

# Single-Pixel Imaging of Spatio-Temporal Flows using Differentiable Latent Dynamics

Sholokhov, Aleksei; Nabi, Saleh; Rapp, Joshua; Brunton, Steven; Kutz, Nathan; Boufounos, Petros  
T.; Mansour, Hassan

TR2024-151 November 02, 2024

## Abstract

Imaging dynamic spatio-temporal flows typically requires high-speed, high-resolution sensors that may be physically or economically prohibitive. Single-pixel imaging (SPI) has emerged as a low-cost acquisition technique where light from a scene is projected through a spatial light modulator onto a single photodiode with a high temporal acquisition rate. The scene is then reconstructed from the temporal samples using computational techniques that leverage prior assumptions on the scene structure. In this paper, we propose to image spatio-temporal flows from incomplete measurements by leveraging scene priors in the form of a reduced-order model (ROM) of the dynamics learned from training data examples. By combining SPI acquisition with the ROM prior implemented as a neural ordinary differential equation, we achieve high-quality image sequence reconstruction with significantly reduced data requirements. Specifically, our approach achieves similar performance levels to leading methods despite using one to two orders of magnitude fewer samples. We demonstrate superior reconstruction at low sampling rates for simulated trajectories governed by Burgers' equation, Kolmogorov flow, and turbulent plumes emulating gas leaks.

*IEEE Transactions on Computational Imaging 2024*



# Single-Pixel Imaging of Spatio-Temporal Flows using Differentiable Latent Dynamics

Aleksei Sholokhov, Saleh Nabi, Joshua Rapp *Member, IEEE*, Steven L. Brunton, *Senior Member, IEEE*, J. Nathan Kutz, *Senior Member, IEEE*, Petros T. Boufounos, *Fellow, IEEE*, and Hassan Mansour, *Senior Member, IEEE*

**Abstract**—Imaging dynamic spatio-temporal flows typically requires high-speed, high-resolution sensors that may be physically or economically prohibitive. Single-pixel imaging (SPI) has emerged as a low-cost acquisition technique where light from a scene is projected through a spatial light modulator onto a single photodiode with a high temporal acquisition rate. The scene is then reconstructed from the temporal samples using computational techniques that leverage prior assumptions on the scene structure. In this paper, we propose to image spatio-temporal flows from incomplete measurements by leveraging scene priors in the form of a reduced-order model (ROM) of the dynamics learned from training data examples. By combining SPI acquisition with the ROM prior implemented as a neural ordinary differential equation, we achieve high-quality image sequence reconstruction with significantly reduced data requirements. Specifically, our approach achieves similar performance levels to leading methods despite using one to two orders of magnitude fewer samples. We demonstrate superior reconstruction at low sampling rates for simulated trajectories governed by Burgers’ equation, Kolmogorov flow, and turbulent plumes emulating gas leaks.

**Index Terms**—Single-pixel imaging, high-dimensional dynamical systems, reduced-order model, neural ODE.

## I. INTRODUCTION

Low-cost cameras with high pixel counts have enabled digital imaging to become ubiquitous. However, traditional pixel arrays are often too slow or too expensive for imaging applications that require high acquisition speed or detection at wavelengths outside of the visible spectrum. Imaging systems that raster scan the illumination or detection pixel-by-pixel enable image formation from a cost-effective single-pixel detector, but the scanning process may likewise be too slow to capture dynamic scenes. For example, methane gas leak monitoring is important for reducing greenhouse gas emissions, but methane is only detectable in the mid-infrared (MIR) wavelengths. Unfortunately, MIR array detectors come with a significant cost, and capturing spatio-temporal dynamics

requires high-frame rate acquisition [1]. Other applications that require expensive hardware for high-speed imaging include microfluidics [2] and light propagation [3].

For applications in which high-speed, high-resolution sensors are physically or economically prohibitive, a candidate solution is single-pixel imaging (SPI), a technique that uses a single detector element and time-varying spatial light modulation to capture multiplexed measurements of light intensity [4]. SPI acquisition is often combined with computational reconstruction techniques that use prior information on the structure of the scene content to recover a full image despite using fewer modulation patterns than the target pixel resolution. The challenge then is how to mathematically describe prior information to match the target scene. Initial approaches to single-pixel imaging often assumed sparsity of the scene in the Fourier or wavelet basis or in the spatial gradient domain (i.e., total variation) [4]–[7]. Recent approaches have adopted deep learning (DL) architectures for SPI reconstruction and showed that using deep neural networks can dramatically reduce the sampling ratio and offer near-real-time performance [8]–[11]. However, these methods use frame-by-frame processing, thus requiring a relatively large number of temporal samples per frame (SPF) to ensure satisfactory reconstruction of a full video sequence.

This work aims to reconstruct high-quality videos of dynamical systems from a small number of samples per frame by learning priors on the dynamics that can act as a regularization across the entire sequence of frames. Specifically, we consider videos corresponding to complex spatio-temporal dynamical systems described by partial differential equations (PDEs), such as 1D inviscid Burger flow, 2D turbulent flow, and 3D buoyancy-driven flows governed by the Navier–Stokes equation as represented by smoke plumes. For such videos, the ‘ground truth’ refers to the case where the dynamics evolve on an appropriate infinite dimensional manifold. In particular, we show how a reduced-order model (ROM) can be used to regularize compressively-sampled single-pixel image/video reconstruction of a physical phenomenon, even when the number of samples per frame is as low as one. Prior work has shown that data-driven ROMs are highly effective at modeling complex temporal dynamics at a low computational cost, see for example [12].

A. Sholokhov is with the Department of Applied Mathematics, University of Washington, Seattle, WA, USA. He performed this work while he was an intern at MERL.

S. Nabi was with Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02139, USA. He is now with Schneider Electric, US Research Team, 800 Federal St, Andover, MA 01810, USA.

J. Rapp, P. Boufounos, and H. Mansour are with Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02139, USA.

S. Brunton is with the Department of Mechanical Engineering, University of Washington, Seattle, WA, USA.

N. Kutz is with the Department of Applied Mathematics, University of Washington, Seattle, WA, USA.

Manuscript received September, 2023.

We call our method SPI-NODE<sup>1</sup> because we combine SPI acquisition with a reconstruction algorithm that uses a neural ordinary differential equation operator (Neural ODE, or NODE) as regularization. We first pre-train a ROM to compress simulated gas flows into a low-dimensional latent space representation and calculate the system’s evolution within the latent space using Neural ODEs [13], [14]. Next, we efficiently solve an ODE-regularized compressively-sampled SPI reconstruction problem using the adjoint sensitivity method to obtain the necessary derivatives following the technique proposed in [13]. The reconstruction procedure minimizes the measurement mismatch between the samples recorded by the SPI detector and the synthesized measurements obtained by sampling the reconstructed trajectory using the compressive measurement operator. The ODE-based regularization provides a strong signal prior and allows for the reduction in the number of samples per frame required for successful reconstruction below the requirements of current state-of-the-art methods. While we focus our exploration on the design of an efficient learning-assisted reconstruction algorithm for SPI, the proposed architecture is also amenable to optimizing the modulation patterns as has become common practice in several recent works on deep learning-assisted SPI, e.g., [9].

Our main contributions are as follows. We first describe single-pixel acquisition of signals governed by dynamical systems. We then introduce a data-driven ROM prior on the temporal dynamics as regularization for the SPI inverse problem. While the performance of SPI algorithms with inadequate priors breaks down as the number of samples per frame decreases, we demonstrate that our SPI-NODE regularization achieves good results for up to two orders of magnitude fewer samples per frame. Finally, we show that this performance extends to a complex real-world dataset of turbulent smoke ascending in an ambient with a different density. We also show the applicability limits of our method. Using a synthetic experiment we show that the reconstruction gains persist when mild-to-moderate noise is applied but currently do not extend to high-noise regimes. In particular, we show that our method performs on par with its strongest competitor when the training dataset is insufficient for obtaining a good ROM. Such observations illustrate the necessity for developing more robust reconstruction methods and for training less data-intensive ROMs.

This paper is organized as follows. In Section II we provide the background on compressive sensing, SPI, and Neural ODEs. In Section III, we describe the acquisition system and our proposed architecture with an appropriate training loss to train a ROM, and present the recovery algorithm with our proposed regularizer. In Section IV, a hierarchy of experiments are provided to demonstrate the feasibility and assessment of our algorithm compared to state-of-the-art methods. Finally, we conclude the paper in Section V.

<sup>1</sup>An abbreviated version of this work has been submitted to ICASSP 2024 introducing the SPI-NODE model as a ROM formulation for single pixel imaging and demonstrating the performance on Burger’s equation and ScalarFlow data.

## II. BACKGROUND

### A. Compressive Sensing

Compressive sensing (CS) emerged in the mid-2000s as an alternative signal acquisition regime to classical Shannon-Nyquist sampling [15]–[18]. In the classical regime, a band-limited signal is sampled at  $N$  specific points in time (and/or space) so that the signal can be reconstructed from the acquired samples using a simple sinc interpolation. On the other hand, CS proposes to acquire a signal using a small number of inner products between the signal and general non-adaptive sampling vectors. If the signal to be acquired has a  $K$ -sparse—or more generally, parsimonious—representation in some transform domain, then the signal can be guaranteed to be recovered using as few as  $\mathcal{O}(K \log(N/K))$  non-adaptive measurements, which is much smaller than the  $N$  samples that are required by sample-then-compress approaches such as JPEG and JPEG-2000 [19]. Recent works have extended the applicability of this framework to general signal priors that include complicated image denoisers [20], [21] as well as deep signal priors [22], [23].

### B. Single-Pixel Imaging

1) *Advantages:* SPI evolved from early observations of the benefits of using multiplexed illumination instead of raster scanning [24] and the possibility of combining multiplexing with a single pixel detector to form an image [25]. Further SPI development was largely motivated as a practical application of compressive sensing theory and efficient reconstruction algorithms [4]. In the ensuing years, many iterations of SPI have been explored with different hardware architectures, reconstruction algorithms, and imaging applications [26]–[28]. SPI has multiple advantages over traditional imaging approaches that employ 2D sensor arrays. First, since all reflected light is focused onto a single detector, SPI has better sensitivity under dim lighting conditions and may enjoy a broader spectral range [27]. Second, SPI moves the cost of acquisition and processing from the encoding/compression stage to the decoding/reconstruction stage. For instance, the modulation pattern does not need to be scene-dependent to ensure accurate reconstruction: [18] showed that modulation patterns can be drawn as independent identically distributed (i.i.d.) samples from a uniform Bernoulli distribution. In contrast, sample-then-compress approaches [19] require an expensive projection of multiple single-pixel snapshots onto a chosen basis, such as Fourier or top- $K$  singular vectors, during their encoding step. In applications where sensors need to be small and low-power, it is often preferable to perform the encoding optically, while decoding can occur offline on high-performance computers. The simplicity of encoding keeps the computational requirements of SPI sensors (and thus their price) low, inviting their application in multispectral imaging [29], [30], optical encryption [31]–[33], remote sensing [34], object tracking [35], and many other fields [36]–[38].

2) *Hardware:* There are two main hardware configurations for SPI. The initial demonstration of SPI flood-illuminated a scene and used a programmable spatial light modulator (SLM) to control the pattern of light reaching the single-pixel



detector at each sample. An alternative configuration instead uses structured illumination to project different patterns of light onto the scene, and all of the reflected light is focused onto the detector. Mathematically, both configurations can be considered equivalent. The structured illumination approach is closely related to computational ghost imaging [39], which showed that image reconstruction based on correlations with patterned illumination was not limited to a quantum phenomenon [40].

A variety of hardware methods exist for performing spatial light modulation, including diffusers, liquid crystal devices (LCDs), and rotating masks, but the most common SLMs for SPI are digital micro-mirror devices (DMDs) [41], [42], which orient their mirrors to reflect light either towards (+1) or away from (0) the detector. DMDs are usually preferred for their high frame rate, broad spectral range, high spatial resolution, and low cost [28].

3) *Reconstruction Algorithms*: Since every measurement of a SPI system contains highly-mixed information about the observed image frame, one needs to sample a large number of such measurements to allow for faithful reconstruction. Thus, many recent works propose ways to reduce the number of required samples per frame (SPF). One approach uses signal-dependent control of the modulation pattern to ensure that each single-pixel measurement contains as much information as possible, e.g., by focusing the DMD mirrors on the selected regions of the frame [43]–[45]. Although highly effective, this approach cannot be applied to detectors that do not offer fine-grained control of their mirrors. An alternative line of work has focused on creating more sophisticated reconstruction techniques that are agnostic to the sampling patterns by improving the efficiency of image recovery using classical [5]–[7], [46] as well as learned priors [8]–[11] for image recovery. Our work follows this second approach by learning priors about the dynamics of scenes described by ODEs.

### C. Neural ODEs

A Neural Ordinary Differential Equation (Neural ODE or NODE) model [13] is a differentiable ODE solver. It utilizes the adjoint sensitivity method [47] to back-propagate the gradients through the solution. The emergence of GPU-compatible implementations of differentiable ODE solvers such as `torchdiffeq` [13] and `DiffRax` [48] caused a broad adoption of Neural ODE models in many applications, including for system identification [49], [50], classification [51], [52] medical imaging [53], [54], reduced-order modelling [14], [55], and modelling the flow of air [56] and fluids [57], to name a few. Neural ODEs have demonstrated superior ability to model highly non-linear dynamics compared to linear models especially when the dimensionality of the space over which the dynamics evolve is small. Thus, this technique is especially useful in applications where the size of the latent space dimension needs to be restricted due to their operation in a low measurement regime [58]–[61]. Consequently, we demonstrate in this paper that compact-but-accurate reduced-order models, such as the ones obtained by using autoencoder and Neural ODEs of Figure 1, are essential

for enabling single-pixel imaging systems operating in the low sample-per-frame regime.

### D. Gas Sensing

In backscatter absorption gas imaging (BAGI), the field of view is illuminated by a laser whose wavelength corresponds to the absorption band of a gas of interest [62]. Whereas the background simply scatters the light back toward the detector, a plume of gas attenuates the illumination, appearing as a dark cloud in the image. For good detection sensitivity, many BAGI implementations use a uniform background such as a retro-reflector, so the background appears bright in contrast to the absorptive gas plume. The absorption bands of many gases of interest are in the infrared range, and due to the expense of detectors sensitive to those wavelengths, most BAGI systems use single-pixel detectors and raster scanning to form an image [63]–[65], although some systems based on short-wave infrared (SWIR) focal plane arrays have also been developed [66].

Gibson et al. demonstrated methane gas BAGI monitoring using a SWIR-sensitive single-pixel camera [1]. However, to achieve real-time acquisition with simple reconstruction, the image resolution was limited to  $16 \times 16$  pixels from a full set of 256 mask patterns without compressive reconstruction. In this work, we show that higher-resolution results can be achieved for similar hardware by taking advantage of redundancy in the spatial and temporal dimensions. Although we do not collect any experimental data in this work, we emulate measurements that could have been acquired using the sensing hardware of Gibson et al. We also note that the sensing apparatus in [1] could be modified for other wavelengths, such as MIR, where methane has its strongest absorption. In addition to using a laser source and detector at  $3 \mu\text{m}$  as in [64], a standard DMD can be made MIR-compatible by replacing the sealing window [67].

## III. METHOD

The subsequent subsections outline the mathematical architecture for SPI-NODE. Specifically, each step in the SPI-NODE algorithmic pipeline is detailed. Also, refer to Appendix A for more details on training and architecture.

### A. Acquisition System

Let  $\mathbf{x}(t) \in \mathcal{X} \subseteq \mathbb{R}^n$ ,  $t \in [0, T]$  denote temporal snapshots of a high-dimensional spatio-temporal system indexed in time by  $t$ , e.g. the desired video recording of a scene that we wish to observe. In many applications we cannot record  $\mathbf{x}(t)$  at full resolution in real time due, for example, to a prohibitive cost of the required hardware or lack of high-dimensional sensors. Instead, we collect measurements from  $p$  detectors, where each detector provides a linear combination of the light emitted by  $\mathbf{x}(t)$  at time  $t$ :

$$\mathbf{y}(t) = \mathbf{A}\mathbf{x}(t), \quad \mathbf{A} \in \mathbb{R}^{p \times n}, \quad t \in [0, T]. \quad (1)$$

In particular, we consider a single pixel camera setup where for every time instance  $t$ , a vector of  $p$  acquisitions  $\mathbf{y} \in \mathbb{R}^p$

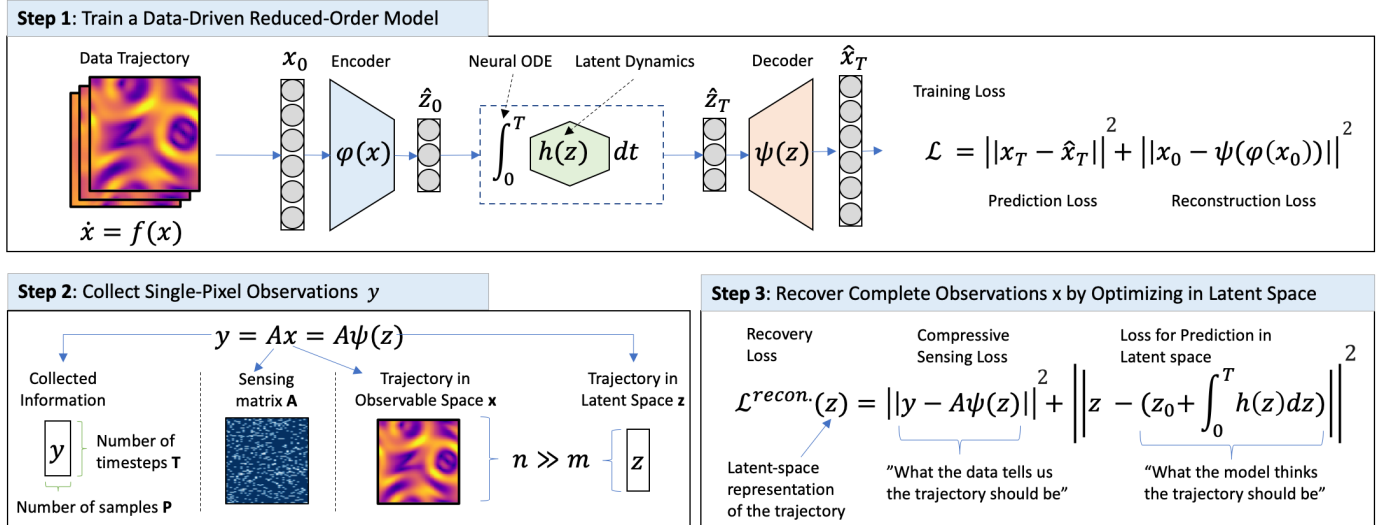


Fig. 1: Illustration of SPI-NODE – the proposed single-pixel imaging (SPI) system that leverages Neural ODEs (NODE) as a reduced-order model to allow for reconstruction of spatio-temporal flows from low sample-per-frame rates. Step 1 presents a pictorial illustration of the NODE training methodology where the spatio-temporal dynamics are learned for a low dimensional latent space representation that is reachable through an autoencoder. Step 2 shows the SPI acquisition process. Step 3 shows the objective function of the reconstruction algorithm that solves for the latent space variables that are then decoded to produce the full resolution image of the flow. Figure adapted from [14].

are obtained by a high sampling-rate photo-detector using the projection matrix  $\mathbf{A}$ . Note that this assumes the dynamical system is pseudo-stationary during the acquisition of each frame. The rows of the matrix  $\mathbf{A}$  correspond to a binary mask pattern that can be optically encoded, for example, using a DMD [41], [42]. In this setup, the number of mirrors of the DMD array represent the pixel resolution of the desired reconstructed image and the angles of the mirrors represent the binary weights. The light from a scene lands on a DMD that orients each mirror to reflect light either towards (+1) or away from (0) the detector according to a predetermined pattern. Light reflected from the DMD towards the detector is then focused by a lens onto the single-pixel detector. Finally, the detector integrates the signal over a short time and converts the measured intensity into a voltage that is digitized by an A/D converter [4]. The modulation pattern (i.e., the mirror angles) is also recorded in the detector’s memory. Fig. 2 illustrates a single-pixel imaging setup where a gas plume is imaged using a DMD array and an MIR photo-detector.

Several existing techniques attempt to reconstruct  $\mathbf{x}(t)$  directly by inverting equation (1) under mild regularity priors; examples include differential ghost imaging (DGI) [68], spatial total variation regularization (TVR) [26], spatio-temporal total variation regularization (3D-TV) [6], [7], and Fourier-domain regularized inversion (FDRI) [69]. Alternatively, one can recover  $\mathbf{x}(t)$  by recovering its representation  $\mathbf{z}(t) \in \mathcal{Z} \subseteq \mathbb{R}^m$  on a low-dimensional manifold ( $m \ll n$ ). Specifically, given an invertible mapping  $\psi(\mathbf{z}) : \mathcal{Z} \rightarrow \mathcal{X}$ , we replace the problem of solving for a vector  $\mathbf{x}$  in high-dimensional space using  $p$  linear observations with solving for a low-dimensional vector  $\mathbf{z}$  from  $p$  non-linear observations:

$$\mathbf{y}(t) = \mathbf{A}\psi(\mathbf{z}(t)). \quad (2)$$




Fig. 2: Schematic of an SPI setup for gas leak monitoring. An MIR source flood-illuminates the region of interest. The light is absorbed by the gas plume but reflected from the backscatter surface. The reflected light is focused with lens L1 onto the DMD, which modulates the light by angling the micro-mirrors to reflect some of the light towards lens L2, which is focused onto the single-pixel MIR detector. An image is reconstructed from the measured intensity value at the detector, the known DMD pattern, and learned priors about gas dynamics.

We call the space  $\mathcal{X}$  an observable space,  $\mathcal{Z}$  a latent space, and the mapping  $\psi(\mathbf{z})$  a decoder. Often a suitable  $\psi(\mathbf{z})$  is not known in advance, in which case an approximation  $\psi_\theta(\mathbf{z})$  is trained using a dataset of full resolution images  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ , where  $\theta$  denotes the parameters of the model, e.g. the weights of an autoencoder network. However, if  $\mathbf{y}(t)$  is insufficient for a unique reconstruction of  $\mathbf{x}(t)$  then it will remain insufficient for reconstructing  $\mathbf{z}(t)$ . In order to reduce the sufficient number of measurements that can uniquely reconstruct  $\mathbf{x}(t)$ , and by proxy  $\mathbf{z}(t)$ , additional structural sparsity or priors may

be introduced, such as joint training of the sampling matrices and the decoder [11]. In this work we take a different route and supplement the encoder  $\psi(z)$  with prior knowledge of the dynamics of  $\mathbf{z}(t)$  on the manifold over time.

### B. Reduced-Order Model with Non-Linear Latent Dynamics

Let  $\mathbf{x}(t)$  be modelled as an autonomous dynamical system on a finite space  $\mathcal{X} \subseteq \mathbb{R}^n$ :

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t)). \quad (3)$$

Considering the high dimensionality of  $\mathbf{x}(t)$ , it is often expensive to directly use the relationship (3) for predicting the behaviour of the system even when  $\mathbf{f}(\mathbf{x})$  is known. For example, if  $\mathbf{x}(t)$  is a flow of 2D liquid then  $\mathbf{f}(\mathbf{x})$  may be a discretization of Navier–Stokes equations, in which case running a simulation of (3) may take months on a cluster. However, a variety of works provide both theoretical [70] and practical [71], [72] evidence that many physical systems evolve on a lower-dimensional manifold  $\mathcal{Z}$ . In that space, the dynamics evolve according to a (generally unknown) function  $\mathbf{h}(\mathbf{z})$ :

$$\frac{d}{dt}\mathbf{z}(t) = \mathbf{h}(\mathbf{z}(t)). \quad (4)$$

Thus, one can predict the dynamics of the system  $\mathbf{x}$  at a future time  $T$  by projecting the initial condition  $\mathbf{x}(0)$  into the latent space, performing an integration there, and mapping the resulting trajectory back to the observable space:

$$\begin{aligned} \mathbf{z}(0) &= \psi^{-1}(\mathbf{x}(0)) \\ \mathbf{z}(T) &= \mathbf{z}(0) + \int_0^T \mathbf{h}(\mathbf{z}(t))dt \\ \mathbf{x}(T) &= \psi(\mathbf{z}(T)). \end{aligned} \quad (5)$$

When  $m \ll n$  we refer to the triplet  $(\psi, \psi^{-1}, \mathbf{h})$  as a reduced-order model (ROM) of  $\mathbf{f}$ . It is often the case that for a given system  $\mathbf{f}$  there exists no reduced-order representation  $(\psi, \psi^{-1}, \mathbf{h})$  such that the relation (5) holds exactly. In this case, we seek an approximate ROM  $(\psi_{\theta^*}, \phi_{\theta^*}, h_{\theta^*})$  that minimizes the difference between the data  $\mathbf{x}(t)$  and the prediction  $\hat{\mathbf{x}}(t)$  over a chosen class of models  $(\psi_{\theta}, \phi_{\theta}, h_{\theta})$  parameterized by  $\theta$ .

1) *ROM Architecture*: In this work we model  $\psi$ ,  $\psi^{-1}$ , and  $\mathbf{h}$  with neural networks  $\psi_{\theta}$ ,  $\phi_{\theta}$ , and  $h_{\theta}$ , respectively. Specifically, the pair  $(\psi, \psi^{-1})$  is modelled with an auto-encoder  $(\psi_{\theta}, \phi_{\theta})$ , and  $\mathbf{h}$  is modelled with a fully-connected network  $h_{\theta}$ , as illustrated in the top panel of Figure 1. All considered models share this architecture; however, the exact constitution of the networks  $\psi_{\theta}$ ,  $\phi_{\theta}$ , and  $h_{\theta}$  are problem-dependent and discussed more closely in the related sections.

2) *Training Loss*: Similar to prior works [59], [73], [74], we define a *training loss*  $\mathcal{L}$  as a sum of reconstruction and prediction losses. The former ensures that  $\phi_{\theta}$  and  $\psi_{\theta}$  are inverse mappings of each other, whereas the latter matches the model’s predictions to the available data. Formally, for a given set of trajectories  $\mathbf{X}_i$ ,  $i \in \{1, \dots, N\}$ , where each trajectory  $\mathbf{X}_i \in \mathbb{R}^{n \times K}$  is a set of  $K$  snapshots  $\mathbf{X}_i(t_j) \in \mathbb{R}^n$

of the system, for  $K$  time-steps,  $t_j$ ,  $j \in \{0, \dots, K-1\}$ , the loss function  $\mathcal{L}_{\theta}^{\text{data}}$  is defined as:

$$\begin{aligned} \mathcal{L}(\theta) &= \sum_{i=1}^N \left[ \sum_{j=0}^{K-1} \|\mathbf{X}_i(t_j) - \psi_{\theta}(\phi_{\theta}(\mathbf{X}_i(t_j)))\|^2 \right. \\ &\quad \left. + \sum_{j=1}^{K-1} \left\| \psi_{\theta} \left( \phi_{\theta}(\mathbf{X}_i(t_0)) + \int_{t_0}^{t_j} h_{\theta}(\mathbf{z}(t))dt \right) - \mathbf{X}_i(t_j) \right\|^2 \right]. \end{aligned} \quad (6)$$

Note that  $\mathbf{z}$  is the latent space variable defined in Eq. (2). To obtain a ROM  $(\psi_{\theta^*}, \phi_{\theta^*}, h_{\theta^*})$ , we minimize the loss above. We note that each trajectory  $\mathbf{X}_i$  may be captured over its own time-frame and use a distinct, possibly non-uniform step-size, in which case the loss function should be modified accordingly<sup>2</sup>. To simplify the notation, without loss of generality, in the rest of the paper we assume that all trajectories were recorded over the same time-frame with an equal and uniform step-size  $\Delta t$ .

### C. Recovery Algorithm

We use the approximate ROM  $(\psi_{\theta^*}, \phi_{\theta^*}, h_{\theta^*})$  above to regularize the latent-space dynamics of the single-pixel imaging reconstructions. Specifically, we obtain the reconstruction  $\mathbf{x}^* = \psi(\mathbf{z}^*)$  by minimizing the following loss with respect to  $\mathbf{Z} = [\mathbf{z}_0, \dots, \mathbf{z}_{K-1}] \in \mathbb{R}^{m \times K}$ :

$$\begin{aligned} \mathcal{L}_{\theta^*}^{\text{recon}}(\mathbf{z}) &= \sum_{j=0}^{K-1} \|y_j - A\psi_{\theta^*}(\mathbf{z}_j)\|_2^2 \\ &\quad + \lambda \left\| \mathbf{z}_0 + \int_{t_0}^{t_j} h_{\theta^*}(\mathbf{z})dz - \mathbf{z}_j \right\|_2^2, \end{aligned} \quad (7)$$

where the regularization parameter  $\lambda$  controls the tradeoff between the adherence to the latent dynamics  $h_{\theta^*}$  or to the SPI measurement mismatch during the signal reconstruction phase. We minimize the loss (7) using a gradient-based technique, with the gradients obtained using automatic differentiation frameworks.

## IV. EXPERIMENTS

### A. Burgers’ Equation

We first study the performance of our framework on flows governed by Burgers’ equation with  $[-\pi, \pi]$ -periodic boundary conditions:

$$\begin{aligned} u_t + uu_x &= \nu u_{xx} \\ u(-\pi, t) &= u(\pi, t), \quad \forall t \in [0, T] \end{aligned} \quad (8)$$

where  $u_t$  represents the first partial derivative with respect to time, and  $u_x$  and  $u_{xx}$  represent the first and second partial spatial derivatives, respectively.

<sup>2</sup>The implementation is affected only in evaluating the integral in (6). This part is handled by the `torchdiffeq` [75] library, which supports non-uniform time-frames within a batch.

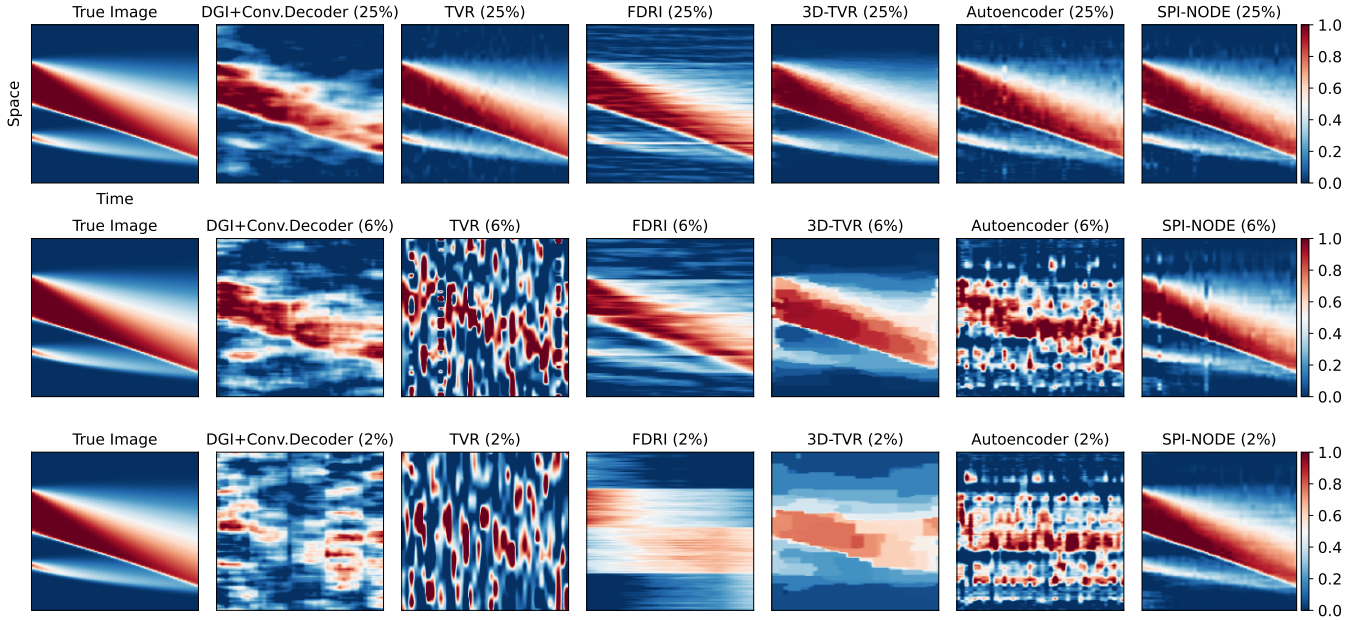


Fig. 3: Example reconstruction of a trajectory of Burgers’ equation for three SPF rates: 25%, 6%, and 2%. All algorithms complete the reconstruction faithfully when SPF is large. However, all algorithms except our proposed method, fail to recover the image due to the insufficient SPF rate ( $2/128 \lesssim 2\%$ ). In contrast, SPI-NODE achieves nearly the same reconstruction quality as in with SPF rate of 25%.

1) *ROM Training*: To obtain a training dataset for the ROM we replicate the experimental setup from [14, Section 3]. Namely, we generate 1024 trajectories on a discretized spatial domain  $[-\pi, \pi]$  with 128 grid-points. To generate a diverse set of initial conditions we sum the first 10 harmonic terms with random coefficients:

$$u(x, 0) = \frac{1}{10} \sum_{k=1}^{10} a_k \cos(kx) + b_k \sin((k+1)x), \quad (9)$$

where  $a_k, b_k \sim \mathcal{N}(0, 1)$ . We solve Equation 8 for  $t \in [0, 2]$  with  $\Delta t = 0.1$  using a spectral solver [76].

2) *Architecture*: We define  $\phi_\theta$  and  $\psi_\theta$  to be fully-connected neural networks with 3 hidden layers, each 512 neurons wide, and ReLU activations after all but the output layers. We set the observable space dimension  $n = 128$  and the latent space dimension  $m = 16$ . The network  $h_\theta$  is a fully-connected network with 3 layers, each 256 neurons wide, and ReLU activations after all but the output layers. See more details of the training procedure in Appendix A.

3) *SPI Recovery*: For the recovery phase, we generate 128 trajectories with “bump” initial conditions—a smooth approximation of a bump with two opposing steeply-curved sigmoids:

$$u(x, 0) = \frac{1}{1 + \exp(-k(x-a))} - \frac{1}{1 + \exp(-k(x-b))}, \quad (10)$$

where  $a < b$  are sampled uniformly in  $[-\pi, \pi]$  and  $k = 20$ . We choose this shape to ensure that the training and sensing trajectories are sufficiently different to prevent memorization effects. We choose compressive sensing matrices  $\mathbf{A}_t$  to be

binary  $\{0, 1\}$  matrices with each row having 64 non-zero components out of 128 sampled uniformly.

4) *Performance Evaluation*: We compare the performance of our SPI-NODE regularization approach against five alternatives: an autoencoder-enhanced variation of digital ghost imaging [68], [77] (DGI+Conv. Decoder, [11]), spatial total variation regularization (TVR, [26]), spatio-temporal total variation regularization (3D-TVR, [6], [7]), Fourier-domain regularized inversion (FDRI, [78]), and an approach which uses an autoencoder (AE, [22]).

We measure reconstruction accuracy with three commonly-used metrics, treating the Burgers’ trajectories as 2D-images in spatial and in temporal domains. For two 2D images  $a$  and  $b$ , the normalized root mean-squared error (NRMSE) is defined as the Frobenius norm of the residuals for every snapshot divided by the Frobenius norm of the true snapshot:

$$\text{NRMSE}(a, b) = \|a - b\|_F / \|a\|_F. \quad (11)$$

The peak signal-to-noise ratio (PSNR) is defined as the logarithm of the maximal value from the true image:

$$\text{PSNR}(a, b) = 20 \log \left( \frac{\max_{i,j} (|a[i, j]|)}{\|a - b\|_2} \right). \quad (12)$$

Finally, the structural similarity index measure (SSIM) [79] is defined as a product of relative luminance  $l(a, b)$ , contrast  $c(a, b)$ , and structure  $s(a, b)$

$$\begin{aligned} \text{SSIM}(a, b) &= l(a, b) \cdot c(a, b) \cdot s(a, b) \\ &= \left( \frac{2\mu_a\mu_b + c_1}{\mu_a^2 + \mu_b^2 + c_1} \right) \left( \frac{2\sigma_a\sigma_b + c_2}{\sigma_a^2 + \sigma_b^2 + c_2} \right) \left( \frac{\sigma_{ab} + c_2/2}{\sigma_a\sigma_b + c_2/2} \right), \end{aligned} \quad (13)$$




Fig. 4: Median PSNR, MSE, and SSIM with 95% confidence intervals. SPI-NODE achieves the best performance across all metrics at low SPF rates (3% or less).

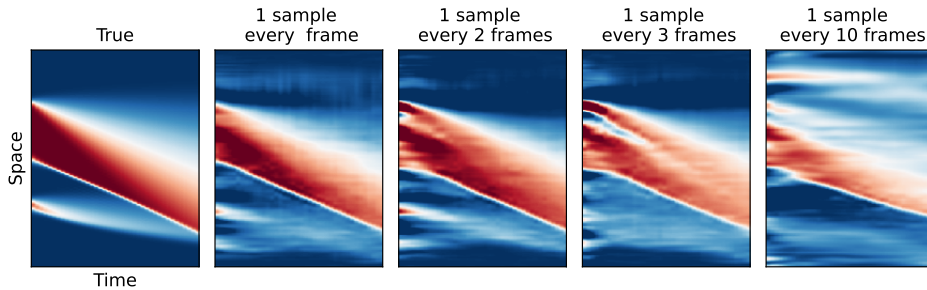


Fig. 5: An example SPI reconstruction with SPI-NODE for the regime when one sample is taken every  $K$  frames.

where, for a picture  $x$ ,  $\mu_x$  is the pixel sample mean,  $\sigma_x$  is the standard deviation,  $\sigma_{xy}$  is the cross-covariance, and  $c_i$  are the constants which stabilize the division when the denominator approaches 0. Typically,  $c_i$  are set to be proportional to the square of the dynamic range of the pixel values, e.g.  $c_i = (k_i \times 2^{\text{bits-per-pixel}} - 1)^2$ , with  $k_1 = 0.01$  and  $k_2 = 0.03$  being default choices. Finally, the reconstruction performance is evaluated as a function of the samples-per-frame (SPF) rate, i.e., the number of SPI measurements divided by the number of grid points  $n = 128$  in every frame that is recovered.

5) *Results:* Figure 3 provides an example of reconstruction at three SPF values. We notice that the DGI, TVR, and Autoencoder approaches do not reconstruct images faithfully until the SPF approaches 25 – 50% of the number of pixels in a snapshot. The FDRI and 3D-TVR methods have priors that are better suited to the Burgers’ trajectories, and their performance declines more gracefully as the SPF decreases. In contrast, SPI-NODE regularization is able to reconstruct most trajectories with high accuracy using a very small number of samples per frame, i.e., as low as 2% of pixels.

Figure 4 presents quantitative results of simulations comparing the median PSNR, MSE, and SSIM for each method over different numbers of samples per frame taken. We observe that SPI-NODE regularization achieves  $\approx 25\%$  higher PSNR and up to  $\approx 33\%$  higher SSIM relative to FDRI when SPF is low. It also achieves  $\approx 70\%$  better PSNR and more than  $2\times$  higher SSIM relative to the Autoencoder-based algorithms. While 3D-TVR is shown to perform better than all other methods at 6% SPF rate and above, SPI-NODE has the best performance at low SPF rate. The slight degradation in the performance of SPI-NODE at the 6% SPF rate and above can be attributed

to the suboptimal choice of regularization parameter  $\lambda$ , which was fixed for all SPFs and was set to be more favorable to the low SPF regime. Superior performance could likely be achieved with SPI-NODE by weighting the data fidelity more heavily as the SPF rate increases.

The success of SPI-NODE is due to the reduced-order model serving as a strong prior for the dynamics across time: the model reconstructs the trajectory as a whole instead of reconstructing every snapshot separately. Of the comparison algorithms, only 3D-TVR likewise takes advantage of the entire measurement sequence together. Still, SPI-NODE outperforms 3D-TVR when the SPF rate is small because it is a stronger prior on the dynamics of the system than the total variation prior. In fact, a successful reconstruction is possible even when the number of samples per frame is below one, i.e., one sample is taken every  $K$  frames, as illustrated by Figure 5 (the aggregated results are presented on Figure 10 in Appendix B). Such a situation arises when a slow SPI sensor is trying to capture fast-changing dynamics. In Appendix C, we further explore the robustness of the SPI-NODE regularization approach to additive noise.

## B. Kolmogorov Flow

We next study the performance of our framework on Kolmogorov flow, first-described by [80] and later analyzed by [81], [82]. Such flows are characterized by high intrinsic dimensionality and are not amenable to traditional order-reduction methods through projection of the governing equations. Therefore, it serves as a representative example to






Fig. 6: An example of a reconstruction of Kolmogorov with three different methods using 8 samples per frame of  $66 \times 66$  pixels, an SPF rate of 0.18%.




Fig. 7: SPI reconstruction performance using three types of regularizers. Only our proposed SPI-NODE regularization provides a strong enough prior for reconstruction with an SPF rate as low as 0.18%.

showcase our method. The following equation describes the behavior of a 2D velocity field  $\mathbf{u}(x, y, t)$ :

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{\text{Re}} \nabla^2 \mathbf{u} + f \quad (14)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (15)$$

where  $p$  is a 2D pressure field,  $f = A \sin(ky)x$  represents the driving force with amplitude  $A$  and wavenumber  $k$ , and  $\text{Re}$  is the Reynolds number. In all our experiments,  $k = 4$  and  $\text{Re} = 40$ ; we chose this setup to match the one from [83], which identified that this combination of hyper-parameters leads to the occurrence of extreme instability events, making the long-term prediction of such flow truly challenging.

To train a ROM that predicts  $u(x, y, t)$ , we generate 8192 data trajectories as solutions of (14), by adopting the spectral solver and altering the initial conditions, from [83]<sup>3</sup>. To avoid the initial transient behavior, each solution was first simulated from  $t = 0$  to  $t = 100$  and then snapshots were recorded from  $t = 100$  to  $t = 110$ , with the time-step  $\Delta t = 0.5$ .

As in Section III, the ROM consists of an encoder  $\phi(u)$ , a decoder  $\psi(z)$ , and the latent dynamics  $h(z)$ . The pair  $\phi(u)$  and  $\psi(z)$  were represented by a convolutional autoencoder, to capture spatial patterns, and the function  $h(z)$  was represented by a fully-connected network. The ROM was trained by minimizing the data-driven loss (6) until the prediction MSE on a

<sup>3</sup>[https://github.com/zhong1wan/data-assisted/blob/master/Kolmogorov/kol2d\\_odd.py](https://github.com/zhong1wan/data-assisted/blob/master/Kolmogorov/kol2d_odd.py)

holdout set stopped improving, which took about 60 epochs. See Appendix A for more details about the architecture and training. The sensing step using SPI is consistent with the one described in the previous section.

A sample result of the reconstruction algorithm is displayed in Figure 6. For the image recovery step, we acquire 8 samples per frame where each frame consists of  $66 \times 66 = 4356$  pixels, which yields a SPF rate of 0.18%. The columns represent different timesteps. The first row represents the true flow, the next four rows show the images recovered by our method, FDRI, Autoencoder, and 3D-TVR methods, respectively. We observe that the SPI-NODE reconstruction faithfully recovers many of the details of the flow, whereas the other algorithms recover more blurry images.

The aggregated results in Figure 7 show that other methods require more than 2% SPF rate for a faithful recovery of the signal, whereas our method can provide a decent reconstruction with SPF as low as 0.18%, a difference of one order of magnitude in the required data intake.

### C. Application to Real Data: Reconstruction of a Gas Plume

In this section we use our technique to reconstruct videos of gas leaks from two similar real-world datasets: ScalarFlow [84] and GasVid [85]. From the full-scale videos, we emulate the measurements that would be recorded by a single-pixel camera. The physical problem at hand is an ascending turbulent plume in an environment of higher density. The mixing associated with such buoyancy-driven flows is a complex and high-dimensional phenomenon and the optimization or control of such a system is a tedious task [86], [87]. We find that our method performs better than FDRI for the former dataset and similarly well for the latter. This comparison illustrates the necessary amount and quality of the data for a ROM to serve as a good reconstruction prior.

1) *ScalarFlow: Smoke Machine Plume Data:* The dataset consists of 3D reconstructions of density and velocity of 104 smoke plumes, each 150 frames long. We assemble our dataset by taking the front views (2D) of the reconstructions and reducing the spatial resolution to  $320 \times 192$  pixels, for computational and memory limitations. We then split the trajectories into non-overlapping sets: train (92), validation (5), test (5), and reconstruction (2) respectively. We use the first three sets to train, fine-tune, and select a final ROM, respectively. We use a similar architecture of the ROM as in Section IV-B with the details outlined in Appendix A. After obtaining the ROM, we use the fourth set of data for SPI reconstruction. For our simulated SPI experiments, we sample 2 to 128 single-pixel observations per frame which corresponds to SPF rates ranging from  $3 \times 10^{-3}\%$  to 0.2%.

The results are presented in Figure 8(a), the mean NRMSE, PSNR, and SSIM are displayed in Figure 8(b). We see that SPI-NODE consistently outperforms FDRI by all three metrics within the given ranges of SPF rates. For example, in Figure 8(a) we see that FDRI reconstructs a blurry image with 32 samples per frame (SPF rate of 0.05%), whereas our method is able to reconstruct considerably more accurate detail.

2) *GasVid: Methane Leaks Data:* We use a subset of the GasVid dataset [85], which consists of 10 methane gas leak videos of 10 to 20 seconds each. We split the videos into 1-second-long intervals, where each interval contains 10 steps:  $T = 1s$ ,  $\Delta t = 0.1s$ . To emulate a BAGI measurement with a uniform background, we next remove the background from each video using the Gaussian mixture-based background/foreground segmentation (MOG2) algorithm by [88], [89] implemented in OpenCV library [90]. Finally, we split all intervals into four non-overlapping batches: train (390), validation (50), test (49), and reconstruction (2). We use the same architecture of ROM as in Section IV-B.

The example reconstructions are shown in Figure 9(a) for which we sample 512 single-pixel observations per video frame of size  $240 \times 320$  pixels, corresponding to an SPF rate of 0.66%. The aggregated performance metrics – mean PSNR, NRMSE and SSIM – are displayed in Figure 9(b) for SPF rates ranging from 0.003% to just below 1%.

We see that SPI-NODE algorithm was able to faithfully reconstruct the general trajectories of the gas plumes. Although the limited training dataset and noise contamination do not allow the model to capture the details of turbulent flows, the method is still able to recover the altitude, direction, and spread of the plumes correctly, which are the most important aspects for practical applications in real-world gas monitoring applications. However, SPI-NODE and FDRI show indistinguishable average performance in terms of SSIM and NRMSE. It is interesting to also note that the FDRI reconstruction better captures the intensity of the smoke cloud, which could be an important characteristic to extract in some applications that require a gas density estimation.

3) *Discussion:* GasVid and ScalarFlow datasets are different in three aspects that determine the success of ROM-based SPI reconstruction. First, the background subtraction struggles to separate grey smoke from a dynamic grey background in GasVid, which leads to phantom clouds of smoke in the processed data. ScalarFlow, on the other hand, is noise-free. Second, the trajectories in GasVid are affected by hidden variables such as the direction of the wind, its strength, and the velocity of the leak. These variables make it hard to train a ROM that predicts the trajectory accurately using only the initial condition. In contrast, the ScalarFlow setup excludes ambient air movement to a reasonable degree and maintains constant smoke generation parameters. Finally, the number of training frames in GasVid (3900) was more than three times smaller than that of ScalarFlow (12750). Together, these three factors explain how the training procedure obtains a ROM that is a useful prior for the ScalarFlow dataset and why the same procedure fails to find a prior for GasVid that is more useful than the sparse-frequencies prior of FDRI.

## V. CONCLUSION

In this paper, we introduced a new type of regularization that can be used for single-pixel imaging of scenes with spatio-temporal dynamics governed by an ODE. By training a reduced-order model with neural ODE, we learn an efficient representation of spatio-temporal flows that can act as a strong

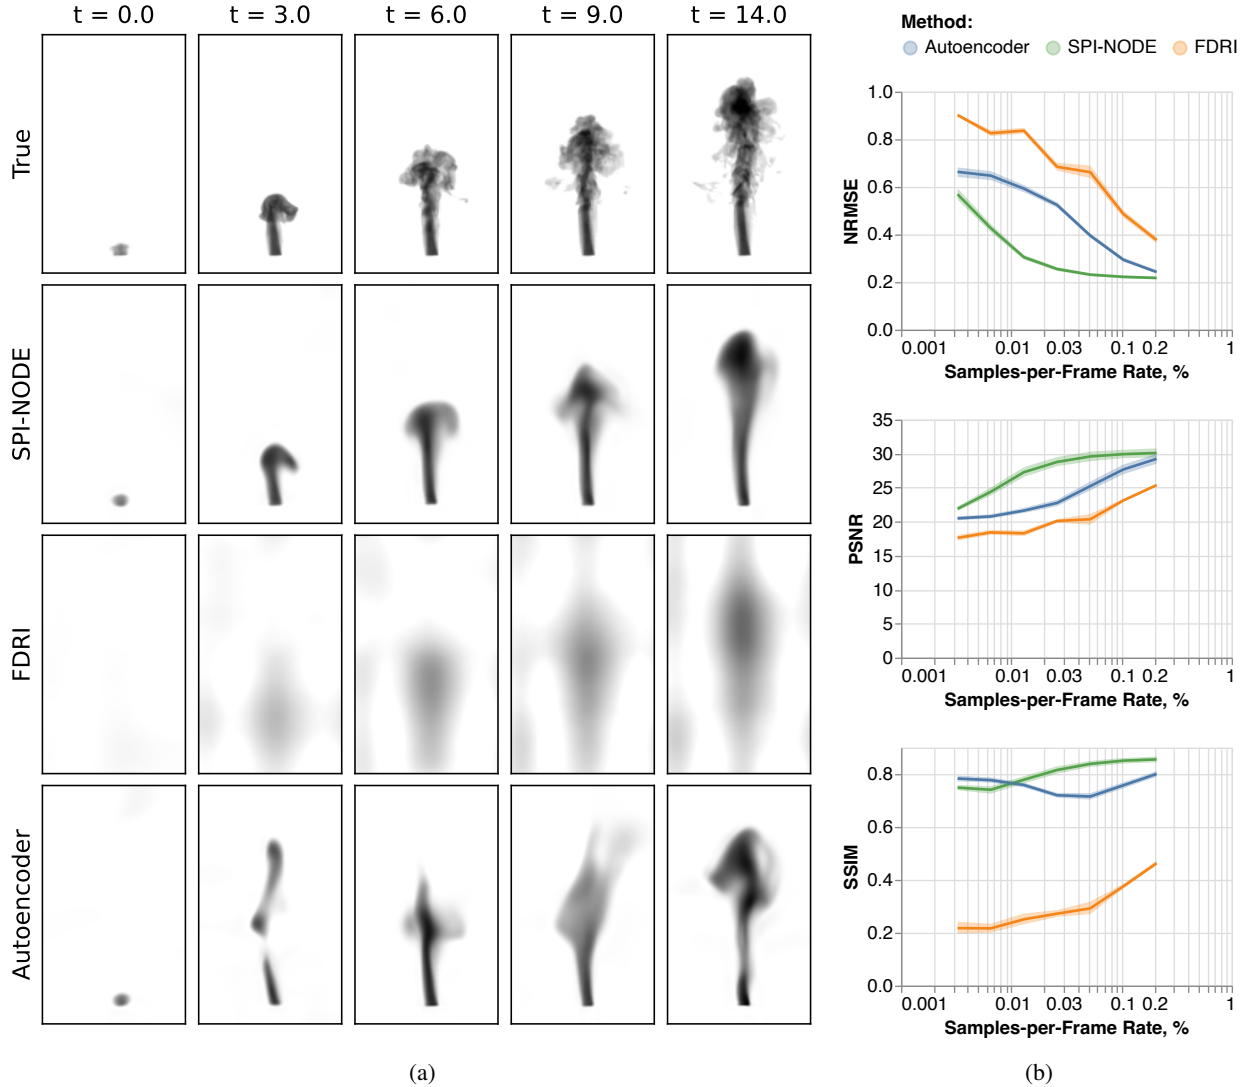


Fig. 8: SPI reconstructions of gas plumes using ScalarFlow Dataset. The larger, cleaner dataset of ScalarFlow results in a more useful ROM prior than the GasVid data. The left panel shows SPI reconstructions by three algorithm using 16 samples per frame (0.026% SPF rate). Right panel presents mean NRMSE, PSNR, and SSIM with 95% confidence intervals.

prior for SPI reconstruction. SPI-NODE is a good fit for settings when large amounts of clean and representative data are available, and when the application in question necessitates a low sample-per-frame ratio (e.g.,  $< 1\%$ ). In such a setting, our method leverages the ROM to compensate for low data intake with a strong data-driven prior and yields superior reconstruction quality. Our approach to SPI regularization has promise for applications such as gas leak monitoring that are limited by the cost of mid-infrared pixel arrays and the slow speed of raster-scanned methods. Further applications in which our regularization approach could help SPI for high speed dynamical systems could also include gas monitoring for HVAC [91] and weather-robust sensing [92], as well as imaging for microfluidics [2] or light propagation [3]. The low SPF rates that we achieve can ease constraints on data transfer and storage while allowing recovery of high-resolution images that retain important aspects of the flows.

Importantly, the technique developed shows that SPI-based methods can be significantly enhanced by using appropriate priors for the image reconstruction. Although the priors considered here are physics-based in the sense that the data is assumed to be generated by differential or partial differential equations, a broader set of regularizations can be critically enabling based upon how expert knowledge of the system can inform the prior. Finally, we emphasize that the success of the proposed approach is dependant on the ability to learn a representative model of the target spatio-temporal flow. That implies having sufficient training data or knowing the exact physics to properly learn the ROM. On the algorithmic side, it remains an open question to find the optimal regularization parameter from the measurements alone that will achieve the best reconstruction.



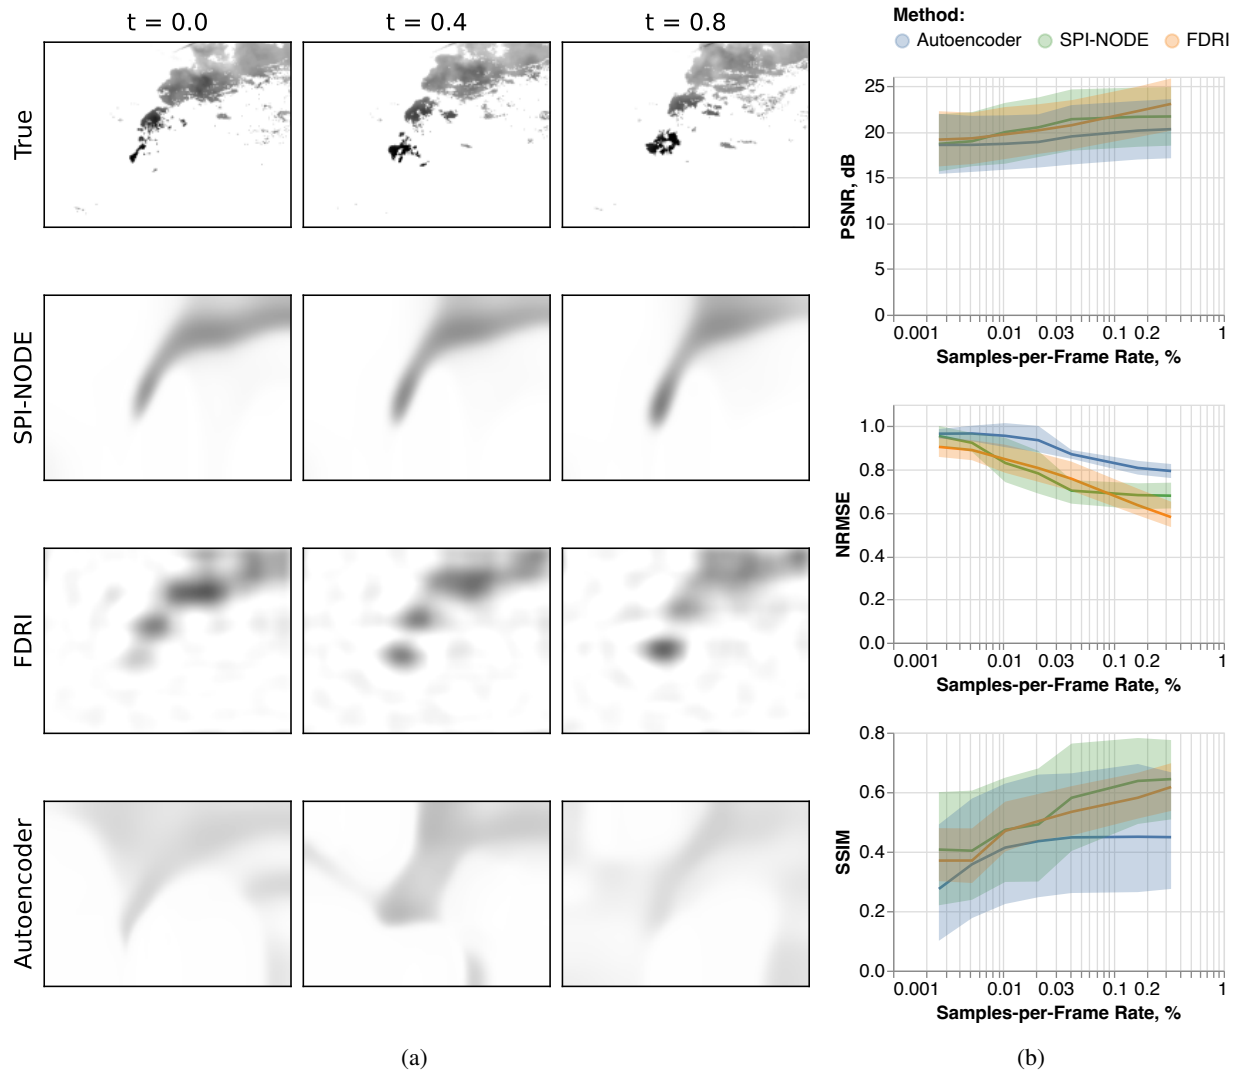


Fig. 9: SPI reconstructions of gas plumes using a subset of GasVid dataset. The size of training data and the noise in it limit the ROM to be a useful prior which leads to quantitatively indistinguishable performance. The left panel shows SPI reconstructions by three algorithm using 16 samples per frame (0.026% SPF rate). Right panel presents mean PNSR, NRMSE, and SSIM with 95% confidence intervals.

#### ACKNOWLEDGEMENTS

AS, SLB, and JNK acknowledge support from the National Science Foundation AI Institute in Dynamic Systems (grant number 2112085). JNK further acknowledges support from the Air Force Office of Scientific Research (FA9550-19-1-0011). JR, PB, and HM are supported exclusively by MERL.

#### APPENDIX A DETAILS OF TRAINING AND ARCHITECTURES

Table I contains details on the architecture and training of the SPI-NODE models. Asterisk-marked number of epochs stands for "stopped when the prediction on the development set plateaued" and refers to the actual number of epochs performed before the training was stopped. Table II describes the architecture of one block of the autoencoder for the

Kolmogorov Flow, GasVid, and ScalarFlow datasets. We used circular padding for the former and same padding for the rest.

All experiments were computed using an Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20 GHz equipped with a Tesla K80 GPU. The computer had Linux 4.15 installed as an OS.

#### APPENDIX B BURGER'S EQUATION: AGGREGATE RESULTS FOR SUB-1 SPF REGIME

In Figure 10 we present the aggregated results for recovering Burgers' solutions under SPF rates less than one. Since no algorithm can recover an image using zero samples per frame, we took reconstructions obtained for 1 SPF and then used linear interpolation to fill the gaps for methods that reconstruct the dynamics on frame-by-frame basis (i.e., for all but




Fig. 10: Aggregate results for SPI-NODE reconstructions when the number of samples per frame is below 1.

TABLE I: Architecture and training details of SPI-NODE models

	Burgers' Equation	Kolmogorov Flow	ScalarFlow	GasVid
Section	IV-A	IV-B	IV-C1	IV-C2
<i>Autoencoder</i>				
Type	FC	Conv.	Conv.	Conv.
Num. Blocks	3	4	6	6
Block Type	FC+ReLU	Table II	Table II	Table II
Activation Fn.	ReLU	SILU	SILU	SILU
Skip Conn.	-	-	+	-
Spectral Norm.	-	+	+	+
<i>Latent Space Dynamics</i>				
Type	FC	FC	FC	FC
Dimension	16	64	64	64
Num. Layers	3	3	2	2
Width	256	256	128	128
Activation Fn.	ReLU	ReLU	ReLU	ReLU
Skip Conn.	-	-	-	-
Spectral Norm.	-	-	-	-
<i>Training</i>				
Batch Size	32	64	16	16
Learning Rate	$10^{-4}$	$10^{-3}$	$10^{-3}$	$10^{-3}$
Epochs	500	61*	75*	320*

TABLE II: Auto-encoder block architecture

Step	Input Size	Action
<i>Encoder</i>		
1	(c, w, h)	Padding (1)
2	(c, w+2, h+2)	2D Convolution 3x3, Stride 1, Padding 0
3	(2c, w, h)	SILU Activation
4	(2c, w, h)	Average Pooling 2D (2x2)
5	(2c, w/2, h/2)	Next block (go to step 1)
<i>Decoder</i>		
1	(c, w, h)	2x Bilinear Interpolation
2	(c, 2w, 2h)	Padding (1)
3	(c, 2w+2, 2h+2)	2D Deconvolution 3x3, Stride 1, Padding 0
4	(c/2, 2w+2, 2h+2)	Spectral Normalization
5	(c/2, 2w+2, 2h+2)	SILU Activation (except for the output layer)
6	(c/2, 2w+2, 2h+2)	Cropping (2)
6	(c/2, 2w, 2h)	Next block (go to step 1)

PINODE). We see that PINODE, on average, is capable of reconstructing trajectories with up to one sample per 3 frames, beyond which its quantitative performance becomes indistinguishable from that of other algorithms. Interestingly, the performance of the Autoencoder-based method slightly improves as the gap between the sampled frames increases. We posit that this behavior results from the temporal smoothing effect of linear interpolation. In contrast, when temporal interpolation is not used—as in Figures 3 and 6—the autoencoder-based

method could yield temporally non-smooth reconstructions since the recovered latent space vectors of adjacent frames could be substantially far apart so long as their decoding matches the SPI measurements. Temporal linear interpolation prevents such large displacement and leads to smoother images and thus higher SSIM scores.

## APPENDIX C

### BURGERS' EQUATION: ROBUSTNESS TO NOISE

We also study the robustness of SPI-NODE to noise in the SPI measurements  $y$  and compare the reconstruction performance of all algorithms. Figure 11 presents an aggregated comparison of reconstruction of solutions of Burgers' equations under the noisy setting for SPF = 12.5%; the setup is equivalent to the one for SPF = 8 in Figure 4. Namely, we plot performance metrics (PSNR, MSE, SSIM) of reconstructions from different models against Signal-to-Noise Ratio (SNR) in the vector of compressive sensing single-pixel observations  $y$ . The plots show that while the performance of SPI-NODE is superior at high SNRs, it begins to suffer earlier than other methods at lower SNRs.

## REFERENCES

- G. M. Gibson, B. Sun, M. P. Edgar, D. B. Phillips, N. Hempler, G. T. Maker, G. P. A. Malcolm, and M. J. Padgett, "Real-time imaging of methane gas leaks using a single-pixel camera," *Optics Express*, vol. 25, no. 4, pp. 2998–3005, Feb. 2017.
- K. Goda, K. K. Tsia, and B. Jalali, "Serial time-encoded amplified imaging for real-time observation of fast dynamic phenomena," *Nature*, vol. 458, no. 7242, pp. 1145–1149, Apr. 2009.
- A. Veltén, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, and R. Raskar, "Femto-photography: Capturing and visualizing the propagation of light," *ACM Transactions on Graphics*, vol. 32, no. 4, July 2013.
- M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 83–91, 2008.
- O. Katz, Y. Bromberg, and Y. Silberberg, "Compressive ghost imaging," *Applied Physics Letters*, vol. 95, no. 13, pp. 131110, 2009.
- T. Goldstein, L. Xu, K. F. Kelly, and R. Baraniuk, "The STOne transform: Multi-resolution image enhancement and compressive video," *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 5581–5593, 2015.
- Y. Le Montagner, E. Angelini, and J.-C. Olivo-Marin, "Video reconstruction using compressed sensing measurements and 3d total variation regularization for bio-imaging applications," in *2012 19th IEEE International Conference on Image Processing*, Orlando, FL, USA, Sept. 2012, pp. 917–920.
- M. Lyu, W. Wang, H. Wang, H. Wang, G. Li, N. Chen, and G. Situ, "Deep-learning-based ghost imaging," *Scientific Reports*, vol. 7, no. 1, pp. 17865, 2017.




Fig. 11: Performance of reconstruction of solutions of Burgers' equations under noise for  $\text{SPF}=12.5\%$ . We plot the performance metrics of different models against Signal-to-Noise Ratio (SNR) in the vector of compressive sensing single-pixel observations  $y$ . The performance of SPI-NODE remains superior until all models start reconstructing the scene very poorly.

- [9] C. F. Higham, R. Murray-Smith, M. J. Padgett, and M. P. Edgar, "Deep learning for real-time single-pixel video," *Scientific Reports*, vol. 8, no. 1, pp. 2369, 2018.
- [10] F. Wang, H. Wang, H. Wang, G. Li, and G. Situ, "Learning from simulation: An end-to-end deep-learning approach for computational ghost imaging," *Optics Express*, vol. 27, no. 18, pp. 25560–25572, 2019.
- [11] F. Wang, C. Wang, C. Deng, S. Han, and G. Situ, "Single-pixel imaging using physics enhanced deep learning," *Photonics Research*, vol. 10, no. 1, pp. 104–110, 2022.
- [12] P. Wu, J. Sun, X. Chang, W. Zhang, R. Arcucci, Y. Guo, and C. C. Pain, "Data-driven reduced order model with temporal convolutional neural network," *Computer Methods in Applied Mechanics and Engineering*, vol. 360, pp. 112766, 2020.
- [13] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Advances in Neural Information Processing Systems*, 2018, pp. 6571–6583.
- [14] A. Sholokhov, Y. Liu, H. Mansour, and S. Nabi, "Physics-informed neural ODE (PINODE): embedding physics into models using collocation points," *Scientific Reports*, vol. 13, no. 1, pp. 10166, June 2023.
- [15] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [16] E. J. Candès, "Compressive sampling," in *Proceedings of the International Congress of Mathematicians*. Madrid, Spain, 2006, vol. 3, pp. 1433–1452.
- [17] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [18] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, pp. 253–263, 2008.
- [19] A. Skodras, C. Christopoulos, and T. Ebrahimi, "The JPEG 2000 still image compression standard," *IEEE Signal Processing Magazine*, vol. 18, no. 5, pp. 36–58, 2001.
- [20] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *IEEE Global Conference on Signal and Information Processing*, 2013, pp. 945–948.
- [21] S. Sreehari, S. V. Venkatakrishnan, B. Wohlberg, G. T. Buzzard, L. F. Drummy, J. P. Simmons, and C. A. Bouman, "Plug-and-play priors for bright field electron tomography and sparse interpolation," *IEEE Transactions on Computational Imaging*, vol. 2, no. 4, pp. 408–423, 2016.
- [22] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," in *International Conference on Machine Learning*, 2017, pp. 537–546.
- [23] Z. Wu and R. Zhang, "Learning physics by data for the motion of a sphere falling in a non-Newtonian fluid," *Communications in Nonlinear Science and Numerical Simulation*, vol. 67, pp. 577–593, 2019.
- [24] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur, "A theory of multiplexed illumination," in *IEEE International Conference on Computer Vision*, 2003, vol. 2, pp. 808–815.
- [25] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch, "Dual photography," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 745–755, 2005.
- [26] L. Bian, J. Suo, Q. Dai, and F. Chen, "Experimental comparison of single-pixel imaging algorithms," *Journal of the Optical Society of America A*, vol. 35, no. 1, pp. 78–87, 2018.
- [27] M. P. Edgar, G. M. Gibson, and M. J. Padgett, "Principles and prospects for single-pixel imaging," *Nature Photonics*, vol. 13, no. 1, pp. 13–20, Jan. 2019.
- [28] G. M. Gibson, S. D. Johnson, and M. J. Padgett, "Single-pixel imaging 12 years on: A review," *Optics Express*, vol. 28, no. 19, pp. 28190–28208, 2020.
- [29] L. Bian, J. Suo, G. Situ, Z. Li, J. Fan, F. Chen, and Q. Dai, "Multispectral imaging using a single bucket detector," *Scientific Reports*, vol. 6, no. 1, pp. 1–7, 2016.
- [30] Z. Li, J. Suo, X. Hu, C. Deng, J. Fan, and Q. Dai, "Efficient single-pixel multispectral imaging via non-mechanical spatio-spectral modulation," *Scientific Reports*, vol. 7, no. 1, pp. 41435, 2017.
- [31] S. Jiao, J. Feng, Y. Gao, T. Lei, and X. Yuan, "Visual cryptography in single-pixel imaging," *Opt. Express*, vol. 28, no. 5, pp. 7301–7313, Mar 2020.
- [32] P. Clemente, V. Durán, V. Torres-Company, E. Tajahuerce, and J. Lancis, "Optical encryption based on computational ghost imaging," *Opt. Lett.*, vol. 35, no. 14, pp. 2391–2393, Jul 2010.
- [33] W. Chen and X. Chen, "Ghost imaging for three-dimensional optical security," *Applied Physics Letters*, vol. 103, no. 22, pp. 221106, 2013.
- [34] C. Zhao, W. Gong, M. Chen, E. Li, H. Wang, W. Xu, and S. Han, "Ghost imaging lidar via sparsity constraints," *Applied Physics Letters*, vol. 101, no. 14, pp. 141123, 2012.
- [35] E. Li, Z. Bo, M. Chen, W. Gong, and S. Han, "Ghost imaging of a moving target with an unknown constant speed," *Applied Physics Letters*, vol. 104, no. 25, pp. 251120, 2014.
- [36] B. Sun, M. P. Edgar, R. Bowman, L. E. Vittert, S. Welsh, A. Bowman, and M. J. Padgett, "3D computational imaging with single-pixel detectors," *Science*, vol. 340, no. 6134, pp. 844–847, 2013.
- [37] P. Zhang, W. Gong, X. Shen, and S. Han, "Correlated imaging through atmospheric turbulence," *Physical Review A*, vol. 82, no. 3, pp. 033817, 2010.
- [38] J. Cheng, "Ghost imaging through turbulent atmosphere," *Optics Express*, vol. 17, no. 10, pp. 7916–7921, 2009.
- [39] J. H. Shapiro, "Computational ghost imaging," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 78, no. 6, Dec. 2008.
- [40] B. I. Erkmen and J. H. Shapiro, "Ghost imaging: from quantum to classical to computational," *Advances in Optics and Photonics*, vol. 2, no. 4, pp. 405, Dec. 2010.
- [41] J. B. Sampsell, "An overview of Texas Instruments digital micromirror device (DMD) and its application to projection displays," in *Society for Information Display International Symposium Digest of Technical Papers*, 1993, vol. XXIV, pp. 1012–1015.
- [42] L. Hornbeck, "Current status of the digital micromirror device (DMD) for projection television applications," in *Proceedings of IEEE International Electron Devices Meeting*, Washington, DC, USA, 1993, pp. 381–384.
- [43] Z. Zhang, X. Wang, G. Zheng, and J. Zhong, "Fast Fourier single-pixel imaging via binary illumination," *Scientific Reports*, vol. 7, no. 1, pp. 12029, 2017.
- [44] M.-J. Sun, L.-T. Meng, M. P. Edgar, M. J. Padgett, and N. Radwell, "A Russian Dolls ordering of the Hadamard basis for compressive single-pixel imaging," *Scientific Reports*, vol. 7, no. 1, pp. 3464, 2017.
- [45] Z.-H. Xu, W. Chen, J. Penuelas, M. Padgett, and M.-J. Sun, "1000 fps computational ghost imaging using LED-based structured illumination," *Optics Express*, vol. 26, no. 3, pp. 2427–2434, 2018.

- [46] J. Suo, L. Bian, F. Chen, and Q. Dai, "Signal-dependent noise removal for color videos using temporal and cross-channel priors," *Journal of Visual Communication and Image Representation*, vol. 36, pp. 130–141, 2016.
- [47] L. S. Pontryagin, *Mathematical Theory of Optimal Processes*, Routledge, 2018.
- [48] P. Kidger, *On Neural Differential Equations*, Ph.D. thesis, University of Oxford, 2021.
- [49] A. Quaglino, M. Gallieri, J. Masci, and J. Koutník, "SNODE: Spectral discretization of neural ODEs for system identification," arXiv:1906.07038, 2019.
- [50] X. He, Y. Choi, W. D. Fries, J. L. Belof, and J.-S. Chen, "gLaSDI: Parametric physics-informed greedy latent space dynamics identification," *Journal of Computational Physics*, p. 112267, 2023.
- [51] D. Ruiz-Balet and E. Zuazua, "Neural ODE control for classification, approximation, and transport," *SIAM Review*, vol. 65, no. 3, pp. 735–773, 2023.
- [52] M. Zakwan, L. Xu, and G. Ferrari-Trecate, "Robust classification using contractive Hamiltonian neural ODEs," *IEEE Control Systems Letters*, vol. 7, pp. 145–150, 2022.
- [53] J. Xu, E. Z. Chen, X. Chen, T. Chen, and S. Sun, "Multi-scale neural ODEs for 3D medical image registration," in *Medical Image Computing and Computer Assisted Intervention*, Strasbourg, France, 2021, Springer, pp. 213–223.
- [54] J. Ru, B. Lu, B. Chen, J. Shi, G. Chen, M. Wang, Z. Pan, Y. Lin, Z. Gao, J. Zhou, et al., "Attention guided neural ODE network for breast tumor segmentation in medical images," *Computers in Biology and Medicine*, vol. 159, pp. 106884, 2023.
- [55] A. J. Linot and M. D. Graham, "Data-driven reduced-order modeling of spatiotemporal chaos with neural ordinary differential equations," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 32, no. 7, 2022.
- [56] J. M. Dhadphale, V. R. Unni, A. Saha, and R. Sujith, "Neural ODE to model and prognose thermoacoustic instability," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 32, no. 1, 2022.
- [57] F. Zhou and L. Li, "Forecasting reservoir inflow via recurrent neural ODEs," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021, vol. 35, pp. 15025–15032.
- [58] K. Lee and K. T. Carlberg, "Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders," *Journal of Computational Physics*, vol. 404, pp. 108973, 2020.
- [59] C. Gin, B. Lusch, S. L. Brunton, and J. N. Kutz, "Deep learning models for global coordinate transformations that linearise PDEs," *European Journal of Applied Mathematics*, vol. 32, no. 3, pp. 515–539, 2021.
- [60] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, "Data-driven discovery of coordinates and governing equations," *Proceedings of the National Academy of Sciences*, vol. 116, no. 45, pp. 22445–22451, 2019.
- [61] B. Kim, V. C. Azevedo, N. Thuerey, T. Kim, M. Gross, and B. Solenthaler, "Deep fluids: A generative network for parameterized fluid simulations," in *Computer Graphics Forum*. Wiley Online Library, 2019, vol. 38, pp. 59–70.
- [62] T. G. McRae and T. J. Kulp, "Backscatter absorption gas imaging: a new technique for gas visualization," *Applied Optics*, vol. 32, no. 21, pp. 4037–4050, July 1993.
- [63] T. J. Kulp, P. Powers, R. Kennedy, and U.-B. Goers, "Development of a pulsed backscatter-absorption gas-imaging system and its application to the visualization of natural gas leaks," *Applied Optics*, vol. 37, no. 18, pp. 3912–3922, June 1998.
- [64] D. J. M. Stothard, M. H. Dunn, and C. F. Rae, "Hyperspectral imaging of gases with a continuous-wave pump-enhanced optical parametric oscillator," *Optics Express*, vol. 12, no. 5, pp. 947–955, Mar. 2004.
- [65] G. Gibson, B. van Well, J. Hodgkinson, R. Pride, R. Strzoda, S. Murray, S. Bishon, and M. Padgett, "Imaging of methane gas using a scanning, open-path laser system," *New Journal of Physics*, vol. 8, no. 2, pp. 26, Feb. 2006.
- [66] K. J. Nutt, N. Hempler, G. T. Maker, G. P. A. Malcolm, M. J. Padgett, and G. M. Gibson, "Developing a portable gas imaging camera using highly tunable active-illumination and computer vision," *Optics Express*, vol. 28, no. 13, pp. 18566–18576, June 2020.
- [67] P. Gattinger, I. Zorin, A. Ebner, C. Rankl, and M. Brandstetter, "Mid-infrared DMD-based spectral-coding spectroscopy with a supercontinuum laser source," *Optics Express*, vol. 30, no. 4, pp. 6440–6449, Feb. 2022.
- [68] F. Ferri, D. Magatti, L. Lugiato, and A. Gatti, "Differential ghost imaging," *Physical Review Letters*, vol. 104, no. 25, pp. 253603, 2010.
- [69] A. Pastuszczak, R. Stojek, P. Wróbel, and R. Kotyński, "Differential real-time single-pixel imaging with Fourier domain regularization: applications to VIS-IR imaging and polarization imaging," *Optics Express*, vol. 29, no. 17, pp. 26685–26700, 2021.
- [70] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, 2012.
- [71] B. R. Noack, M. Morzynski, and G. Tadmor, *Reduced-Order Modelling for Flow Control*, vol. 528, Springer Science & Business Media, 2011.
- [72] B. Chen, K. Huang, S. Raghupathi, I. Chandratreya, Q. Du, and H. Lipson, "Discovering state variables hidden in experimental data," arXiv:2112.10755, 2021.
- [73] N. Takeishi, Y. Kawahara, and T. Yairi, "Learning Koopman invariant subspaces for dynamic mode decomposition," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [74] J. Morton, F. D. Witherden, and M. J. Kochenderfer, "Deep variational Koopman models: Inferring Koopman observations for uncertainty-aware dynamics modeling and control," arXiv:1902.09742, 2019.
- [75] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [76] L. N. Trefethen, *Spectral Methods in MATLAB*, SIAM, 2000.
- [77] W. Gong and S. Han, "A method to improve the visibility of ghost images obtained by thermal light," *Physics Letters A*, vol. 374, no. 8, pp. 1005–1008, 2010.
- [78] K. M. Czajkowski, A. Pastuszczak, and R. Kotyński, "Real-time single-pixel video imaging with Fourier domain regularization," *Optics Express*, vol. 26, no. 16, pp. 20009–20022, 2018.
- [79] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Processing Magazine*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [80] V. Dovzhenko, A. Obukhov, and V. Ponomarev, "Generation of vortices in axisymmetric shear flow," *Akademiia Nauk SSSR Izvestiia Mekhanika Zhidkosti i Gaza*, vol. 17, pp. 27–36, 1981.
- [81] J. Tithof, B. Suri, R. K. Pallantla, R. O. Grigoriev, and M. F. Schatz, "Bifurcations in a quasi-two-dimensional Kolmogorov-like flow," *Journal of Fluid Mechanics*, vol. 828, pp. 837–866, 2017.
- [82] B. Suri, J. Tithof, R. O. Grigoriev, and M. F. Schatz, "Forecasting fluid flows using the geometry of turbulence," *Physical Review Letters*, vol. 118, no. 11, pp. 114501, 2017.
- [83] Z. Y. Wan, P. Vlachas, P. Koumoutsakos, and T. Sapsis, "Data-assisted reduced-order modeling of extreme events in complex dynamical systems," *PLOS One*, vol. 13, no. 5, 2018.
- [84] M.-L. Eckert, K. Um, and N. Thuerey, "ScalarFlow: a large-scale volumetric data set of real-world scalar transport flows for computer animation and machine learning," *ACM Transactions on Graphics*, vol. 38, no. 6, pp. 239:1–16, Nov. 2019.
- [85] J. Wang, L. P. Tchapmi, A. P. Ravikumar, M. McGuire, C. S. Bell, D. Zimmerle, S. Savarese, and A. R. Brandt, "Machine vision for natural gas methane emissions detection using an infrared camera," *Applied Energy*, vol. 257, pp. 113998, 2020.
- [86] S. Nabi, P. Grover, and C. Caulfield, "Nonlinear optimal control strategies for buoyancy-driven flows in the built environment," *Computers & Fluids*, vol. 194, pp. 104313, 2019.
- [87] S. Nabi, P. Grover, and C. Caulfield, "Robust preconditioned one-shot methods and direct-adjoint-looping for optimizing Reynolds-averaged turbulent flows," *Computers & Fluids*, vol. 238, pp. 105390, 2022.
- [88] Z. Zivkovic and F. van der Heijden, "Efficient adaptive density estimation per image pixel for the task of background subtraction," *Pattern Recognition Letters*, vol. 27, no. 7, pp. 773–780, 2006.
- [89] Z. Zivkovic, "Improved adaptive Gaussian mixture model for background subtraction," in *Proceedings of the International Conference on Pattern Recognition (ICPR)*, 2004, vol. 2, pp. 28–31.
- [90] G. Bradski, "The OpenCV Library," *Dr. Dobbs's Journal of Software Tools*, 2000.
- [91] E. Hagenaaars, A. Pandharipande, A. Murthy, and G. Leus, "Single-pixel thermopile infrared sensing for people counting," *IEEE Sensors Journal*, vol. 21, no. 4, pp. 4866–4873, Feb. 2021.
- [92] Y. Zhang, G. M. Gibson, M. P. Edgar, G. Hammond, and M. J. Padgett, "Dual-band single-pixel telescope," *Optics Express*, vol. 28, no. 12, pp. 18180–18188, June 2020.





**Aleksei Sholokhov** received the B.S. degree in mathematics from the Moscow Institute of Physics and Technology, Moscow, Russia, in 2018, with a focus on optimization methods, and his PhD in Applied Mathematics from the University of Washington, Seattle, WA, USA, in 2023. He is currently a Machine Learning Engineer with Stripe, Inc.



**Saleh Nabi** received the B.Sc. from K. N. Toosi University of Technology, Tehran, Iran, in 2005, the M.Sc. from Isfahan University of Technology, Isfahan, Iran, in 2008, and the Ph.D. from University of Alberta, Canada, in 2013, all in Mechanical Engineering. Currently serving as a Principal AI Researcher at Schneider Electric, Saleh previously held the position of Principal Research Scientist at Mitsubishi Electric Research Labs (MERL). His research interests are at the intersection of fluid mechanics, scientific machine learning, dynamical

systems, and optimal control in complex systems. His current research involves hybrid methods using traditional tools along with deep learning-based methods for efficient and robust control and estimation of PDEs with applications to HVACs, li-ion batteries, and energy storage systems.



**Joshua Rapp** (S'12-M'21, IEEE) received the B.S. degree in electrical engineering from Tufts University, Medford, MA, USA in 2014 and the M.S. and Ph.D. degrees in electrical engineering from Boston University, Boston, MA, USA in 2016 and 2020, respectively.

He was a Doctoral Fellow at Draper, Cambridge, MA, USA from 2015 to 2019 and a Postdoctoral Researcher with the Department of Electrical Engineering at Stanford University, Stanford, CA, USA from 2020 to 2021. He is currently a Research

Scientist with Mitsubishi Electric Research Laboratories, Cambridge, MA, USA. His research interests include computational imaging, optical sensing, lidar, and statistical signal processing.

Dr. Rapp is a member of Optica, Eta Kappa Nu, and Tau Beta Pi. He was the winner of the 2021 IEEE Signal Processing Society Best PhD Dissertation award and the 2020 Best EE Dissertation Award from Boston University. His work also received a Best Student Paper Award at the IEEE International Conference on Image Processing in 2018, a Young Author Best Paper Award from the IEEE Signal Processing Society in 2020 and the Best Student Poster award at the Conference on Lasers and Electro-Optics in 2021.



**Steven L. Brunton** (Senior Member, IEEE) received the B.S. degree in mathematics from the California Institute of Technology, Pasadena, CA, USA, in 2006, with a focus on control and dynamical systems, and the Ph.D. degree in mechanical and aerospace engineering from Princeton University, Princeton, NJ, USA, in 2012. He is currently Professor of Mechanical Engineering and a Data Science Fellow with the eScience Institute, University of Washington.



**J. Nathan Kutz** (Senior Member, IEEE) received the B.S. degree in physics and mathematics from the University of Washington, Seattle, WA, USA, in 1990, and the Ph.D. degree in applied mathematics from Northwestern University, Evanston, IL, USA, in 1994. He is currently Director of the AI Institute in Dynamics Systems, a Professor of applied mathematics and electrical engineering, and a Senior Data Science Fellow with the eScience Institute, University of Washington.



**Petros T. Boufounos** (S'02, M'06, SM'13, F'22) is a Distinguished Research Scientist, a Deputy Director and the Computational Sensing Senior Team Leader at Mitsubishi Electric Research Laboratories (MERL). Dr. Boufounos completed his undergraduate and graduate studies at MIT. He received the S.B. degree in Economics in 2000, the S.B. and M.Eng. degrees in Electrical Engineering and Computer Science (EECS) in 2002, and the Sc.D. degree in EECS in 2006. Between September 2006 and December 2008, he was a postdoctoral associate

with the Digital Signal Processing Group at Rice University. Dr. Boufounos joined MERL in January 2009, where he has been heading the Computational Sensing Team since 2016.

Dr. Boufounos' immediate research focus includes signal acquisition and processing, computational sensing, inverse problems, quantization, and data representations. He is also interested in how signal acquisition interacts with other fields that use sensing extensively, such as machine learning, robotics, and dynamical system theory. He has over 40 patents granted and more than 10 pending, and more than 100 peer reviewed journal and conference publications in these topics. Dr. Boufounos was the general co-chair of the ICASSP 2023 organizing committee and is currently a regional director-at-large in the IEEE Signal Processing Society's Board of Governors. He has also served as an Area Editor and a Senior Area Editor for the IEEE Signal Processing Letters, an AE for IEEE Transactions on Computational Imaging, and as a member of the SigPort editorial board and the IEEE Signal Processing Society Theory and Methods technical committee. Dr. Boufounos is an IEEE Fellow and an IEEE SPS Distinguished Lecturer for 2019-2020.



**Hassan Mansour** (S'99-M'09-SM'17) received the B.E. degree in computer and communications engineering from the American University of Beirut, Beirut, Lebanon, in 2003, and the M.A.Sc. degree in electrical and computer engineering and the Ph.D. degree in electrical and computer engineering from The University of British Columbia, Vancouver, BC, Canada, in 2005 and 2009, respectively. Between January 2010 and January 2013, he was a Postdoctoral Research Fellow with the Department of Computer Science, the Mathematics Department,

and the Department of Earth, Ocean, and Atmospheric Sciences, The University of British Columbia. He is a Senior Principal Research Scientist and Computational Sensing Team Leader with Mitsubishi Electric Research Laboratories, Cambridge, MA, USA. His research interests are in inverse problems, machine learning, compressed sensing, sparse signal reconstruction, image enhancement, and scalable video compression and transmission. His current research is focused on the design of efficient acquisition schemes and reconstruction algorithms for natural images, radar sensing, video analytics, and seismic imaging. Dr. Mansour is a member of the IEEE Signal Processing Society. He has also served on the Computational Imaging Technical Committee and the Sensor Array and Multichannel Technical Committee. He was an Associate Editor for the IEEE Transactions on Signal Processing between 2018 and 2022.