measurement. In the EHT 2017 setting, we fit the model on the first 1.6 hours forecast the subsequent 1.4 hours. Figure 4b shows the per-frame normalized image-space L_2 error $\|\hat{I}_t - I_t\|_2^2 / \|I_t\|_2^2$, with the extrapolation region shaded in gray. NeuralDMD continues to produce coherent motion throughout the forecast horizon, whereas the vanilla MLP rapidly diverges once observations are no longer available. Since StarWarps estimates frames only within observation interval, it does not support forecasting.

Interpretability. Beyond reconstruction quality, NeuralDMD recovers spatial modes $\{w_j(x,y)\}$ together with a continuous-time spectrum $\{\Omega_j = \alpha_j + i\omega_j\}$. The modes capture dominant spatial structures, while the spectrum characterizes their temporal evolution through decay rates and oscillation frequencies. In the black hole setting, for example, the learned frequencies provide direct estimates of characteristic dynamical timescales, such as the angular velocity of rotating structures (examples provided in the Supp. S5).

5.4 Robustness and Ablations

Table 3 ablates coordinate-network backbones and positional-encoding bandwidth on EHT 2017 reconstructions. Here, L denotes the number of sinusoidal positional-encoding frequencies used for coordinate inputs; we keep the number of layers and channels fixed across all INRs. For each method we report a low- and high-bandwidth setting $(L_{\text{low}} \mid L_{\text{high}})$, and results across a wider range of bandwidths can be found in the Supp. S5.7. Reported values are mean±std over 5 random initializations and measure reconstruction quality on recovered image frames (not forecasting). Across all backbone and bandwidth settings, NeuralDMD achieves the best overall reconstruction quality, indicating that the dynamical inductive bias provides benefits beyond the choice of INR architecture alone. We also ablate the parameterization of (Ω, b) by directly optimizing $\{\Omega_j, b_j\}$ as learnable parameters rather than neural network outputs, while keeping the same mode network and measurement loss ("Direct Ω/b"). Directly optimizing $\{\Omega_j, b_j\}$ yields the weakest performance across all metrics, highlighting that over-parameterizing these quantities as neural network outputs is crucial for effective optimization. We further evaluate the robustness of NeuralDMD under varying noise levels and observation sparsity, demonstrating strong reconstruction performance in challenging measurement regimes (Supp.S5).

5.5 Limitations: Breakdown of the Linear Dynamics Assumption

A key assumption of NeuralDMD and other DMD-based methods is that the observed dynamics are sufficiently linear over the observation window [63, 72]. We study the limits of this assumption using Navier–Stokes cylinder-flow simulations [60], generated with the `jax-fluids` library [5], whose nonlinearity increases with Reynolds number. The PDE for this systems is

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \frac{1}{\mathrm{Re}} \nabla^2 \mathbf{u}, \ \nabla \cdot \mathbf{u} = 0, \tag{11}$$

where $\mathbf{u}$ is the fluid velocity field, p is the pressure, and Re is the Reynolds number. This system is nonlinear due to the convective term $(\mathbf{u} \cdot \nabla)\mathbf{u}$. As Re increases

Table 3: EHT 2017 ablations. Reconstruction metrics for different coordinate-network backbones, positional-encoding bandwidths L, and the direct optimization baseline for (Ω, b).

Method	$L_{\text{low}} \mid L_{\text{high}}$	RMSE↓	PSNR↑	SSIM↑	LPIPS↓
SIREN	3	0.072±0.011	23.0±1.2	0.58±0.032	0.29±0.026
	10	0.075±0.0058	23.0±0.68	0.56±0.019	0.32±0.011
MFN	3	0.13±0.0095	18.0±0.68	0.54±0.0087	0.30±0.0036
	10	0.15±0.025	17.0±1.6	0.53±0.020	0.32±0.022
FFM	256	0.093±0.0037	21.0±0.33	0.42±0.0061	0.66±0.018
	1024	0.094±0.0047	21.0±0.42	0.47±0.031	0.59±0.017
Vanilla MLP	3	0.051±0.0029	26.0±0.55	0.81±0.013	0.22±0.0040
	10	0.17±0.12	17.0±6.0	0.63±0.084	0.28±0.042
Direct Ω/b	2	0.13±0.0039	18.0±0.26	0.30±0.011	0.41±0.029
	10	0.19±0.0043	14.0±0.20	0.18±0.0002	0.63±0.013
NeuralDMD (ours)	**2**	**0.0349±0.010**	**29.18±1.1**	**0.6557±0.099**	**0.1905±0.0803**
	10	**0.066±0.0033**	**24.0±0.47**	**0.54±0.0021**	**0.31±0.040**

the flow transitions from smooth laminar behavior to increasingly turbulent dynamics, providing a controlled way to study the breakdown of low-dimensional linear representations.

We evaluate NeuralDMD on two cylinder-flow regimes at Re= 200 and Re= 5200, which exhibit increasing levels of nonlinear dynamical complexity. For reconstructions, we use 120 and 150 modes, respectively. Figure 5 reports PSNR versus normalized time, where $t_{\text{norm}} \in [0, 1]$ is the data-fitting window and $t_{\text{norm}} > 1$ is extrapolation. We compare against a *persistence* [76] baseline, which uses the previous frame during the observation window ($\hat{I}_t = I_{t-1}$), and the final observed frame for extrapolation. We define the failure point as when NeuralDMD matches the persistence PSNR.

While NeuralDMD achieves high fidelity during fitting (PSNR ~ 60 for Re= 200 and ~ 40 for Re= 5200), its forecasting accuracy deteriorates as nonlinearity increases. For both Reynolds numbers, NeuralDMD becomes comparable to the persistence baseline at longer horizons, around $t_{\text{norm}} \approx 1.6$, indicating the eventual breakdown of the low-dimensional linear approximation. At Re= 5200, a vanilla MLP can attain higher PSNR during fitting, but with less reliable extrapolation, consistent with its lack of an explicit dynamical prior.

6 Discussion

We introduce NeuralDMD, a simulation-free and interpretable framework for recovering spatio-temporal dynamics from sparse, noisy, and indirect observations while enabling forecasting beyond the observation window. Although we demonstrated the method using image domain and sparse Fourier measurements, the framework naturally extends to arbitrary differentiable observation operators, including dynamic acquisition processes whose measurements evolve on the

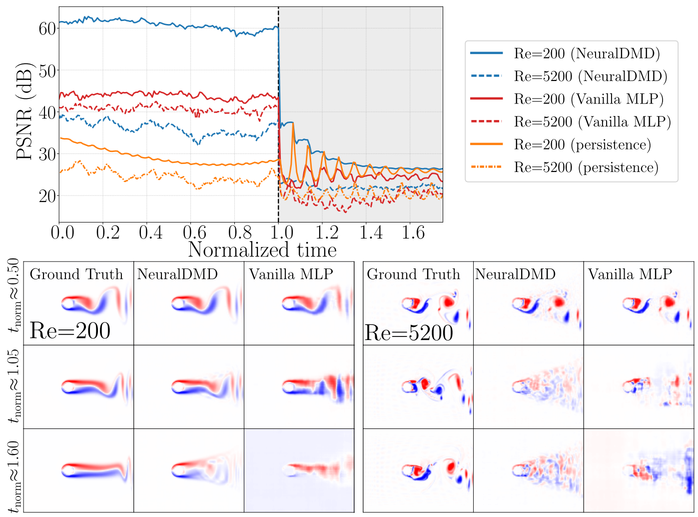

Fig. 5: top: PSNR vs. normalized time for Navier-Stokes cylinder flow simulations at Re= 5200 and Re= 200. The white and gray regions denotes the observation and forecasting regions respectively. **bottom:** Ground truth and NeuralDMD reconstructions at $t_{\mathrm{norm}} = 0.5,\ 1.05,\ 1.60$. Error grows more rapidly for the highly non-linear case of Re= 5200 than Re= 200 (the same fact is reflected in the PSNR plot).

same timescale as the underlying scene. By combining neural representations with the inductive bias of dynamical systems, NeuralDMD consistently outperforms conventional neural representations. A key feature of NeuralDMD is its explicit low-rank dynamical structure. Although the modeled dynamics are linear, many complex systems admit accurate low-dimensional approximations over finite time windows, yielding interpretable modes and spectra that capture their dominant spatio-temporal behavior.

More broadly, we believe NeuralDMD is particularly relevant for scientific imaging problems where the governing physics remains uncertain and faithful simulations are unavailable. In these settings, learning dynamics directly from observations allows the data to inform, refine, or challenge existing physical models rather than constraining inference to the behavior pre-encoded in simulations. We view this as an important step toward imaging algorithms that not only reconstruct phenomena but also help reveal the underlying physics and enable scientific discovery.

Acknowledgements. We thank S. Nousias, A. Broderick, M. Shalaby, A. Fuentes, M. Foschi, S. Rashidi, R. Dahale, D. Posselt, and M. Als for helpful comments, and E. Schnetter and D. Lang for their provision of computational resources. A. Levis's work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). This work was also supported by the Ontario Research Fund – Research Excellence under Project Number RE012-045.

References

1. Askham, T., Kutz, J.N.: Variable projection methods for an optimized dynamic mode decomposition (2017), https://arxiv.org/abs/1704.02343
2. Attal, B., O'Toole, M.: Towards mixed-state coded diffraction imaging. IEEE Transactions on Pattern Analysis and Machine Intelligence **47**(8), 6122–6133 (2025). https://doi.org/10.1109/TPAMI.2023.3340902
3. Azizzadenesheli, K., Kovachki, N., Li, Z., Liu-Schiaffini, M., Kossaifi, J., Anandkumar, A.: Neural operators for accelerating scientific simulations and design. Nature Reviews Physics **6**(5), 320–328 (2024)
4. Bannister, R.N.: A review of operational methods of variational and ensemble-variational data assimilation. Quarterly Journal of the Royal Meteorological Society **143**(703), 607–633 (2017). https://doi.org/https://doi.org/10.1002/qj.2982, https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.2982
5. Bezgin, D.A., Buhendwa, A.B., Adams, N.A.: Jax-fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. Computer Physics Communications (2022)
6. Bouman, K.L., Johnson, M.D., Dalca, A.V., Chael, A.A., Roelofs, F., Doeleman, S.S., Freeman, W.T.: Reconstructing video from interferometric measurements of time-varying sources (2018), https://arxiv.org/abs/1711.01357
7. Cao, A., Johnson, J.: Hexplane: A fast representation for dynamic scenes (2023), https://arxiv.org/abs/2301.09632
8. Cao, R., Divekar, N., Nuñez, J., Upadhyayula, S., Waller, L.: Neural space–time model for dynamic multi-shot imaging. Nature Methods **21**, 2336–2341 (09 2024). https://doi.org/10.1038/s41592-024-02417-0
9. Carrassi, A., Bocquet, M., Demaeyer, J., Grudzien, C., Raanes, P.N., Vannitsem, S., Chrust, M.: Data assimilation in the geosciences: An overview of methods, issues, and perspectives. WIREs Climate Change **9**(5), e535 (2018). https://doi.org/10.1002/wcc.535
10. Chael, A.: eht-imaging (May 2022). https://doi.org/10.5281/zenodo.7226661, https://doi.org/10.5281/zenodo.7226661
11. Chan, A., Vasconcelos, N.: Layered dynamic textures. In: Weiss, Y., Schölkopf, B., Platt, J. (eds.) Advances in Neural Information Processing Systems. vol. 18. MIT Press (2005), https://proceedings.neurips.cc/paper_files/paper/2005/file/2bc8ae25856bc2a6a1333d1331a3b7a6-Paper.pdf
12. Chan, A., Vasconcelos, N.: Variational layered dynamic textures. pp. 1062–1069 (06 2009). https://doi.org/10.1109/CVPRW.2009.5206556
13. Chang, Y., Benchekroun, O., Chiaramonte, M.M., Chen, P.Y., Grinspun, E.: Neural representation of shape-dependent laplacian eigenfunctions (2024), https://arxiv.org/abs/2408.10099
14. Chen, B., Huang, K., Raghupathi, S., Chandratreya, I., Du, Q., Lipson, H.: Automated discovery of fundamental variables hidden in experimental data. Nature Computational Science **2**(7), 433–442 (Jul 2022). https://doi.org/10.1038/s43588-022-00281-6, https://doi.org/10.1038/s43588-022-00281-6
15. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations (2019), https://arxiv.org/abs/1806.07366
16. Dey, S.: Dynamic mode decomposition and koopman theory (2022), https://arxiv.org/abs/2211.07561
17. Dhruv, V., Prather, B., Wong, G., Gammie, C.F.: A survey of general relativistic magnetohydrodynamic models for black hole accretion systems (2025), https://arxiv.org/abs/2411.12647

18. Doretto, G., Chiuso, A., Wu, Y., Soatto, S.: Dynamic textures. International Journal of Computer Vision **51**, 91–109 (02 2003). `https://doi.org/10.1023/A:1021669406132`
19. Event Horizon Telescope Collaboration, et al.: First Sagittarius A* Event Horizon Telescope Results. III: Imaging of the Galactic Center Supermassive Black Hole. The Astrophysical Journal Letters **930**(L14) (2023). `https://doi.org/10.3847/2041-8213/ac6429`
20. Feng, B.T., Chael, A.A., Bromley, D., Levis, A., Freeman, W.T., Bouman, K.L.: Dynamic black-hole emission tomography with physics-informed neural fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12511–12521 (June 2026)
21. Fridovich-Keil, S., Meanti, G., Warburg, F., Recht, B., Kanazawa, A.: K-planes: Explicit radiance fields in space, time, and appearance (2023), `https://arxiv.org/abs/2301.10241`
22. Gammie, C.F., McKinney, J.C., Toth, G.: Harm: A numerical scheme for general relativistic magnetohydrodynamics. The Astrophysical Journal **589**(1), 444–457 (May 2003). `https://doi.org/10.1086/374594`, `http://dx.doi.org/10.1086/374594`
23. Gao, Q., Xu, Q., Cao, Z., Mildenhall, B., Ma, W., Chen, L., Tang, D., Neumann, U.: Gaussianflow: Splatting gaussian dynamics for 4d content creation (2024), `https://arxiv.org/abs/2403.12365`
24. Garcia, A.C., van Gemert, J., Brinks, D., Tömen, N.: Learning physics from video: Unsupervised physical parameter estimation for continuous dynamical systems (2025), `https://arxiv.org/abs/2410.01376`
25. Guangcong, Chen, Z., Loy, C.C., Liu, Z.: Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
26. Hofherr, F., Koestler, L., Bernard, F., Cremers, D.: Neural implicit representations for physical parameter inference from a single video. In: 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2092–2102 (2023). `https://doi.org/10.1109/WACV56688.2023.00213`
27. Hsieh, J.T., Liu, B., Huang, D.A., Fei-Fei, L., Niebles, J.C.: Learning to Decompose and Disentangle Representations for Video Prediction. arXiv e-prints arXiv:1806.04166 (Jun 2018). `https://doi.org/10.48550/arXiv.1806.04166`
28. Issaoun, S., Johnson, M.D., Blackburn, L., Broderick, A., Tiede, P., Wielgus, M., Doeleman, S.S., Falcke, H., Akiyama, K., Bower, G.C., Brinkerink, C.D., Chael, A., Cho, I., Gómez, J.L., Hernández-Gómez, A., Hughes, D., Kino, M., Krichbaum, T.P., Liuzzo, E., Loinard, L., Markoff, S., Marrone, D.P., Mizuno, Y., Moran, J.M., Pidopryhora, Y., Ros, E., Rygl, K., Shen, Z.Q., Wagner, J.: Persistent non-gaussian structure in the image of sagittarius a* at 86 ghz. The Astrophysical Journal **915**(2), 99 (jul 2021). `https://doi.org/10.3847/1538-4357/ac00b0`, `https://dx.doi.org/10.3847/1538-4357/ac00b0`
29. Jaques, M., Burke, M., Hospedales, T.: Physics-as-inverse-graphics: Unsupervised physical parameter estimation from video (2020), `https://arxiv.org/abs/1905.11169`
30. Johnson, M.D., Narayan, R., Psaltis, D., Blackburn, L., Kovalev, Y.Y., Gwinn, C.R., Zhao, G.Y., Bower, G.C., Moran, J.M., Kino, M., Kramer, M., Akiyama, K., Dexter, J., Broderick, A.E., Sironi, L.: The scattering and intrinsic structure of sagittarius a* at radio wavelengths. The Astrophysical Journal **865**(2), 104 (sep 2018). `https://doi.org/10.3847/1538-4357/aadcff`, `https://dx.doi.org/10.3847/1538-4357/aadcff`

31. Kim, Y., Dean, S.: Sparse-to-field reconstruction via stochastic neural dynamic mode decomposition (2025), https://arxiv.org/abs/2511.20612
32. Knollmüller, J., Arras, P., Enßlin, T.: Resolving horizon-scale dynamics of sagittarius a* (2023), https://arxiv.org/abs/2310.16889
33. Kossaifi, J., Kovachki, N., Azizzadenesheli, K., Anandkumar, A.: Multi-grid tensorized fourier neural operator for high-resolution pdes (2023), https://arxiv.org/abs/2310.00120
34. Kovachki, N., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., Anandkumar, A.: Neural operator: learning maps between function spaces with applications to pdes. J. Mach. Learn. Res. **24**(1) (Jan 2023)
35. Krishnapriyan, A.S., Gholami, A., Zhe, S., Kirby, R.M., Mahoney, M.W.: Characterizing possible failure modes in physics-informed neural networks. arXiv e-prints arXiv:2109.01050 (Sep 2021). https://doi.org/10.48550/arXiv.2109.01050
36. Kutz, J.N., Brunton, S.L., Brunton, B.W., Proctor, J.L.: Dynamic Mode Decomposition. Society for Industrial and Applied Mathematics, Philadelphia, PA (2016). https://doi.org/10.1137/1.9781611974508, https://epubs.siam.org/doi/abs/10.1137/1.9781611974508
37. Kutz, J.N., Brunton, S.L., Brunton, B.W., Proctor, J.L.: Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems. Software, Environments, Tools, SIAM (2016). https://doi.org/10.1137/1.9781611974508
38. Levis, A., Lee, D., Tropp, J.A., Gammie, C.F., Bouman, K.L.: Inference of black hole fluid-dynamics from sparse interferometric measurements. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2340–2349 (October 2021)
39. Levis, A., Srinivasan, P.P., Chael, A.A., Ng, R., Bouman, K.L.: Gravitationally lensed black hole emission tomography. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 19841–19850 (2022)
40. Li, Z., Wang, Q., Cole, F., Tucker, R., Snavely, N.: Dynibar: Neural dynamic image-based rendering (2023), https://arxiv.org/abs/2211.11082
41. Lin, E.Y.H., Wang, Z., Lin, R., Miau, D., Kainz, F., Chen, J., Zhang, X., Lindell, D.B., Kutulakos, K.N.: Learning lens blur fields. IEEE Transactions on Pattern Analysis and Machine Intelligence **47**(9), 7383–7395 (Sep 2025). https://doi.org/10.1109/tpami.2025.3578587, http://dx.doi.org/10.1109/TPAMI.2025.3578587
42. Lin, E.Y.H., Wang, Z., Lin, R., Miau, D., Kainz, F., Chen, J., Zhang, X., Lindell, D.B., Kutulakos, K.N.: Learning lens blur fields. IEEE Transactions on Pattern Analysis and Machine Intelligence p. 1–12 (2025). https://doi.org/10.1109/tpami.2025.3578587, http://dx.doi.org/10.1109/TPAMI.2025.3578587
43. Lorenc, A.C.: Analysis methods for numerical weather prediction. Quarterly Journal of the Royal Meteorological Society **112**(474), 1177–1194 (1986). https://doi.org/10.1002/qj.49711247414
44. Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E.: Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. Nature Machine Intelligence **3**(3), 218–229 (Mar 2021). https://doi.org/10.1038/s42256-021-00302-5, http://dx.doi.org/10.1038/s42256-021-00302-5
45. Medeiros, L., Psaltis, D., Lauer, T.R., Özel, F.: Principal-component interferometric modeling (primo), an algorithm for eht data. i. reconstructing images from simulated eht observations. The Astrophysical Journal **943**(2), 144 (Feb 2023). https://doi.org/10.3847/1538-4357/acaa9a, http://dx.doi.org/10.3847/1538-4357/acaa9a

46. Mezić, I.: Analysis of fluid flows via spectral properties of the koopman operator. Annual Review of Fluid Mechanics **45**, 357–378 (2013). https://doi.org/10.1146/annurev-fluid-011212-140652
47. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis (2020), https://arxiv.org/abs/2003.08934
48. Moscibrodzka, M., Gammie, C.F.: ipole - semianalytic scheme for relativistic polarized radiative transport (2017), https://arxiv.org/abs/1712.03057
49. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics **41**(4), 1–15 (Jul 2022). https://doi.org/10.1145/3528223.3530127, http://dx.doi.org/10.1145/3528223.3530127
50. Pan, S., Brunton, S.L., Kutz, J.N.: Neural Implicit Flow: a mesh-agnostic dimensionality reduction paradigm of spatio-temporal data. arXiv e-prints arXiv:2204.03216 (Apr 2022). https://doi.org/10.48550/arXiv.2204.03216
51. Pesce, D.W., Blackburn, L., Chaves, R., Doeleman, S.S., Freeman, M., Issaoun, S., Johnson, M.D., Lindahl, G., Natarajan, I., Paine, S.N., Palumbo, D.C.M., Roelofs, F., Tiede, P.: ngehtsim (Dec 2023), https://github.com/Smithsonian/ngehtsim
52. Petrini, L., Cagnetta, F., Vanden-Eijnden, E., Wyart, M.: Learning sparse features can lead to overfitting in neural networks*. Journal of Statistical Mechanics: Theory and Experiment **2023**(11), 114003 (nov 2023). https://doi.org/10.1088/1742-5468/ad01b9, https://doi.org/10.1088/1742-5468/ad01b9
53. Previtali, D., Valceschini, N., Mazzoleni, M., Previdi, F.: Identification of dynamic textures using dynamic mode decomposition. IFAC-PapersOnLine **53**(2), 2423–2428 (2020). https://doi.org/https://doi.org/10.1016/j.ifacol.2020.12.045, https://www.sciencedirect.com/science/article/pii/S2405896320302974, 21st IFAC World Congress
54. Proctor, J.L., Brunton, S.L., Kutz, J.N.: Dynamic mode decomposition with control. SIAM Journal on Applied Dynamical Systems **15**(1), 142–161 (2016). https://doi.org/10.1137/15M1013857
55. Proctor, J.L., Brunton, S.L., Kutz, J.N.: Generalizing koopman theory to allow for inputs and control. SIAM Journal on Applied Dynamical Systems **17**(1), 909–930 (2018). https://doi.org/10.1137/16M1062296
56. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes (2020), https://arxiv.org/abs/2011.13961
57. Ren, J., Xie, K., Mirzaei, A., Liang, H., Zeng, X., Kreis, K., Liu, Z., Torralba, A., Fidler, S., Kim, S.W., Ling, H.: L4gm: Large 4d gaussian reconstruction model (2024), https://arxiv.org/abs/2406.10324
58. Salvador Jasin, D., Vava, R.: Svd-rom: Svd-based reduced-order modeling for large-scale data (Mar 2026). https://doi.org/10.5281/zenodo.19311360, https://doi.org/10.5281/zenodo.19311360
59. Sashidhar, D., Kutz, J.N.: Bagging, optimized dynamic mode decomposition for robust, stable forecasting with spatial and temporal uncertainty quantification. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **380**(2229) (Jun 2022). https://doi.org/10.1098/rsta.2021.0199, http://dx.doi.org/10.1098/rsta.2021.0199
60. Schäfer, M., Turek, S.: Benchmark computations of laminar flow around a cylinder. In: Flow Simulation with High-Performance Computers II (1996)
61. Schmid, P.J.: Dynamic mode decomposition of numerical and experimental data. Journal of Fluid Mechanics **656**, 5–28 (2010). https://doi.org/10.1017/S0022112010001217

62. SCHMID, P.J.: Dynamic mode decomposition of numerical and experimental data. Journal of Fluid Mechanics **656**, 5–28 (2010). https://doi.org/10.1017/S0022112010001217
63. Schmid, P.J.: Dynamic mode decomposition of numerical and experimental data. Journal of Fluid Mechanics **656**, 5–28 (2010)
64. Serrano, L., Boudec, L.L., Koupaï, A.K., Wang, T.X., Yin, Y., Vittaut, J.N., Gallinari, P.: Operator learning with neural fields: Tackling pdes on general geometries (2023), https://arxiv.org/abs/2306.07266
65. Shekarforoush, S., Lindell, D.B., Fleet, D.J., Brubaker, M.A.: Residual multiplicative filter networks for multiscale reconstruction (2022), https://arxiv.org/abs/2206.00746
66. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions (2020), https://arxiv.org/abs/2006.09661
67. Sodickson, D.K., Manning, W.J.: Simultaneous acquisition of spatial harmonics (smash): Fast imaging with radiofrequency coil arrays. Magnetic Resonance in Medicine **38**(4), 591–603 (1997). https://doi.org/10.1002/mrm.1910380414
68. Sun, M., Wang, K., Mishra, Y.N., Qiu, S., Heidrich, W.: Space-time fourier ptychography for in vivo quantitative phase imaging. Optica **11**(9), 1250–1260 (2024). https://doi.org/10.1364/OPTICA.531646
69. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains (2020), https://arxiv.org/abs/2006.10739
70. The Event Horizon Telescope Collaboration: First sagittarius a* event horizon telescope results. i. the shadow of the supermassive black hole in the center of the milky way. The Astrophysical Journal Letters **930**(2), L12 (may 2022). https://doi.org/10.3847/2041-8213/ac6674, https://dx.doi.org/10.3847/2041-8213/ac6674
71. Thompson, A.R., Moran, J.M., Swenson, G.W.: Very-Long-Baseline Interferometry, pp. 391–483. Springer International Publishing, Cham (2017)
72. Tu, J.H., Rowley, C.W., Luchtenburg, D.M., Brunton, S.L., Kutz, J.N.: On dynamic mode decomposition: Theory and applications. Journal of Computational Dynamics **1**(2), 391–421 (2014). https://doi.org/10.3934/jcd.2014.1.391
73. Tu, R., SaraerToosi, A., Conroy, N.S., Pekhimenko, G., Levis, A.: Bhcast: Unlocking black hole plasma dynamics from a single blurry image with long-term forecasting. CVPR (2026)
74. Weiss, S., Maier, R., Cremers, D., Westermann, R., Thuerey, N.: Correspondence-free material reconstruction using sparse surface constraints. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4685–4694 (2020). https://doi.org/10.1109/CVPR42600.2020.00474
75. Wen, G., Li, Z., Azizzadenesheli, K., Anandkumar, A., Benson, S.M.: U-fno – an enhanced fourier neural operator-based deep-learning model for multiphase flow (2022), https://arxiv.org/abs/2109.03697
76. Wilks, D.S.: Statistical Methods in the Atmospheric Sciences, International Geophysics, vol. 100. Academic Press, Amsterdam, 3 edn. (2011)
77. World Meteorological Organization (WMO): OSCAR/Surface: Observing Systems Capability Analysis and Review Tool. https://oscar.wmo.int/surface/, accessed: Nov. 2025

78. World Meteorological Organization (WMO): Guide to the WMO Integrated Global Observing System (WIGOS). Tech. Rep. WMO-No. 1165, World Meteorological Organization, Geneva, Switzerland (2019), https://library.wmo.int/records/item/55696-guide-to-the-wmo-integrated-global-observing-system
79. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering (2024), https://arxiv.org/abs/2310.08528
80. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond (2022), https://arxiv.org/abs/2111.11426
81. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. Computer Graphics Forum (2022). https://doi.org/10.1111/cgf.14505
82. Zhang, C., Cherniavskii, D., Zadaianchuk, A., Tragoudaras, A., Vozikis, A., Nijdam, T., Prinzhorn, D.W.E., Bodracska, M., Sebe, N., Gavves, E.: Morpheus: Benchmarking physical reasoning of video generative models with real physical experiments (2025), https://arxiv.org/abs/2504.02918
83. Zhang, H., Rowley, C.W., Deem, E.A., Cattafesta, L.N.: Online dynamic mode decomposition for time-varying systems. SIAM Journal on Applied Dynamical Systems **18**(3), 1586–1609 (2019). https://doi.org/10.1137/18M1192329
84. Zhang, W., Liu, Y.S., Han, Z.: Neural signed distance function inference through splatting 3d gaussians pulled on zero-level set (2024), https://arxiv.org/abs/2410.14189
85. Zhao, B., Levis, A., Connor, L., Srinivasan, P.P., Bouman, K.: Revealing the 3d cosmic web through gravitationally constrained neural fields. In: The Thirteenth International Conference on Learning Representations
86. Zhong, E.D., Bepler, T., Berger, B., Davis, J.H.: CryoDRGN: reconstruction of heterogeneous cryo-EM structures using neural networks. Nature Methods **18**, 176–185 (2021). https://doi.org/10.1038/s41592-020-01049-4, https://www.nature.com/articles/s41592-020-01049-4