



MOX-Report No. 37/2023

**Latent Dynamics Networks (LDNets): learning the intrinsic dynamics
of spatio-temporal processes**

Regazzoni, F.; Pagani, S.; Salvador, M.; Dede', L.; Quarteroni, A.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Latent Dynamics Networks (LDNets): learning the intrinsic dynamics of spatio-temporal processes

Francesco Regazzoni^{1,*}, Stefano Pagani¹, Matteo Salvador^{1,2}, Luca Dede¹, Alfio Quarteroni^{1,3}

¹ MOX, Department of Mathematics, Politecnico di Milano, Milan, Italy

² Institute for Computational and Mathematical Engineering, Stanford University, California, USA

³ École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland (*Professor Emeritus*)

* *Corresponding author* (francesco.regazzoni@polimi.it)

Abstract

Predicting the evolution of systems that exhibit spatio-temporal dynamics in response to external stimuli is a key enabling technology fostering scientific innovation. Traditional equations-based approaches leverage first principles to yield predictions through the numerical approximation of high-dimensional systems of differential equations, thus calling for large-scale parallel computing platforms and requiring large computational costs. Data-driven approaches, instead, enable the description of systems evolution in low-dimensional latent spaces, by leveraging dimensionality reduction and deep learning algorithms. We propose a novel architecture, named Latent Dynamics Network (LDNet), which is able to discover low-dimensional intrinsic dynamics of possibly non-Markovian dynamical systems, thus predicting the time evolution of space-dependent fields in response to external inputs. Unlike popular approaches, in which the latent representation of the solution manifold is learned by means of auto-encoders that map a high-dimensional discretization of the system state into itself, LDNets automatically discover a low-dimensional manifold while learning the latent dynamics, without ever operating in the high-dimensional space. Furthermore, LDNets are meshless algorithms that do not reconstruct the output on a predetermined grid of points, but rather at any point of the domain, thus enabling weight-sharing across query-points. These features make LDNets lightweight and easy-to-train, with excellent accuracy and generalization properties, even in time-extrapolation regimes. We validate our method on several test cases and we show that, for a challenging highly-nonlinear problem, LDNets outperform state-of-the-art methods in terms of accuracy (normalized error 5 times smaller), by employing a dramatically smaller number of trainable parameters (more than 10 times fewer).

1 Introduction

Mathematical models based on differential equations, such as Partial Differential Equations (PDEs) and Stochastic Differential Equations (SDEs), can yield quantitative predictions of the evolution of space-dependent quantities of interest in response to external stimuli. Pivotal examples are given by fluid dynamics and turbulence [46], wave propagation phenomena [42], the deformation of solid bodies and biological tissues [57], molecular dynamics [14], price evolution of financial assets [61], epidemiology [34]. However, the development of traditional modeling-and-simulation approaches carry several mathematical and computational challenges. Model development requires a deep understanding of the physical processes, the adoption of physics first principles or empirical rules, and their translation into mathematical terms. The values of parameters and of boundary and initial conditions required to close the model are often unknown, increasing the intrinsic dimensionality of the solution space. Finally, the computational cost that accompanies the (possibly many-query) numerical approximation of such mathematical models may be prohibitive and hinder their use in relevant applications [48, 52].

In recent years, we are witnessing the introduction of a new paradigm, namely data-driven modeling [6, 15, 22, 47, 59, 62], as opposed to traditional physics-based modeling, enabled by recent advances in optimization, high-performance computing, GPU-based hardware, artificial neural networks (NNs) and Machine/Deep Learning in general. Data-driven modeling methods hold promise in overcoming the limitations of traditional

physics-based models, either as a replacement for them or in synergy with them [2]. On the one hand, data-driven techniques are employed to learn a model directly from experimental data [15, 62]. On the other hand, instead, they are used to build a surrogate for a high-fidelity model – the latter being typically based on the numerical approximation of systems of differential equations – from a dataset of precomputed high-fidelity simulation snapshots [2, 6]. This paradigm is successful in many-query contexts, that is when the computational resources spent in the offline phase (generation of the training data and construction of the data-driven surrogate model) are repaid by a large number of evaluations of the trained model (online phase), as is the case of sensitivity analysis, parameter estimation and uncertainty quantification. Another case of interest is when real-time responses are needed, like, e.g., in clinical scenarios [49].

Several methods have been recently proposed for automatically learning the dynamics of systems exhibiting spatio-temporal behavior [25, 44, 59, 64]. Typically, these methods discretize the space-dependent output field into a high-dimensional vector (e.g., by point-wise evaluation on a grid, or by expansion with respect to a Finite Element basis or to a Fourier basis) and then compress it by means of dimensionality reduction techniques, based e.g. on proper orthogonal decomposition (POD) of a set of snapshots [29, 30, 52, 63], on fully connected auto-encoders, or on convolutional auto-encoders [19, 26, 35, 41, 44, 64]. The underlying assumption is that the dynamics can be represented by a limited number of state variables, called latent variables, whose time evolution is learned either through NNs with recurrent structure (such as RNNs [39], LSTMs [41, 64] or ODE-Nets [24]), dynamic mode decomposition [19], SINDy [17, 59], fully-connected NNs (FCNNs) [26], or DeepONets [44].

When a high-fidelity model is available, there are also techniques for building reduced-order models by exploiting knowledge of the equations [4, 8–10, 16, 31, 51]. These latter methods are however intrusive, unlike the formers, which learn a model in a data-driven manner using only a dataset of input-output pairs. Intrusive techniques are typically based on projecting the high-fidelity model into a low-dimensional space, obtained by POD or by greedy algorithms. In the case of nonlinear models, however, such techniques require special arrangements, such as the (discrete) empirical interpolation method [7, 21, 23], but this entails a difficult trade-off between accuracy and computational cost. Furthermore, many problems feature a slow decay Kolmogorov n -width, an index of the amenability of the solution manifold to be approximated by an n -dimensional linear subspace [13, 20]. In many cases of interest, such as advection-dominated problems or high Reynolds number flow equations, POD-based methods achieve reasonable accuracy only for high values of n [35]. This limits their use in practical applications.

In this paper, we present a novel family of NNs, called Latent Dynamics Networks (LDNets), that can effectively learn, in a data-driven manner, the temporal dynamics of space-dependent fields and predict their evolution for unseen time-dependent input signals and unseen scalar parameters. LDNets automatically discover a compact encoding of the system state in terms of (typically a few) latent scalar variables. Remarkably, the latent representation is learned without the need of using an auto-encoder to explicitly compress a high-dimensional discretization of the system state. Furthermore, LDNets are based on an intrinsically space-dependent reconstruction of the output fields. Indeed, instead of yielding a discrete representation of the fields (e.g. point values on a spatial mesh), LDNets are able to generate output fields defined in any point of space, in a *meshless* manner. As a consequence, the (typically high-dimensional) discrete representation of the output is never explicitly constructed. These features make the training of LDNets extremely lightweight, and boost their generalization ability even in the presence of few training samples and even in time-extrapolation regimes, that is for longer time horizons than those seen during training.

2 Methods

We introduce the notation used throughout this paper and we present the proposed method.

2.1 Notation

We denote by $\mathbf{y}: \Omega \times [0, T] \rightarrow \mathbb{R}^{d_y}$ an output field we aim to predict, where $\Omega \subset \mathbb{R}^d$ is the space domain and $T > 0$ is the final time. The evolution of \mathbf{y} is driven by the input $\mathbf{u}: [0, T] \rightarrow \mathbb{R}^{d_u}$, that is, a set of time-dependent signals or constant parameters. Our goal is to unveil, starting from data, the laws underlying the dependence of \mathbf{y} on \mathbf{u} .

More precisely, we denote input signals as $\mathbf{u}: [0, T] \rightarrow U$, taking values in the set $U \subseteq \mathbb{R}^{d_u}$, and we denote by $\mathcal{U} \subseteq \{\mathbf{u}: [0, T] \rightarrow U\}$ the set of admissible input signals. Then, we denote by $\mathbf{y}: \Omega \times [0, T] \rightarrow Y$ the output (space-time dependent) field, with values in $Y \subseteq \mathbb{R}^{d_y}$. Finally, we denote by $\mathcal{Y} \subseteq \{\mathbf{y}: \Omega \times [0, T] \rightarrow Y\}$ the space of possible outputs. We assume that the map $\mathbf{u} \mapsto \mathbf{y}$ is well defined (i.e. the output \mathbf{y} is unambiguously determined by the input \mathbf{u}) and it is consistent with the arrow of time (i.e. $\mathbf{y}(\mathbf{x}, t)$ depends on $\mathbf{u}(s)|_{s \in [0, t]}$ but not on $\mathbf{u}(s)|_{s \in (t, T]}$).

An important (albeit not exclusive) example is the case when the dynamics we aim to learn underlies a differential model, that is when the map $\mathbf{u} \mapsto \mathbf{y}$ is defined as the composition of an observation operator and the solution map $\mathbf{u} \mapsto \mathbf{z}$ of a partial differential equation (PDE) in the form of

$$\begin{cases} \partial_t \mathbf{z}(\mathbf{x}, t) = \mathcal{F}(\mathbf{z}(\mathbf{x}, t), \mathbf{u}(t)) & \text{in } \Omega \times (0, T] \\ \mathbf{y}(\mathbf{x}, t) = \mathcal{G}(\mathbf{z}(\mathbf{x}, t), \mathbf{x}) & \text{in } \Omega \times (0, T] \\ \mathbf{z}(\mathbf{x}, 0) = \mathbf{z}_0(\mathbf{x}) & \text{in } \Omega, \end{cases} \quad (1)$$

where $\mathbf{z} \in \mathcal{Z} \subseteq \{\mathbf{z}: \Omega \times [0, T] \rightarrow Z\}$, with $Z \subseteq \mathbb{R}^{d_z}$, is the state variable (typically, \mathcal{Z} is a Sobolev space). Here, $\mathbf{z}_0: \Omega \rightarrow Z$ is the initial state, \mathcal{F} is a differential operator, and \mathcal{G} is the observation operator. In particular, LDNets can be also used for the model-order reduction of (1), by passing through data generated via a numerical approximation of (1), e.g. by the Finite Element method, called full-order model (FOM). Meaningful examples are provided in Sec. 3. Nonetheless, in this paper neither knowledge nor even the existence of a model such as (1) is required: the training of an LDNet only requires input-output pairs.

Remark 1. The case when the output field \mathbf{y} is determined not only by some time-dependent inputs \mathbf{u} , but also by some inputs that are time-independent (typically called *parameters*) is a special case of the one considered here. Still, to keep the notation compact, we will use the same symbol \mathbf{u} to collectively denote time-dependent inputs (i.e. signals) and time-constant inputs (i.e. parameters).

Remark 2. The FOM (1) is an autonomous system. The non-autonomous case can be recovered as a special case by setting $\mathbf{u}(t) \cdot \mathbf{e}_k = t$ for some k , where \mathbf{e}_k is the k -th element of the canonical base of \mathbb{R}^{d_u} .

2.2 Training data

The training data are collected by considering a finite number of realizations of the map $\mathbf{u} \mapsto \mathbf{y}$, each one referred to as a *training sample*. For each training sample $i \in \mathcal{S}_{\text{train}}$, we collect the following discrete observations:

- $\mathbf{u}_i(\tau)$, for $\tau \in \mathcal{S}^i$;
- $\mathbf{y}_i(\boldsymbol{\xi}, \tau)$, for $\tau \in \mathcal{T}^i$, $\boldsymbol{\xi} \in \mathcal{P}_\tau^i$;

where $\mathcal{S}^i \subset [0, T]$, $\mathcal{T}^i \subset [0, T]$ and $\mathcal{P}_\tau^i \subset \Omega$ are discrete sets of observations. We remark that the observation times and points can be either shared among samples (i.e. $\mathcal{S}^i \equiv \mathcal{S}$, $\mathcal{T}^i \equiv \mathcal{T}$ and $\mathcal{P}_\tau^i \equiv \mathcal{P}$ for any i and for any τ) or be different from one sample to another.

Our goal is to learn the map $\mathbf{u} \mapsto \mathbf{y}$, that is to infer the output $\mathbf{y}(\mathbf{x}, t)$ corresponding to inputs $\mathbf{u}(t)$ outside the training set.

2.3 LDNets

An LDNet consists of two sub-networks, $\mathcal{NN}_{\text{dyn}}$ and $\mathcal{NN}_{\text{rec}}$, that is two FCNNs with trainable parameters \mathbf{w}_{dyn} and \mathbf{w}_{rec} , respectively (see Fig. 1). The first NN, namely $\mathcal{NN}_{\text{dyn}}$, evolves the dynamics of the latent variables $\mathbf{s}(t) \in \mathbb{R}^{d_s}$ according to the differential equation

$$\dot{\mathbf{s}}(t) = \mathcal{NN}_{\text{dyn}}(\mathbf{s}(t), \mathbf{u}(t); \mathbf{w}_{\text{dyn}}) \quad \text{in } (0, T] \quad (2)$$

with a prescribed arbitrary initial condition (in our numerical test, we set $\mathbf{s}(0) = \mathbf{0}$). The inputs of $\mathcal{NN}_{\text{dyn}}$ are the latent states $\mathbf{s}(t)$ and the input signal $\mathbf{u}(t)$ at the current time t . Instead, the second NN, $\mathcal{NN}_{\text{rec}}$, is used to reconstruct $\tilde{\mathbf{y}}$, an approximation of the output field \mathbf{y} at any time $t \in [0, T]$ and at any query point $\mathbf{x} \in \Omega$:

$$\tilde{\mathbf{y}}(\mathbf{x}, t) = \mathcal{NN}_{\text{rec}}(\mathbf{s}(t), \mathbf{u}(t), \mathbf{x}; \mathbf{w}_{\text{rec}}) \quad \text{in } \Omega \times (0, T] \quad (3)$$

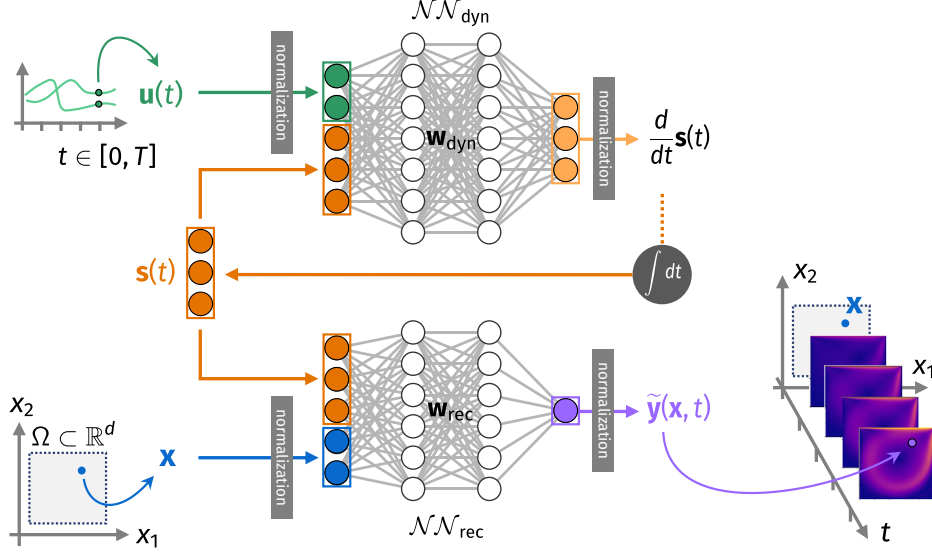


Figure 1: **LDNet architecture.** The network $\mathcal{NN}_{\text{dyn}}$ receives the input $\mathbf{u}(t)$ and the latent state $\mathbf{s}(t)$ and returns the time derivative of the latent state, thus defining its dynamics. The network $\mathcal{NN}_{\text{rec}}$, instead, is evaluated only when an estimate of the output field \mathbf{y} is sought. More precisely, an approximation of $\mathbf{y}(\mathbf{x}, t)$ is recovered by giving as an input to $\mathcal{NN}_{\text{rec}}$ the latent state at time t and the query space coordinate $\mathbf{x} \in \Omega$. In general, the reconstruction network $\mathcal{NN}_{\text{rec}}$ might take as an input $\mathbf{u}(t)$ as well (see e.g. Sec. 3, Test Case 2); for simplicity, in the figure we represent the special case when $\mathcal{NN}_{\text{rec}}$ does not depend on $\mathbf{u}(t)$.

We remark that the reconstruction network $\mathcal{NN}_{\text{rec}}$ is independently queried for every point $\mathbf{x} \in \Omega$ for which the solution is sought.

Hence, the LDNet defines a map from a time-dependent input signal $\mathbf{u} \in \mathcal{U}$ to a space-time dependent field $\tilde{\mathbf{y}} \in \mathcal{Y}$ through the solution of the following system of ordinary differential equations (ODEs):

$$\begin{cases} \dot{\mathbf{s}}(t) = \mathcal{NN}_{\text{dyn}}(\mathbf{s}(t), \mathbf{u}(t); \mathbf{w}_{\text{dyn}}) & \text{in } (0, T] \\ \mathbf{s}(0) = \mathbf{0} \\ \tilde{\mathbf{y}}(\mathbf{x}, t) = \mathcal{NN}_{\text{rec}}(\mathbf{s}(t), \mathbf{u}(t), \mathbf{x}; \mathbf{w}_{\text{rec}}) & \text{for } \mathbf{x} \in \Omega \text{ and } t \in [0, T], \end{cases} \quad (4)$$

where $\mathbf{s}(t) \in \mathbb{R}^{d_s}$ is the vector of latent states. The number of latent states d_s is set by the user, and should be regarded as an hyperparameter. We remark that, thanks to the hidden nature of $\mathbf{s}(t)$, we can assume without loss of generality the initial condition $\mathbf{s}(0) = \mathbf{0}$ (see [56] for a discussion on this topic in a similar framework). In this work, we always consider hyperbolic tangent (tanh) activation functions.

Remark 3. The formulation (4) is the most general one. A special case is the one where $\mathcal{NN}_{\text{rec}}$ does not depend on $\mathbf{u}(t)$. Whether or not to include the latter dependency in $\mathcal{NN}_{\text{rec}}$ is an architectural choice that shall be regarded as a hyperparameter, possibly subject to selection via cross-validation. In many cases, however, the choice can be driven by the *physics* of the underlying process. Specifically, we will leave an explicit dependency whenever the output $\mathbf{y}(\mathbf{x}, t)$ depends on the input $\mathbf{u}(t)$ instantaneously. The case where the dependency is neglected is the one that we will mostly consider in our test cases, expect for Test Case 2, in which we will allow $\mathcal{NN}_{\text{rec}}$ to depend on $\mathbf{u}(t)$ in a direct way.

In practice, the ODE system (4) is discretized by a suitable numerical method. In this work, we employ a Forward Euler scheme with a uniform time step size Δt , but other schemes could be considered as well (e.g. time-adaptive Runge-Kutta schemes [53]). In case the observation times \mathcal{T}^i do not coincide with the discrete times $k\Delta t$, for $k = 1, \dots$, we perform a re-sampling of \mathbf{u} through a piecewise linear interpolation. Similarly, to evaluate the predicted output $\tilde{\mathbf{y}}$ in correspondence of the observation times $\tau \in \mathcal{T}^i$, we interpolate the discrete solution of $\mathbf{s}(t)$ at the time instants τ .

We denote with the symbol $\mathcal{RN}_{\text{dyn}}$ (to evoke its recurrent nature) the operator mapping the time series of inputs $\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{T}^i}$ associated with a given sample i to the latent state \mathbf{s}_i evolution. More precisely, we

have, for any sample i and at any time $t \in [0, T]$:

$$\mathbf{s}_i(t) = \mathcal{RNN}_{\text{dyn}}(\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{I}^i}, t; \mathbf{w}_{\text{dyn}})$$

With this notation, the LDNet output $\tilde{\mathbf{y}}_i(\mathbf{x}, t)$ is the result of the composition of $\mathcal{NN}_{\text{rec}}$ with $\mathcal{RNN}_{\text{dyn}}$:

$$\tilde{\mathbf{y}}_i(\mathbf{x}, t) = \mathcal{NN}_{\text{rec}}(\mathcal{RNN}_{\text{dyn}}(\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{I}^i}, t; \mathbf{w}_{\text{dyn}}), \mathbf{u}_i(t), \mathbf{x}; \mathbf{w}_{\text{rec}}). \quad (5)$$

To train the LDNet, we define the loss function:

$$\mathcal{L}(\mathbf{w}_{\text{dyn}}, \mathbf{w}_{\text{rec}}) = \sum_{i \in \mathcal{S}_{\text{train}}} \sum_{\tau \in \mathcal{I}^i} \sum_{\xi \in \mathcal{I}_{\tau}^i} \mathcal{E}(\tilde{\mathbf{y}}_i(\xi, \tau), \mathbf{y}_i(\xi, \tau)) + \alpha_{\text{dyn}} \mathcal{R}(\mathbf{w}_{\text{dyn}}) + \alpha_{\text{rec}} \mathcal{R}(\mathbf{w}_{\text{rec}}),$$

where the symbol \sum denotes the average operator (that is the sum over a set divided by the cardinality of the set), and where $\tilde{\mathbf{y}}_i(\xi, \tau)$ are the outputs of the LDNet associated with the trainable parameters \mathbf{w}_{dyn} and \mathbf{w}_{rec} as defined in (5). The discrepancy metric \mathcal{E} is typically defined as

$$\mathcal{E}(\tilde{\mathbf{y}}, \mathbf{y}) = \frac{\|\tilde{\mathbf{y}} - \mathbf{y}\|^2}{y_{\text{norm}}^2} \quad (6)$$

with y_{norm} being a normalization factor defined from case to case and where $\|\cdot\|$ denotes the euclidean norm. The first term of \mathcal{L} represents therefore the normalized mean square error between observations and LDNet predictions. Moreover, to mitigate overfitting, suitable regularization terms on the NN weights could be introduced, with weighting factors α_{dyn} and α_{rec} . In this work, we define \mathcal{R} as the mean of the squares of the NN weights (yielding the so-called L^2 -regularization or Tikhonov regularization).

Remark 4. The quadratic discrepancy metric (6), while being the most natural choice, is not the unique one. For instance, it can be replaced by goal-oriented metrics (an example is given in Test Case 2).

Training an LDNet consists in employing suitable optimization methods to approximate the solution of the following non-convex minimization problem:

$$(\mathbf{w}_{\text{dyn}}^*, \mathbf{w}_{\text{rec}}^*) = \underset{\mathbf{w}_{\text{dyn}}, \mathbf{w}_{\text{rec}}}{\operatorname{argmin}} \mathcal{L}(\mathbf{w}_{\text{dyn}}, \mathbf{w}_{\text{rec}}).$$

The two NNs are simultaneously trained. Thanks to the simultaneous end-to-end training of the two NNs, the latent space is discovered at the same time as learning the dynamics of the system. This generalizes the approach presented in [56] for the case of time signals as outputs.

2.4 Normalization layers

In order to facilitate training, we normalize the inputs and the outputs of the NNs. Specifically:

- We normalize the signals \mathbf{u} , the output fields \mathbf{y} and the space variables \mathbf{x} , so that each entry approximately spans the interval $[-1, 1]$. We normalize each entry independently of the others. More precisely we normalize each scalar variable α through the affine transformation $\tilde{\alpha} = (\alpha - \alpha_0)/\alpha_w$ where α_0 is a reference value and α_w is a reference width. To define α_0 and α_w , we follow two different strategies.

1. If the variable takes values in a bounded interval $[\alpha_{\min}, \alpha_{\max}]$, we set

$$\begin{aligned} \alpha_0 &= (\alpha_{\min} + \alpha_{\max})/2, \\ \alpha_w &= (\alpha_{\max} - \alpha_{\min})/2. \end{aligned}$$

2. If the variable is sampled from a distribution with unbounded support (e.g., when α is normally distributed), we set α_0 equal to the sample mean and α_w equal to three times the sample standard deviation.

- We also normalize the time variable, by dividing the time steps by a characteristic time scale Δt_{ref} . The normalization constant Δt_{ref} impacts the output of $\mathcal{NN}_{\text{dyn}}$, that is dimensionally proportional to the inverse of time. Since finding a good value for Δt_{ref} is in general not straightforward, we typically consider it as a hyperparameter, tuned through a suitable automatic algorithm (see Section. 2.8).

- We do not normalize the latent states \mathbf{s} , since their distribution is not known before training. Indeed, when hyperparameters are well tuned, the training algorithm tends to generate models that produce latent states with approximately normalized values.

In practice, normalization can be achieved either by modifying the training data accordingly, or by embedding the two NNs between two normalization layers (namely, one input layer and one output layer) each. Formally, the second approach consists in defining $\mathcal{NN}_{\text{dyn}}$ and $\mathcal{NN}_{\text{rec}}$ as follows, where we $\widetilde{\mathcal{NN}}_{\text{dyn}}$ and $\widetilde{\mathcal{NN}}_{\text{rec}}$ are two FCNNs:

$$\begin{aligned}\mathcal{NN}_{\text{dyn}}(\mathbf{s}, \mathbf{u}; \mathbf{w}_{\text{dyn}}) &= \Delta t_{\text{ref}}^{-1} \widetilde{\mathcal{NN}}_{\text{dyn}}(\mathbf{s}, (\mathbf{u} - \mathbf{u}_0) \oslash \mathbf{u}_w; \mathbf{w}_{\text{dyn}}) \\ \mathcal{NN}_{\text{rec}}(\mathbf{s}, \mathbf{u}, \mathbf{x}; \mathbf{w}_{\text{rec}}) &= \mathbf{y}_0 + \mathbf{y}_w \odot \widetilde{\mathcal{NN}}_{\text{rec}}(\mathbf{s}, (\mathbf{u} - \mathbf{u}_0) \oslash \mathbf{u}_w, (\mathbf{x} - \mathbf{x}_0) \oslash \mathbf{x}_w; \mathbf{w}_{\text{rec}})\end{aligned}$$

where \odot and \oslash denote the Hadamard (i.e. element-wise) product and division, respectively.

2.5 Imposing a-priori physical knowledge

The architecture of LDNets reflects certain features of the physics they are meant to capture. With respect to the space variable, the representation is continuous, unlike methods that reconstruct a discretized solution thus losing the correspondence between neighboring points. With respect to the time variable, the dynamics is driven by a system of differential equations which makes LDNets consistent with the arrow of time (i.e., with the causality principle [56]). These features make it natural to introduce a-priori physical knowledge in the construction and training of LDNets. In this regard, we distinguish between weak imposition and strong imposition.

Weak imposition consists of introducing physics-informed terms [54] into the loss function, aimed at promoting solutions that satisfy certain requirements (such as irrotationality of a velocity field, to make an example). In this paper we do not show examples in this regard, but simply highlight that the continuous representation of the output field used by LDNets makes the introduction of such terms very straightforward through the use of automatic differentiation.

Strong imposition, on the other hand, consists of modifying the architecture of the LDNet components in order to obtain models that automatically satisfy certain properties [5, 38]. In what follows, we provide two examples of how this can be applied to ensure both temporal (acting on $\mathcal{NN}_{\text{dyn}}$) and spatial (acting on $\mathcal{NN}_{\text{rec}}$) properties.

2.5.1 Equilibrium configuration imposition

In many real-life applications, data are collected starting from an equilibrium configuration. This entails that the initial state should be an equilibrium for the latent dynamics as well, in virtue of the interpretation of \mathbf{s} as a compact encoding of the full-order system state. Therefore, we define the right-hand side of the latent state evolution equation as follows, where $\widetilde{\mathcal{NN}}_{\text{dyn}}$ is a trainable FCNN and where $\mathbf{u}_{\text{eq}} \in U$ is the input at equilibrium:

$$\mathcal{NN}_{\text{dyn}}(\mathbf{s}, \mathbf{u}; \mathbf{w}_{\text{dyn}}) = \widetilde{\mathcal{NN}}_{\text{dyn}}(\mathbf{s}, \mathbf{u}; \mathbf{w}_{\text{dyn}}) - \widetilde{\mathcal{NN}}_{\text{dyn}}(\mathbf{0}, \mathbf{u}_{\text{eq}}; \mathbf{w}_{\text{dyn}})$$

As a consequence, the initial state $\mathbf{s} = \mathbf{0}$ of the model is an equilibrium for any choice of the trainable parameters \mathbf{w}_{dyn} .

2.5.2 Prescribed solution in subsets of the domain (e.g. Dirichlet boundary conditions)

The evolution of the output field is often unknown except on a subset of the domain Ω , such as for example a portion Γ_D of its boundary $\partial\Omega$. This happens, e.g., when there is a FOM that features a Dirichlet boundary condition like

$$\mathbf{y}(\mathbf{x}, t) = \mathbf{y}_D(\mathbf{x}) \quad \text{on } \Gamma_D. \quad (7)$$

In this case, the solution is constrained to satisfy (7) by defining $\mathcal{NN}_{\text{rec}}$ as

$$\mathcal{NN}_{\text{rec}}(\mathbf{s}, \mathbf{u}, \mathbf{x}; \mathbf{w}_{\text{rec}}) = \mathbf{y}_{\text{lift}}(\mathbf{x}) + \widetilde{\mathcal{NN}}_{\text{rec}}(\mathbf{s}, \mathbf{u}, \mathbf{x}; \mathbf{w}_{\text{rec}})\psi(\mathbf{x})$$

where $\widetilde{\mathcal{NN}}_{\text{rec}}$ is a trainable FCNN, \mathbf{y}_{lift} is the lifting of the boundary datum, that is an extension of \mathbf{y}_D to the whole domain Ω , and $\psi: \Omega \rightarrow \mathbb{R}$ is a mask, that is a smooth function such that $\psi(\mathbf{x}) = 0$ if and only if $\mathbf{x} \in \Gamma_D$. See [58] for further details and [12] for a general approach to construct the mask ψ based on approximate distance functions.

2.6 Error metrics

To evaluate the generalization accuracy of a trained LDNet, we test it on unseen data, that is on samples belonging to a test set denoted by $\mathcal{S}_{\text{test}}$. Specifically, we employ two metrics.

Normalized root mean square error (NRMSE) is obtained as the square root of the mean of the squares of the normalized errors obtained on the testing set:

$$\text{NRMSE} = \sqrt{\sum_{i \in \mathcal{S}_{\text{test}}} \sum_{\tau \in \mathcal{T}^i} \sum_{\xi \in \mathcal{P}_{\tau}^i} \frac{\|\tilde{\mathbf{y}}_i(\xi, \tau) - \mathbf{y}_i(\xi, \tau)\|^2}{y_{\text{norm}}^2}}.$$

Pearson dissimilarity ($1 - \rho$) is defined from the Pearson correlation coefficient ρ :

$$\rho = \frac{\sum_{i \in \mathcal{S}_{\text{test}}} \sum_{\tau \in \mathcal{T}^i} \sum_{\xi \in \mathcal{P}_{\tau}^i} (\tilde{\mathbf{y}}_i(\xi, \tau) - \bar{\tilde{\mathbf{y}}}) \cdot (\mathbf{y}_i(\xi, \tau) - \bar{\mathbf{y}})}{\sqrt{\sum_{i \in \mathcal{S}_{\text{test}}} \sum_{\tau \in \mathcal{T}^i} \sum_{\xi \in \mathcal{P}_{\tau}^i \|\tilde{\mathbf{y}}_i(\xi, \tau) - \bar{\tilde{\mathbf{y}}}\|^2 \sum_{i \in \mathcal{S}_{\text{test}}} \sum_{\tau \in \mathcal{T}^i} \sum_{\xi \in \mathcal{P}_{\tau}^i \|\mathbf{y}_i(\xi, \tau) - \bar{\mathbf{y}}\|^2}}$$

where we denote the average outputs as

$$\begin{aligned} \bar{\mathbf{y}} &= \sum_{i \in \mathcal{S}_{\text{test}}} \sum_{\tau \in \mathcal{T}^i} \sum_{\xi \in \mathcal{P}_{\tau}^i} \mathbf{y}_i(\xi, \tau) \\ \bar{\tilde{\mathbf{y}}} &= \sum_{i \in \mathcal{S}_{\text{test}}} \sum_{\tau \in \mathcal{T}^i} \sum_{\xi \in \mathcal{P}_{\tau}^i} \tilde{\mathbf{y}}_i(\xi, \tau) \end{aligned}$$

We remark that both metrics (RMSE and $1 - \rho$) are robust with respect to multiplicative rescaling of the outputs. For both of them, the smaller the value of the metric, the higher the accuracy of the predictions.

2.7 Training algorithm

To train the LDNet, we employ a two stage strategy. First, we perform a limited number of epochs (typically, a few hundreds) with the Adam optimizer [33], starting with a learning rate of 10^{-2} . Then, we switch to a second-order accurate optimizer, namely BFGS [28]. BFGS is more accurate than Adam, but more prone to get stuck in local minima, which is why it is useful to precede it with some Adam iterations, which provide a good initial guess.

To evaluate the gradient of the loss function with respect to the trainable parameters, we combine back-propagation-through-time for $\mathcal{RN}_{\text{dyn}}$ with back-propagation for $\mathcal{NN}_{\text{rec}}$ [28]. To initialize the parameters of the two NNs, we employ a Glorot uniform strategy for weights and zero values for the biases [28].

2.8 Hyperparameters tuning algorithms

The hyperparameters of the proposed method are the number of layers and neurons of $\mathcal{NN}_{\text{dyn}}$ and $\mathcal{NN}_{\text{rec}}$, the L^2 regularization weights α_{dyn} and α_{rec} , the normalization time constant Δt_{ref} and, whenever necessary in the different test cases, the number of latent states d_s . To automatically tune them, we employ the Tree-structured Parzen Estimator (TPE) Bayesian algorithm [1, 11]. The hyperparameters search space is defined as an hypercube, with a log-uniform sampling. We perform K-fold cross validation while monitoring the value of the discrepancy metric in Eq. (6). We also employ the Asynchronous Successive Halving (ASHA) scheduler to early terminate hyperparameters configurations that are either bad or not promising [36, 37].

We simultaneously train multiple NNs associated to different hyperparameters settings on a supercomputer endowed with several CPUs via Message Passing Interface (MPI). Each NN exploits Open Multi-Processing (OpenMP) for Hyper-Threading, which allows for a speed-up in the computationally-intensive tensorial operations involved during the training phase. For the implementation, we rely on the Ray Python distributed framework [43].

3 Results

We demonstrate the effectiveness of LDNets through several test cases. First, we consider a linear PDE model to analyze the ability of LDNets to extract a compact latent representation of models that are progressively less amenable to reduction. Then, we consider the time-dependent version of a benchmark problem in fluid dynamics. Finally, we compare LDNets with state-of-the-art methods in a challenging task, that is, learning the dynamics of the Aliev-Panfilov model [3], a highly non-linear excitation-propagation PDE model used in the field of cardiac electrophysiology modeling. We focus on synthetically generated data obtained by numerical approximation of differential models, thus allowing us to test LDNet predictions against ground-truth results.

3.1 Test Case 1: Advection-Diffusion-Reaction equation

We consider the linear advection-diffusion-reaction (ADR) equation on the interval $\Omega = (-1, 1)$:

$$\frac{\partial z(x, t)}{\partial t} - \mu_1 \frac{\partial^2 z(x, t)}{\partial x^2} - \mu_2 \frac{\partial z(x, t)}{\partial x} + \mu_3 z(x, t) = f(x, t) \quad x \in (-1, 1), t \in (0, T]. \quad (8)$$

This PDE is widely used, e.g., to describe the concentration $z(x, t)$ of a substance dissolved in a channel [53]. The constant parameters μ_1 , μ_2 and μ_3 respectively represent diffusion, advection and reaction coefficients, while the forcing term $f(x, t)$ is a prescribed external source, defined as $f(x, t) = A(t) \cos(2\pi F(t)x - P(t))$, that is a sine wave with amplitude $A(t)$, frequency $F(t)$ and phase $P(t)$. We consider an initial condition $z(x, 0) = z_0(x)$ and periodic boundary conditions.

To generate the training dataset, we employ a high-fidelity FFT-based solver on 101 equally spaced grid points, combined with an adaptive-time integration scheme for stiff problems [18, 50]. Then, we subsample the time domain in 100 equally distributed intervals. We challenge LDNets in predicting the space-time evolution of the target variable $\mathbf{y}(x, t) = z(x, t)$ by considering three cases of increasing complexity (Test Cases 1a, 1b, 1c), in which the input \mathbf{u} is associated either to the parameters μ_1 , μ_2 and μ_3 , or to the forcing term $f(x, t)$. In all the test cases, we define y_{norm} as the difference between maximum and minimum value taken by the output on the whole training set. In all the cases presented below, unless otherwise stated, we set $\Delta t = 5 \cdot 10^{-2}$.

3.1.1 Sampling of inputs

In order to generate training and testing data, we need to sample the space of inputs. This calls for a probability distribution on the latter space. For the parameters (μ_1 , μ_2 and μ_3), we take for simplicity a uniform distribution on a suitable hypercube. For the time-dependent signals ($A(t)$, $F(t)$ and $P(t)$), instead, probability distributions on function spaces are needed. For $A(t)$ and $P(t)$, we consider a Gaussian Process distribution [55] with mean μ and with the following covariance kernel

$$K(t_1, t_2) = \sigma^2 \exp \left[-\frac{(t_1 - t_2)^2}{2\tau^2} \right],$$

characterized by standard deviation σ and characteristic time-scale τ . For what concerns $F(t)$ instead, in order to let it vary within a bounded set (f_{\min}, f_{\max}) , we define it as

$$F(t) = \frac{1}{2} \left[f_{\min} + f_{\max} + (f_{\max} - f_{\min}) \tanh \left(\frac{3}{5} \gamma(t) \right) \right],$$

with $\gamma(t)$ sampled from a Gaussian Process with mean $\mu = 0$, standard deviation $\sigma = 1$ and prescribed characteristic time-scale τ . The values of μ , σ and τ are indicated below for each test case.

3.1.2 Test Case 1a: finite latent dimension, constant parameters

First, we consider $z_0(x) = \cos(\pi x)$ and $f \equiv 0$. We aim at predicting the evolution of $z(x, t)$, depending on the constant parameters $\mathbf{u}(t) \equiv (\mu_1, \mu_2, \mu_3)$. Due to the linearity of the equation, the solution is, at any time t , a sine wave with period 2, and can be thus unambiguously identified by two scalars (namely, the wave amplitude and phase, or equivalently, the real and imaginary part of the Fourier transform at frequency 0.5). In other terms, the intrinsic dimension of the solution manifold is strictly equal to 2. This provides therefore an ideal testbed for the capability of LDNets to recognize and learn a low-dimensional encoding of the system state from data.

To generate training and testing data, we employ a Monte Carlo sampling of the hypercube defined by the bounds $(\mu_1, \mu_2, \mu_3) \in [0, 0.05] \times [-0.1, 0.1] \times [0, 0.01]$. We consider 100 training samples and 500 testing samples. We select the hyperparameters according to the algorithm presented in Sec. 2.8. The ranges of hyperparameter values used in the tuning algorithm are reported in Tab. 1, in the row *tuning*. The selected values are instead reported in the row *final*.

In Tab. 2 we report the training and testing accuracy metrics obtained by training an LDNet with $d_s = 2$ latent variables. In the table we show the accuracy achieved after 500, 5000 or 50000 epochs of BFGS (in all the cases, we first run 200 epochs of Adam). The LDNet, trained on 100 samples, achieves an excellent accuracy when tested on 500 unseen samples. Indeed, the NRMSE is $1.88 \cdot 10^{-5}$ on the testing set, against a training NRMSE of $1.81 \cdot 10^{-5}$. Pearson dissimilarity is $3.30 \cdot 10^{-9}$ on the testing set and $3.00 \cdot 10^{-9}$ on the training set. The very small differences in the accuracy metrics between training and testing sets provide evidence that the trained LDNet reproduces the FOM dynamics with great fidelity and without overfitting, that is with remarkably good generalization capabilities.

3.1.3 Test Case 1b: finite latent dimension, time-dependent inputs

We now consider the case of time-dependent inputs, with a forcing term $f(x, t) = A(t) \cos(\pi x - P(t))$, where $\mathbf{u}(t) = (A(t), P(t))$ represents an input signal that can vary in time within a bounded set. We thus fix the values of the parameters to $\mu_1 = 0.05$, $\mu_2 = 0$ and $\mu_3 = 0.002$. Similarly to Test Case 1a, the solution manifold has dimension 2 (thanks to the equation being linear and to the forcing term having constant frequency), but learning the dynamics becomes more challenging due to the presence of time-dependent inputs. We let the amplitude $A(t)$ and the phase $P(t)$ vary in time, as described in Sec. 3.1.1. Specifically, for $A(t)$ we set $\mu = 2/5$, $\sigma = 2/15$ and $\tau = 1$, while for $P(t)$ we set $\mu = 0$, $\sigma = 4/3$ and $\tau = 1$.

First, we consider 100 training samples and we set $\Delta t = 0.05$. We tune the hyperparameters starting from the ranges indicated in Tab. 1, and we select the optimal values reported in the same table. The LDNets predictions $\tilde{\mathbf{y}}$ for some test samples are displayed against reference outputs \mathbf{y} in Fig. 2.

Now, we perform two tests. First, by keeping $\Delta t = 0.05$, we vary the number of training samples in the set $\{25, 50, 100, 200, 400\}$. Then, we let Δt vary in the set $\{0.2, 0.1, 0.05, 0.02\}$, by keeping the number of training samples equal to 100. In both the cases, we do not vary the hyperparameters reported in Tab. 1. As desirable, the accuracy of predictions improves as the number of training samples increases and as the time discretization step size is reduced (Fig. 3b).

Test Case 1b provides an ideal testbed to assess the capabilities of LDNets in discovering a compact representation of the FOM state. The state of (8) indeed evolves on a two-dimensional manifold, being the state fully characterized, at any time t , by two scalars. In other terms, provided that the forcing term is defined as in Test Case 1b, the model (8) has an intrinsic latent dimension equal to 2. At any time t , in fact, $z(\cdot, t)$ is a sine wave with frequency 0.5. Among the infinitely many equivalent parametrization, one is given by the Fourier transform of $z(\cdot, t)$ at frequency 0.5, that is determined by its real and imaginary part, respectively denoted by $Re(\hat{z}(0.5))$ and $Im(\hat{z}(0.5))$.

Our results show that, during the training process of an LDNet, the algorithm discovers a compact representation of the solution field $z(\cdot, t)$, represented by the two latent variables $s_1(t)$ and $s_2(t)$. We now investigate whether there is a relationship between the pairs (s_1, s_2) and $(Re(\hat{z}(0.5)), Im(\hat{z}(0.5)))$. With this goal, we train four different LDNets, starting from a different random initializations of the trainable parameters. Then, we evaluate the trained LDNets on 24 test samples, and we collect the trajectories in the latent space (s_1, s_2) . Finally, we plot these trajectories by displaying each point with a color that depends on the corresponding value of $Re(\hat{z}(0.5))$ computed from the reference solution, and we repeat the same procedure by considering the values of $Im(\hat{z}(0.5))$. The results are shown in Fig. 4. We notice that

Test cases	$\mathcal{NN}_{\text{dyn}}$				Hyperparameters			$\alpha_{\text{dyn}}, \alpha_{\text{rec}}$	Trainable parameters		
	layers	neurons	layers	$\mathcal{NN}_{\text{rec}}$ neurons	Δt_{ref}	$\mathcal{NN}_{\text{dyn}}$ # param.	$\mathcal{NN}_{\text{rec}}$ # param.				
Test Case 1a											
tuning	2	3 – 30	2	3 – 15	$[10^{-1}, 10^1]$		0				
final	2	9	2	11	0.5		0		164	188	
Test Case 1b											
tuning	2	3 – 30	2	3 – 15	$[10^{-1}, 10^1]$		$[10^{-6}, 10^{-1}]$				
final	2	10	2	7	2.3		10^{-5}		182	92	
Test Case 1c											
tuning	1 – 3	5 – 20	1,2	4 – 15	$[10^{-1}, 10^1]$		$[10^{-5}, 10^{-1}]$				
fine tuning	2	twice $\mathcal{NN}_{\text{rec}}$	2	4 – 9	8		$[2 \cdot 10^{-4}, 2 \cdot 10^{-3}]$				
f_{max}											
d_s											
0.5	2	10	2	5	8		$1.95 \cdot 10^{-3}$		192	56	
0.5	3	16	2	8	8		$2.10 \cdot 10^{-4}$		435	121	
0.5	4	16	2	8	8		$2.10 \cdot 10^{-4}$		468	129	
0.5	5	16	2	8	8		$2.00 \cdot 10^{-4}$		501	137	
1	2	12	2	6	8		$1.05 \cdot 10^{-3}$		254	73	
1	3	14	2	7	8		$4.20 \cdot 10^{-4}$		353	99	
1	4	16	2	8	8		$2.70 \cdot 10^{-4}$		468	129	
1	5	16	2	8	8		$3.50 \cdot 10^{-4}$		501	137	
2	2	14	2	7	8		$6.60 \cdot 10^{-4}$		324	92	
2	3	16	2	8	8		$3.90 \cdot 10^{-4}$		435	121	
2	4	16	2	8	8		$2.20 \cdot 10^{-4}$		468	129	
2	5	16	2	8	8		$2.00 \cdot 10^{-4}$		501	137	

Table 1: Test Case 1: hyperparameters ranges and selected values. We remark that, in the fine tuning stage of test Case 1c, we select twice as many neurons for $\mathcal{NN}_{\text{dyn}}$ as for $\mathcal{NN}_{\text{rec}}$, in order to reduce the number of independent hyperparameters. See text for details.

BFGS epochs	500	5000	50000
Training time	8m 5s	1h 5m	10h 8m
$\text{NRMSE}_{\text{train}}$	$1.02 \cdot 10^{-3}$	$7.08 \cdot 10^{-5}$	$1.81 \cdot 10^{-5}$
$\text{NRMSE}_{\text{test}}$	$1.19 \cdot 10^{-3}$	$7.23 \cdot 10^{-5}$	$1.88 \cdot 10^{-5}$
$1 - \rho_{\text{train}}$	$9.42 \cdot 10^{-6}$	$4.58 \cdot 10^{-8}$	$3.00 \cdot 10^{-9}$
$1 - \rho_{\text{test}}$	$1.33 \cdot 10^{-5}$	$4.87 \cdot 10^{-8}$	$3.30 \cdot 10^{-9}$

Table 2: Test Case 1a: training and test accuracy metrics for LDNets trained with an increasing number of BFGS training epochs (500, 5000 and 50000). Training time refer to a single-CPU standard laptop.




Figure 2: Test Case 1b. 100 testing samples, comparing the reference outputs \mathbf{y} (left) with the LDNet predictions $\tilde{\mathbf{y}}$ (right). For each sample, the horizontal axis refers to space, and the vertical axis refers to time.

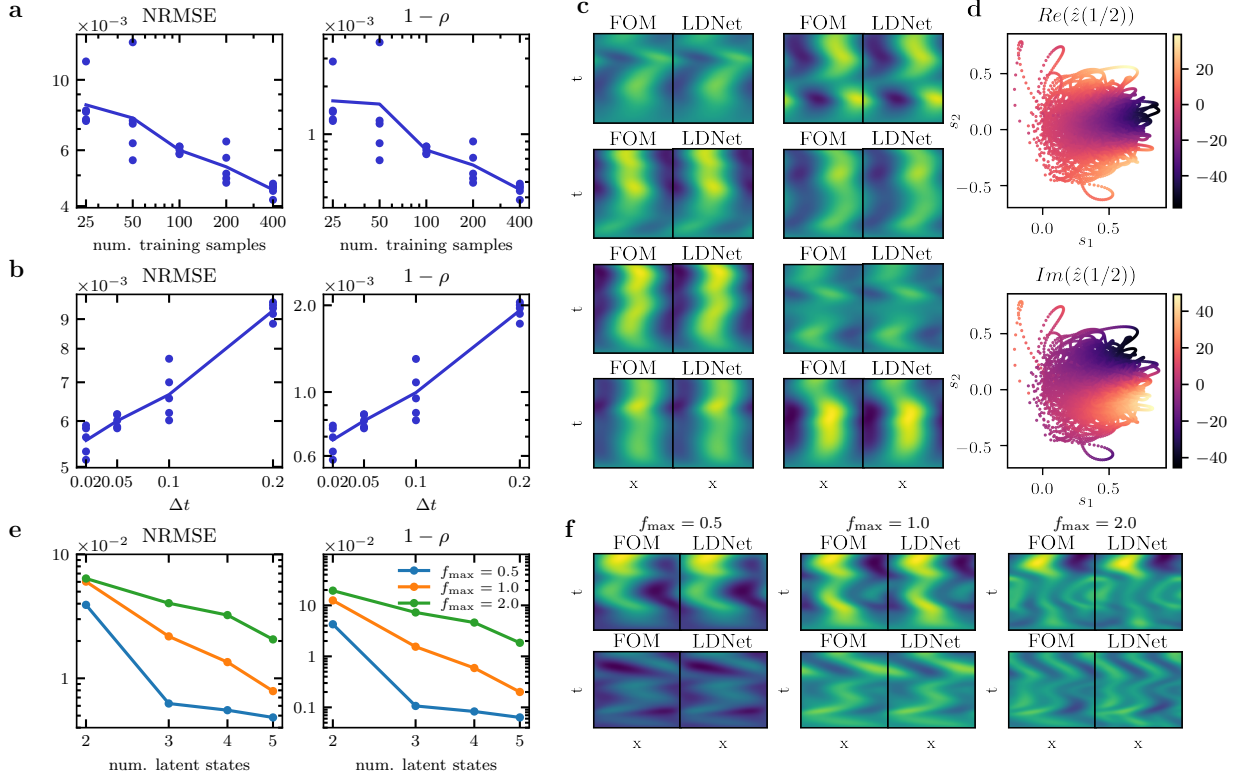


Figure 3: **Results of Test Case 1.** (a)-(b): Testing accuracy of Test Case 1b, as a function of the number of training samples (with $\Delta t = 0.05$) and of Δt (with 100 training samples). For each setting we run 5 training runs with random weights initialization. Each dot corresponds to a training run, while the solid line is the geometric mean. (c): FOM against LDNet predictions on 8 testing samples for Test Case 1b. The abscissa corresponds to space and the ordinate to time. (d): Mapping from the latent space trajectories and the Fourier space coefficients of the FOM solution for the testing samples of Test Case 1b. (e): Testing accuracy of Test Case 1c as a function of the number of latent states and of the maximum input frequency f_{\max} . (f): FOM against LDNet predictions on 2 testing samples for Test Case 1c, obtained by employing 5 latent states, for different maximum input frequencies (reported above the figure).




Figure 4: Test Case 1b. Trajectories in the latent space (s_1, s_2) of 24 testing samples obtained with four different LDNets, by starting from as many different initial guesses for the trainable parameters (each of them corresponding to different columns). In the first (respectively, second) row, each point in the latent space is colored according to the corresponding value of $Re(\hat{z}(0.5))$ (respectively, $Im(\hat{z}(0.5))$).

trajectories significantly differ among the four LDNets. This is not surprising, as latent states are internal variables in the LDNet, hidden within the input-output relationship. Nonetheless, a common pattern emerges in the connection between the latent states and the Fourier coefficients: each of the four trained LDNets underlies a well-defined relation between (s_1, s_2) and $(Re(\hat{z}(0.5)), Im(\hat{z}(0.5)))$. As a matter of fact, each LDNet discovers a different compact encoding for the FOM state, each of which underlying a relationship with the Fourier coefficients of the solution. In other terms, despite not being explicitly instructed to do that, the LDNet *discovers* an operator that is equivalent to the Fourier transform of the state. At the same time, the reconstruction NN ($\mathcal{NN}_{\text{rec}}$) *discovers* the inverse operator, as it is able to reconstruct the function $z(\cdot, t)$ from the two scalars (s_1, s_2) . Remarkably, this is obtained in a fully data-driven manner, without explicitly using any prior Fourier-based feature extraction.

3.1.4 Test Case 1c: infinite latent dimension

Finally, we consider a forcing term $f(x, t) = A(t) \cos(2\pi F(t)x - P(t))$, where $\mathbf{u}(t) = (A(t), F(t), P(t))$ is the time-dependent input signal. The forcing frequency $F(t)$ varies within an interval $[0.25, f_{\max}]$. Hence, the solution manifold of (8) has a potentially infinite dimension, being z the superimposition of a continuum of frequencies. Still, the results show that LDNets are able to discover effective low-dimensional encodings of the state. Specifically, to sample $F(t)$ we set $\tau = 1$, $f_{\min} = 0.25$ and f_{\max} as indicated below; for $A(t)$ we set $\mu = 1$, $\sigma = 1/3$ and $\tau = 1$; for $P(t)$ we set $\mu = 0$, $\sigma = 4/3$ and $\tau = 1$.

In this test case, we are interested in studying the impact of the number of latent states d_s on the LDNet accuracy, in three increasingly challenging cases, namely by setting $f_{\max} = 0.5$, $f_{\max} = 1$ and $f_{\max} = 2$ (see Fig. 5). In all the cases, we consider $d_s \in \{2, 3, 4, 5\}$. For each combination of f_{\max} and d_s , we retune the hyperparameters, in order to compare the (in principle) best accuracy attainable. With this purpose, we first run a preliminary hyperparameters tuning step, by considering a wide range of values (see Tab. 1, row *tuning*). Then, we shrink the variability by selecting the most (generally with respect to different combinations) promising area of the hyperparameter space, and we perform a fine tuning for each combination of f_{\max} and d_s independently (row *fine tuning*). The selected hyperparameters are listed in Tab. 1.

First, we set $f_{\max} = 0.5$ and we train LDNets for increasing number of latent states, from 2 to 5.

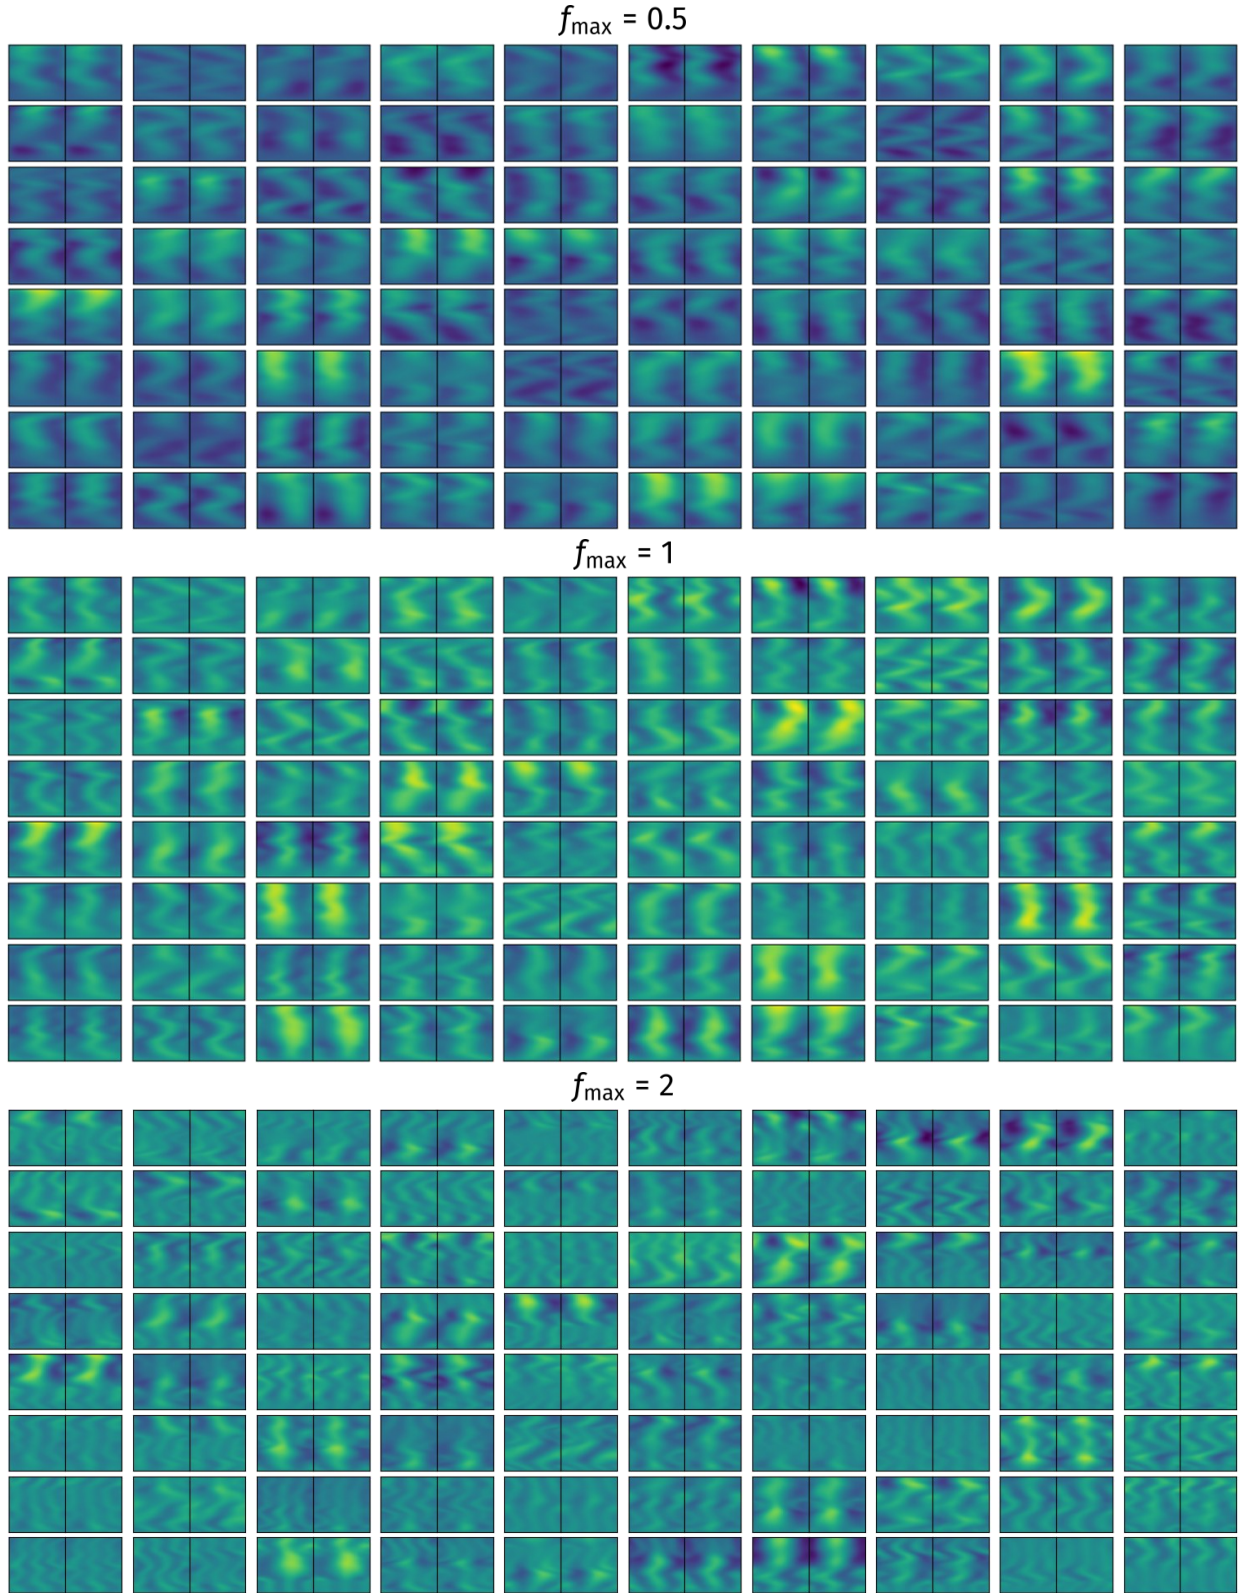


Figure 5: Test Case 1c. 80 testing samples for each f_{\max} value considered in this work (namely 0.5, 1 and 2), comparing the reference outputs \mathbf{y} (left) with the LDNet predictions $\tilde{\mathbf{y}}$ (right), for $d_s = 5$. For each sample, the horizontal axis refers to space, and the vertical axis refers to time.

Remarkably, as the number of latent states increases, LDNets discover more effective encodings, that reflect in an increasing prediction accuracy (Fig. 3e, blue line), with a saturation due to the relatively small training set size (100 samples). By increasing f_{\max} , the FOM state gets less prone to be represented by a compact encoding; still, as the number of latent states increases, the accuracy is enhanced.

3.2 Test Case 2: Unsteady Navier-Stokes

The 2D lid-driven cavity is a well-known benchmark problem in fluid dynamics [66], which may exhibit a wide range of flow patterns and vortex structures when increasing the Reynolds number. We challenge LDNets in learning an unsteady version of the lid-driven cavity problem, where the velocity prescribed on the top portion of the boundary (Γ_{top}) is a time-dependent input $u(t)$ (see Fig. 6a). In particular, as reported in Sec. 3.1.1, we consider a Gaussian Process Distribution with mean $\mu = 0$, standard deviation $\sigma = 5$ and characteristic time-scale $\tau = 5$. During the simulations, the Reynolds number varies over time by reaching peaks of nearly 1500. This problem is challenging also because of discontinuities in the velocity field at the two top corners. The goal here is to predict the velocity field (that is, we set $\mathbf{y}(\mathbf{x}, t) = \mathbf{v}(\mathbf{x}, t)$) for each prescribed $u(t)$:

$$\begin{aligned} \rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla) \mathbf{v} - \mu \Delta \mathbf{v} + \nabla p &= \mathbf{0} & \mathbf{x} \in \Omega, t \in (0, T], \\ \nabla \cdot \mathbf{v} &= 0 & \mathbf{x} \in \Omega, t \in (0, T], \\ \mathbf{v} &= u(t) \mathbf{e}_x & \mathbf{x} \in \Gamma_{\text{top}}, t \in (0, T], \\ \mathbf{v} &= \mathbf{0} & \mathbf{x} \in \partial\Omega \setminus \Gamma_{\text{top}}, t \in (0, T], \\ \mathbf{v} &= \mathbf{0} & \mathbf{x} \in \Omega, t = 0, \end{aligned} \tag{9}$$

where the dependence of the velocity \mathbf{v} and pressure p on space and time is understood. As shown in [58], a simple quadratic loss function is not adequate for capturing small vortex structures, because of their small impact, compared to medium- and large-scale structures, to the loss function. Therefore, we use the following goal-oriented metric, where we denote by \mathbf{v} and $\hat{\mathbf{v}}$ the reference and predicted velocities, respectively:

$$\mathcal{E}(\mathbf{v}, \hat{\mathbf{v}}) = \frac{\|\mathbf{v} - \hat{\mathbf{v}}\|^2}{v_{\text{norm}}^2} + \gamma \left\| \frac{\mathbf{v}}{\epsilon + \|\mathbf{v}\|} - \frac{\hat{\mathbf{v}}}{\epsilon + \|\hat{\mathbf{v}}\|} \right\|^2 \tag{10}$$

with hyperparameters γ and $\epsilon \ll 1$, and where v_{norm} is a reference velocity magnitude. Specifically, we pick $\gamma = 1 \cdot 10^{-1}$ and $\epsilon = 1 \cdot 10^{-4}$. The second term of the metric (10) allows to match the flow direction, even in the regions of small flow magnitude.

We generate training data through a FEM-based solver of (9), on a 100×100 triangular grid, accounting for nearly 91K degrees of freedom. We employ Taylor-Hood elements in the FEM solver, i.e. \mathbb{P}_2 Finite Elements for \mathbf{v} and \mathbb{P}_1 Finite Elements for p . We consider a semi-implicit time discretization with $\Delta t = 2 \cdot 10^{-1}$. The training and validation datasets consist of 80 simulations with $T = 20$, where 200 points are uniformly sampled in $\Omega = (0, 1)^2$. Regarding the testing data, we run 200 simulations with $T = 40$, a time span that is twice as long as the one of the training set, and we uniformly take 400 points in space from the domain $\Omega = (0, 1)^2$. To train LDNets, we take 100 evenly distributed snapshots in time, and we randomly take 200 points in space for each time-step.

We train three LDNets, by increasing the number of latent states from 1 to 5 and 10. Specifically, we fix the number of latent states and we perform hyperparameters tuning (see Sec. 2.8) by monitoring the following goal-oriented loss function (10). We define in Tab. 3 the initial ranges and final values of the hyperparameters for different numbers of latent states (1, 5 and 10).

We report in Tab. 4 the NRMSE and Pearson dissimilarity values on the testing set for $d_s = 1$, $d_s = 5$ and $d_s = 10$. We notice that the size of both \mathcal{N}_{dyn} and \mathcal{N}_{rec} architectures increases while the testing error of the LDNets decreases with respect to the number of latent states. Similarly to Test Case 1, as the number of latent states increases the LDNets are more and more efficient in discovering an effective compact representation of the system state and thus providing reliable predictions (Figs. 6b and 6c). Still, small NRMSEs are attained even with a small number of latent states.

Furthermore, we challenge the trained LDNets in predicting the flow evolution even on a longer time horizon than that considered in the training dataset (specifically, twice as long), that is for $t \in (20, 40]$,

Test cases	Hyperparameters					Trainable parameters		
	$\mathcal{NN}_{\text{dyn}}$ layers	neurons	$\mathcal{NN}_{\text{rec}}$ layers	neurons	Δt_{ref}	$\alpha_{\text{dyn}}, \alpha_{\text{rec}}$	$\mathcal{NN}_{\text{dyn}}$ # param.	$\mathcal{NN}_{\text{rec}}$ # param.
Test Case $d_s = 1$								
tuning	1 – 6	5 – 35	1 – 6	5 – 35	$[10^{-1}, 10^1]$	0		
final	2	7	4	24	5.4	0	85	1'970
Test Case $d_s = 5$								
tuning	1 – 6	5 – 35	1 – 6	5 – 35	$[10^{-1}, 10^1]$	0		
final	2	27	4	33	8.6	0	1'085	3'731
Test Case $d_s = 10$								
tuning	1 – 6	5 – 35	1 – 6	5 – 35	$[10^{-1}, 10^1]$	0		
final	3	19	5	28	9.2	0	1'188	3'698

Table 3: Test Case 2: hyperparameters ranges and selected values. See text for details.

Number of latent states	1	5	10
$\text{NRMSE}_{\text{test}} (0 < t < 40)$	$4.08 \cdot 10^{-3}$	$1.94 \cdot 10^{-3}$	$1.39 \cdot 10^{-3}$
$\text{NRMSE}_{\text{test}} (0 < t < 20)$	$4.10 \cdot 10^{-3}$	$1.93 \cdot 10^{-3}$	$1.34 \cdot 10^{-3}$
$\text{NRMSE}_{\text{test}} (20 < t < 40)$	$4.07 \cdot 10^{-3}$	$1.95 \cdot 10^{-3}$	$1.44 \cdot 10^{-3}$
$1 - \rho_{\text{test}} (0 < t < 40)$	$5.14 \cdot 10^{-3}$	$1.16 \cdot 10^{-3}$	$6.01 \cdot 10^{-4}$
$1 - \rho_{\text{test}} (0 < t < 20)$	$5.48 \cdot 10^{-3}$	$1.23 \cdot 10^{-3}$	$6.00 \cdot 10^{-4}$
$1 - \rho_{\text{test}} (20 < t < 40)$	$4.81 \cdot 10^{-3}$	$1.11 \cdot 10^{-3}$	$6.02 \cdot 10^{-4}$

Table 4: Test Case 2: test accuracy metrics for LDNets trained with an increasing number of latent states (1, 5 and 10).

when the training samples are generated for $t \in [0, 20]$. Remarkably, we observe a negligible propagation of the approximation error along the prolonged time frame, making the trained LDNets reliable also for time-extrapolation (Figs. 6b and 6d, Tab. 4). This is a remarkable achievement, considering that the dynamical system at hand does not present a periodic or quasi-periodic regime.

We depict in Fig. 7 the time evolution of the streamlines associated with a specific sample belonging to the testing set, along with the corresponding input signal $u(t)$. We see that the approximation provided by LDNets with 10 latent variables has an excellent agreement with the FOM.

3.3 Test Case 3: Aliev-Panfilov electrophysiology model

We consider the Aliev-Panfilov (AP) model [3], a nonlinear system of PDEs describing the propagation of the electrical potential $z(x, t)$ in an excitable tissue. The AP model envisages a recovery variable $w(x, t)$ that tracks the refractoriness of the tissue by modulating the repolarization phase. The model, supplemented with homogeneous Neumann boundary conditions (encoding electrical insulation) and zero initial conditions for both the variables, reads

$$\begin{aligned}
\frac{\partial z}{\partial t} - D \frac{\partial^2 z}{\partial x^2} &= Kz(1-z)(z-\alpha) - zw + I_{\text{stim}} & x \in (0, L), t \in (0, T], \\
\frac{\partial w}{\partial t} &= \left(\gamma + \frac{\mu_1 w}{\mu_2 + z} \right) (-w - Kz(z-b-1)) & x \in (0, L), t \in (0, T], \\
\frac{\partial z(0, t)}{\partial x} &= \frac{\partial z(L, t)}{\partial x} = 0 & t \in (0, T], \\
z(x, 0) &= w(x, 0) = 0, & x \in (0, L).
\end{aligned} \tag{11}$$

with parameters $D = 0.1 \text{ mm ms}^{-1}$, $K = 8$, $\alpha = 0.1$, $\gamma = 0.02$, $\mu_1 = 0.2$, $\mu_2 = 0.3$, $b = 0.15$, $L = 100 \text{ mm}$, $T = 500 \text{ ms}$. Note that z is a non-dimensional potential (which can be mapped to its physiological values by the relationship $(100z - 80) \text{ mV}$) and the model is rescaled with respect to the time constant $\tau = 12.9 \text{ ms}$ (for further details, see [27]).

The excitation-propagation process is triggered by an external stimulus $I_{\text{stim}}(x, t)$, applied at two stimulation points, respectively located at $x = 1/4L$ and $x = 3/4L$, and consisting of square impulses, to mimic the action of a (natural or artificial) pacemaker. The input of the model is given by $\mathbf{u}(t) = (I_{\text{stim}}(x_1^{\text{stim}}, t), I_{\text{stim}}(x_2^{\text{stim}}, t))$. To generate the training samples, we randomly trigger the applied stimuli, either in correspondence of x_1^{stim} , x_2^{stim} or both points, by randomly picking the stimulation times. For the approximation of the AP model, we employ the finite difference method both in space and time, on a regular grid with 800 points in space and 10^5 time steps. Then, we subsample the space-time grid by retaining 100 points in space and 500 time instants.

The AP model solution features the fast-slow dynamics of a cardiac action potential (steep depolarization fronts followed by slow repolarization of the electrical potential to its resting value) and the wavefront propagation in space generating collisions of waves from different stimulation points. These features make this problem a challenging test case for comparing the proposed method against popular approaches to learning space-time dynamics of complex systems.

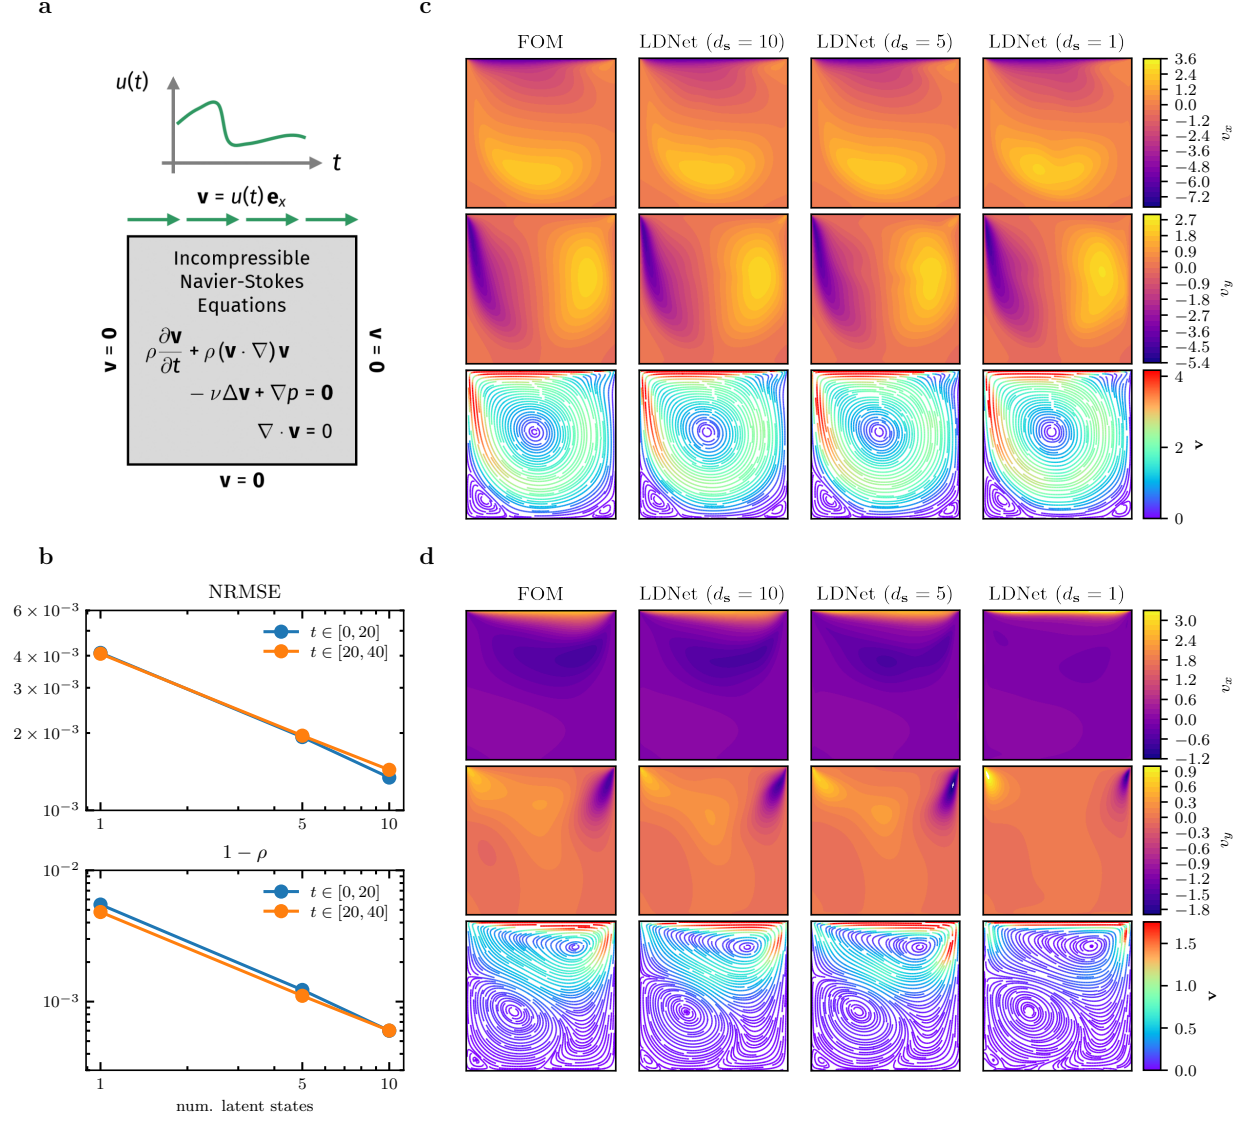
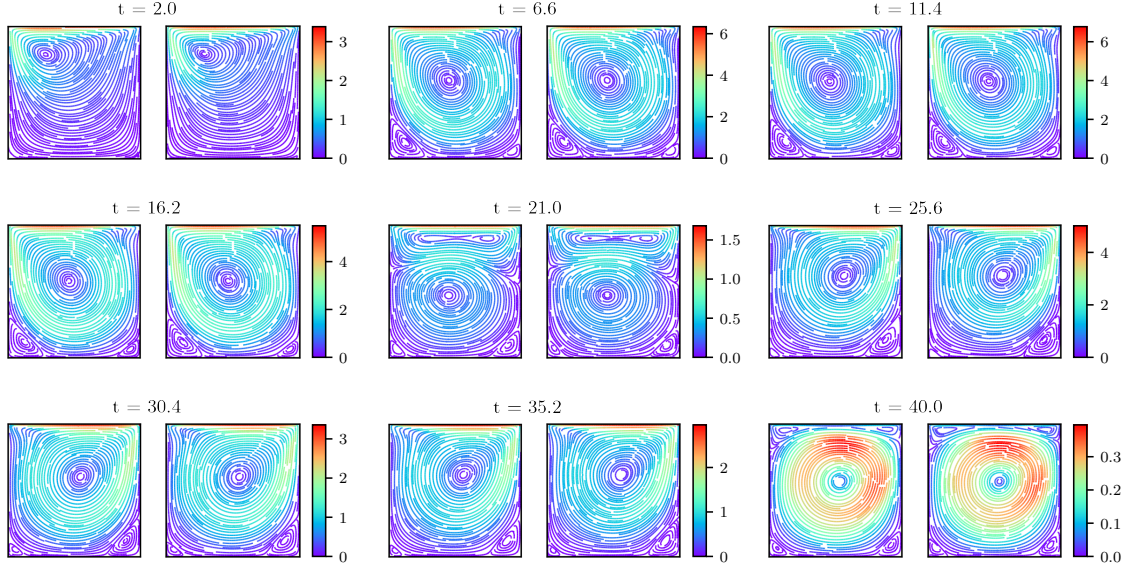


Figure 6: Test Case 2. (a): Computational domain and equations of the FOM. (b): Error metrics (NRMSE and Pearson dissimilarity) of LDNets for different number of latent variables (1, 5 and 10). The training dataset consists of 80 simulations with $T = 20$, while the test dataset comprises 200 simulations with $T = 40$. The blue lines refer to the testing error obtained in the interval $t \in [0, 20]$ (that is the same interval seen during training), while orange lines refer to the testing error in the interval $t \in [20, 40]$. (c): A snapshot of the velocity field within the interval $t \in [0, 20]$ (interpolation interval) of a testing sample. (d): A snapshot of the velocity field within the interval $t \in [20, 40]$ (extrapolation interval) of a testing sample.

a



b




Figure 7: Test Case 2. **(a)**: Streamlines of the velocity field \mathbf{v} for one testing sample over time ($t \in [0, 40]$). Horizontal and vertical axis refer to space in the domain $\Omega = (0, 1)^2$. For each subplot, the left one represents the FOM while the right one represents the LDNet approximation for $d_s = 10$. **(b)**: Input signal $u(t)$ over time applied to the top portion of the boundary Γ_{top} .

	Hyperparameters							
	d_s	$\mathcal{NN}_{\text{enc}}, \mathcal{NN}_{\text{dec}}$		$\mathcal{NN}_{\text{dyn}}$		Δt_{ref}	$\alpha_{\text{enc}}, \alpha_{\text{dec}}$	α_{dyn}
		layers	neurons	layers	neurons			
AE/ODE								
tuning	4 – 12	1 – 4	4 – 80	1 – 10	4 – 40	$[10^1, 10^3]$	$[10^{-5}, 10^{-2}]$	$[10^{-5}, 10^{-1}]$
final	12	1	75	4	38	$1.80 \cdot 10^1$	$6.89 \cdot 10^{-3}$	$2.80 \cdot 10^{-2}$
AE/LSTM								
tuning	4 – 12	1 – 4	4 – 80				$[10^{-5}, 10^{-2}]$	$[10^{-5}, 10^0]$
final	12	1	75				$6.89 \cdot 10^{-3}$	$4.81 \cdot 10^{-1}$

Table 5: Test Case 3: hyperparameters ranges and selected values for the AE/ODE and AE/LSTM models. For the encoder (respectively, the decoder) the number of neurons refers to the first (respectively, last) hidden layer. In the other layers, the number of neurons is linearly varied to connect the first (respectively, last) hidden layer to the layer of latent variables.

	Hyperparameters					
	d_s	$\mathcal{NN}_{\text{dyn}}$		$\mathcal{NN}_{\text{rec}}$		Δt_{ref}
		layers	neurons	layers	neurons	$\alpha_{\text{dyn}}, \alpha_{\text{rec}}$
LDNet						
tuning	4 – 12	1 – 3	4 – 15	1 – 5	4 – 20	$[10^1, 10^3]$
final	12	1	5	5	17	$2.05 \cdot 10^2$

Table 6: Test Case 3: hyperparameters ranges and selected values for the LDNet.

We compare LDNets with state-of-the-art approaches in which dimensionality reduction is achieved by training an auto-encoder (AE) on a discrete representation of the output $z(\cdot, t)$. Once trained, the encoder is employed to compute the trajectories of the latent states throughout the training set, and the dynamics in the latent space is learned either through an ODE-Net [24] or an LSTM [32]. We denote the resulting models by AE/ODE and AE/LSTM, respectively. Then, we further train the NN that tracks the dynamics of the latent states simultaneously with the decoder, that is in an end-to-end (e2e) fashion, and we denote the resulting models by AE/ODE-e2e and AE/LSTM-e2e, respectively. Furthermore, we benchmark LDNets against a classical method of model-order reduction of PDE models, namely the POD-DEIM method [45, 60]. These methods are described in detail in Appendix A.

We challenge LDNets and the above-mentioned methods in the task of predicting the space-time dynamics of the target value $\mathbf{y}(x, t) = z(x, t)$, given the time series of impulses in the two stimulation points. In order to ensure a fair comparison, for all the methods that require a choice of hyperparameters we use an algorithm for their automatic tuning, as described in Sec. 2.8. The ranges used for tuning and the final hyperparameter values are reported in Tab. 5 and 6. In this test case, we employ the technique described in Sec. 2.5.1 to impose in a strong manner the equilibrium condition of the initial state when $\mathbf{u}(t) = \mathbf{0}$. To achieve a significant dimensionality reduction, we set a maximum number of 12 latent variables both for auto-encoder-based methods and for LDNets. We notice that the hyperparameter tuning algorithm selects the maximum number of latent states (i.e. $d_s = 12$) for all the methods. For the POD-DEIM method, we test different number d_s of POD modes (i.e., basis functions) for the state and the nonlinear term, ranging from 12 to 60.

In Tab. 7 we report training and testing errors obtained for the different methods, along with the number of trainable parameters. The results of this comparison are also reported in Figs. 8–9.

We observe that the results obtained with the POD-DEIM method using 12 modes are unsatisfactory, compared to the other methods, confirming the importance of adopting nonlinear dimensional reduction techniques for this class of traveling-front problems. Indeed, due to the presence of traveling fronts, this problem features a slow decay of the Kolmogorov n -width [52], that reflects in a poor accuracy of the electrical potential reconstruction given by the POD-DEIM method when 12 modes are used. Then, we consider higher number of modes, namely 24, 36, 48 and 60. We report in Fig. 10, Fig. 11 and Tab. 7

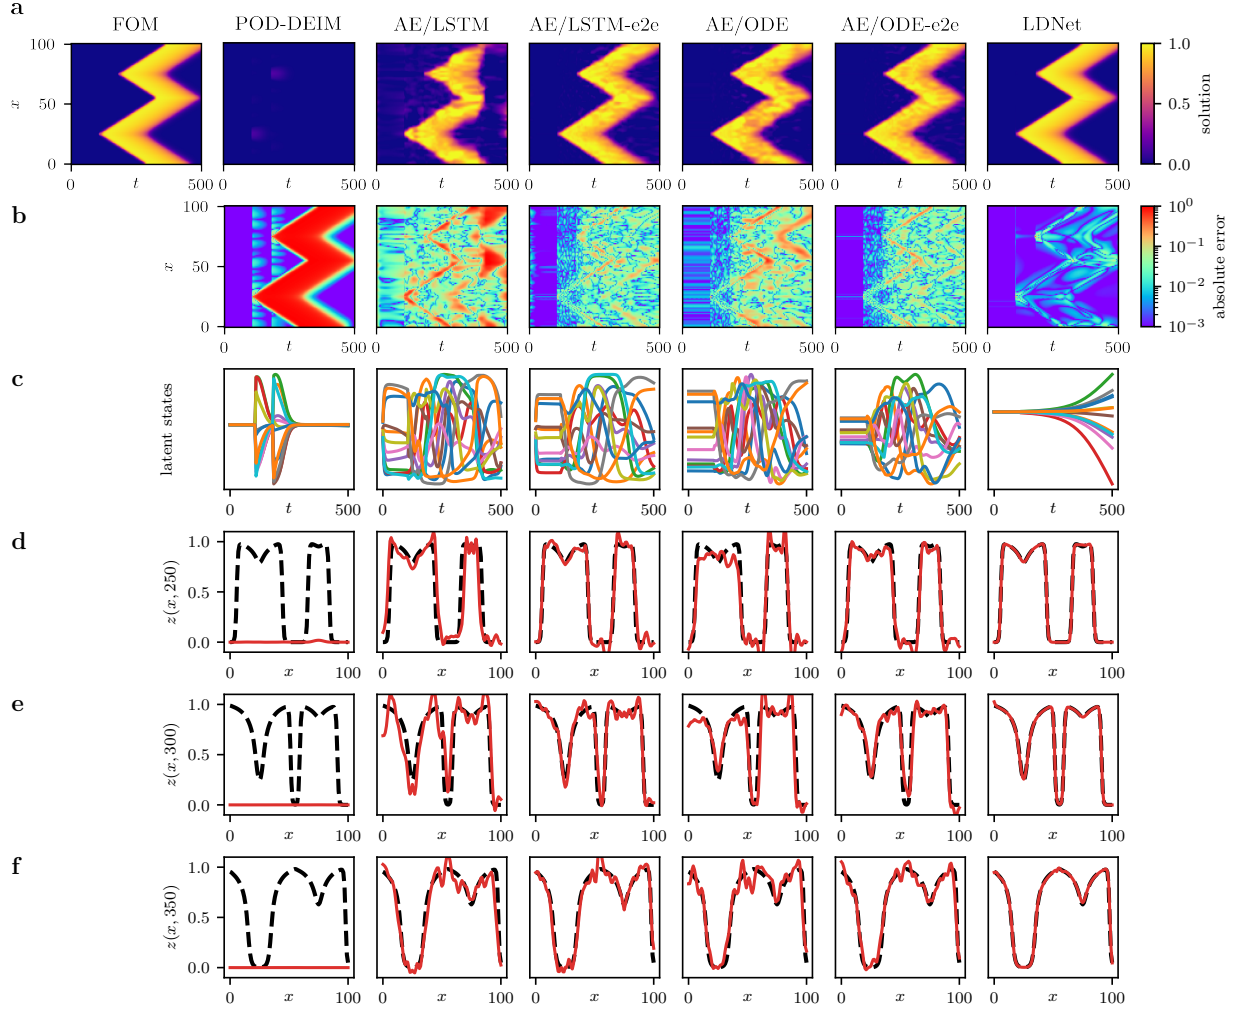


Figure 8: **Test Case 3: Methods comparison.** We compare the results obtained with different methods for a sample belonging to the test dataset. The left-most column reports the FOM solution of the AP model (the abscissa denotes time, the ordinate denotes space). For each method we report: **(a)** the space-time solution; **(b)** the space-time error with respect to the FOM solution; **(c)** the time-evolution of the 12 latent variables; **(d)-(e)-(f)** three snapshots of the space-dependent output field at $t = 250, 300$ and 350 , in which we compare the predicted solution (red solid line) with the FOM solution (black dashed line).

	NRMSE		Number of trainable parameters			
	training	testing	$\mathcal{NN}_{\text{enc}}$	$\mathcal{NN}_{\text{dec}}, \mathcal{NN}_{\text{rec}}$	$\mathcal{RNN}_{\text{dyn}}$	total
POD-DEIM ($d_s = 12$)	$4.05 \cdot 10^{-1}$	$3.92 \cdot 10^{-1}$				
POD-DEIM ($d_s = 24$)	$3.59 \cdot 10^{-1}$	$3.47 \cdot 10^{-1}$				
POD-DEIM ($d_s = 36$)	$1.71 \cdot 10^{-1}$	$1.62 \cdot 10^{-1}$				
POD-DEIM ($d_s = 48$)	$7.48 \cdot 10^{-2}$	$7.57 \cdot 10^{-2}$				
POD-DEIM ($d_s = 60$)	$2.97 \cdot 10^{-2}$	$2.90 \cdot 10^{-2}$				
AE/LSTM	$1.90 \cdot 10^{-1}$	$1.98 \cdot 10^{-1}$	8'562	8'651	720	17'933
AE/LSTM-e2e	$2.05 \cdot 10^{-2}$	$5.87 \cdot 10^{-2}$	8'562	8'651	720	17'933
AE/ODE	$2.09 \cdot 10^{-2}$	$4.58 \cdot 10^{-2}$	8'562	8'651	5'484	22'697
AE/ODE-e2e	$1.78 \cdot 10^{-2}$	$3.37 \cdot 10^{-2}$	8'562	8'651	5'484	22'697
LDNet	$7.09 \cdot 10^{-3}$	$7.37 \cdot 10^{-3}$	0	1'480	228	1'708

Table 7: Test Case 3: Training and testing errors obtained with the different methods and number of trainable parameters.

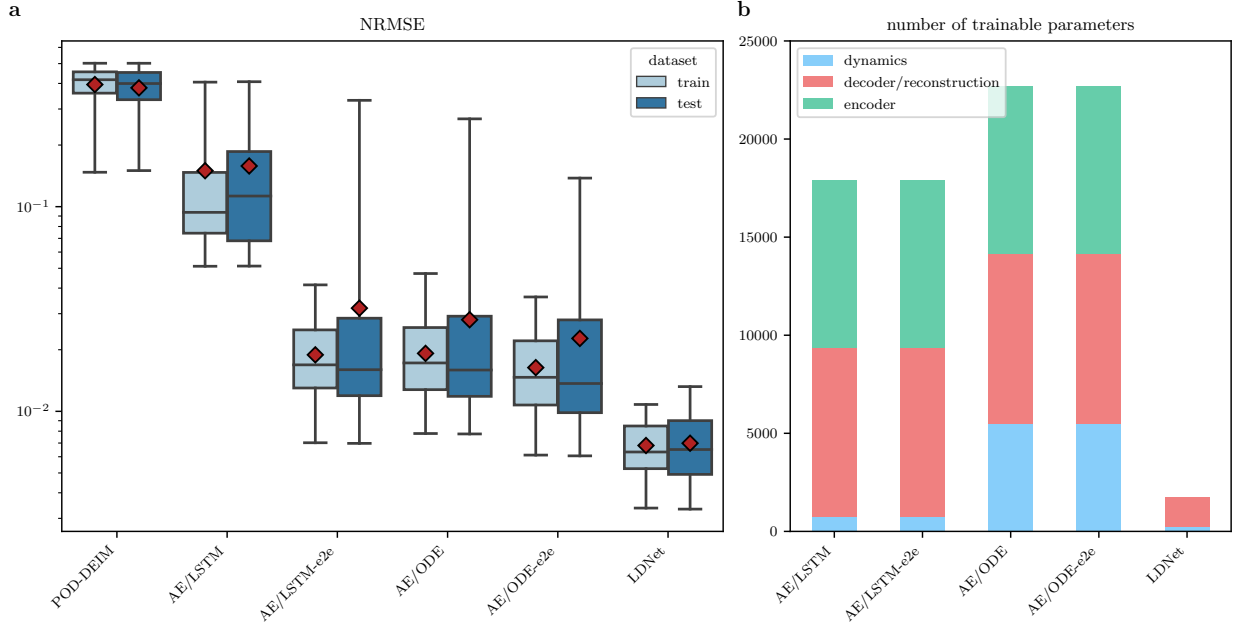


Figure 9: **Results of Test Case 3.** (a): Boxplots of the distribution of the testing (blue) and training (light blue) errors obtained with each method. The red diamonds represent the average error on each dataset. (b): Number of trainable parameters of each method. The bin *encoder* is present only for auto-encoder-based methods, but not for LDNets. The bin *dynamics* refers to the NN that evolves the latent states. The POD-DEIM method is not included, as it does not envisage a training stage.

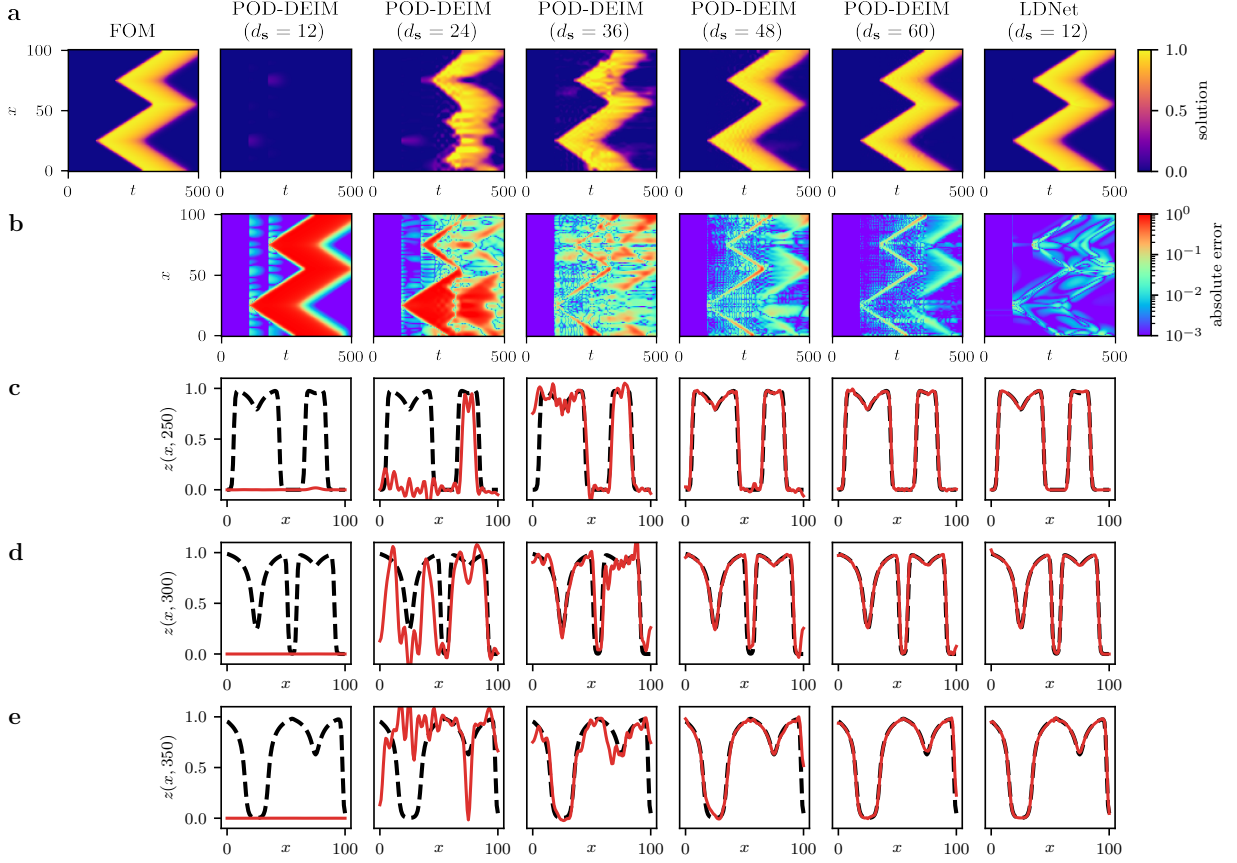


Figure 10: We compare the results obtained with the POD-DEIM method, for an increasing number of considered modes (reported in the titles), against the results obtained with our proposed method. The figure shows the predictions obtained for a sample belonging to the test dataset. The left-most column reports the FOM solution of the AP model (the abscissa denotes time, the ordinate denotes space). **(a)** the space-time solution; **(b)** the space-time error with respect to the FOM solution; **(c)-(d)-(e)** three snapshots of the space-dependent output field at $t = 250, 300$ and 350 , in which we compare the predicted solution (red solid line) with the FOM solution (black dashed line).

the results. By increasing the number of considered modes, the accuracy increases. With 60 modes, e.g., the POD-DEIM method is more accurate than auto-encoder-based methods with 12 latent variables, while LDNets with 12 latent variables are still more accurate. Nonetheless, we observe that an increase of number of modes in the POD-DEIM method is accompanied by a significant increase in the computational cost, since the computational complexity scales with the cube of d_s . In addition, the POD-DEIM model is solved on the same temporal discretization as the high-fidelity one, which represents a strong limitation with respect to the other methods described in this paper.

A better accuracy is achieved by both auto-encoder-based methods and by LDNets, thanks to their ability to express a nonlinear relationship between the latent states and the solution. Still, LDNet outperforms the other methods, with a testing NRMSE equal to $7 \cdot 10^{-3}$ and with a much lower overfitting. The testing NRMSE of auto-encoder-based methods is nearly 5 times larger than with LDNets or more. Remarkably, our method achieves better accuracy with significantly fewer trainable parameters: auto-encoder-based methods require more than tenfold the number of parameters. This testifies to the good architectural design of LDNets.




Figure 11: We compare the results obtained with the POD-DEIM method, for an increasing number of considered modes, against the results obtained with our proposed method. Specifically, we report boxplots of the distribution of the testing (blue) and training (light blue) errors obtained with each method. The red diamonds represent the average error on each dataset.

Discussion

We have presented LDNets, a novel class of NNs that learn in a data-driven manner the evolution of systems exhibiting spatio-temporal dynamics in response to external inputs. An LDNet is trained in a supervised way from observations of input-output pairs, which can either come from experimental measurements or be synthetically generated through the numerical approximation of mathematical models of which one seeks a surrogate or reduced-order model. This latter case is the one considered in this manuscript to demonstrate the capabilities of the proposed method.

LDNets provide a paradigm-shift from state-of-the-art methods based on dimensionality reduction (e.g., exploiting POD or auto-encoders) of a high-dimensional discretization of the system state. Specifically, LDNets automatically discover a compact representation of the system state, without necessitating the explicit construction of an encoder. This enables the training algorithm to select a compact representation of the state that is functional not only in reconstructing the space-dependent field for each time instant, but also in predicting its dynamics; an auto-encoder, conversely, when trained, extracts features on a purely statistical basis, being agnostic of the importance of each feature in determining the evolution of the system.

Unlike standard approaches that reconstruct a high-dimensional discretization of the output, corresponding e.g. to evaluations at the vertices of a computational mesh, our approach is in this sense meshless. The reconstruction NN is indeed queried for each point in space independently. This design principle gives LDNets several benefits. First, the meshless nature of LDNets combined with the automatic discovery of the latent space allows them to operate in a low-dimensional space without ever going through a high-dimensional discretization, as auto-encoder-based methods do. This makes LDNets very lightweight structures, easy to train, and not prone to overfitting. The LDNet architecture enables the sharing of the trainable parameters needed to evaluate the solution at different points (that is, the same weights are employed regardless of the query point). The low overfitting of LDNets is thus not surprising, as weight-sharing is often the key of good generalization properties of many architectures, such as CNNs and RNNs [28]. Second, it provides a continuous representation of the output, and, thus, allows for additional and possibly physics-informed terms to be introduced into the loss function [54], opening up countless possibilities for extending the purely

black-box method proposed in this paper to grey-box approaches. Third, the loss function can be defined by stochastically varying the points in space at which the error is evaluated (see Test Case 2), thus lightening the computational burden associated with training. Note that this is not possible when the model returns the entire batch of observations. This aspect also opens up to multiple developments, such as stochastic, minibatch-based training algorithms, or even adaptive refinements of the evaluation points, by sampling more densely where the error is larger.

The time-dynamics of LDNets is based on a recurrent architecture that is consistent, by construction, with the arrow of time. This differentiates LDNets from other approaches in which time is seen as a parameter [26], or approaches, based e.g. on DeepONets, that take as input the entire time-history of $\mathbf{u}(t)$ with a fixed length [40, 44]. The latter approaches do not easily allow for predictions over time frames longer than those used during training, or allow for time-extrapolation only in periodic or quasi-periodic problems [65]. LDNets, on the other hand, allow predictions for arbitrarily long times. We remark that the reliability of time-extrapolation is constrained by the characteristics of the problem at hand and the available training data. For example, if the system is characterized by a divergent behavior such that, as time progresses, the state enters regions increasingly distant from the initial condition, then the reliability of the predictions is not guaranteed in time-extrapolation regimes. When the system state remains bounded, however, the predictions of LDNets are significantly accurate even in time-extrapolation regimes, as showcased in Test Case 2.

LDNets represent, as proved by the results of this work, an innovative tool capable of learning spatio-temporal dynamics with great accuracy and by using a remarkably small number of trainable parameters. They are able to discover, simultaneously with the system dynamics, compact representations of the system state, as shown in Test Case 1 where the Fourier transform of a sinusoidal signal is automatically discovered. Once trained, LDNets provide predictions for unseen inputs with negligible computational effort (order of milliseconds for the considered Test Cases). LDNets provide a new flexible and powerful tool for data-driven emulators that is open to a wide range of variations in the definition of the loss function (like, e.g., including physics-informed terms), in the training strategies, and, finally, in the NN architectures. The comparison with state-of-the-art methods on a challenging problem, such as predicting the excitation-propagation pattern of a biological tissue in response to external stimuli, highlights the full potential of LDNets, which outperform the accuracy of existing methods while still using a significantly lighter architecture.

Acknowledgements

FR, SP and LD are members of the INdAM research group GNCS. This project has been partially supported by the INdAM GNCS Project CUP_E55F22000270001. LD acknowledges the support of the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program (Italy) within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence).

Appendices

A Alternative methods

In this appendix, we describe alternative methods to LDNets, whose performance is compared in Sec. 2. These methods are based on a space discretization of the solution field.

We consider hence a space-discretization operator $\mathcal{A}: \mathcal{Y} \rightarrow \mathbb{R}^{N_h}$, with $N_h \gg 1$. The operator \mathcal{A} typically consists of a point-wise evaluation of \mathbf{y} on a grid of points (e.g. the nodes of a computational mesh), or of an expansion with respect to a Finite Element basis or to a Fourier basis. The subscript h refers to the characteristic size associated with the discretization (e.g., the mesh element size for a mesh-based discretization; the sampling period for a discretization based on the discrete Fourier transform). Hence, we write $\mathbf{Y}(t) = \mathcal{A}(\mathbf{y}(\cdot, t)) \in \mathbb{R}^{N_h}$. The space-discretization operator is typically accompanied by a discrete-to-continuous operator $\mathcal{A}': \mathbb{R}^{N_h} \rightarrow \mathcal{Y}$, such that $\mathbf{y}(\cdot, t) \simeq \mathcal{A}'(\mathbf{Y}(t))$.

Let us consider, for simplicity, the case when the discretization operator is associated with the evaluation

of the output field on a grid of points $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \dots, \boldsymbol{\xi}_{N_{\text{nodes}}}$. Then we have

$$\mathbf{Y}_i(\tau) = (\mathbf{y}_i(\boldsymbol{\xi}_1, \tau), \mathbf{y}_i(\boldsymbol{\xi}_2, \tau), \dots, \mathbf{y}_i(\boldsymbol{\xi}_{N_{\text{nodes}}}, \tau)) \quad \text{for } i \in \mathcal{S}_{\text{train}}, \tau \in \mathcal{T}^i$$

In this case, we have $N_h = d_{\mathbf{y}} N_{\text{nodes}}$. We are thus restricting ourselves to the case when the training dataset is such that $\mathcal{P}_{\tau}^i \equiv \mathcal{P} := (\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \dots, \boldsymbol{\xi}_{N_{\text{nodes}}})$. If, instead, observation points vary between samples or between time steps, interpolation would be required.

A.1 Projection-based reduced-order models

Projection-based reduced-order models (ROMs) allow the efficient simulation of complex spatio-temporal systems like (1) for many different queries of the input $\mathbf{u}(t)$ [4, 8–10, 31, 51]. They leverage Galerkin or Petrov-Galerkin projection of the PDE system onto problem-specific linear manifolds to reduce the complexity of high-fidelity discretization, deliberately rich in degrees of freedom for the sake of accurately approximating the PDE solution. Here, to ease the notation, we only consider the case in which the model output coincides with the FOM state ($\mathbf{y} \equiv \mathbf{z}$), that is the observation operator is the identity function. The case when the operator differs from the identity is a straightforward generalization of the one considered here.

Projection-based ROMs are intrusive, as they require manipulation of the discretization of FOM (1). Their construction first involves dimensionality reduction techniques applied to pre-computed samples of the discretized FOM state:

$$\mathbf{Z}(t) = \mathcal{A}(\mathbf{z}(\cdot, t)) \in \mathbb{R}^{N_h},$$

collected by numerically approximating the space-discrete counterpart of the FOM, for different samples of the input $\mathbf{u}(t)$ in the training set:

$$\begin{cases} \frac{d}{dt} \mathbf{Z}(t) = \mathbf{F}(\mathbf{Z}(t), \mathbf{u}(t)) & \text{for } t \in (0, T] \\ \mathbf{Z}(t) = \mathbf{Z}_0 := \mathcal{A}(\mathbf{z}_0), \end{cases} \quad (12)$$

where \mathbf{F} is a suitable discretization of the differential operator \mathcal{F} . Solution snapshots are stored into a matrix

$$Z = \{\mathcal{A}(\mathbf{z}_i(\cdot, \tau))\}_{i \in \mathcal{S}_{\text{train}}}^{\tau \in \mathcal{T}^i} \in \mathbb{R}^{N_h \times N_{\text{snapshots}}}$$

from which $d_{\mathbf{s}} \ll N_h$ basis functions $\boldsymbol{\phi}_k$, $k = 1, \dots, d_{\mathbf{s}}$ are extracted by proper orthogonal decomposition (POD), or other linear dimensionality reduction techniques, like reduced-basis greedy algorithm [31, 52]. This allows expressing the discretized state compactly as

$$\mathbf{Z}(t) \approx V \mathbf{s}(t) = \sum_{k=1}^{d_{\mathbf{s}}} s_k(t) \boldsymbol{\phi}_k, \quad (13)$$

where $V = [\boldsymbol{\phi}_1 | \dots | \boldsymbol{\phi}_{d_{\mathbf{s}}}]$ is the transformation matrix, collecting the basis functions in its columns, and $\mathbf{s}(t) = [s_1(t), \dots, s_{d_{\mathbf{s}}}(t)]^T$ is the reduced state, a vector made of the coefficients s_k associated with each basis function. In this context, the solution manifold associated with the latent space is thus the column space of the matrix V .

Given a test manifold spanned by $d_{\mathbf{s}}$ basis functions $\{\boldsymbol{\psi}_k\}_{k=1}^{d_{\mathbf{s}}}$ (collected into a matrix $W = [\boldsymbol{\psi}_1 | \dots | \boldsymbol{\psi}_{d_{\mathbf{s}}}]$), the projection of the dynamical system (12) generates a system of $d_{\mathbf{s}}$ equations for the reduced state of the form:

$$\begin{cases} \frac{d}{dt} \mathbf{s}(t) = W^T \mathbf{F}(V \mathbf{s}(t), \mathbf{u}(t)) & \text{for } t \in (0, T] \\ \mathbf{s}(0) = \mathbf{s}_0 := V^T \mathbf{Z}_0, \end{cases} \quad (14)$$

In this formulation, computational efficiency hinges on the form of the discretized operator \mathbf{F} . In the linear and affine cases, the dependence on the solution can be separated from the operator, generating a set or pre-computable terms that ensure computational efficiency in the resolution.

In the case nonlinearities are present in the model, instead, \mathbf{F} requires the reduced state to be projected back onto the state space for its evaluation, before being projected onto the low-dimensional dynamics

of (14). This double projection makes the computational costs of (14) comparable to the ones of (12). Hyper-reduction techniques based on interpolation and linear dimensionality reduction techniques provide a computationally efficient alternative for evaluating nonlinearities [23, 31, 52]. In our test case, we first separate the nonlinear part from the linear one:

$$\mathbf{F}(\mathbf{s}(t), \mathbf{u}(t)) = \mathbf{F}_{nl}(\mathbf{s}(t), \mathbf{u}(t)) + \mathbf{F}_l(\mathbf{s}(t), \mathbf{u}(t)),$$

and we then rely on the discrete version of the empirical interpolation method (DEIM), which employs an interpolation-based projection of the nonlinearity onto the span of sparsely sampled basis functions, precomputed with POD [23]. Specifically, the non-linear term \mathbf{F}_{nl} is approximated by

$$\mathbf{F}_{nl}(\mathbf{s}(t), \mathbf{u}(t)) \approx U(P^T U)^{-1} P^T \mathbf{F}_{nl}(\mathbf{s}(t), \mathbf{u}(t)),$$

being U the transformation matrix, collecting POD basis functions of the nonlinear term, and P a sparse matrix that samples \mathbf{F}_{nl} on a subset of interpolation indices. Here, we denote this model-order reduction method as POD-DEIM method.

The success and greatest limitation of projection-based ROM lies in being intrusively linked to high-fidelity discretization via a linear subspace. This, on the one hand, ensures the ROM consistency, with a reduced state directly mappable to the system state thanks to (13), and convergence to the FOM as the dimensionality of the reduced manifold increases. On the other hand, if the solution of the parametric problem entails large variability, as in the case of advection-dominated problems, the accuracy-efficiency balance is compromised by the curse of dimensionality.

A.2 Methods based on auto-encoders

For the reasons above, many researchers have shifted to the development of reduced-order models based on non-linear dimensionality reduction techniques, such as NN-based auto-encoders [19, 26, 35, 41, 44, 64]. In this framework, to reduce the dimensionality of the discretized output field \mathbf{Y} , we train an auto-encoder, with latent code dimension d_s , by considering the following minimization problem:

$$\begin{aligned} (\mathbf{w}_{\text{enc}}^*, \mathbf{w}_{\text{dec}}^*) = \underset{\mathbf{w}_{\text{enc}}, \mathbf{w}_{\text{dec}}}{\operatorname{argmin}} \quad & \sum_{i \in \mathcal{S}_{\text{train}}} \sum_{\tau \in \mathcal{T}^i} \|\mathcal{NN}_{\text{dec}}(\mathcal{NN}_{\text{enc}}(\mathbf{Y}_i(\tau); \mathbf{w}_{\text{enc}}); \mathbf{w}_{\text{dec}}) - \mathbf{Y}_i(\tau)\|_{\mathbf{Y}}^2 \\ & + \alpha_{\text{enc}} \mathcal{R}(\mathbf{w}_{\text{enc}}) + \alpha_{\text{dec}} \mathcal{R}(\mathbf{w}_{\text{dec}}) \end{aligned}$$

where $\mathcal{NN}_{\text{enc}}$ and $\mathcal{NN}_{\text{dec}}$ (with trainable parameters \mathbf{w}_{enc} and \mathbf{w}_{dec} , respectively) are the encoder and the decoder, respectively. For normalization purposes, to evaluate distances in the discrete space we employ the rescaled euclidean norm $\|\cdot\|_{\mathbf{Y}}^2 := \|\cdot\|^2 / (N_{\text{nodes}} y_{\text{norm}}^2)$.

Once the auto-encoder is trained, i.e. the parameters $\mathbf{w}_{\text{enc}}^*$ and $\mathbf{w}_{\text{dec}}^*$ are available, we compute the latent codes associated with each training sample $i \in \mathcal{S}_{\text{train}}$ and each observation time $\tau \in \mathcal{T}^i$, that we denote as

$$\hat{\mathbf{s}}_i(\tau) = \mathcal{NN}_{\text{enc}}(\mathbf{Y}_i(\tau); \mathbf{w}_{\text{enc}}^*) \in \mathbb{R}^{d_s}.$$

At this stage, we train a second model (denoted by $\mathcal{NN}_{\text{dyn}}$, with trainable parameters \mathbf{w}_{dyn}) to predict the dynamics of the latent codes $\hat{\mathbf{s}}_i(t)$ as a function of the inputs $\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{T}^i}$:

$$\mathbf{w}_{\text{dyn}}^* = \underset{\mathbf{w}_{\text{dyn}}}{\operatorname{argmin}} \quad \sum_{i \in \mathcal{S}_{\text{train}}} \sum_{t \in \mathcal{T}^i} \|\mathcal{NN}_{\text{dyn}}(\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{T}^i}, t; \mathbf{w}_{\text{dyn}}) - \hat{\mathbf{s}}_i(t)\|^2 + \alpha_{\text{dyn}} \mathcal{R}(\mathbf{w}_{\text{dyn}})$$

In this work, we consider two different architectures for $\mathcal{NN}_{\text{dyn}}$, namely an ODE-Net [24] and an LSTM [32]. In both cases, we employ normalization layers as described in Sec. 2.4.

Once both the auto-encoder $\mathcal{NN}_{\text{dec}} \circ \mathcal{NN}_{\text{enc}}$ and the dynamics network $\mathcal{NN}_{\text{dyn}}$ have been trained, the predictions on the test set are obtained as follows, for $i \in \mathcal{S}_{\text{test}}$ and $t \in \mathcal{T}^i$:

$$\tilde{\mathbf{Y}}_i(t) = \mathcal{NN}_{\text{dec}}(\mathcal{NN}_{\text{dyn}}(\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{T}^i}, t; \mathbf{w}_{\text{dyn}}^*); \mathbf{w}_{\text{dec}}^*),$$

and

$$\tilde{\mathbf{y}}_i(\cdot, t) = \mathcal{A}'(\mathcal{NN}_{\text{dec}}(\mathcal{NN}_{\text{dyn}}(\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{T}^i}, t; \mathbf{w}_{\text{dyn}}^*); \mathbf{w}_{\text{dec}}^*)).$$

Here, we denote this method as AE/ODE and AE/LSTM, depending on the architecture employed for $\mathcal{RN}_{\text{dyn}}$.

We remark that the encoder \mathcal{N}_{enc} is only instrumental to recover labeled data $\hat{\mathbf{s}}_i(\tau)$ needed to train $\mathcal{RN}_{\text{dyn}}$ and it is not used to predict the output during the testing phase. In other words, only $\mathcal{RN}_{\text{dyn}}$ and \mathcal{N}_{dec} are retained in the testing phase. The latter observation suggests a third training stage in which, starting from the pre-trained values $(\mathbf{w}_{\text{dyn}}^*, \mathbf{w}_{\text{dec}}^*)$, we further train the $\mathcal{RN}_{\text{dyn}}$ and \mathcal{N}_{dec} in a simultaneous manner, that is in an end-to-end (e2e) way:

$$(\mathbf{w}_{\text{dyn}}^{**}, \mathbf{w}_{\text{dec}}^{**}) = \underset{\mathbf{w}_{\text{dyn}}, \mathbf{w}_{\text{dec}}}{\operatorname{argmin}} \sum_{i \in \mathcal{S}_{\text{train}}} \sum_{t \in \mathcal{T}^i} \|\mathcal{N}_{\text{dec}}(\mathcal{RN}_{\text{dyn}}(\{\mathbf{u}_i(\tau)\}_{\tau \in \mathcal{T}^i}, t; \mathbf{w}_{\text{dyn}}); \mathbf{w}_{\text{dec}}) - \mathbf{Y}_i(t)\|_{\mathbf{Y}}^2 \\ + \alpha_{\text{dyn}} \mathcal{R}(\mathbf{w}_{\text{dyn}}) + \alpha_{\text{dec}} \mathcal{R}(\mathbf{w}_{\text{dec}})$$

To denote the models obtained after this third training stage, we employ the abbreviation AE/ODE-e2e and AE/LSTM-e2e.

References

- [1] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. “Optuna: A Next-generation Hyperparameter Optimization Framework”. In: *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2019.
- [2] M. Alber, A. Buganza Tepole, W. R. Cannon, S. De, S. Dura-Bernal, K. Garikipati, G. Karniadakis, W. W. Lytton, P. Perdikaris, L. Petzold, et al. “Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences”. In: *NPJ digital medicine* 2.1 (2019), p. 115.
- [3] R. R. Aliev and A. V. Panfilov. “A simple two-variable model of cardiac excitation”. In: *Chaos, Solitons & Fractals* 7.3 (1996), pp. 293–301.
- [4] A. C. Antoulas. *Approximation of large-scale dynamical systems*. Vol. 6. Siam, 2005.
- [5] F. As’ad, P. Avery, and C. Farhat. “A mechanics-informed artificial neural network approach in data-driven constitutive modeling”. In: *International Journal for Numerical Methods in Engineering* 123.12 (2022), pp. 2738–2759.
- [6] Y. Bar-Sinai, S. Hoyer, J. Hickey, and M. P. Brenner. “Learning data-driven discretizations for partial differential equations”. In: *Proceedings of the National Academy of Sciences* 116.31 (2019), pp. 15344–15349.
- [7] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. “An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations”. In: *Comptes Rendus Mathématique* 339.9 (2004), pp. 667–672.
- [8] P. Benner, S. Griwet-Talocia, A. Quarteroni, G. Rozza, W. Schilders, and L. M. Silveira. *Model Order Reduction - Volume 2: Snapshot-Based Methods and Algorithms*. Berlin, Boston: De Gruyter, 2021.
- [9] P. Benner, S. Gugercin, and K. Willcox. “A survey of projection-based model reduction methods for parametric dynamical systems”. In: *SIAM Review* 57.4 (2015), pp. 483–531.
- [10] P. Benner, V. Mehrmann, and D. C. Sorensen. *Dimension reduction of large-scale systems*. Vol. 35. Springer, 2005.
- [11] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl. “Algorithms for hyper-parameter optimization”. In: *Advances in neural information processing systems* 24 (2011).
- [12] S. Berrone, C. Canuto, M. Pintore, and N. Sukumar. “Enforcing Dirichlet boundary conditions in physics-informed neural networks and variational physics-informed neural networks”. In: *arXiv preprint arXiv:2210.14795* (2022).
- [13] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk. “Convergence rates for greedy algorithms in reduced basis methods”. In: *SIAM journal on mathematical analysis* 43.3 (2011), pp. 1457–1472.

- [14] G. A. Bird. “Molecular gas dynamics and the direct simulation of gas flows”. In: *Molecular gas dynamics and the direct simulation of gas flows* (1994).
- [15] J. Bongard and H. Lipson. “Automated reverse engineering of nonlinear dynamical systems”. In: *Proceedings of the National Academy of Sciences* 104.24 (2007), pp. 9943–9948.
- [16] J. Bruna, B. Peherstorfer, and E. Vanden-Eijnden. “Neural galerkin scheme with active learning for high-dimensional evolution equations”. In: *arXiv preprint arXiv:2203.01360* (2022).
- [17] S. L. Brunton, J. L. Proctor, and J. N. Kutz. “Discovering governing equations from data by sparse identification of nonlinear dynamical systems”. In: *Proceedings of the National Academy of Sciences* (2016), p. 201517384.
- [18] S. L. Brunton and J. N. Kutz. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press, 2022.
- [19] S. L. Brunton, B. R. Noack, and P. Koumoutsakos. “Machine learning for fluid mechanics”. In: *Annual review of fluid mechanics* 52 (2020), pp. 477–508.
- [20] A. Buffa, Y. Maday, A. T. Patera, C. Prud’homme, and G. Turinici. “A priori convergence of the greedy algorithm for the parametrized reduced basis method”. In: *ESAIM: Mathematical modelling and numerical analysis* 46.3 (2012), pp. 595–603.
- [21] C. Canuto, T. Tonn, and K. Urban. “A posteriori error analysis of the reduced basis method for non-affine parametrized nonlinear PDEs”. In: *SIAM Journal on Numerical Analysis* 47.3 (2009), pp. 2001–2022.
- [22] M. Cenedese, J. Axås, B. Bäuerlein, K. Avila, and G. Haller. “Data-driven modeling and prediction of non-linearizable dynamics via spectral submanifolds”. In: *Nature communications* 13.1 (2022), pp. 1–13.
- [23] S. Chaturantabut and D. C. Sorensen. “Nonlinear model reduction via discrete empirical interpolation”. In: *SIAM Journal on Scientific Computing* 32.5 (2010), pp. 2737–2764.
- [24] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. “Neural ordinary differential equations”. In: *Advances in neural information processing systems (NeurIPS 2018 Proceedings)* 31 (2018).
- [25] D. Floryan and M. D. Graham. “Data-driven discovery of intrinsic dynamics”. In: *Nature Machine Intelligence* (2022), pp. 1–8.
- [26] S. Fresca, L. Dede’, and A. Manzoni. “A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized PDEs”. In: *Journal of Scientific Computing* 87.2 (2021), pp. 1–36.
- [27] S. Göktepe and E. Kuhl. “Computational modeling of cardiac electrophysiology: a novel finite element approach”. In: *International journal for numerical methods in engineering* 79.2 (2009), pp. 156–178.
- [28] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio. *Deep learning*. Vol. 1. 2. MIT press Cambridge, 2016.
- [29] M. Guo and J. S. Hesthaven. “Data-driven reduced order modeling for time-dependent problems”. In: *Computer methods in applied mechanics and engineering* 345 (2019), pp. 75–99.
- [30] J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified reduced basis methods for parametrized partial differential equations*. Springer, 2016.
- [31] J. S. Hesthaven, C. Pagliantini, and G. Rozza. “Reduced basis methods for time-dependent problems”. In: *Acta Numerica* 31 (2022), pp. 265–345.
- [32] S. Hochreiter and J. Schmidhuber. “Long short-term memory”. In: *Neural Computation* 9.8 (1997), pp. 1735–1780.
- [33] D. P. Kingma and J. Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [34] S. Lai, N. W. Ruktanonchai, L. Zhou, O. Prosper, W. Luo, J. R. Floyd, A. Wesolowski, M. Santillana, C. Zhang, X. Du, et al. “Effect of non-pharmaceutical interventions to contain COVID-19 in China”. In: *nature* 585.7825 (2020), pp. 410–413.

- [35] K. Lee and K. T. Carlberg. “Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders”. In: *Journal of Computational Physics* 404 (2020), p. 108973.
- [36] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar. “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization”. In: *Journal of Machine Learning Research* 18.1 (2017), pp. 6765–6816.
- [37] L. Li, K. Jamieson, A. Rostamizadeh, E. Gonina, J. Ben-Tzur, M. Hardt, B. Recht, and A. Talwalkar. “A System for Massively Parallel Hyperparameter Tuning”. In: *arXiv preprint arXiv:1810.05934* (2020).
- [38] K. Linka and E. Kuhl. “A new family of Constitutive Artificial Neural Networks towards automated model discovery”. In: *Computer Methods in Applied Mechanics and Engineering* 403 (2023), p. 115731.
- [39] Y. Liu, J. N. Kutz, and S. L. Brunton. “Hierarchical deep learning of multiscale differential equation time-steppers”. In: *Philosophical Transactions of the Royal Society A* 380.2229 (2022), p. 20210200.
- [40] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. “Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators”. In: *Nature Machine Intelligence* 3.3 (2021), pp. 218–229.
- [41] R. Maulik, B. Lusch, and P. Balaprakash. “Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders”. In: *Physics of Fluids* 33.3 (2021), p. 037106.
- [42] J. J. Monaghan. “Simulating free surface flows with SPH”. In: *Journal of computational physics* 110.2 (1994), pp. 399–406.
- [43] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul, M. I. Jordan, and I. Stoica. “Ray: A Distributed Framework for Emerging AI Applications”. In: *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation*. 2018, pp. 561–577.
- [44] V. Oommen, K. Shukla, S. Goswami, R. Dingreville, and G. E. Karniadakis. “Learning two-phase microstructure evolution using neural operators and autoencoder architectures”. In: *arXiv preprint arXiv:2204.07230* (2022).
- [45] S. Pagani, A. Manzoni, and A. Quarteroni. “Numerical approximation of parametrized problems in cardiac electrophysiology by a local reduced basis method”. In: *Computer Methods in Applied Mechanics and Engineering* 340 (2018), pp. 530–558.
- [46] A. T. Patera. “A spectral element method for fluid dynamics: laminar flow in a channel expansion”. In: *Journal of computational Physics* 54.3 (1984), pp. 468–488.
- [47] B. Peherstorfer, S. Gugercin, and K. Willcox. “Data-Driven Reduced Model Construction with Time-Domain Loewner Models”. In: *SIAM Journal on Scientific Computing* 39.5 (2017), A2152–A2178.
- [48] B. Peherstorfer and K. Willcox. “Dynamic data-driven reduced-order models”. In: *Computer Methods in Applied Mechanics and Engineering* 291 (2015), pp. 21–41.
- [49] M. Peirlinck, F. S. Costabal, J. Yao, J. Guccione, S. Tripathy, Y. Wang, D. Ozturk, P. Segars, T. Morrison, S. Levine, et al. “Precision medicine in human heart modeling: Perspectives, challenges, and opportunities”. In: *Biomechanics and modeling in mechanobiology* 20 (2021), pp. 803–831.
- [50] L. Petzold. “Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations”. In: *SIAM journal on scientific and statistical computing* 4.1 (1983), pp. 136–148.
- [51] C. Prud’Homme, D. V. Rovas, K. Veroy, L. Machiels, Y. Maday, A. T. Patera, and G. Turinici. “Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bound methods”. In: *Journal of Fluids Engineering* 124.1 (2002), pp. 70–80.
- [52] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced basis methods for partial differential equations: an introduction*. Vol. 92. Springer, 2015.
- [53] A. Quarteroni and A. Valli. *Numerical approximation of partial differential equations*. Vol. 23. Springer Science & Business Media, 2008.
- [54] M. Raissi, P. Perdikaris, and G. E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (2019), pp. 686–707.

- [55] C. E. Rasmussen. “Gaussian processes in machine learning”. In: *Summer school on machine learning*. Springer, 2003, pp. 63–71.
- [56] F. Regazzoni, L. Dedè, and A. Quarteroni. “Machine learning for fast and reliable solution of time-dependent differential equations”. In: *Journal of Computational Physics* 397 (2019), p. 108852.
- [57] F. Regazzoni, M. Salvador, P. Africa, M. Fedele, L. Dedè, and A. Quarteroni. “A cardiac electromechanical model coupled with a lumped-parameter model for closed-loop blood circulation”. In: *Journal of Computational Physics* 457 (2022), p. 111083.
- [58] F. Regazzoni, S. Pagani, and A. Quarteroni. “Universal Solution Manifold Networks (USM-Nets): Non-Intrusive Mesh-Free Surrogate Models for Problems in Variable Domains”. In: *Journal of Biomechanical Engineering* 144.12 (Sept. 2022). 121004.
- [59] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz. “Data-driven discovery of partial differential equations”. In: *Science advances* 3.4 (2017), e1602614.
- [60] N. D. Santo, A. Manzoni, S. Pagani, and A. Quarteroni. “Reduced-order modeling for applications to the cardiovascular system”. In: *Model Order Reduction - Volume 3 Applications*. Ed. by P. Benner, S. Grivet-Talocia, A. Quarteroni, G. Rozza, W. Schilders, and L. M. Silveira. Berlin, Boston: De Gruyter, 2021, pp. 251–278.
- [61] E. Scalas, R. Gorenflo, and F. Mainardi. “Fractional calculus and continuous-time finance”. In: *Physica A: Statistical Mechanics and its Applications* 284.1-4 (2000), pp. 376–384.
- [62] M. Schmidt and H. Lipson. “Distilling free-form natural laws from experimental data”. In: *science* 324.5923 (2009), pp. 81–85.
- [63] L. Sirovich. “Turbulence and the dynamics of coherent structures. I. Coherent structures”. In: *Quarterly of Applied Mathematics* 45.3 (1987), pp. 561–571.
- [64] P. R. Vlachas, G. Arampatzis, C. Uhler, and P. Koumoutsakos. “Multiscale simulations of complex systems by learning their effective dynamics”. In: *Nature Machine Intelligence* 4.4 (2022), pp. 359–366.
- [65] M. Zhu, H. Zhang, A. Jiao, G. E. Karniadakis, and L. Lu. “Reliable extrapolation of deep neural operators informed by physics or sparse observations”. In: *arXiv preprint arXiv:2212.06347* (2022).
- [66] O. C. Zienkiewicz and R. L. Taylor. *The finite element method for solid and structural mechanics*. Elsevier, 2005.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 34/2023** Caldana, M.; Antonietti, P. F.; Dede', L.
A Deep Learning algorithm to accelerate Algebraic Multigrid methods in Finite Element solvers of 3D elliptic PDEs
- 35/2023** Ferrari, L.; Manzi, G.; Micheletti, A.; Nicolussi, F.; Salini, S.
Pandemic Data Quality Modelling: A Bayesian Approach
- 32/2023** Gambarini, M.; Ciaramella, G.; Miglio, E.; Vanzan, T.
Robust optimization of control parameters for WEC arrays using stochastic methods
- 31/2023** Orlando, G.
Assessing ChatGPT for coding finite element methods
- 30/2023** Antonietti, P. F.; Bonizzoni, F.; Verani, M.
A DG-VEM method for the dissipative wave equation
- 29/2023** Carbonaro, D.; Mezzadri, F.; Ferro, N.; De Nisco, G.; Audenino, A.L.; Gallo, D.; Chiastra, C.; Morbiducci, U.; Perotto, S.
Design of innovative self-expandable femoral stents using inverse homogenization topology optimization
- 28/2023** Zingaro, A.; Vergara, C.; Dede', L.; Regazzoni, F.; Quarteroni, A.
A comprehensive mathematical model for myocardial perfusion
- Fumagalli, A.; Panzeri, L.; Formaggia, L.; Scotti, A.; Arosio, D.
A mixed-dimensional model for direct current simulations in presence of a thin high-resistivity liner
- 27/2023** Beirao da Vega, L.; Canuto, C.; Nochetto, R.H.; Vacca, G.; Verani, M.
Adaptive VEM for variable data: convergence and optimality