


Extracting Interpretable Physical Parameters from Spatiotemporal Systems Using Unsupervised Learning

Peter Y. Lu^{1,*}, Samuel Kim², and Marin Soljačić¹

¹*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

 (Received 27 August 2019; revised 6 May 2020; accepted 21 July 2020; published 11 September 2020)

Experimental data are often affected by uncontrolled variables that make analysis and interpretation difficult. For spatiotemporal systems, this problem is further exacerbated by their intricate dynamics. Modern machine learning methods are particularly well suited for analyzing and modeling complex datasets, but to be effective in science, the result needs to be interpretable. We demonstrate an unsupervised learning technique for extracting interpretable physical parameters from noisy spatiotemporal data and for building a transferable model of the system. In particular, we implement a physics-informed architecture based on variational autoencoders that is designed for analyzing systems governed by partial differential equations. The architecture is trained end to end and extracts latent parameters that parametrize the dynamics of a learned predictive model for the system. To test our method, we train our model on simulated data from a variety of partial differential equations with varying dynamical parameters that act as uncontrolled variables. Numerical experiments show that our method can accurately identify relevant parameters and extract them from raw and even noisy spatiotemporal data (tested with roughly 10% added noise). These extracted parameters correlate well (linearly with $R^2 > 0.95$) with the ground truth physical parameters used to generate the datasets. We then apply this method to nonlinear fiber propagation data, generated by an *ab initio* simulation, to demonstrate its capabilities on a more realistic dataset. Our method for discovering interpretable latent parameters in spatiotemporal systems will allow us to better analyze and understand real-world phenomena and datasets, which often have unknown and uncontrolled variables that alter the system dynamics and cause varying behaviors that are difficult to disentangle.

DOI: [10.1103/PhysRevX.10.031056](https://doi.org/10.1103/PhysRevX.10.031056)

Subject Areas: Computational Physics,
Interdisciplinary Physics,
Nonlinear Dynamics

I. INTRODUCTION

Physics has traditionally relied upon human ingenuity to identify key variables, discover physical laws, and model dynamical systems. With the recent explosion of available data coupled with advances in machine learning, fundamentally new methods of discovery are now possible. However, a major issue with applying these novel techniques to scientific and industrial applications is their interpretability: neural networks and deep learning are often seen as inherently black box methods. To make progress, we must incorporate scientific domain knowledge into our network architecture design and algorithms

without sacrificing the flexibility provided by deep learning models [1–3]. In this work, we show that we can leverage unsupervised learning techniques in a physics-informed architecture to build models that learn to both identify relevant interpretable parameters and perform prediction. Because relevant parameters are necessary for predictive success, the two tasks of extracting parameters and creating a predictive model are closely linked, and we exploit this relationship to do both using a single architecture.

We focus our attention on spatiotemporal systems with dynamics governed by partial differential equations (PDEs). These systems are ubiquitous in nature and include physical phenomena in fluid dynamics and electromagnetism. Recently, there has been significant interest in the data-driven analysis and discovery of PDEs: e.g., explicitly identifying PDEs using sparse linear regression with a library of possible terms [4], using a convolutional architecture with symbolic regression to identify PDEs [5,6], and representing PDE solutions as neural networks to solve and identify PDEs [7–9]. However, previous works on PDE

*lup@mit.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

discovery and parameter extraction often assume the entire dataset is governed by the same dynamics and also explicitly provide the key dynamical parameters (with potentially unknown values). In more complex scenarios, we may have limited control over the systems that we are studying and yet still want to model them and extract relevant physical features from their dynamics. If we attempt to study such systems by naively training a predictive model, we are likely to fail in one of two ways: first, a single explicit PDE model will be unable to capture the variations in the dynamics caused by uncontrolled variables in the data, and second, a generic deep learning method for time-series prediction such as long short-term-memory-based models [10,11] will not be interpretable or provide any physical insight and may also result in unphysical solutions at later times due to overfitting. To avoid these problems and gain a better understanding of the physical system, we must first identify important parameters or variables that are uncontrolled and that change in the raw data, producing varying dynamics. Recent work on learning parametric PDEs has taken steps toward addressing this issue [12]. We use an unsupervised learning method to automate the process of determining the relevant parameters that control the system dynamics and constructing a predictive model—all without requiring information on the form of the governing PDE.

We propose a model architecture [Fig. 1(a)] based on variational autoencoders (VAEs) [13]. VAEs are widely used for dimensionality reduction and unsupervised learning tasks [14] and have been shown to be effective for studying a wide variety of physical phenomena, e.g., discovering simple representations of systems in classical and quantum mechanics [15], modeling protein folding and molecular dynamics [16,17], and identifying condensed matter phase transitions [18,19]. In terms of interpretability, the VAE architecture and its derivatives have also been shown to disentangle independent factors of variation [20–22]. The choice of a VAE-based architecture is motivated both by their prior success in extracting useful representations and by their strong theoretical foundation (Appendix E). Prior works have also applied other methods of parameter identification, such as using principal component analysis as a postprocessing step with a standard autoencoder [23]. This, however, relies heavily on the poorly understood implicit regularization of the neural network architecture rather than the explicit regularization of a VAE, and, in our experience, VAEs produce more consistent and interpretable results.

Our architecture consists of an encoder [Fig. 1(b)] that extracts physical parameters characterizing the system dynamics and a decoder [Fig. 1(c)] that acts as a predictive model and propagates an initial condition forward in time given the extracted parameters. This differs from a traditional VAE due to the additional initial condition provided to the decoder, allowing the encoder to focus on extracting

latent parameters that parametrize the dynamics of the system rather than the physical state. Our architecture can be thought of as a conditional VAE [24], although only the decoder is conditional. While similar architectures have been recently proposed for physical systems such as interacting particles [23] and moving objects [25], our model is specifically designed to study spatiotemporal phenomena, which have a continuous set of degrees of freedom.

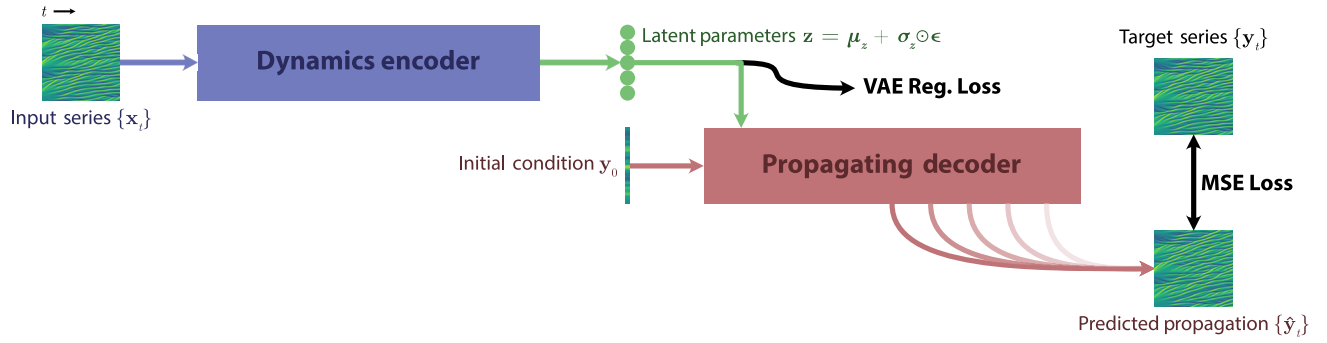
To take advantage of the spatiotemporal structure of PDE-governed systems, we use convolutional layers—commonly employed in image recognition tasks [26,27] to efficiently represent local features—in both the encoder and decoder portions of our architecture. The translation invariance of the convolutions allows us to train on small patches of data and then evaluate on larger systems with arbitrary boundary conditions. In the decoder, the convolutional layers are placed in a recurrent architecture to represent time propagation—analogue to a PDE solver with a finite difference approximation [5]. In addition, our architecture efficiently parametrizes PDE propagation by dynamically generating the convolutional kernels and biases of the decoder using the extracted latent parameters from the encoder. In this way, the latent parameters directly control the local propagation of the physical states in the decoder, resulting in more stable model predictions and a more physical encoding of the dynamics. These architecture choices provide the key physics-informed inductive biases that enhance the interpretability of the extracted parameters and ensure a physically reasonable predictive model.

To demonstrate the capabilities of this approach, we test our method on simulated data from PDE models for chaotic wave dynamics, optical nonlinearities, and convection and diffusion (Sec. III). These numerical experiments show that our method can accurately identify and extract relevant physical parameters that characterize variations in the observed dynamics of a spatiotemporal system (Sec. IV A), while at the same time construct a flexible and transferable predictive model (Sec. IV B). We further show that the parameter extraction is robust to noisy data and can still be effective for chaotic systems where accurate prediction is difficult. Finally, we apply this method to nonlinear optical fiber propagation using data generated from an *ab initio* electromagnetic simulation to test the model on a more realistic dataset (Sec. V). The goal of our approach is to provide an additional tool for studying complex spatiotemporal systems when there are unknown and uncontrolled variables present.


II. MODEL ARCHITECTURE

Our model (Fig. 1) has an encoder-decoder architecture based on a variational autoencoder [13,20]. Given a dataset of time series from a spatiotemporal system, the *dynamics encoder* (DE) extracts latent parameters which parametrize the varying dynamics in the dataset. These latent parameters are then used by the *propagating decoder* (PD) to

(a) Architecture overview



(b) Dynamics encoder



(c) Propagating decoder




FIG. 1. VAE-based model architecture. (a) The architecture consists of the dynamics encoder (DE) and the propagating decoder (PD) with kernels and biases given by a latent-to-kernel network. (b) The DE extracts the latent distribution parameters μ_z and σ_z from the input series $\{x_i\}_{i=0}^{T_x}$ using dilated convolutions and inverse-variance weighted (IVW) averaging. During training, the latent parameters are sampled from the extracted distribution $z \sim \mathcal{N}(\mu_z, \sigma_z^2)$. (c) The PD then uses a fully connected latent-to-kernel network to map the latent parameters z to the kernels and biases of the dynamic convolutional layers, which are used in a recurrent fashion to predict the propagation of the system $\{\hat{y}_i\}_{i=1}^{T_y}$ starting from the initial condition $\hat{y}_0 = y_0$. The model is trained end to end using the mean-squared error (MSE) loss between the predicted propagation series $\{\hat{y}_i\}_{i=1}^{T_y}$ and target series $\{y_i\}_{i=1}^{T_y}$ along with VAE regularization. Time-series limits are dropped in the labels for conciseness.

simulate the system given an initial condition and boundary conditions. During training, the model is optimized to match the output of the PD to a time-series example from the dataset. The goal of the VAE architecture is to allow the

PD to push the DE to extract useful and informative latent parameters.

For training, the network requires time-series data that are grouped in pairs: the *input series* $\{x_i\}_{i=0}^{T_x}$ is the input to

the DE, and the *target series* $\{\mathbf{y}_t\}_{t=0}^{T_y}$ provides the initial condition \mathbf{y}_0 and the training targets $\{\mathbf{y}_t\}_{t=1}^{T_y}$ for the PD. Each pair of time series $(\{\mathbf{x}_t\}_{t=0}^{T_x}, \{\mathbf{y}_t\}_{t=0}^{T_y})$ must follow the same dynamics and thus correspond to the same latent parameters. We can construct such a dataset from the raw data by cropping each original time series to produce a pair of smaller time series. This cropping can be performed randomly in both the time and space dimensions, which allows the network to train on a reduced system size while still making use of all the available data. In our examples, we also choose to crop dynamically during training, akin to data augmentation methods used in image recognition [27].

In detail, the DE network takes the full input series $\{\mathbf{x}_t\}_{t=0}^{T_x}$ and outputs a mean $\mu_{z|\mathbf{x}}$ and a variance $\sigma_{z|\mathbf{x}}^2$, which we will henceforth refer to as μ_z and σ_z^2 for compactness, representing a normal distribution for each latent parameter z . During training, each z is sampled from its corresponding distribution $\mathcal{N}(\mu_z, \sigma_z^2)$ using the VAE reparametrization trick: $z = \mu_z + \sigma_z \epsilon$, where $\epsilon \sim \mathcal{N}(0, 1)$ is independently sampled for every training example during each training step. During evaluation, we simply take $z = \mu_z$. These parameters z along with an initial condition—the first state \mathbf{y}_0 in the target series—are then used by the PD network to predict the future propagation of the system. The predicted propagation series $\{\hat{\mathbf{y}}_t\}_{t=1}^T$ produced by the PD can be computed up to an arbitrary future time T , where $T = T_y$ during training to match the target series. By providing the PD with an initial condition, we allow the DE to focus on encoding parameters that characterize the dynamics of the data rather than encoding a particular state of the system. This is further reinforced by training on randomly cropped pairs of time series as well as by the VAE regularization term (Appendix E).

The full architecture is trained end to end using a mean-squared error loss between the predicted propagation series $\{\hat{\mathbf{y}}_t\}_{t=1}^{T_y}$ from the PD and the target series $\{\mathbf{y}_t\}_{t=1}^{T_y}$. We also add the VAE regularization loss—the Kullback-Leibler divergence term $D_{\text{KL}}[\mathcal{N}(\mu_z, \sigma_z^2) \|\mathcal{N}(0, 1)]$ —which encourages each latent parameter distribution $\mathcal{N}(\mu_z, \sigma_z^2)$ generated by the DE to approach the standard normal prior distribution $\mathcal{N}(0, 1)$. The total loss function is given by

$$\mathcal{L} = \frac{1}{T_y} \sum_{t=1}^{T_y} (\mathbf{y}_t - \hat{\mathbf{y}}_t)^2 \quad (1)$$

$$+ \beta \sum_z D_{\text{KL}}[\mathcal{N}(\mu_z, \sigma_z^2) \|\mathcal{N}(0, 1)], \quad (2)$$

where T_y is the length of the target series without the initial \mathbf{y}_0 , and β is a regularization hyperparameter which we tune for each dataset. The β parameter is key to learning disentangled representations [20,21]. By using the VAE sampling method and regularizer, we compel the

model to learn independent and interpretable latent parameters (Appendix E). For additional training details, see Appendix B.

The source code for our implementation is available in Ref. [28].

A. Dynamics encoder

The dynamics encoder network is designed to take advantage of existing symmetry or structure in the time-series data. We implement the DE as a deep convolutional network in both the time and space dimensions to allow the network to efficiently extract relevant features. To ensure the DE can handle arbitrary system sizes and time-series lengths, the architecture only contains convolutional layers with a weighted average applied at the output to obtain the latent parameters. The weights for the final averaging are also learned by the network and interpreted as variances so that the overall variance can also be computed. In this way, the network is able to focus on areas of the input series that are most important for estimating the latent parameters, akin to a visual attention mechanism [29].

Explicitly, we first compute local quantities,

$$\mu(t, \mathbf{r}) = f_{\text{DE}, \mu}(\{\mathbf{x}_{t'}\}_{t'=0}^{T_x}), \quad (3)$$

$$\log \sigma^2(t, \mathbf{r}) = f_{\text{DE}, \sigma^2}(\{\mathbf{x}_{t'}\}_{t'=0}^{T_x}), \quad (4)$$

where $f_{\text{DE}, \mu}$ and f_{DE, σ^2} are multilayer convolutional networks in space and time (see Appendix A for details). Then, instead of using a fully connected layer to compute the final mean μ_z and variance σ_z^2 for each latent parameter, we combine the local quantities by performing an inverse-variance weighted average using weights given by

$$w(t, \mathbf{r}) = \sigma^{-2}(t, \mathbf{r}) / \sum_{t, \mathbf{r}} \sigma^{-2}(t, \mathbf{r}) \quad (5)$$

to obtain

$$\mu_z = \sum_{t, \mathbf{r}} w(t, \mathbf{r}) \mu(t, \mathbf{r}), \quad (6)$$

$$\sigma_z^2 = C^d / \sum_{t, \mathbf{r}} \sigma^{-2}(t, \mathbf{r}), \quad (7)$$

where C is a constant chosen to correct for the correlations between nearby points and d is the total number of time and space dimensions of the input. This averaging serves two purposes: it allows the DE to scale to arbitrary system sizes and geometries, and it improves the parameter extraction by placing greater emphasis on regions of high confidence. Assuming that nonoverlapping patches should be treated as independent while overlapping patches are increasingly

correlated, we take $C = 31$ to be the linear size of the receptive field of the convolutional networks $f_{\text{DE},\mu}$ and f_{DE,σ^2} .

B. Propagating decoder

The propagating decoder network is designed as a predictive model for spatiotemporal systems. We structure the PD as a multilayered convolutional network f_{PD} (see Appendix A for details) with a residual skip connection that maps a state $\hat{\mathbf{y}}_t$ to the next state in the time series $\hat{\mathbf{y}}_{t+1}$. Thus, each propagation step is given by

$$\hat{\mathbf{y}}_{t+1} = \hat{\mathbf{y}}_t + f_{\text{PD}}(\hat{\mathbf{y}}_t). \quad (8)$$

To generate the predicted propagation series $\{\hat{\mathbf{y}}_t\}_{t=1}^{T_y}$ for comparison with the target series $\{\mathbf{y}_t\}_{t=1}^{T_y}$, we begin with the initial condition $\hat{\mathbf{y}}_0 = \mathbf{y}_0$ and then recursively apply the PD to propagate $\hat{\mathbf{y}}_t \rightarrow \hat{\mathbf{y}}_{t+1}$, forming a recurrent network. The PD acts as a physics simulator or, in this case, a PDE integrator with explicit time stepping. This architecture reflects both the spatial and temporal structure of PDE-governed systems and incorporates boundary conditions by properly padding $\hat{\mathbf{y}}_t$ at each time step before applying the convolutional layers. For example, to use periodic boundary conditions during evaluation, we apply periodic padding at each time step. During training, we treat the edges of the target series—cropped from a full training example—as a boundary condition by using the spatial boundary of each \mathbf{y}_t in the target series to pad the corresponding state $\hat{\mathbf{y}}_t$ before propagation.

Unlike the convolutional layers of the DE, the kernels and biases for f_{PD} are not directly trained. Instead, the kernel weights and biases are a function of the latent parameters \mathbf{z} . This type of layer is known as a dynamic convolution [30] or a cross-convolution [25]. Each convolutional kernel and corresponding bias is constructed by a separate fully connected *latent-to-kernel* network that maps the latent parameters to each kernel or bias, forming a multiplicative connection in the PD. Thus, we can interpret the PD convolutional kernels and biases as encoding the dynamics of the system parametrized by \mathbf{z} .

III. SIMULATED PDE DATASETS

To study the ability of our architecture to perform parameter extraction, we generate simulated datasets of spatiotemporal systems that have spatially uniform, time-independent local dynamics in a box with periodic boundary conditions; i.e., we consider PDEs of the form

$$\frac{\partial \mathbf{u}(t, \mathbf{r})}{\partial t} = F[\mathbf{u}, \nabla \mathbf{u}, \mathbf{u}^2, (\mathbf{u} \cdot \nabla) \mathbf{u}, \dots], \quad (9)$$

where F is a general space- and time-independent, nonlinear local operator acting on \mathbf{u} . This allows us to design

an optimized, physics-informed model architecture. We test our model on a variety of spatiotemporal systems by creating the following three datasets that cover linear, nonlinear, and chaotic dynamics as well as giving both 1D and 2D examples. For details on the generation of the simulated datasets, see Appendix D.

A. 1D Kuramoto-Sivashinsky dataset

The Kuramoto-Sivashinsky equation,

$$\frac{\partial u}{\partial t} = -\gamma \partial_x^4 u - \partial_x^2 u - u \partial_x u, \quad (10)$$

is a nonlinear scalar wave equation with a viscosity damping parameter γ . This is a key example of a chaotic PDE [31] due to the instability caused by the negative second derivative term and was originally derived to model laminar flame fronts [32,33]. The 1D Kuramoto-Sivashinsky dataset has a training set with 5000 examples and a test set with 10 000 examples.

B. 1D Nonlinear Schrödinger dataset

The nonlinear Schrödinger equation,

$$i \frac{\partial \psi}{\partial t} = -\frac{1}{2} \partial_x^2 \psi + \kappa |\psi|^2 \psi, \quad (11)$$

is a complex scalar wave equation with a cubic nonlinearity controlled by the coefficient κ . In our data, we represent $\psi = u_1 + iu_2$ as a real two-component vector $\mathbf{u} = (u_1, u_2)$. This equation can be used to model the evolution of wave packets in nonlinear optics and is known to exhibit soliton solutions [34]. The 1D nonlinear Schrödinger dataset has a training set with 5000 examples and a test set with 10 000 examples.

C. 2D Convection-diffusion dataset

The 2D convection-diffusion equation,

$$\frac{\partial u}{\partial t} = D \nabla^2 u - \mathbf{v} \cdot \nabla u, \quad (12)$$

is a linear scalar wave equation consisting of a diffusion term with constant D and a velocity-dependent convection term with velocity field \mathbf{v} . The equation describes a diffusing quantity that is also affected by the flow or drift of the system, e.g., dye diffusing in a moving fluid. We consider the case of a constant velocity field. The 2D convection-diffusion dataset has a training set with 1000 examples and a test set with 1000 examples.

IV. NUMERICAL EXPERIMENTS

We perform numerical experiments by training the model on both the original noiseless datasets and the datasets with added $\sigma = 0.1$ Gaussian noise—corresponding to 10% noise

relative to the initial conditions. Then, we evaluate the trained models on the full size noiseless test set examples (no cropping). By also training on noisy datasets, we test the robustness of our method and show the effect of noise on the extracted parameters and prediction performance.

A. Parameter extraction

During training, the model will only use a minimal set of latent parameters to encode the variation in the dynamics and will align each latent parameter in this subspace with an independent factor of variation due to the VAE regularization [20,21]. Intuitively, the regularization encourages each latent parameter to independently collapse to a noninformative prior $\mathcal{N}(0, 1)$, and so the model prefers to minimize its use of the latent parameters and maintain their independence (Appendix E). Therefore, the number of latent parameters provided to the model is not critical as long as it is greater than the number of independent factors of variation. In our experiments, we allow the model to use five latent parameters. Because the 1D datasets have only one varying physical parameter and the 2D dataset has three varying physical parameters, the trained model will only make use of one or three latent parameters, respectively, and the rest will collapse to the prior.

We can determine the number of relevant latent parameters and empirically verify this claim by examining the statistics of the extracted distribution parameters μ_z and σ_z^2 from the dynamics encoder for each dataset. A latent parameter that is useful to the propagating decoder for predicting the target series will have high variance in μ_z and a low mean σ_z^2 , implying that the extracted parameter is precise and informative. A parameter which has collapsed to the prior and is noninformative will have low variance in μ_z and high mean σ_z^2 . These statistics indeed show that our model can correctly determine the number of relevant parameters for each dataset (Fig. 2). In real applications, we will not have access to the ground truth physical parameters, so we must rely on these metrics to identify the relevant parameters extracted by the model.

To evaluate the performance of our parameter extraction method, we compare the extracted latent parameters from the model with the true physical parameters used to generate our simulated datasets: the viscosity damping parameter γ for the 1D Kuramoto-Sivashinsky dataset, the nonlinearity parameter κ for the 1D nonlinear Schrödinger dataset, and the diffusion constant D and drift velocity components v_x, v_y for the 2D convection-diffusion dataset. Because these simulated physical parameters are drawn from normal distributions (Appendix D), we expect the relevant latent parameters—which have prior distribution $\mathcal{N}(0, 1)$ —to be linearly related to the true parameters (Appendix E). For real experimental systems, this is also a reasonable assumption for uncontrolled variables because natural parameters tend to be normally distributed due to the central limit theorem. We assess the quality of the

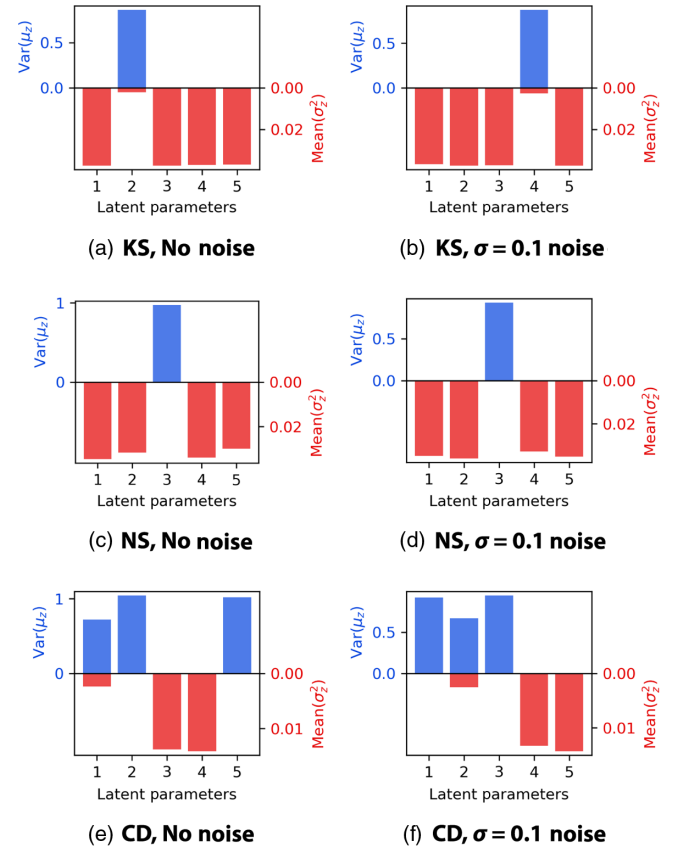


FIG. 2. Identification of relevant latent parameters: the variance of μ_z (blue) and mean of σ_z^2 (red) for the five latent parameters in the models trained on the (a),(b) 1D Kuramoto-Sivashinsky (KS), (c),(d) 1D nonlinear Schrödinger (NS), and (e),(f) 2D convection-diffusion (CD) datasets both with and without added noise. In each case, the model has correctly identified the number of relevant parameters (one for the Kuramoto-Sivashinsky and nonlinear Schrödinger datasets, and three for the convection-diffusion dataset), which are characterized by high variance in μ_z and a low mean σ_z^2 . These relevant latent parameters correspond to interpretable physical parameters that parametrize the dynamics of the system. The other latent parameters with near zero variance in μ_z and high mean σ_z^2 have collapsed to the prior and are noninformative. Note that while one would expect these collapsed parameters to have $\sigma_z^2 = 1$, the actual extracted σ_z^2 for the collapsed noninformative parameters is less than one. This is an artifact of evaluating the model on a larger system size and longer time series than the cropped patches used during training (see Appendix F for details).

extracted parameters by linearly fitting the relevant latent parameters with the ground truth physical parameters to obtain parameter predictions and R^2 correlation coefficients. Our numerical experiments show excellent parameter extraction on all three datasets (Fig. 3) with $R^2 > 0.95$ for all parameters (Table I) and no degradation in performance with added Gaussian $\sigma = 0.1$ noise. However, we do observe some nonlinear behavior at the edges of the parameter range likely due to data sparsity in those regions.

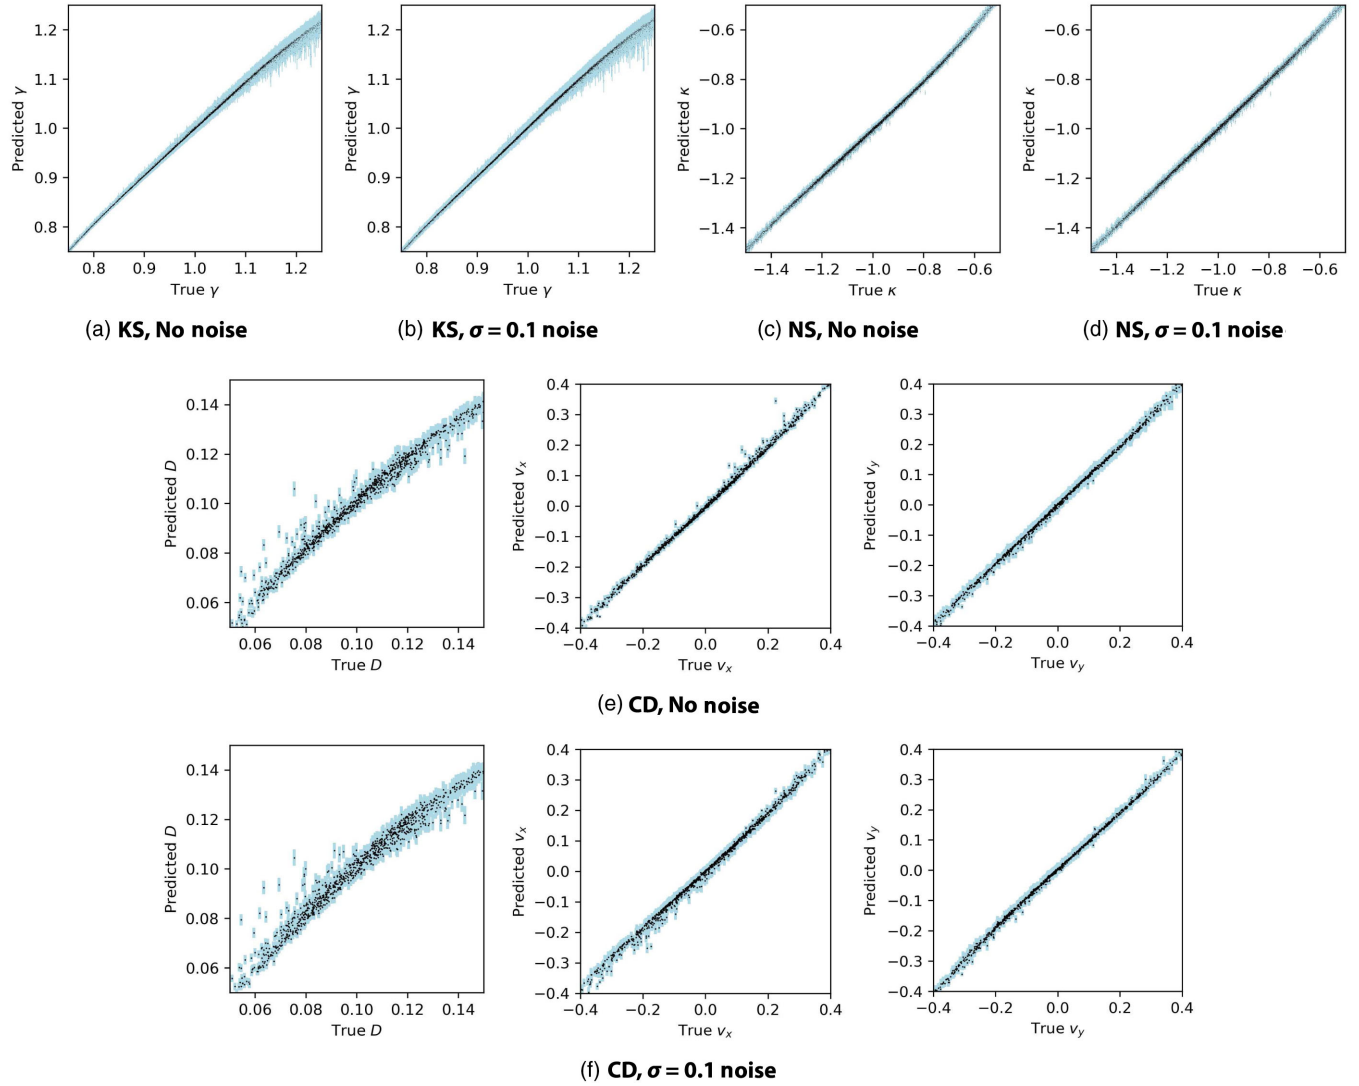


FIG. 3. Predicted physical parameters from a linear fit with the relevant latent parameter (Fig. 2) versus the ground truth physical parameters from the (a),(b) 1D Kuramoto-Sivashinsky, (c),(d) 1D nonlinear Schrödinger, and (e),(f) 2D convection-diffusion datasets. Because the drift velocity \mathbf{v} from the CD dataset has an inherent rotational symmetry, they are encoded in a two-dimensional latent subspace (Table II), so we instead show the predicted drift velocity components v_x , v_y from a multivariate linear regression in the subspace of two relevant latent parameters, which extracts the linear combination of latent parameters that correspond to v_x and v_y . The light blue shaded bars are the 95% confidence intervals produced by the models. Results are shown for models trained on (a),(c),(e) the original noiseless datasets as well as (b),(d),(f) the datasets with added $\sigma = 0.1$ Gaussian noise.

Looking more closely at the results for the three-parameter convection-diffusion dataset, we see that the trained model correctly extracts three relevant latent parameters: one latent parameter corresponds to the diffusion constant and the remaining two-dimensional latent subspace corresponds to the drift velocity vector (Table II). In particular, the model learns a rotated representation of the drift velocity vector as a two-dimensional latent subspace due to the inherent rotational symmetry of the dynamics (Appendix F), so we can recover the v_x , v_y components of the drift velocity by performing a multivariate linear fit [Figs. 3(e) and 3(f)]. The successful separation of diffusion from drift velocity in the extracted

parameters demonstrates our model's ability to distinguish distinct and interpretable factors of variation in the dynamics of the system.

B. Prediction performance

In addition to testing parameter extraction, we evaluate the prediction performance of the trained models on their corresponding test sets. Because of training speed and stability considerations, our models are initially trained with a PD architecture containing only 16 hidden channels (Appendix A). To show the potential for further model refinement, we fix the weights of the DEs trained on the original noiseless datasets and then train additional

TABLE I. R^2 correlation coefficients from linear fits of the relevant latent parameters (Fig. 2) with the ground truth physical parameters for each dataset—both with and without added noise. For the three-parameter convection-diffusion dataset, the diffusion constant D is fit with a corresponding extracted latent parameter, while the drift velocity components v_x , v_y are fit with a corresponding two-dimensional subspace of the latent parameters due to the inherent rotational symmetry.

Dataset	Parameter	No noise	$\sigma = 0.1$ noise
Kuramoto-Sivashinsky	γ	0.993	0.995
Nonlinear Schrödinger	κ	0.997	0.998
Convection-diffusion	D	0.963	0.959
Convection-diffusion	v_x	0.997	0.994
Convection-diffusion	v_y	0.998	0.996

TABLE II. R^2 correlation coefficients from individual linear fits of the 2D convection-diffusion dataset parameters with each relevant latent parameter (LP). High correlations are in bold font, emphasizing the interpretability of the learned latent parameters as either corresponding to the diffusion constant D or the drift velocity components v_x , v_y . The drift velocity is matched with two latent parameters that form a two-dimensional latent subspace corresponding to the velocity vector.

Parameter	No noise			$\sigma = 0.1$ noise		
	LP 1	LP 2	LP 5	LP 1	LP 2	LP 3
D	0.963	0.000	0.003	0.003	0.959	0.001
v_x	0.000	0.205	0.766	0.395	0.006	0.554
v_y	0.001	0.818	0.205	0.568	0.000	0.473

PDs, each with an expanded 64 hidden channels. These refined predictive models perform better than the original predictive models used during end-to-end training. For

comparison, the datasets are also evaluated with a stiff exponential integrator for semilinear differential equations [exponential time-differencing fourth-order Runge-Kutta (ETDRK4) [35]] using a second-order finite difference discretization on the same time and space meshes provided in the datasets. Although this integrator is the same one used during dataset generation, the time step and mesh size are set to match the available data to provide a reasonable baseline. During dataset generation, the solution is obtained starting with the exact form of the initial condition and converged using much finer mesh sizes (Appendix D). Also, note that using a nonstiff integrator fails on many of the examples in the datasets, so a stiff integrator is required.

The models trained on the 1D Kuramoto-Sivashinsky and 1D nonlinear Schrödinger datasets both perform reasonably when compared with the traditional finite difference method [Figs. 4(a) and 4(b)], with the model trained on the Kuramoto-Sivashinsky dataset maintaining a higher accuracy than its traditional counterpart. The prediction error of the 2D convection-diffusion model is dominated by the uncertainty in the parameter extraction, so the prediction performance is comparable to a finite difference exponential integrator with similar noise in the PDE parameters [Fig. 4(c)]. For the models trained on datasets with added $\sigma = 0.1$ noise, we see some negative impact on prediction performance (Fig. 4) but no effect on the parameter extraction quality (Tables I and II).

The refined models, trained on the noiseless datasets, demonstrate that the PD—the predictive network—can be improved independently of the DE—the parameter extraction network. Moreover, the solutions generated by these models remain stable and physically reasonable well beyond the number of time steps propagated during training, suggesting that the models have indeed learned




FIG. 4. The root mean square prediction error (RMSE) during evaluation at each propagation time step, averaged over the corresponding test set, for models trained without noise (blue line), trained with $\sigma = 0.1$ Gaussian noise (orange line), and refined by fixing the DE and training with 64 hidden channels in the PD network on the noiseless datasets (green line). Shown for comparison is an evaluation by a second-order finite difference (FD) discretization integrated with the ETDRK4 method [35] (dashed red line), which reduces to an exact exponential integrator for the 2D convection-diffusion system. Also, since the prediction accuracy on 2D convection-diffusion is dominated by parameter extraction uncertainty, we include for comparison the solver with additional parameter noise: $\sigma = 0.005$ for D , $\sigma = 0.01$ for v_x , v_y (dashed purple line). The black vertical dotted line denotes the length of the target series T_y used for training each set of models, i.e., the maximum number of time steps propagated by the model during training.

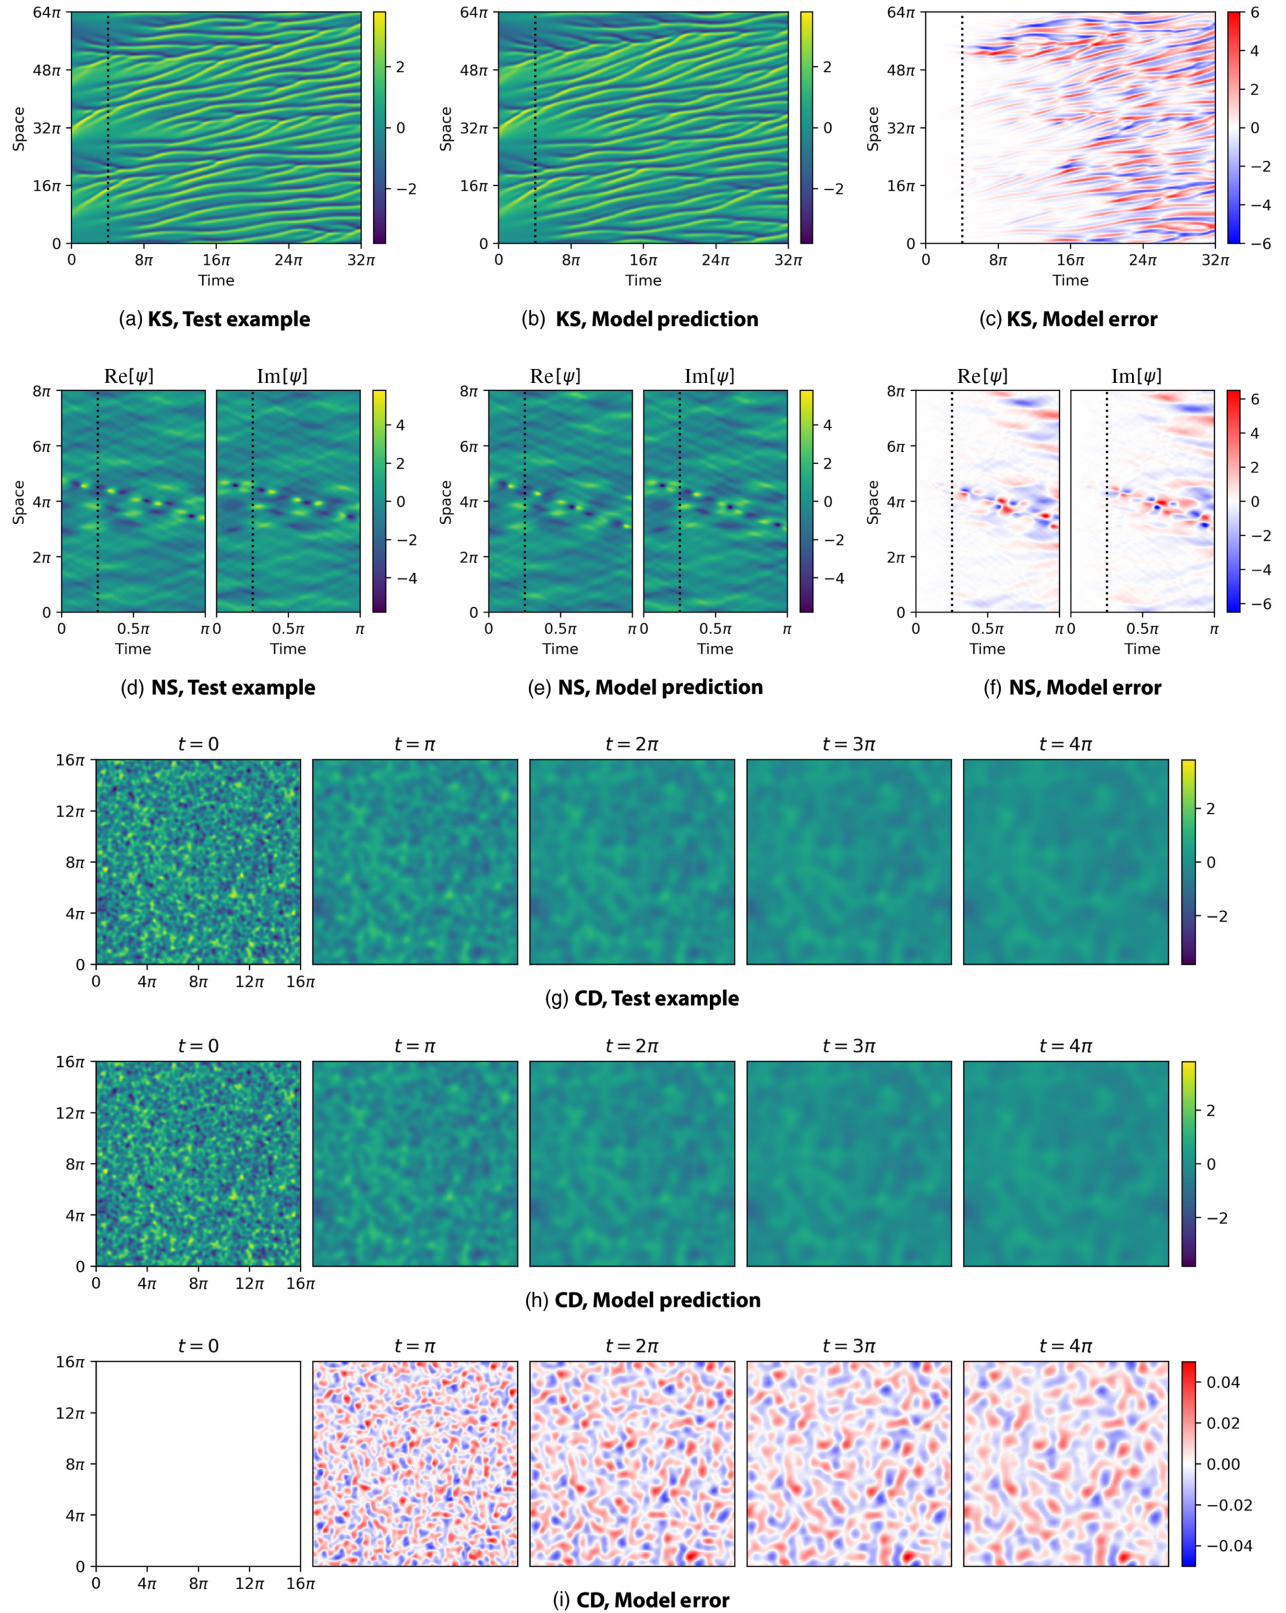


FIG. 5. Time evolution of (a),(d),(g) examples from the 1D Kuramoto-Sivashinsky, 1D nonlinear Schrödinger, and 2D convection-diffusion test sets with periodic boundary conditions, (b),(e),(h) the predicted propagation of the refined model given the initial condition at time $t = 0$, and (c),(f),(i) the model prediction error. The black vertical dotted line denotes the maximum amount of time propagated by the model during training, corresponding to the length of the target series. For the CD dataset, the maximum amount of time propagated by model during training is $t = \pi$.

physically meaningful propagators of the PDE-governed systems (Fig. 5).

V. APPLICATION TO NONLINEAR FIBER PROPAGATION

To demonstrate our method applied on a more complex and realistic dataset, we use MEEP [36], a finite difference time-domain electromagnetic simulator, to model pulse propagation through an optical fiber with a Kerr nonlinearity. These simulations model Maxwell's equations exactly, with no approximation apart from the discretization, and often reproduce real experiments point by point [37]. The simulated fiber consists of a two-layer core and the surrounding cladding with (relative) permittivity $\varepsilon = 1$ [cross section shown in Fig. 6(a)]. We generate a dataset with 200 different geometries by varying the size of the two-layer core through r_i and r_o as well as the corresponding permittivities ε_i and ε_o , representing uncontrolled experimental variables related to fabrication (see Appendix D for details). Then, we excite a randomly

generated pulse in each fiber and train our model on the flux-normalized amplitude $A(x, t)$ of the resulting pulse propagation.

There is no exact first-order PDE describing the evolution of this amplitude. However, in the slowly varying envelope approximation, $A(x, t)$ is governed by an effective nonlinear Schrödinger equation [38],

$$\frac{\partial A}{\partial x} = i \sum_{n=1}^{\infty} \frac{i^n k_n}{n!} \partial_t^n A + i\gamma |A|^2 A, \quad (13)$$

where the dispersion coefficients $k_n = \partial_\omega^n k(\omega)|_{\omega=\omega_0}$ can be computed from the dispersion relation $\omega(k)$ at a carrier frequency $2\pi\omega_0$, and the nonlinearity parameter γ is related to the Kerr nonlinearity (Appendix D) and the shape of the propagating mode in the fiber [39]. Because of the form of this effective equation, we choose our propagation variable to be the distance x rather than the time t ; i.e., our model will predict $A(x, t)$ given an initial pulse $A_0(x=0, t)$ by propagating forward in the x direction. We obtain the ground truth dispersion coefficients using MPB [40], a

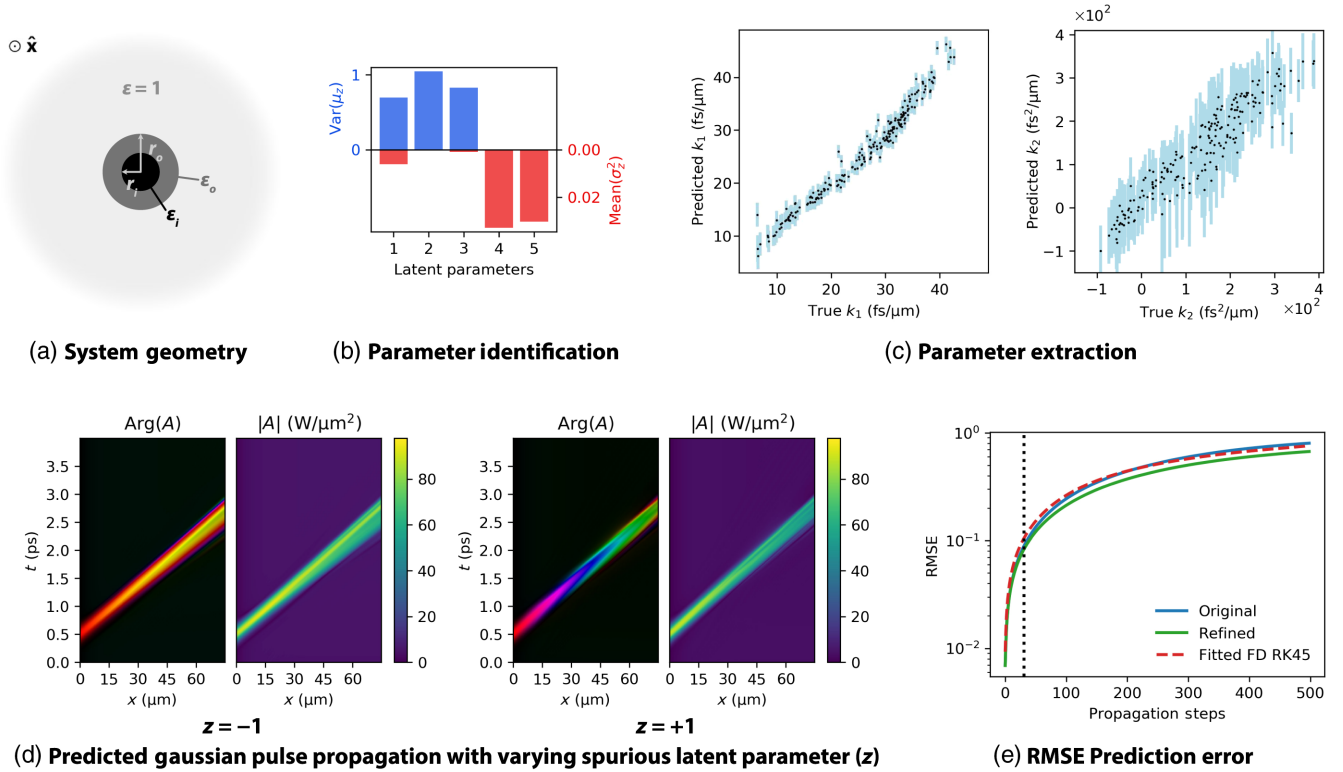


FIG. 6. Analysis of the model trained on the nonlinear fiber propagation dataset. (a) The cross section of the fiber shows the cladding with (relative) permittivity $\varepsilon = 1$, the inner core with radius r_i and permittivity ε_i , and the outer core with radius r_o and permittivity ε_o . (b) Three relevant latent parameters are identified by the trained model. (c) A linear fit of the latent parameters predicts the true k_1 with $R^2 = 0.966$ and k_2 with $R^2 = 0.863$, corresponding to group velocity and second-order dispersion. (d) Varying the remaining spurious latent parameter z while propagating a Gaussian pulse using the trained propagating decoder, we observe that the spurious parameter represents a phase velocity. The plot of $\text{Arg}(A)$ shows the phase of the pulse mapped to the color hue while the amplitude is mapped to lightness. (e) The root mean squared prediction error at each propagation step of the originally trained PD (blue) and a further refined PD (green) are comparable to an explicit solution (second-order finite difference, adaptive fourth-order Runge-Kutta method) of the effective equation (13) with fitted parameters (up to fourth-order dispersion) for each individual test set example (dashed red).

frequency domain electromagnetic eigenmode solver, for comparison with the extracted parameters from the model.

The trained model identified and extracted three relevant latent parameters [Fig. 6(b)]. Two independent directions in the latent parameter space (primarily, latent parameters 2 and 3) correspond to the group velocity $k_1 = 1/v_g$ with $R^2 = 0.966$ and second-order dispersion k_2 with $R^2 = 0.863$ [Fig. 6(c)]. Note that the extracted parameters are not the geometry parameters used to generate the dataset, but rather parameters relevant to the effective propagation of the pulse. These two latent parameters also capture variations in higher-order dispersion terms, which are correlated with the group velocity and second-order dispersion in the dataset. As a result, although important and present in the effective equation (13), the higher-order terms are already well correlated with the existing latent parameters, so no additional parameters are required to characterize the dynamics. See Appendix I for more parameter analysis details.

In addition to the parameters corresponding to k_1 and k_2 , the model extracted another relevant latent direction (primarily, latent parameter 1), which is independent and orthogonal to the previous two. This seemingly spurious latent parameter does not correspond to a term in the effective equation (13); instead, it represents a spurious phase velocity, which is the result of imperfect preprocessing (Appendix D). We discovered this correspondence by varying the spurious parameter, while leaving all other latent parameters fixed, and observing the effect on the model predictions [Fig. 6(d)]. This uncontrolled variable was successfully extracted by the model and subsequently identified, demonstrating the process by which unknown extracted parameters can be understood and interpreted.

We evaluate the prediction performance of the trained propagating decoder as well as a refined version of the PD (with an expanded 64 hidden channels) by comparing the prediction error of our models with an explicit solution to the effective equation (13) as a baseline. The effective equation parameters are determined by a linear fit of the finite difference derivatives computed from the dataset. For each example in the test set, we fit the dispersion parameters k_1 , k_2 , k_3 , k_4 (i.e., up to fourth-order dispersion) and the nonlinearity parameter γ . The effective equation is then integrated using a fourth-order Runge-Kutta method with adaptive step size and second-order finite difference discretization. This approach is a simplified version (with the terms preidentified) of explicit PDE identification methods, such as SINDy [4]. Our predictive models achieve similar prediction performance when compared with this explicit effective equation baseline [Fig. 6(e)].

VI. DISCUSSION

We developed a general unsupervised learning method for extracting unknown dynamical parameters from noisy spatiotemporal data *without* detailed knowledge of the

underlying system or the exact form of the governing PDE. While we do not explicitly extract the governing PDE, our method provides a set of tunable relevant parameters, which characterize the system dynamics in independent and physically interpretable ways, coupled with a highly transferable predictive model. This is often enough to provide significant insight into the physics of the system: for example, by examining the effect of varying an unidentified relevant parameter on the predictions of the model, we can disentangle the parameter from other effects (parametrized by the other relevant parameters) and interpret it independently. This is precisely how we identified the spurious phase velocity parameter extracted by the model trained on nonlinear fiber propagation [Fig. 6(d)]. One potential complication of interpreting the parameters extracted using our method is that each parameter must represent an independent factor of variation in the dynamics of the dataset. This means that if features of the dynamics are highly correlated in the underlying dataset, they will be parametrized by the same parameter; e.g., the parameters extracted from the nonlinear fiber propagation dataset corresponding to group velocity and second-order dispersion also capture higher-order dispersion terms (Appendix I).

The flexibility and robustness of our model comes from using a generic physics-informed neural network model for nonlinear PDEs. The interpretability of the resulting extracted parameters is a result of the variational autoencoder training and regularization (Appendix E) as well as the inductive biases imposed by our physics-informed network design. By using appropriate spatial averaging in the dynamics encoder and dynamic convolutions in the propagating decoder, we ensure that both the parameter extraction and the propagation prediction from our model are physically motivated and generalizable to arbitrary system sizes and geometries. The dynamic convolutions, in particular, are an important physical inductive bias for encouraging the model to learn latent parameters which govern the propagation dynamics. As a result, the learned parameter-to-kernel mappings in the trained predictive model are fully transferable, which we can demonstrate by evaluating the predictive model *without retraining* on a different set boundary conditions (see Appendix H).

Our strategy for modeling spatiotemporal systems is to retain the expressiveness of a neural network model while imposing general constraints—such as locality through using convolutional layers—to help the network learn more efficiently. For particular applications, this could also include spatial symmetries [41–43], e.g., properly transforming fluid flow vectors, as well as additional symmetries of the internal degrees of freedom, e.g., the global phase of the nonlinear Schrödinger equation. These architecture-based constraints encourage the model to learn physically relevant representations and can be tailored to individual applications, allowing us to incorporate domain

knowledge into the model. This also lets us use datasets that are much smaller than is traditionally required by deep learning methods. The model trained on nonlinear fiber propagation used only 200 examples (Sec. V), and we have been able to successfully train models on the Kuramoto-Sivashinsky dataset with as few as 10 examples (Appendix G). Additionally, the model's robustness to noise is a powerful feature of deep learning methods and provides a promising avenue for studying dynamical systems in a data-driven fashion [44].

The primary challenges associated with applying our current implementation involve setting hyperparameters to improve training stability (see Appendix C for more details) and choosing the β regularization hyperparameter, which controls parameter extraction (Appendix E). The choice of β can be somewhat ambiguous, with very high values resulting in no relevant parameters and very low values failing to enforce independence (for our choices, see Appendix B). This trade-off and other related issues are a very active area of research, and parameter extraction methods will continue to improve following the rapid advances in unsupervised learning and disentangling representations, e.g., through a deeper theoretical understanding of the β -VAE [21,45,46] and alternative formulations [22,47]. Physics-informed inductive biases, however, will remain the key ingredient for ensuring the representations are interpretable [48].

The predictive model will also likely achieve better accuracies by using more sophisticated architectures, such as echo state networks which have been shown to perform extremely well on even chaotic PDEs [49], or by explicitly incorporating differentiable PDE solvers with gradients computed by the adjoint method [50]. An echo state network or other alternative decoder architecture would have to be adapted to retain the transferability of our current PD. Using a differentiable PDE solver would allow the PD network to focus on encoding the PDE rather than also learning a stable integration method, and thus may improve the interpretability of the model.

Our unsupervised learning method is also highly complementary with the significant body of work applying machine learning methods for more accurate predictions of specific physical systems, such as multiscale hydrodynamic systems [51] and turbulence modeling [52–55]. These methods often combine a known physics model with a machine learned correction or parametrize an unknown part of the physics model using neural networks. Such machine learning–physics hybrid models can be adapted into decoders for our unsupervised learning method, with the extracted relevant parameters representing independent variations of and providing insight into the machine learned portions of the model. Our current method can also be adapted in the future for a more general class of spatio-temporal systems by incorporating spatially inhomogeneous latent parameters and will also be able to use

data with incomplete physical state observations by inferring missing information.

The ultimate goal of this work is to provide additional insight into complex spatiotemporal dynamics using a data-driven approach. Our method is an example of a new machine learning tool for studying the physics of spatio-temporal systems with an emphasis on interpretability.

ACKNOWLEDGMENTS

We would like to acknowledge useful discussions with Jason Fleischer, Rumen Dangovski, and Li Jing. This research is sponsored in part by the Army Research Office and was accomplished under Cooperative Agreement No. W911NF-18-2-0048. This work is further supported in part by the U.S. Department of Defense through the National Defense Science and Engineering Graduate Fellowship (NDSEG) Program and by the MIT–SenseTime Alliance on Artificial Intelligence. This material is also based in part upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00111890042.

APPENDIX A: MODEL IMPLEMENTATION DETAILS

For the 1D datasets, the dynamics encoder (DE) uses 2D convolutions with output channel sizes (4,16,64,64,5), linear kernel sizes (3,3,3,3,1), and dilation factors (1,2,4,8,1). For the 2D dataset, the DE uses 3D convolutions with output channel sizes (8,64,64,64,5) with the same kernel sizes and dilation factors. The $f_{\text{DE},\mu}$ and f_{DE,σ^2} networks share the same convolution weights for the first four layers and have distinct final layers to produce μ and $\log \sigma^2$. The final output channel size is determined by the number of latent parameters; in our tests, we use five latent parameters.

For the 1D datasets, the propagating decoder (PD) architecture uses three 1D dynamic convolutional layers with output channel sizes (16, 16, data channel size), linear kernel size 5, and periodic padding. For the 2D datasets, the PD uses three 2D dynamic convolutional layers with the same output channel sizes, kernels, and padding. The refined models increase the number of hidden channels in the PD from 16 to 64, resulting in output channel sizes (64, 64, data channel size).

The latent-to-kernel network, which maps the latent parameters to each kernel or bias in the PD, consists of two fully connected layers, i.e., one hidden layer. For the 1D datasets, the hidden layers have size ($4 \times \text{input channel size} \times \text{output channel size}$) for kernels and ($4 \times \text{output channel size}$) for biases, where the input and output channel sizes refer to the corresponding dynamic convolution in the PD. For the 2D dataset, the hidden layers have size ($16 \times \text{input channel size} \times \text{output channel size}$) for kernels and ($4 \times \text{output channel size}$) for biases.

Our architecture uses rectified linear unit activations throughout except for the unactivated output layers of the DE, PD, and latent-to-kernel networks. The output of the PD convolutional network f_{PD} uses a tanh activation with a learnable scaling parameter $[x \mapsto \lambda \tanh(x/\lambda)]$ with learnable parameter λ initialized to 1] to stabilize the recurrent architecture. The network f_{PD} also has a fixed multiplicative prefactor set to 10^{-6} to improve the initial training stability. For the nonlinear fiber propagation dataset, we add Gaussian noise with $\sigma = 10^{-2}$ between propagation steps in the PD network for improved prediction stability.

Our model is implemented using PyTorch v1.1 [56], and the source code is available in Ref. [28].

APPENDIX B: TRAINING DETAILS

All models are trained using batch size 50 and the Adam optimizer [57] with learning rate 10^{-3} . Models for the 1D datasets and the noisy 2D convection-diffusion dataset were trained for 2000 epochs; the model for the noiseless 2D convection-diffusion dataset was trained for 4000 epochs; and the corresponding refined models were trained for 2000 epochs. The VAE regularization hyperparameter is set to $\beta = 0.02$ for the 1D datasets and $\beta = 10^{-4}$ for the 2D convection-diffusion dataset. The model for the nonlinear fiber propagation application was trained for 40 000 epochs due to small size of the dataset and significant data augmentation; the corresponding refined model was trained for 20 000 epochs; and the VAE regularization hyperparameter was set to $\beta = 7 \times 10^{-4}$. During validation, we choose β such that we obtain the maximum number of relevant latent parameters while still maintaining statistical independence among the parameters as well as a clean separation between the relevant and irrelevant (i.e., collapsed to the prior) parameters. All hyperparameter tuning is done using the training set for validation.

For the 1D Kuramoto-Sivashinsky dataset, we train the model using a random 64×94 crop—in the time and space dimensions, respectively—for the input series and another random 64×76 crop for the target series. For the 1D nonlinear Schrödinger dataset, we train using a 64×94 crop for the input series and a 32×76 crop for the target series. For the 2D convection-diffusion dataset, we train using a $45 \times 62 \times 62$ crop—in the one time and two space dimensions, respectively—for the input series and a $16 \times 44 \times 44$ crop for the target series. For the nonlinear fiber propagation dataset, we train using a 128×158 crop for the input series and a 32×76 crop for the target series. During evaluation, we always use the full size time series from the test set for both the input and target series.

APPENDIX C: NOTES ON TRAINING STABILITY

We find our model to be particularly sensitive to the architecture of the propagating decoder: with larger, more complex networks and more propagation steps during

training resulting in increasing instability. The dynamics encoder does influence stability, but the effect is more indirect through its interaction with the PD and is not very architecture sensitive. This instability is likely related to the problem of vanishing and exploding gradients seen often in recurrent architectures, which is mitigated using gating mechanisms like in long short-term memory networks [10,11] or by explicitly using unitary norm-preserving matrices [58]. Importantly, this is only a significant problem when training the model end to end using both the DE and PD; when we fix the DE weights, we are able to further refine a more complex PD model without instability (Sec. IV B). This also does *not* affect the prediction performance and stability of the model during evaluation, which generalizes well past the number of time steps propagated during training (Fig. 4). We currently implement a learnable gating mechanism (Appendix A) that significantly stabilizes the network, but further work is required to fully address this issue.

APPENDIX D: DATASET GENERATION DETAILS

For each time-series example in the 1D Kuramoto-Sivashinsky dataset, we sample the viscosity damping parameter γ from a truncated normal distribution ($\mu = 1$, $\sigma = 0.125$, and truncation interval $[0.5, 1.5]$). We then use the ETDRK4 integrator [35] to generate each time series to within a local relative error of 10^{-3} . Each time series consists of a uniform time mesh with 256 points for a total time $T = 32\pi$ and a space mesh with $M = 256$ points for an $L = 64\pi$ unit cell. These are produced by solving on a finer time and space mesh to ensure convergence and then resampling to the dataset mesh sizes. Each initial state is generated from independently sampled, normally distributed Fourier components with a Gaussian band-limiting envelope of varying widths (uniformly sampled in the interval $[8\pi/L, M\pi/4L]$) and then normalized to unit variance.

For each time-series example in the 1D nonlinear Schrödinger dataset, we sample the nonlinearity coefficient κ from a truncated normal distribution ($\mu = -1$, $\sigma = 0.25$, and truncation interval $[-2, 0]$). We then use the ETDRK4 integrator [35] to generate each time series to within a local relative error of 10^{-3} . Each time series consists of a uniform time mesh with 256 points for a total time $T = \pi$ and a space mesh with $M = 256$ points for an $L = 8\pi$ unit cell. These are produced by solving on a finer time and space mesh to ensure convergence and then resampling to the dataset mesh sizes. The initial states are generated in an analogous manner to the Kuramoto-Sivashinsky dataset.

For each time-series example in the 2D convection-diffusion dataset, we vary the parameters by sampling the diffusion constant D from a truncated normal distribution ($\mu = 0.1$, $\sigma = 0.025$, and truncation interval $[0, 0.2]$), and each velocity component v_x, v_y from a normal distribution

($\mu = 0$, $\sigma = 0.2$). Because the convection-diffusion equation is linear, we use the exact solution to generate the dataset. Each time series consists of a uniform time mesh with 64 points for a total time $T = 4\pi$ and an $M \times M = 256 \times 256$ space mesh for an $L \times L = 16\pi \times 16\pi$ unit cell. Each initial state is generated from independently sampled, normally distributed Fourier components with a Gaussian band-limiting envelope of varying widths (uniformly sampled in the interval $[16\pi/L, M\pi/2L]$) and then normalized to unit variance.

The nonlinear fiber propagation dataset aims to roughly model a set of highly dispersive nonlinear optical fibers with variations in geometry due to fabrication. Each of the 200 examples in the nonlinear fiber propagation dataset corresponds to a randomly generated fiber geometry excited with a randomly generated pulse. Specifically, each fiber geometry [Fig. 6(a)] corresponds to a set of geometry parameters sampled from normal distributions: inner core radius r_i ($\mu = 75$ nm, $\sigma = 3$ nm), outer core radius r_o ($\mu = 150$ nm, $\sigma = 7.5$ nm), inner core (relative) permittivity ε_i ($\mu = 30$, $\sigma = 2$), and outer core permittivity ε_o ($\mu = 8$, $\sigma = 1$). There is an overall fixed Kerr nonlinearity that corresponds to a nonlinear refractive index $n_2 = (3.375 \times 10^{-4} \text{ } \mu\text{m}^2 \text{ W}^{-1})/\varepsilon$, where ε is the material permittivity. The excited pulse is generated from independently sampled, normally distributed frequency components with a Gaussian band-limiting envelope centered on a carrier frequency $f = 200$ THz ($\lambda = 1.5 \text{ } \mu\text{m}$) and with a width of $f/20$. This pulse is slowly turned on with a sigmoid of width $20/f$, and, for the test set, the pulse is also turned off at half the total simulation time, allowing the pulse to fully propagate through the fiber. During the simulation, the amplitude $A(x, t)$ of the electric field at the center of fiber is recorded and later normalized to the total flux passing through the fiber. The space and time Fourier components of each resulting dataset example $A(x, t)$ are shifted to remove the carrier frequency and the peak wave number, resulting in a slowly varying envelope. Then, the final dataset is normalized again so that the amplitude at the initial point $x = 0$ has, on average, unit variance over the whole dataset. Each example consists of an x -direction mesh with 500 points for a propagation length of $75 \text{ } \mu\text{m}$ and a uniform t -direction mesh with 800 points for a total time 4.00 ps (training set) or 1000 points for a total time 5.00 ps (test set).

The dataset generation scripts are available in Ref. [28].

APPENDIX E: UNDERSTANDING THE EFFECTS OF VAE REGULARIZATION

Using VAE [13] or β -VAE [20] regularization in our model provides three main benefits for learning physically interpretable representations: the regularization encourages the model to minimize use of the latent parameters, enforces independence among the learned latent parameters, and matches the marginal latent distribution to a

standard normal prior. We can explicitly see these effects by decomposing the data-averaged VAE regularization term in the following way [22,59]:

$$\mathbb{E}_{p_D(\mathbf{x})}\{D_{\text{KL}}[q(\mathbf{z}|\mathbf{x})\|p(\mathbf{z})]\} = D_{\text{KL}}[q(\mathbf{z}, \mathbf{x})\|q(\mathbf{z})p_D(\mathbf{x})] \quad (\text{E1})$$

$$+ D_{\text{KL}}\left(q(\mathbf{z})\|\prod_i q(z_i)\right) \quad (\text{E2})$$

$$+ \sum_i D_{\text{KL}}[q(z_i)\|p(z_i)], \quad (\text{E3})$$

where $p_D(\mathbf{x})$ is the data distribution, $p(\mathbf{z}) = \prod_i p(z_i)$ are the standard normal priors for the latent parameters $\mathbf{z} = (z_1, z_2, \dots, z_i, \dots)$, $q(\mathbf{z}|\mathbf{x})$ is the output distribution of the dynamics encoder, $q(\mathbf{z}, \mathbf{x}) = q(\mathbf{z}|\mathbf{x})p_D(\mathbf{x})$ is the joint distribution of the encoded latent parameters and the data, and $q(\mathbf{z}) = \int d\mathbf{x} q(\mathbf{z}, \mathbf{x})$ and $q(z_i) = \int d\mathbf{x} \prod_{j \neq i} dz_j q(\mathbf{z}, \mathbf{x})$ are the marginal distributions of the latent parameters \mathbf{z} or a single latent parameter z_i , respectively. The three terms in this decomposition correspond directly to the three effects: the first term (E1) represents the mutual information between the latent parameters and the data; the second term (E2) represents the total correlation between the latent parameters; and the third term (E3) consists of KL divergences between the marginal distribution for individual latent parameters and the standard normal prior.

By minimizing the mutual information between the latent space and the data (E1) as well as correlations among the latent parameters (E2), the model is compelled to learn a latent space with minimal information and independent parameters; i.e., the model will use a minimal set of independent relevant latent parameters to capture only the necessary information for better prediction performance. The rest of the unused latent parameters will collapse to the prior. Furthermore, by matching the *marginal* latent parameter distributions $q(z_i)$ to the standard normal priors $p(z_i)$ (E3), the VAE regularizer encourages a linear relationship between the relevant learned latent parameters and the true physical parameters *if* the physical parameters are normally distributed in the data. Even if a physical parameter z_{phys} is non-normally distributed, the VAE regularization will still compel the model to learn a monotonic relationship between z_{phys} and a corresponding latent parameter z given by

$$z_{\text{phys}} = \pm \text{CDF}_{p(z_{\text{phys}})}^{-1} \circ \text{CDF}_{p(z)}^{-1}(z), \quad (\text{E4})$$

where $\text{CDF}_{p(\cdot)}$ is the cumulative distribution function for the probability distribution $p(\cdot)$. One caveat—in addition to ambiguities introduced by symmetries of the physical parameters—is that the relationship may not be monotonic

for physical parameter distributions which have support on a topologically distinct space from the real line, e.g., a uniformly distributed periodic angle parameter. However, the result may still be interpretable; e.g., an angle parameter may be encoded as a ring in a two-dimensional latent subspace.

Although this decomposition is suggestive of the effects of VAE regularization, the study of the performance of VAE-based models and the relative importance and model dependence of each of these effects is still very much ongoing [21,22,45–47,60]. While training our model, we empirically observe that the latent parameters retain their independence and that their marginal distributions match the standard normal priors, so only an increase in information stored in the latent space is traded for better prediction performance. We believe we can attribute this to the physics-informed inductive biases present in our architecture, which allows our model to achieve its best performance using a minimal set of independent and normally distributed latent parameters.

APPENDIX F: RAW PARAMETER EXTRACTION RESULTS

We can explicitly see the relevant and collapsed latent parameters in the raw data by plotting the latent parameters versus the true physical parameters (Fig. 7). The latent parameters that show a correlation with the true physical parameters also have small variances σ_z^2 and correspond to

the identified relevant latent parameters (Fig. 2), while the remaining latent parameters have collapsed to the prior. Note that for the collapsed parameters, we see variances σ_z^2 that are less than one—the value expected for parameters which have collapsed to the prior distribution $\mathcal{N}(0, 1)$. This is because, to average over systems of different sizes, the model makes the assumption that patches separated far in space or time provide independent estimates of the extracted parameters and computes the total variance accordingly. This assumption is reasonable for relevant parameters but will artificially lower the extracted variances for collapsed parameters. During testing, we choose to evaluate on the full system size resulting in this artifact. If we were to evaluate on smaller patches that match the size of the crops used during training, we would indeed see that the collapsed parameters have $\sigma_z^2 = 1$.

We also note that, for the model trained on the 2D convection-diffusion dataset, the latent parameters associated with the drift velocity \mathbf{v} are not aligned with the v_x, v_y velocity components. This is an expected result due to the inherent ambiguity of choosing a coordinate basis—introduced by the rotational symmetry of the velocity vector—and makes judging the extraction performance more difficult. Instead of examining one latent parameter at a time, we must consider the two-dimensional latent subspace associated with the velocity vector. Taking the two relevant latent parameters that are correlated with the drift velocity (Table II), we can perform a multivariate

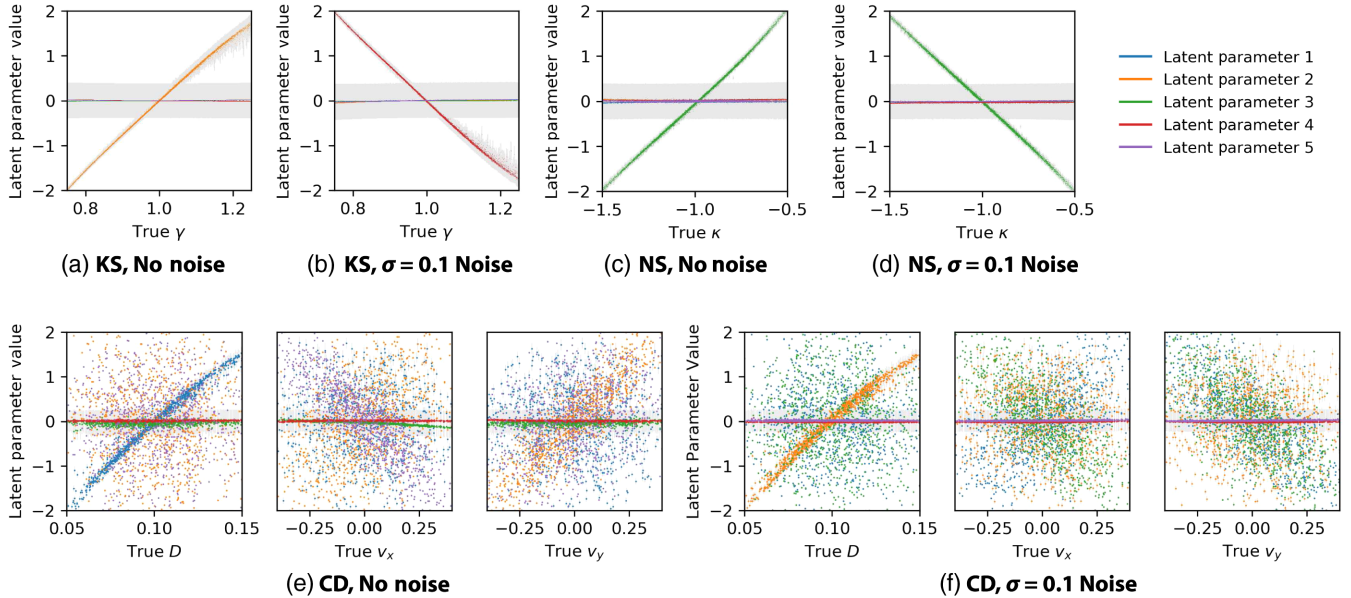


FIG. 7. The five latent parameters in the models trained on the (a),(b) 1D Kuramoto–Sivashinsky, (c),(d) 1D nonlinear Schrödinger, and (e),(f) 2D convection–diffusion datasets both with and without added noise versus the ground truth physical parameters used to generate the datasets. In these plots of the raw extracted latent parameters, we see the direct correspondence between the identified relevant latent parameters in Fig. 2 and the true physical parameters as well as the collapse of the unused latent parameters. Note that the relevant latent parameters corresponding to the drift velocity \mathbf{v} in the convection–diffusion model are not precisely aligned with the two components v_x, v_y due to the inherent rotational symmetry. The gray shaded bars are the 95% confidence intervals ($\pm 1.96\sigma_z$) produced by the model.

linear regression of the velocity components v_x , v_y in this two-dimensional latent subspace to verify that the model has indeed learned a simple rotated representation of the velocity vector [Figs. 3(e) and 3(f)].

APPENDIX G: PERFORMANCE SCALING WITH DATASET SIZE

Because of the significant physics-informed inductive biases in our architecture, our model still achieves usable results even when trained on very small datasets. We test the dataset size dependence of our method using the 1D Kuramoto-Sivashinsky system and find that the model is still able to identify the relevant latent parameter even with a dataset of just 10 examples (Fig. 8). The accuracy and precision of the extracted parameter and the prediction performance do begin to suffer when using such extremely small datasets, but the model is still able to provide some insight into the dynamics of the spatiotemporal system represented by the data.

APPENDIX H: ALTERNATIVE BOUNDARY CONDITIONS

The fully convolutional structure of the propagating decoder means that we are able to evaluate our model on arbitrary geometries and boundary conditions. By training on small crops and evaluating on the full size

examples in the test set (Sec. IV), we have already shown the trained model can be directly evaluated on larger system sizes. To show direct evaluation on an alternative boundary condition, we test the refined predictive model—originally trained on the 1D Kuramoto-Sivashinsky dataset with periodic boundaries—on a new test example generated with Dirichlet hard wall boundary conditions (Fig. 9). In general, we can apply alternative boundary conditions by adjusting the padding scheme of each propagation step in the PD. For Dirichlet boundaries, this corresponds to applying antireflection padding at each propagation step. This preliminary test suggests that we can achieve similar prediction performance using an alternative boundary condition, which the model has never previously seen, and demonstrates the transferability of the learned convolutional kernels.

APPENDIX I: NONLINEAR FIBER PARAMETER ANALYSIS

In the three-dimensional relevant latent space (z_1, z_2, z_3) extracted by the model trained on the nonlinear fiber propagation dataset [Fig. 6(b)], we determine two independent and interpretable directions by a linear fit: $(-0.101, 0.971, 0.218)$ and $(0.477, -0.0303, -0.878)$, which correspond to the group velocity $k_1 = 1/v_g$ and second-order dispersion k_2 , respectively [Fig. 6(c)]. The final direction $(0.882, -0.107, 0.459)$ in the latent

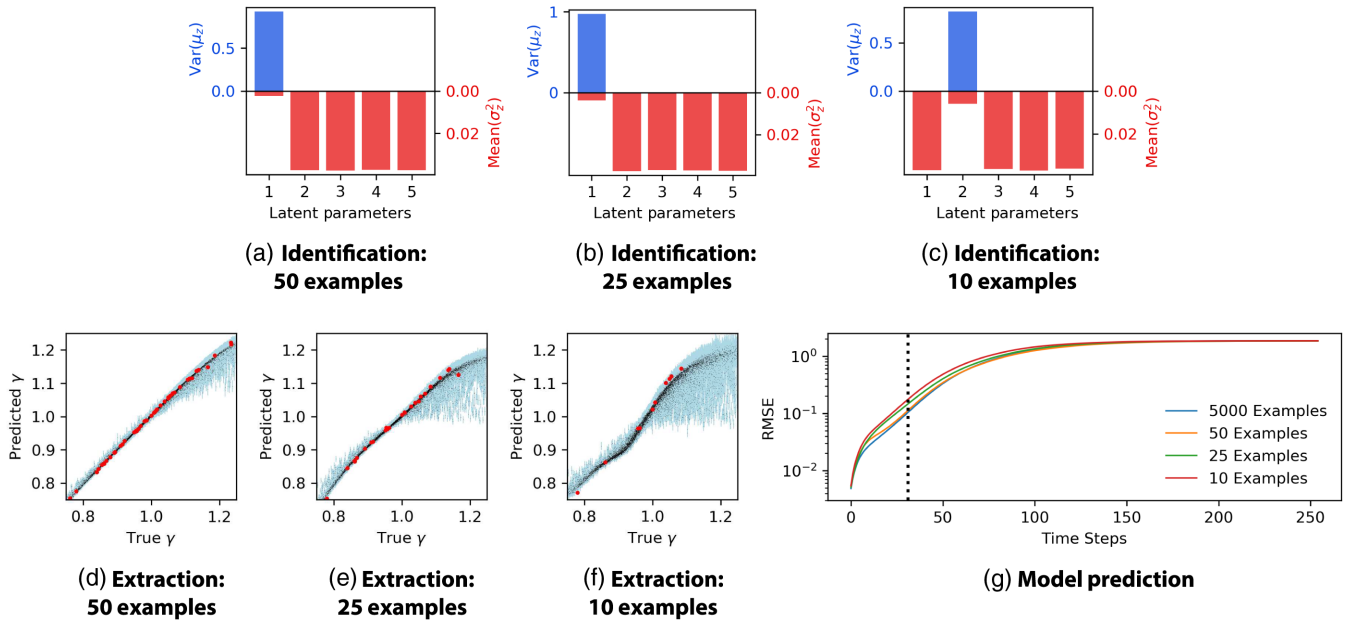


FIG. 8. Parameter extraction and prediction performance for models trained on noiseless 1D Kuramoto-Sivashinsky datasets with 50, 25, and 10 examples to show the effect of dataset size. The (a)–(c) parameter identification and (d)–(f) extraction plots demonstrate the ability of the model to identify and extract a relevant latent parameter using very few examples. In the extraction plots, the small black points show the evaluation on the test set, the large red points show the training examples used for each model, and the light blue shaded bars are the 95% confidence intervals produced by the model. There is also a subtle decrease in prediction performance seen in (g) the RMSE averaged over the 10 000 example test set. The parameter identification and extraction plots for the 5000 example dataset are shown in Figs. 2(a) and 3(a), respectively.




FIG. 9. Time evolution of (a) a 1D Kuramoto-Sivashinsky example with Dirichlet boundary conditions, (b) the predicted propagation using transferred kernels from the refined predictive model given the initial condition at time $t = 0$, and (c) the prediction error of the model. The refined predictive model for the 1D Kuramoto-Sivashinsky dataset was originally trained on data with periodic boundaries and is adapted—without retraining—to use Dirichlet hard wall boundaries instead by adjusting the padding at each propagation time step. The black vertical dotted line denotes the maximum amount of time propagated by the model during the original training, corresponding to the length of the target series.

space—orthogonal to the previous two—is a seemingly spurious relevant parameter unrelated to the parameters of the effective equation (13) and represents a spurious phase velocity [Fig. 6(d)].

For the examples in the nonlinear fiber propagation dataset, higher-order dispersion terms k_n are still significant. However, because these terms are correlated with k_1 and k_2 in the data, the model does not require additional latent parameters to capture their effect. Instead, the existing latent parameters also adjust the higher-order dispersion terms; in other words, the latent parameters each correspond to a dispersion operator that includes higher-order dispersion along with k_1 and k_2 .

-
- [1] A. Karpatne, G. Atluri, J. H. Faghmous, M. Steinbach, A. Banerjee, A. Ganguly, S. Shekhar, N. Samatova, and V. Kumar, *Theory-Guided Data Science: A New Paradigm for Scientific Discovery from Data*, *IEEE Trans. Knowledge Data Eng.* **29**, 2318 (2017).
 - [2] A. Karpatne, W. Watkins, J. Read, and V. Kumar, *Physics-Guided Neural Networks (PGNN): An Application in Lake Temperature Modeling*, *arXiv:1710.11431*.
 - [3] C. Wang, H. Zhai, and Y.-Z. You, *Emergent Schrödinger Equation in an Introspective Machine Learning Architecture*, *Sci. Bull.* **64**, 1228 (2019).
 - [4] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, *Data-Driven Discovery of Partial Differential Equations*, *Sci. Adv.* **3**, e1602614 (2017).
 - [5] Z. Long, Y. Lu, X. Ma, and B. Dong, *PDE-Net: Learning PDEs from Data*, in *Proceedings of the 35th International Conference on Machine Learning, Stockholm, 2018*, edited by J. Dy and A. Krause, Vol. 80, pp. 3208–3216.
 - [6] Z. Long, Y. Lu, and B. Dong, *PDE-Net 2.0: Learning PDEs from Data with a Numeric-Symbolic Hybrid Deep Network*, *J. Comput. Phys.* **399**, 108925 (2019).
 - [7] M. Raissi, P. Perdikaris, and G. Karniadakis, *Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations*, *J. Comput. Phys.* **378**, 686 (2019).
 - [8] J. Berg and K. Nyström, *Data-Driven Discovery of PDEs in Complex Datasets*, *J. Comput. Phys.* **384**, 239 (2019).
 - [9] M. Raissi, *Deep Hidden Physics Models: Deep Learning of Nonlinear Partial Differential Equations*, *J. Mach. Learn. Res.* **19**, 1 (2018), <http://www.jmlr.org/papers/v19/18-046.html>.
 - [10] S. Hochreiter and J. Schmidhuber, *Long Short-Term Memory*, *Neural Comput.* **9**, 1735 (1997).
 - [11] F. A. Gers, J. Schmidhuber, and F. Cummins, *Learning to Forget: Continual Prediction with LSTM*, *Neural Comput.* **12**, 2451 (2000).
 - [12] S. Rudy, A. Alla, S. Brunton, and J. Kutz, *Data-Driven Identification of Parametric Partial Differential Equations*, *SIAM J. Appl. Dyn. Syst.* **18**, 643 (2019).
 - [13] D. P. Kingma and M. Welling, *Auto-Encoding Variational Bayes*, *arXiv:1312.6114*.
 - [14] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016), <http://www.deeplearningbook.org>.
 - [15] R. Iten, T. Metger, H. Wilming, L. del Rio, and R. Renner, *Discovering Physical Concepts with Neural Networks*, *Phys. Rev. Lett.* **124**, 010508 (2020).
 - [16] C. X. Hernández, H. K. Wayment-Steele, M. M. Sultan, B. E. Husic, and V. S. Pande, *Variational Encoding of Complex Dynamics*, *Phys. Rev. E* **97**, 062412 (2018).
 - [17] C. Wehmeyer and F. Noé, *Time-Lagged Autoencoders: Deep Learning of Slow Collective Variables for Molecular Kinetics*, *J. Chem. Phys.* **148**, 241703 (2018).
 - [18] S. J. Wetzel, *Unsupervised Learning of Phase Transitions: From Principal Component Analysis to Variational Autoencoders*, *Phys. Rev. E* **96**, 022140 (2017).
 - [19] M. Cristoforetti, G. Jurman, A. I. Nardelli, and C. Furlanello, *Towards Meaningful Physics from Generative Models*, *arXiv:1705.09524*.
 - [20] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, *beta-VAE: Learning Basic Visual Concepts with a Constrained*

- Variational Framework, in *Proceedings of the International Conference on Learning Representations, 2017, Toulon, France*, <https://openreview.net/forum?id=Sy2fzU9gl>.
- [21] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, *Understanding Disentangling in β -VAE*, [arXiv:1804.03599](https://arxiv.org/abs/1804.03599).
- [22] T. Q. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud, *Isolating Sources of Disentanglement in Variational Autoencoders*, in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2018), pp. 2610–2620.
- [23] D. Zheng, V. Luo, J. Wu, and J. B. Tenenbaum, *Unsupervised Learning of Latent Physical Properties Using Perception-Prediction Networks*, in *Proceedings of the Conference on Uncertainty in Artificial Intelligence, 2018, Monterey, California, USA* (AUAI Press, Corvallis, Oregon, 2018).
- [24] K. Sohn, H. Lee, and X. Yan, *Learning Structured Output Representation Using Deep Conditional Generative Models*, in *Advances in Neural Information Processing Systems 28*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2015) pp. 3483–3491.
- [25] T. Xue, J. Wu, K. Bouman, and W. Freeman, *Visual Dynamics: Stochastic Future Generation via Layered Cross Convolutional Networks*, *IEEE Trans. Pattern Anal. Mach. Intell.* **41**, 2236 (2019).
- [26] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *ImageNet Classification with Deep Convolutional Neural Networks*, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., Red Hook, NY, 2012), pp. 1097–1105.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition*, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, Las Vegas, NV, USA* (IEEE, New York City, NY, 2016).
- [28] See <https://github.com/peterparity/PDE-VAE-pytorch>.
- [29] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, *Show, Attend and Tell: Neural Image Caption Generation with Visual Attention*, in *Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 2015*, edited by F. Bach and D. Blei, Vol. 37, pp. 2048–2057.
- [30] X. Jia, B. De Brabandere, T. Tuytelaars, and L. V. Gool, *Dynamic Filter Networks*, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2016), pp. 667–675.
- [31] P. Manneville, *Liapounov Exponents for the Kuramoto-Sivashinsky Model*, in *Macroscopic Modelling of Turbulent Flows* (Springer, New York, 1985), pp. 319–326.
- [32] Y. Kuramoto, *Diffusion-Induced Chaos in Reaction Systems*, *Prog. Theor. Phys. Suppl.* **64**, 346 (1978).
- [33] G. Sivashinsky, *On Flame Propagation under Conditions of Stoichiometry*, *SIAM J. Appl. Math.* **39**, 67 (1980).
- [34] M. J. Ablowitz, *Nonlinear Dispersive Waves: Asymptotic Analysis and Solitons*, Cambridge Texts in Applied Mathematics (Cambridge University Press, Cambridge, England, 2011).
- [35] A.-K. Kassam and L. N. Trefethen, *Fourth-Order Time-Stepping for Stiff PDEs*, *SIAM J. Sci. Comput.* **26**, 1214 (2005).
- [36] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. Joannopoulos, and S. G. Johnson, *MEEP: A Flexible Free-Software Package for Electromagnetic Simulations by the FDTD Method*, *Comput. Phys. Commun.* **181**, 687 (2010).
- [37] A. Taflov and S. C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method* (Artech House, London, United Kingdom and Norwood, MA, USA, 2005).
- [38] This effective equation is usually transformed to a comoving time coordinate $\tau = t - x/v_g$, where v_g is the group velocity [39]. However, since the group velocity varies among of the dataset examples and because we want to use minimal preprocessing, we keep the original time coordinate t . We also include higher-order dispersion terms which are relevant for our dataset.
- [39] R. W. Boyd, *Nonlinear Optics*, 3rd ed. (Academic Press, Burlington, VT, 2008).
- [40] S. G. Johnson and J. D. Joannopoulos, *Block-Iterative Frequency-Domain Methods for Maxwell's Equations in a Planewave Basis*, *Opt. Express* **8**, 173 (2001).
- [41] T. Cohen and M. Welling, *Group Equivariant Convolutional Networks*, in *Proceedings of the 33rd International Conference on Machine Learning, New York, 2016*, edited by M. F. Balcan and K. Q. Weinberger, Vol. 48, pp. 2990–2999.
- [42] T. S. Cohen and M. Welling, *Steerable CNNs*, in *Proceedings of the International Conference on Learning Representations, 2017, Toulon, France*, <https://openreview.net/forum?id=rJQKYt5Il>.
- [43] M. Weiler, M. Geiger, M. Welling, W. Boomsma, and T. Cohen, *3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data*, in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2018), pp. 10381–10392.
- [44] S. H. Rudy, J. N. Kutz, and S. L. Brunton, *Deep Learning of Dynamics and Signal-Noise Decomposition with Time-Stepping Constraints*, *J. Comput. Phys.* **396**, 483 (2019).
- [45] A. Alemi, B. Poole, I. Fischer, J. Dillon, R. A. Saurous, and K. Murphy, *Fixing a Broken ELBO*, in *Proceedings of the 35th International Conference on Machine Learning, Stockholm, 2018*, pp. 159–168.
- [46] S. Zhao, J. Song, and S. Ermon, *Towards Deeper Understanding of Variational Autoencoding Models*, [arXiv:1702.08658](https://arxiv.org/abs/1702.08658).
- [47] S. Zhao, J. Song, and S. Ermon, *InfoVAE: Balancing Learning and Inference in Variational Autoencoders*, in *Proceedings of the AAAI Conference on Artificial Intelligence, 2019, Honolulu, HI, USA*, <https://doi.org/10.1609/aaai.v33i01.33015885>.
- [48] F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem, *Challenging Common Assumptions in the Unsupervised Learning of Disentangled*

- Representations*, in *Proceedings of the 36th International Conference on Machine Learning, Long Beach, CA, 2019*, edited by K. Chaudhuri and R. Salakhutdinov, Vol. 97, pp. 4114–4124.
- [49] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, *Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach*, *Phys. Rev. Lett.* **120**, 024102 (2018).
- [50] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, *Neural Ordinary Differential Equations*, in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2018), pp. 6571–6583, <https://papers.nips.cc/paper/7892-neural-ordinary-differential-equations>.
- [51] J. Han, C. Ma, Z. Ma, and W. E, *Uniformly Accurate Machine Learning-Based Hydrodynamic Models for Kinetic Equations*, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 21983 (2019).
- [52] C. Ma, J. Wang, and W. E, *Model Reduction with Memory and the Machine Learning of Dynamical Systems*, *Commun. Comput. Phys.* **25**, 947 (2018).
- [53] J. Ling, A. Kurzawski, and J. Templeton, *Reynolds Averaged Turbulence Modelling Using Deep Neural Networks with Embedded Invariance*, *J. Fluid Mech.* **807**, 155 (2016).
- [54] J.-L. Wu, H. Xiao, and E. Paterson, *Physics-Informed Machine Learning Approach for Augmenting Turbulence Models: A Comprehensive Framework*, *Phys. Rev. Fluids* **3**, 074602 (2018).
- [55] X. I. A. Yang, S. Zafar, J.-X. Wang, and H. Xiao, *Predictive Large-Eddy-Simulation Wall Modeling via Physics-Informed Neural Networks*, *Phys. Rev. Fluids* **4**, 034602 (2019).
- [56] A. Paszke *et al.*, *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2019), pp. 8026–8037, <https://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library>.
- [57] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, in *Proceedings of the International Conference on Learning Representations, San Diego, CA, 2015*, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [58] L. Jing, Y. Shen, T. Dubcek, J. Peurifoy, S. Skirlo, Y. LeCun, M. Tegmark, and M. Soljačić, *Tunable Efficient Unitary Neural Networks (EUNN) and Their Application to RNNs*, in *Proceedings of the 34th International Conference on Machine Learning, Sydney, Australia, 2017*, edited by D. Precup and Y. W. Teh, Vol. 70, pp. 1733–1741.
- [59] M. D. Hoffman and M. J. Johnson, *ELBO Surgery: Yet Another Way to Carve Up the Variational Evidence Lower Bound*, in *Proceedings of the Neural Information Processing Systems Workshop on Advances in Approximate Bayesian Inference, 2016, Barcelona, Spain*.
- [60] M. Rosca, B. Lakshminarayanan, and S. Mohamed, *Distribution Matching in Variational Inference*, [arXiv:1802.06847](https://arxiv.org/abs/1802.06847).