

RESEARCH ARTICLE  

Neural modal ordinary differential equations: Integrating physics-based modeling with neural ordinary differential equations for modeling high-dimensional monitored structures

Zhilu Lai^{1,2,*} , Wei Liu^{3,4}, Xudong Jian^{3,5} , Kiran Bacsa^{3,6}, Limin Sun^{5,7} and Eleni Chatzi^{3,6} 

¹Internet of Things Thrust, Information Hub, HKUST(GZ), Guangzhou, China

²Department of Civil and Environmental Engineering, HKUST, Hong Kong, China

³Future Resilient Systems, Singapore-ETH Centre, Singapore, Singapore

⁴Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore, Singapore

⁵State Key Laboratory for Disaster Reduction in Civil Engineering, Tongji University, Shanghai, China

⁶Department of Civil, Environmental and Geomatic Engineering, ETH-Zürich, Zürich, Switzerland

⁷Shanghai Qizhi Institute, Shanghai, China



*Corresponding author. E-mail: zhilulai@ust.hk

Received: 14 July 2022; **Revised:** 11 October 2022; **Accepted:** 14 October 2022

Keywords: Deep learning; dynamical systems; neural ordinary differential equations; physics-based modeling; physics-informed machine learning

Abstract

The dimension of models derived on the basis of data is commonly restricted by the number of observations, or in the context of monitored systems, sensing nodes. This is particularly true for structural systems, which are typically high-dimensional in nature. In the scope of physics-informed machine learning, this article proposes a framework—termed neural modal ordinary differential equations (Neural Modal ODEs)—to integrate physics-based modeling with deep learning for modeling the dynamics of monitored and high-dimensional engineered systems. In this initiating exploration, we restrict ourselves to linear or mildly nonlinear systems. We propose an architecture that couples a dynamic version of variational autoencoders with physics-informed neural ODEs (Pi-Neural ODEs). An encoder, as a part of the autoencoder, learns the mappings from the first few items of observational data to the initial values of the latent variables, which drive the learning of embedded dynamics via Pi-Neural ODEs, imposing a modal model structure on that latent space. The decoder of the proposed model adopts the eigenmodes derived from an eigenanalysis applied to the linearized portion of a physics-based model: a process implicitly carrying the spatial relationship between degrees-of-freedom (DOFs). The framework is validated on a numerical example, and an experimental dataset of a scaled cable-stayed bridge, where the learned hybrid model is shown to outperform a purely physics-based approach to modeling. We further show the functionality of the proposed scheme within the context of virtual sensing, that is, the recovery of generalized response quantities in unmeasured DOFs from spatially sparse data.

  This research article was awarded an Open Data and Open Materials badge for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Impact Statement

We propose neural modal ordinary differential equations that learn generative dynamical models from spatially sparse sensor data. The proposed method is in the format of dynamical variational autoencoders, and we structure the latent space of the measured data using physics-related features (e.g., modal features), allowing physically interpretable architectures. The delivered models are able to reconstruct the full-field structural response, meaning response in unmeasured locations, given limited sensing locations. We believe this proposed method is helpful and meaningful to the community of structural digital twins, model updating, virtual sensing, and structural health monitoring.

1. Introduction

Physics-based modeling (or first-principles modeling) forms an essential engineering approach to understand and simulate the behavior of structural systems. Often implemented via the use of finite element methods (FEM) (Waisman et al., 2010; Strømmen, 2014), within the context of structural engineering, physics-based modeling is capable of building high-dimensional and high-fidelity models for large and complex civil/mechanical structures. However, such models often suffer from simplified assumptions and approximations, while for the case of monitored operating systems, an established model often fails to reflect a system as is, after possible experience of damaging and deterioration effects. Such limitations can be tackled by means of uncertainty quantification analysis (Sankararaman and Mahadevan, 2013), or more effectively via feedback from monitoring (sensory) data (Farrar and Worden, 2012; Kamariotis et al., 2022). The integration of data with physics-based models or physical laws—*physics-informed machine learning* (Zhu et al., 2019; Willard et al., 2020; Karniadakis et al., 2021; Bae and Koumoutsakos, 2022) has grown into an active research area for modeling physical systems in recent years.

Beyond their exploitation within a broader science and engineering context (Karpatne et al., 2017; Wu et al., 2018; Kashinath et al., 2021), physics-informed machine learning has been specifically applied for learning dynamical systems from either simulated or real-world data. This has been pursued in various ways; for instance, by exploiting the automatic differentiation of neural networks (NNs) to form “custom” activation and loss functions that are tailored to the underlying differential operator (Raissi et al., 2019), by incorporating Lagrangian dynamics into the NN architecture (Cranmer et al., 2020; Roehrl et al., 2020), by imposing the laws of dynamics as constraints to the network (Zhang et al., 2020), or via identification of a sparse set of physics-informative basis functions to establish equations of motion of observed systems (Lai and Nagarajaiah, 2019; Lai et al., 2020). It is further worth noting that a significant tool for fusion lies in the reduction of physics-based models. Notably, Vlachas et al. (2022) propose a combination of a long short-term memory (LSTM) network with an autoencoder (AE), jointly referred to as Learning Effective Dynamics, which can be trained on data from simulations of dynamical systems. In a similar context, applied for reduction of nonlinear structural dynamics, Simpson et al. (2021) combine an LSTM with an AE for delivering fast and accurate simulators of complex high-dimensional structures. In an alternate setting, reduction can efficiently be achieved, while respecting the underlying physics equations, via projection-based methods (Carlberg et al., 2013; Qian et al., 2020; Vlachas et al., 2021). This yields a powerful framework, which can eventually be combined with data, for instance via use of Bayesian filtering as proposed in Tatsis et al. (2022) for the purpose of damage detection and flaw identification. In previous work of part of the authoring team, we delivered hybrid representations that draw from the availability of monitoring data (measurements/observations from the system), which combine a term that reflects our often impartial knowledge of the physics, with a learning term which compensates what our physics representations may not account for, via physics-informed neural ordinary differential equations (ODEs) (Lai et al., 2021) and physics-guided Deep Markov Models (PgDMMs) (Liu et al., 2022).

Learning a dynamical system essentially boils down to learning a governing function (either in parametric or nonparametric form) that describes the evolution of the “system’s state” over time. We summarize the motivation of this article as follows. Firstly, in the context of monitoring, the representation of a dynamical system is restricted by the number of sensing nodes. Compared to a model established by

physics-based modeling, a data-driven model is often a reduced-order model, typically encompassing contributing modes, which considerably sacrifices the true spatial resolution. Due to this, there often exists an inconsistency between the coordinate spaces of the two models, with the high-dimensional physics-based model (such as a FEM) corresponding to spatially dense degrees-of-freedom (DOFs), while a data-driven model often reflects a latent space that is expressed in nonphysical coordinates (Schmid, 2010; Lusch et al., 2018; Simpson et al., 2021). Secondly, the adopted data types are critical to the learning of dynamical systems. If direct measurements of a latent space exist (e.g., in representing structural dynamics, displacement and velocity are considered as such latent variables), it is straightforward to learn the dynamics that are inherent to the extracted data. However, this is not the case in practice, as the measured response (data) is most commonly not a direct measurement of the latent variables; for example, when accelerations are available in the context of vibration-based monitoring (Ou et al., 2021). With these two aspects in mind, in this article, we propose a framework that is capable of integrating high-dimensional physics-based models with machine learning schemes for modeling the dynamics of high-dimensional structural systems, with linear or mildly nonlinear behavior. The term “mildly nonlinear” refers to systems whose response is not significantly different from their linear approximation. Such a discrepancy could be formally quantified using metrics such as the value of the coherence between the input (load) and output (response) signal.

To achieve this, we propose to blend a dynamical version (Girin et al., 2020) of a variational autoencoder (VAE) (Kingma and Welling, 2013), with a projection basis containing the eigenmodes that are derived from the linearization of a physics-based model, termed as neural modal ODEs. We justify these components in the proposed architecture as follows: (a) the majority of the aforementioned projection-based methods, which commonly rely on proper orthogonal decomposition (POD) (Liang et al., 2002), have been applied for the reduction of nonlinear models/simulators (Abgrall and Amsallem, 2016; Amsallem et al., 2015; Balajewicz et al., 2016; Peherstorfer and Willcox, 2016; Marconia et al., 2021; Vlachas et al., 2022). In this case, we rely on the availability of actual measured data but not simulations of full-order models, which may bear with model bias. To this end, the probabilistic version of autoencoders (Hinton and Zemel, 1994), that is, the VAE (Kingma and Welling, 2013), is adopted to learn latent representations from data. Our aim is to devise a *generative model*, which is though inferred from data availability, and not a mere observer. In doing so, we exploit data availability in order to infer the initial values of the latent space, in this way boosting the learning of embedded dynamics. This scheme actually falls in the category of nonintrusive model reduction (Swischuk et al., 2020). In contrast with intrusive model reduction, nonintrusive is data-driven and does not require access to the full-order model. (b) This type of nonintrusive model reduction generally allows for flexibility on the structure of the learned latent space, which need not assume a physically meaningful representation. Since we are interested in monitoring applications, it becomes important to achieve such a physics-based representation, especially for the latent space, since this allows virtual sensing tasks; meaning the inference of structural response in locations that are not directly measured/observed (Vettori et al., 2022). To model and structure the dynamics of the reduced-order models (latent dynamics), we herein adopt our previously developed physics-informed neural ODEs (Pi-Neural ODEs) (Lai et al., 2021) to impose a modal structure, in which, the dynamics are driven by superposing the modal representations derived from physics-based modeling with a residual term learned by NNs. This allows accounting for the portion of physics, which remains unaccounted for. (c) The implemented Pi-Neural ODEs allow for flexibility, as the residual term adaptively accounts for various discrepancies. In this case, this makes up for the fact that our reduction basis exploits linear eigenmodes. If the system exhibits a mild level of nonlinearity, the resulting discrepancy will be accounted for by the imposed NN term in the Pi-Neural ODEs.

We validate the efficacy of the proposed neural modal ODEs on a numerical example, and an experimental dataset derived from a scaled cable-stayed bridge. Based on the results presented in this article, the contribution of the study lies in: (a) establishing a generative modeling approach that integrates physics-based modeling with deep learning to model high-dimensional structural dynamical systems, while retaining the format of an ordinary differential equation; (b) by introducing a physically structured decoder, the model is capable of extrapolating the dynamics to unmeasured DOFs. Such a virtual sensing

scheme can be applied to structures where observations are scarce (Sun et al., 2020); (c) since this is a generative model, it further has the potential of being implemented within the context of model updating.

2. Neural Modal ODEs

We summarize the proposed architecture in the flowchart of Figure 1, which combines an encoder Ψ_{NN} and a decoder Φ_p , with Pi-Neural ODEs (Lai et al., 2021) (Pi-Neural ODEs). The role of the encoder is to perform inference of the initial conditions of the latent variables \mathbf{z}_0 from a handful of observational data of measured DOFs.

The evolution of the dynamics initiating from \mathbf{z}_0 is learned and modeled by means of Pi-Neural ODEs. It assumes that a system can be modeled as a superposition of a physics-based modeling term and a learning-based term, where the latter aims to capture the discrepancy between the physics-based model and the actual system. The physics-informed term in this framework adopts a modal representation derived from the eigenanalysis of the structural matrices of the physics-based model. In the case of a nonlinear system, we rely on the linearized portion of the model.

The prediction of latent quantities $\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_t, \dots, \mathbf{z}_T$ at time step t_0, t_1, \dots, t_T , obtained from the previous step is mapped back to the full-order responses via the decoder, and then to the estimated quantities in the original observation space ($\hat{\mathbf{x}}_0, \hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_t, \dots, \hat{\mathbf{x}}_T$) via a selection matrix \mathbf{E} (each row is a one-hot row vector), selecting corresponding monitored quantities. This is then compared against the actual measurements to minimize the prediction error, which effectuates the training of the proposed model. The decoder is physically structured and also derived from the eigenanalysis of the structural matrices.

In what follows, we offer the details of the formulation of the three outlined components (encoder, Pi-Neural ODEs, and decoder) to the suggested framework.

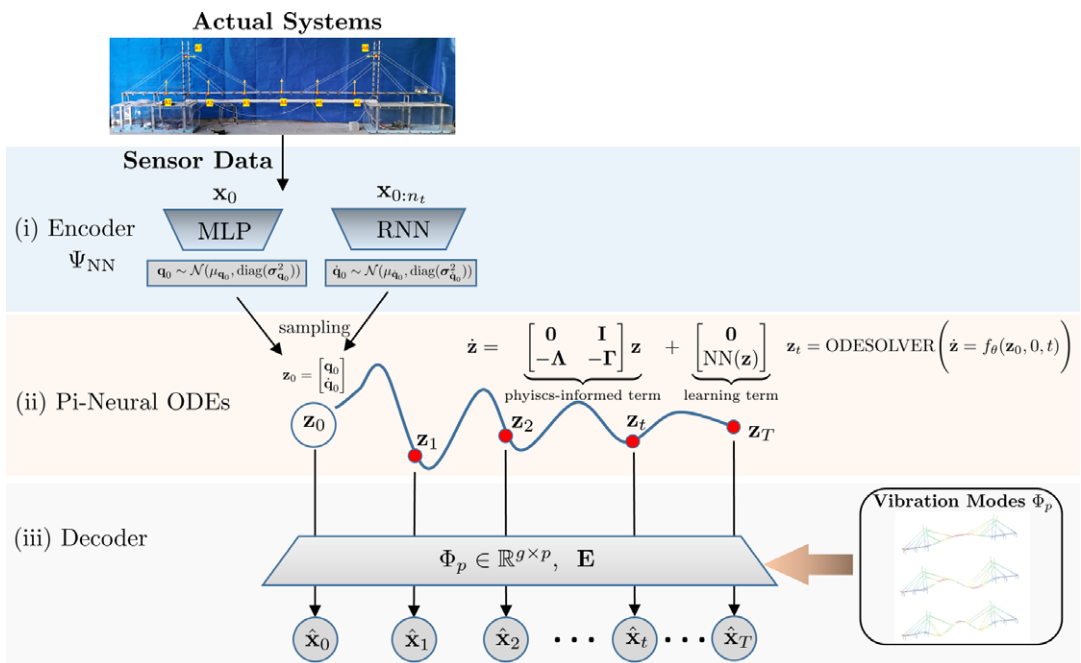


Figure 1. Flow chart of the proposed framework, encompassing an encoder, Pi-Neural ODEs, and a physically structured decoder. The encoder Ψ_{NN} is comprised of a multilayer perceptron (MLP) and a recurrent neural network (RNN).

2.1. Encoder (inference model)

Consider an observation (measurement) dataset $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^N$ with N independent sequences of time series data. Each sequence reflects a multi-DOF time series record, defined as $\mathbf{x}^{(i)} = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T\}^{(i)}$, where the observation vector at time instance t , $\mathbf{x}_t \in \mathbb{R}^m$, reflects m monitored DOFs. When the underlying physics equations are known, the observation \mathbf{x}_t at each time instance t can be assumed to be derived from a corresponding latent (state) variable \mathbf{z}_t , assumed to completely describe the embedded dynamical state. In practice, a common issue is that the latent variables are usually unobserved or only partially observed, via indirect measurements. This limitation is often tackled in prior art via the use of an encoder parameterized by a NN Ψ_{NN} , which is employed to infer the latent variables from observation data.

In delivering such an estimate, we adopt a temporal version (Girin et al., 2020) of the VAE (Kingma and Welling, 2013), that has been implemented in existing literature (Krishnan et al., 2017; Yildiz et al., 2019; Liu et al., 2022). The encoder Ψ_{NN} can be mathematically described as:

$$\Psi_{\text{NN}}(\mathbf{z}_0 | \mathbf{x}_{0:n_t}) = \Psi_{\text{NN}}\left(\begin{bmatrix} \mathbf{q}_0 \\ \dot{\mathbf{q}}_0 \end{bmatrix} | \mathbf{x}_{0:n_t}\right) = \mathcal{N}\left(\begin{bmatrix} \mu_{\mathbf{q}_0} \\ \mu_{\dot{\mathbf{q}}_0} \end{bmatrix}, \begin{bmatrix} \text{diag}(\sigma_{\mathbf{q}_0}^2) & \mathbf{0} \\ \mathbf{0} & \text{diag}(\sigma_{\dot{\mathbf{q}}_0}^2) \end{bmatrix}\right), \tag{1a}$$

where the first few observations from \mathbf{x}_0 to \mathbf{x}_{n_t} (denoted by $\mathbf{x}_{0:n_t}$) are used for inferring \mathbf{z}_0 , that is, \mathbf{z}_0 is conditioned on \mathbf{x}_0 to \mathbf{x}_{n_t} ; the latent variables $\mathbf{z}_t \in \mathbb{R}^{2p}$ are assumed to have dimension of $2p$, and the output of the encoder is intentionally split into $\mathbf{q}_0 \in \mathbb{R}^p$ and $\dot{\mathbf{q}}_0 \in \mathbb{R}^p$ that are corresponding to displacement and velocity states, respectively, that is, $\mathbf{z}_0 = \begin{bmatrix} \mathbf{q}_0 \\ \dot{\mathbf{q}}_0 \end{bmatrix}$. It is further assumed that the inferred state variable \mathbf{z}_0 is a stochastic one, which is in this case essential for reflecting uncertainties, and follows a normal distribu-

tion, of mean value $\begin{bmatrix} \mu_{\mathbf{q}_0} \\ \mu_{\dot{\mathbf{q}}_0} \end{bmatrix}$ and diagonal covariance matrix $\begin{bmatrix} \text{diag}(\sigma_{\mathbf{q}_0}^2) & \mathbf{0} \\ \mathbf{0} & \text{diag}(\sigma_{\dot{\mathbf{q}}_0}^2) \end{bmatrix}$. It should though be

noted that it is common to model uncertainty in structural systems, which are subjected to random environmental influences, using a normal distribution. For most of dynamical VAEs frameworks, which are adopted in the context of modeling dynamical systems with uncertainty, the inherent uncertainties are accounted for via the use of normal distributions, as summarized in the work of Girin et al. (2020).

In practice, Ψ_{NN} is comprised of a feed-forward NN (MLP) and a RNN. We assume that the displacement quantity \mathbf{q}_0 only depends on \mathbf{x}_0 , per the assumption adopted in Yildiz et al. (2019):

$$(\mu_{\mathbf{q}_0}, \sigma_{\mathbf{q}_0}^2) = \text{MLP}(\mathbf{x}_0). \tag{1b}$$

The output of this MLP is a stochastic variable of mean $\mu_{\mathbf{q}_0}$ and variance $\sigma_{\mathbf{q}_0}^2$; the velocity quantity $\dot{\mathbf{q}}_0$ is inferred from the first leading observations $\mathbf{x}_{0:n_t}$, thus a RNN is implemented to take $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{n_t}$ into account:

$$(\mu_{\dot{\mathbf{q}}_0}, \sigma_{\dot{\mathbf{q}}_0}^2) = \text{RNN}(\mathbf{x}_{0:n_t}), \tag{1c}$$

where the output of the RNN is a stochastic variable of mean $\mu_{\dot{\mathbf{q}}_0}$ and variance $\sigma_{\dot{\mathbf{q}}_0}^2$; n_t need not necessarily reflect a large number, larger n_t might dilute the inference of the velocity quantity; for instance, based on empirical trial, in our implementations $n_t = 10$. Once the normal distribution defined in equation (1a) is derived, one can sample \mathbf{z}_0 from this distribution, and use it for computing the evolution of the latent dynamics over time. We use θ_{enc} to denote all the parameters used in the Ψ_{NN} , that is, all the hyperparameters involved in the formulation of the MLP and RNN architectures.

2.2. Modeling latent dynamics via Pi-Neural ODEs

There are generally two strategies in terms of how the temporal dependence between states \mathbf{z} can be modeled. The first strategy is to use a discrete-time model to describe the embedded dynamics, where the

first-order Markovian property is assumed. A popular example, in the deep learning context, can be found in the deep Markov models (Krishnan et al., 2015, 2017; Liu et al., 2022). An alternative lies in adopting continuous models, usually in the form of differential equations, to describe the temporal dependence embedded in the data. The neural ODEs (Chen et al., 2018) form a recently proposed tool that parameterizes the governing differential equations by feed-forward NNs in a continuous format. A specific merit of a continuous modeling approach is that nonequidistant sequential data can be used for training the model. As the Neural ODEs effectively represent a differential equation construct, the trained model can, in turn, be used as a generative model, meaning as a model which can predict the system response given initial conditions or external excitation.

In the previous work of the authors (Lai et al., 2021), we introduced a Pi-Neural ODEs scheme, assuming that a system can be modeled as a superposition of a physics-based modeling term and a learning-based term, where the latter aims to capture the discrepancy between the physics-based model and the actual system. A similar scheme is further discussed in Wagg et al. (2020) for application within the context of digital twinning, as the learning-based term allows for adaptation. The scheme is formally described as follows:

$$\dot{\mathbf{z}} = f_{\theta_{\text{dyn}}}(\mathbf{z}) = f_{\text{phy}}(\mathbf{z}) + f_{\text{NN}}(\mathbf{z}), \tag{2}$$

where $f_{\text{phy}}(\mathbf{z})$ is a physics-based model, which can be built by leveraging the best possible knowledge of the system; $f_{\text{NN}}(\mathbf{z})$ is the learning-based model that is materialized as a NN function of \mathbf{z} . It is noted that the former term $f_{\text{phy}}(\mathbf{z})$ is of a fixed and preassigned structure, while the latter term is adjustable during the process of training the model. The parameter vector θ_{dyn} , reflects the set of hyperparameters involved in the NN representation $f_{\text{NN}}(\mathbf{z})$.

In this article, we adopt this modeling scheme for use within a reduced order modeling (ROM) setting, to model the latent dynamics of a high-dimensional system. We restrict ourselves in this initiating effort to the modeling of linear or mildly nonlinear systems. The mildly nonlinear system we refer to in this article is that the system can be well approximated by the linearization of the system—the first-order Taylor expansion.

In such a case, an approximation of the dynamics can be derived through the solution of an eigenvalue problem of the structural matrices of the physics-based model (in the case of a nonlinear system, we rely on the linearized part), and is reflected in the following decoupled low-dimensional linearized form:

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \ddot{\mathbf{q}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\Lambda & -\Gamma \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix}, \tag{3a}$$

where,

$$\Lambda = \begin{bmatrix} \omega_1^2 & & & \\ & \omega_2^2 & & \\ & & \ddots & \\ & & & \omega_p^2 \end{bmatrix} \quad \Gamma = \begin{bmatrix} 2\xi_1\omega_1 & & & \\ & 2\xi_2\omega_2 & & \\ & & \ddots & \\ & & & 2\xi_p\omega_p \end{bmatrix}, \tag{3b}$$

where Λ and Γ are both diagonal matrices; $\omega_1, \omega_2, \dots, \omega_p$ are the first p leading natural frequencies (the first p maximum frequencies in a descending order) that are retrieved from an eigenanalysis of an a priori available physics-based model; $\xi_1, \xi_2, \dots, \xi_p$ are the corresponding modal damping ratios; $\mathbf{I} \in \mathbb{R}^{p \times p}$ denotes the identity matrix.

Our premise is that the physics-based model in equation (3a) does not fully represent the actual system, which implies that the model-derived modal parameters can be different from the parameters that describe the actual operating system as-is, or that additionally, further to the parameters, the structure of the model is lacking. The latter implies that certain mechanisms are not fully understood and are, thus, modeled inaccurately, for instance, mechanisms related to nonlinearities or damping. To account for such sources

of error or discrepancies, we add a learning-based term to model the dynamics that are unaccounted for, with equation (3a) now defined as:

$$\dot{\mathbf{z}} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{\Lambda} & -\mathbf{\Gamma} \end{bmatrix} \mathbf{z} + \begin{bmatrix} \mathbf{0} \\ \text{NN}(\mathbf{z}) \end{bmatrix} \text{ with } \mathbf{z}(0) = \mathbf{z}_0, \tag{4}$$

where $\mathbf{z} = \begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix}$; NN represents a feed-forward NN that is a function of \mathbf{z} . It is noted that the structure presented in equation (4) has the potential of breaking the fully decoupled structure, which is defined by the first term. This is in fact welcomed since the hypothesis of fully decoupled damping matrices, relating to a Rayleigh viscous damping assumption (Craig and Kurdila, 2006), is a known source of modeling discrepancies for real-world systems (Satake et al., 2003). The learning-based term NN(\mathbf{z}) is thus added to account for possible sources of inconsistency and error. In this physics-informed architecture, during training, the estimated gradients are obtained as the sum of the corresponding gradients derived from the physics-based and learning-based terms. Since the gradients from the physics-based term are fixed, only the gradients of the learning-based term are to be estimated. The combined gradients are restricted in a regime that is closer to the true function’s gradients. Supplementary Appendix further elaborates on the benefit of this physics-informed architecture, which boosts the search for the governing equations close to the actual systems.

The Physics-informed Neural ODE equation (4) governs the evolution of the dynamics. The dynamics of $\mathbf{z}(t)$ can be solved by numerically integrating $\mathbf{z}(t) = \int_{t_0}^t f_{\theta_{\text{dyn}}}(\mathbf{z}) dt$ from t_0 to t given initial conditions \mathbf{z}_0 , with the estimate of the latent state vector $\mathbf{z}(t)$ at each time t offered as:

$$\mathbf{z}(t) = \text{ODESOLVE}(f_{\theta_{\text{dyn}}}, \mathbf{z}_0, t_0, t), \tag{5}$$

where ODESOLVE reflects the chosen numerical integration scheme, with Runge–Kutta methods comprising a typical example of such solvers. The dynamics of the latent state \mathbf{z} , with realization of $\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_t, \dots, \mathbf{z}_T$ (where $\mathbf{z}_t = \begin{bmatrix} \mathbf{q}_t \\ \dot{\mathbf{q}}_t \end{bmatrix}$), are thus computed at each time step, and can be subsequently fed into the decoder model to reconstruct the full field response, as described in what follows.

2.3. Decoder

In the case of a linear dynamical system, the full-order response $\mathbf{x}_t^{\text{full}} \in \mathbb{R}^g$ comprises a modal representation of $\mathbf{x}_t^{\text{full}} \approx \Phi \mathbf{q}_t$ ($\Phi_p \in \mathbb{R}^{g \times p}; \mathbf{q}_t \in \mathbb{R}^p; p \leq g$), where Φ_p is the truncated eigenvector matrix, that is, the leading p columns of full-order eigenvector matrix Φ (corresponding to the largest p eigenvalues).

As illustrated in Figure 1, an estimate of the evolution of the latent state over time $\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_T$ can be obtained by solving the Pi-Neural ODEs via equation (5). It is noted that, within the structural dynamics context, important measurable quantities such as accelerations $\ddot{\mathbf{q}}$ can further be computed on the basis of the governing equation (4): $\ddot{\mathbf{q}} = [-\mathbf{\Lambda} \quad -\mathbf{\Gamma}]\mathbf{z} + \text{NN}(\mathbf{z})$. Thus, beyond the latent states $\mathbf{q}, \dot{\mathbf{q}}$, we can derive further response quantities of interest, such as the acceleration $\ddot{\mathbf{q}}$.

Each response quantity can be respectively emitted to the corresponding full-order response vector (involving all structural DOFs) via the decoder Φ_p ($\mathbb{R}^p \rightarrow \mathbb{R}^g$):

$$\begin{aligned} \text{displacement} : \mathbf{x}_t^{\text{full}} &= \Phi_p(\mathbf{q}_t), \\ \text{velocity} : \dot{\mathbf{x}}_t^{\text{full}} &= \Phi_p(\dot{\mathbf{q}}_t), \\ \text{acceleration} : \ddot{\mathbf{x}}_t^{\text{full}} &= \Phi_p(\ddot{\mathbf{q}}_t), \quad (t = 0, 1, \dots, T), \end{aligned} \tag{6a}$$

where $\mathbf{x}_t^{\text{full}}, \dot{\mathbf{x}}_t^{\text{full}}$, and $\ddot{\mathbf{x}}_t^{\text{full}}$ denote the reconstructed full-order displacement, velocity, and acceleration, respectively. It is noted that further response quantities of interest, such as potentially strains, can be inferred due to availability of a FEM model.

We can only measure a limited of DOFs, $\mathbf{x}_t \in \mathbb{R}^m$ (we use \mathbf{x}_t to denote measured quantities while $\widehat{\mathbf{x}}_t$ denoting the corresponding estimated quantities), via use of appropriate sensors, which form a subset of the full response vector:

$$\widehat{\mathbf{x}}_t = \mathbf{E} \begin{bmatrix} \mathbf{x}_t^{\text{full}} \\ \mathbf{x}_t^{\text{full}} \\ \vdots \\ \mathbf{x}_t^{\text{full}} \end{bmatrix}, \quad (6b)$$

where $\mathbf{E} \in \mathbb{R}^{m \times 3g}$ is a selection matrix (each row is a one-hot row vector), selecting corresponding monitored quantities; $\widehat{\mathbf{x}}_t$ can represent an extended set of the estimated observations, which can correspond to displacement, velocity, acceleration, or further computable response quantities (such as strains). Since we only consider mild nonlinearity, we rely on the observability of the linearized part of the system, where classical observability theory (Kalman, 1960) can be applied to analyze the observability—estimating the full state vector from limited measurements.

The architecture of the proposed framework essentially comprises a sequential version of the VAE, exploiting the presence of an underlying low-dimensional latent representation in the observed dynamics. In the original VAE, the decoder is parameterized by a NN without regularization, which flexibly fits the training data, without necessarily embodying a physical connotation. From an engineering perspective, however, it would be beneficial if the decoder is bestowed with a direct linkage to physical DOFs. One way to achieve this is to seed the modal shape information, computed from physics-based models, which carries within it the spatial information of how each element/node in \mathbf{x} is interconnected. Therefore, we forcibly implement eigenmodes Φ_p as the decoder for emitting the latent variables to the observation space. $\Phi_p = [\phi_1, \phi_2, \dots, \phi_p]$, where each column represents a single eigenmode, can be derived from the structural matrices of the physics-based full order model, for example, a FE model. It is noted that Φ_p is assumed to be time-invariant, thus reflecting an invariant encoding of the spatial relationship between structural DOFs. However, the residual term $\text{NN}(\mathbf{z})$ in the Pi-Neural ODE in equation (4) adaptively accounts for discrepancies that stem from mild nonlinearities, which would also violate the assumption of invariance. We remind that, in Section 2.2, a decoupled structure is adopted as a prior model to encourage the model to mimic the process of a modal decomposition–reconstruction.

It is worth mentioning that, in this framework, the encoder process can be viewed as the transformation from full-order physical coordinates to modal coordinates $\Psi_{\text{NN}}: \mathbf{x} \rightarrow \mathbf{z}$. In real scenarios that involve weakly nonlinear systems, this can be thought of as a “modal-like” coordinate as the learning term $\text{NN}(\mathbf{z})$ can violate the decoupled structure, while the decoder is viewed as the operator which enables the transformation from the modal coordinates’ space to the measured physical coordinates ($\Phi_p: \mathbf{z} \rightarrow \mathbf{x}$).

2.4. Loss function

For the purpose of training the suggested Neural Modal ODE models, which capitalize on the availability of physics information and data, we calculate the measurement prediction error. The model delivers an estimate $\widehat{\mathbf{x}}_{0:T}$ of the measured response quantities $\mathbf{x}_{0:T}$, which in turn allows to minimize the error between the predicted and actual observations, to train the model. The training of the encoder, decoder, and latent dynamic models are performed simultaneously, and the loss function of the framework is given as:

$$\mathcal{L}(\theta; \mathbf{x}) = \mathcal{L}\{\text{DECODER}[\text{ODESOLVE}(f_{\theta_{\text{dyn}}}, \Psi_{\text{NN}}(\mathbf{x}_{0:n_t}), t_0, T)]\}, \quad (7)$$

where $\theta = \theta_{\text{enc}} \cup \theta_{\text{dyn}}$ are all the parameters involved in the deep learning model; $\mathbf{x}_{0:n_t}$ is the first \mathbf{x}_0 to \mathbf{x}_{n_t} data fed into the encoder Ψ_{NN} ; $\mathbf{x}_{0:T}$ is the whole sequence of the data set used for the decoder; the notation DECODER denotes the process given in equations (6a) and (6b).

In the VAE formulation (Kingma and Welling, 2013), the loss function \mathcal{L} is used to maximize a variational lower bound of the data log-likelihood $\log p(\mathbf{x})$; here x is short for $\mathbf{x}_{0:T}$. Using the variational principle with the inference model $\Psi_{\text{NN}}(\mathbf{z}_0 | \mathbf{x}_{0:n_t})$, which is only used to infer the initial condition of \mathbf{z}_0 ,

the evidence lower bound (ELBO) of the data log-likelihood, which is the loss function, is given as follows:

$$\mathcal{L}(\theta; \mathbf{x}) = \sum_{t=0}^T \left\{ \mathbb{E}_{\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})} [\log p(\mathbf{x}_t|\mathbf{z}_t)] - \mathbb{E}_{\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})} [\text{KL}(\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})||p(\mathbf{z}_0))] \right\}, \tag{8}$$

where KL stands for the Kullback–Leibler divergence; a statistical measure that evaluates the closeness of two probability distributions p_1 and p_2 , defined as $\text{KL}(p_1(\mathbf{z})||p_2(\mathbf{z})) := \int p_1(\mathbf{z}) \log \frac{p_1(\mathbf{z})}{p_2(\mathbf{z})} d\mathbf{z}$. In the loss function, the first term $\sum_{t=0}^T \mathbb{E}_{\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})} [\log p(\mathbf{x}_t|\mathbf{z}_t)]$ evaluates the reconstruction accuracy: \mathbf{z}_0 is sampled from the distribution given in equation (1a), and with this given initial condition, one can compute the predicted $\hat{\mathbf{x}}_t \sim \mathcal{N}(\hat{\boldsymbol{\mu}}_t, \hat{\boldsymbol{\Sigma}}_t)$ ($t = 0, 1, \dots, T$) via the latent dynamics model in equation (5) followed by the decoder. Thus, this term can be computed as $\sum_{t=0}^T \log p(\mathbf{x}_t)$ given $\mathbf{z}_0 \sim \Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})$, and $\log p(\mathbf{x}_t)$ has an analytical form when $p(\mathbf{x}_t)$ follows a normal distribution:

$$\log p(\mathbf{x}_t) = -\frac{1}{2} \left[\log |\hat{\boldsymbol{\Sigma}}_t| + (\mathbf{x}_t - \hat{\boldsymbol{\mu}}_t)^T \hat{\boldsymbol{\Sigma}}_t^{-1} (\mathbf{x}_t - \hat{\boldsymbol{\mu}}_t) + d_x \log(2\pi) \right], \tag{9}$$

which is the log-likelihood, and the training of the model is expected to maximize this likelihood given the actual observation data \mathbf{x}_t ; d_x is the dimension of \mathbf{x}_t .

The second term $-\sum_{t=0}^T \mathbb{E}_{\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})} [\text{KL}(\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})||p(\mathbf{z}_0))]$ evaluates the closeness of the inferred initial condition with a prior distribution $p(\mathbf{z}_0)$. In practice, $p(\mathbf{z}_0)$ can be assumed as a normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ if no further prior knowledge is given. The KL terms acts as a penalty term when the inferred initial value is distant from the prior distribution. This term can be alternatively computed as $-\sum_{t=0}^T \text{KL}(\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})||p(\mathbf{z}_0))$ given $\mathbf{z}_0 \sim \Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})$. $\text{KL}(p_1(\mathbf{z})||p_2(\mathbf{z}))$ is described by an analytical formula when both $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$ are normal distributions and $p_2(\mathbf{z}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$:

$$\text{KL}(\Psi_{\text{NN}}(\mathbf{z}_0|\mathbf{x}_{0:n_t})||p(\mathbf{z}_0)) = -\log |\text{diag}(\sigma_{\mathbf{z}_0})| + \frac{\|\sigma_{\mathbf{z}_0}\|^2 + \|\mu_{\mathbf{z}_0}\|^2}{2} - \frac{d_z}{2}, \tag{10}$$

in which, $\sigma_{\mathbf{z}_0} = \begin{bmatrix} \sigma_{q_0} \\ \sigma_{\dot{q}_0} \end{bmatrix}$; $\mu_{\mathbf{z}_0} = \begin{bmatrix} \mu_{q_0} \\ \mu_{\dot{q}_0} \end{bmatrix}$; $|\cdot|$ is the determinant of a matrix; $\|\cdot\|$ is the modulus of a vector; d_z is the dimension of \mathbf{z} .

2.5. Prediction of learned dynamics

The completion of the training process results in the definition of the hyperparameter sets θ_{enc} and θ_{dyn} . This delivers an encoder Ψ_{NN} together with a learned dynamic model $\dot{\mathbf{z}} = f_{\theta_{\text{dyn}}}(\mathbf{z})$, which retains the structure of differential equations. Equations (5) and (6a) can be used for predicting the dynamics given an initial state \mathbf{z}_0 . \mathbf{z}_0 can be either be inferred from the observation dataset via the learned encoder Ψ_{NN} , or—when using the derived model as a generative model—the modeler can assign other specific values for the initial condition \mathbf{z}_0 .

For those readers that are interested in reusing the developed algorithms, a demonstrative implementation in Python, reproducing all steps from Sections 2.1 to 2.5, will be made available at: <https://github.com/zlaidyn/Neural-Modal-ODE-Demo>, including both linear and nonlinear cases of a demonstrative example introduced in the next section.

3. Demonstrative Example of a 4-DOF Structural System

In this section, we implement the proposed framework on a simulated 4-DOF structural systems. The structural system is governed by the following differential equations:

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{C}\dot{\mathbf{x}} + \mathbf{K}\mathbf{x} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ k_n x_n^3 \end{bmatrix} = \mathbf{0}, \tag{11a}$$

where the displacement vector $\mathbf{x} = [x_1, x_2, x_3, x_4]^T$; the mass matrix $\mathbf{M} = \text{diag}(m_1, m_2, m_3, m_4)$, and $m_1 = 1, m_2 = 2, m_3 = 3, m_4 = 4$; the damping matrix $\mathbf{C} = \text{diag}(c_1, c_2, c_3, c_4)$, and $c_1 = c_2 = c_3 = c_4 = 0.1$; and the stiffness matrix:

$$\mathbf{K} = \begin{bmatrix} k_1 + k_2 & -k_2 & 0 & 0 \\ -k_2 & k_2 + k_3 & -k_3 & 0 \\ 0 & -k_3 & k_3 + k_4 & -k_4 \\ 0 & 0 & -k_4 & k_4 \end{bmatrix}, \tag{11b}$$

where $k_1 = 1, k_2 = 2, k_3 = 3, k_4 = 4$. To fully demonstrate the capability of the proposed framework for both linear and nonlinear structural systems, we test three different cases with increasing nonlinearity $k_n = 0.0$ (linear case), 0.5 , and 1.0 , respectively. The linear portion of the three cases are set to be the same and the only variation lies in the coefficient k_n as the nonlinear term.

A total number of 1,000 realizations with randomized initial conditions from a standard normal distribution are generated for each case (the randomization is identical for each case of k_n). As mentioned, we assume that in the here presented application scenarios only a limited subset of the full-order system response quantities are available. In this example, only the displacement of the fourth DOF (x_4) and the accelerations of the first, third, and fourth DOFs ($\ddot{x}_1, \ddot{x}_3, \ddot{x}_4$) are measured. While it is feasible to implement the framework with acceleration measurements only, the accounted displacement of a single DOF is here used to alleviate possible drifting effects that occur in the reconstructed full state. The first n_0 to $n_t = 10$ samples of the sequence are used for the RNN in the encoder to infer the initial latent velocity. As for the decoder Φ_p , we make use of the first $p = 4$ modes obtained via an eigenanalysis of the structural matrices of the physics-based model, thus forming an 8-dimensional latent state. The implementation details are listed in Table 1. The models are trained on the dataset of the first 800 realizations and tested on the remaining 200 realizations.

Figure 2 shows the force-displacement loops of the first DOF of the reference system for different values of the nonlinear coefficient k_n . It indeed reveals that the simulated data delivers different levels of nonlinearity and the measured data are contaminated with noise.

The testing results of an exemplary realization are shown in Figure 3 for all three cases. In this figure, the label ‘‘FEM’’ indicates a linearized model of equation (11a) which is intentionally contaminated with 3% noise. The label ‘‘Hybrid model’’ denotes the proposed framework—neural modal ODEs.

As shown in Figure 3, the ‘‘FEM’’ model approximation does not well approximate the actual response. This is by design since we purposely added noise to the model in order to simulate modeling errors. The corresponding normalized root mean squared error (NRMSE) and R^2 for linear regression between true and predicted responses, both averaged by the dimension 12, are shown in Table 2. It is observed that although the model is recommended for use with linear or mildly nonlinear systems (e.g., $k_n = 0.0, 0.5$), it also performs satisfactorily for the system with relatively stronger nonlinearity ($k_n = 1.0$, which is comparable to the linear stiffness $k_1 = 1.0$). This is due to the adaption ability of the learning-based term, which is supposed to compensate the inaccuracy of the latent dynamics model f_{phy} , as well as to account for the imperfection of

Table 1. Implementation details for the numerical study.

	Encoder		Modeling latent dynamics	Decoder
	RNN ($\mathbf{x}_{0:n_t}$) $n_t = 10$	MLP (\mathbf{x}_0)	NN(\mathbf{z}) in equation (7)	Φ_p
No. of hidden layers	1	2	2	Invariant, $\Phi_p = [\phi_1, \phi_2, \phi_3, \phi_4]$
No. of neurons in each hidden layer	32	128	128	

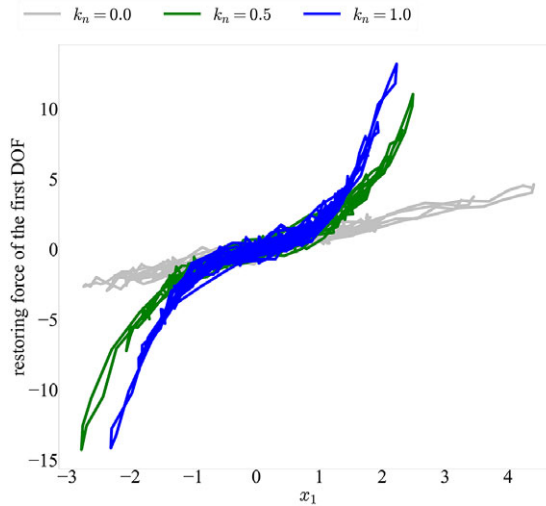


Figure 2. The force-displacement loops of the first DOF of the reference system for different values of the nonlinear coefficient k_n .

the decoder Φ_p . It is also understandable that when the system becomes nonlinear, the assumption that the decoder is invariant does not hold while the responses would become energy-dependent.

For $k_n = 1.0$ (Figure 3c), the recovery performance is not as good as in the other two cases. The recovered response for x_2 , in particular, is not perfectly aligned with the measured data. This implies that the decoder derived from the linear portion is not close to the actual one. Given the limited number of measurements (observations), the model returns a discrepancy with respect to the true model.

4. Illustration on a Model Cable-Stayed Bridge

In this section, the proposed framework is validated on a laboratory-based monitoring dataset derived from a scaled cable-stayed bridge, which was built and tested by the Research Division on Structural Control and Health Monitoring at Tongji University, China.

4.1. Experimental setup and data description

As shown in Figure 4a, this model bridge consists of one 6-m continuous beam, 2 towers, and 16 cables. The beam and towers are made of aluminum alloy, and additional metal weights are attached onto the beam and towers, ensuring that the scaled model's dynamic properties closely approximate those of the real cable-stayed bridge.

Cable-stayed bridges are known for exhibiting geometric nonlinearities. Generally speaking, the nonlinear effects in cable-stayed bridges include: (a) cable sag effects: cables sag because of their self-weight, resulting in variation of their axial stiffness; (b) P-delta effects: the horizontal components of cable forces bend the vertically compressed bridge pylon, introducing additional bending moments. The main girder of a cable-stayed bridge also suffers from the P-delta effects, where the bending girder is compressed by the horizontal components of cable forces; (c) large displacement effects: the displacement of the girder can be large, as the main girder of a cable-stayed bridge is mainly supported by flexible cables; thus, the small deformation assumption and linear beam theory do not apply in this scenario.

In this study, the model cable-stayed bridge exhibits nonlinearity in terms of the P-delta and large displacement effects, but these two effects are mild owing to the relatively small dimension of the scaled model. Further, the cable sag effects are negligible, as the steel cables are light. As a result, this scaled cable-stayed bridge model manifests mild nonlinearity and can be well approximated by the proposed scheme.

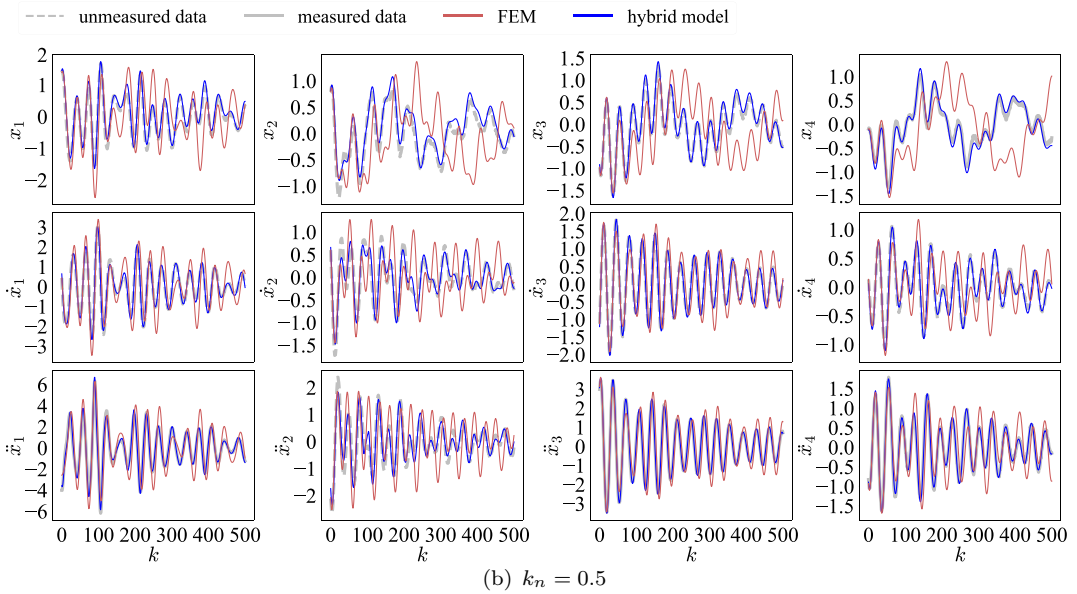
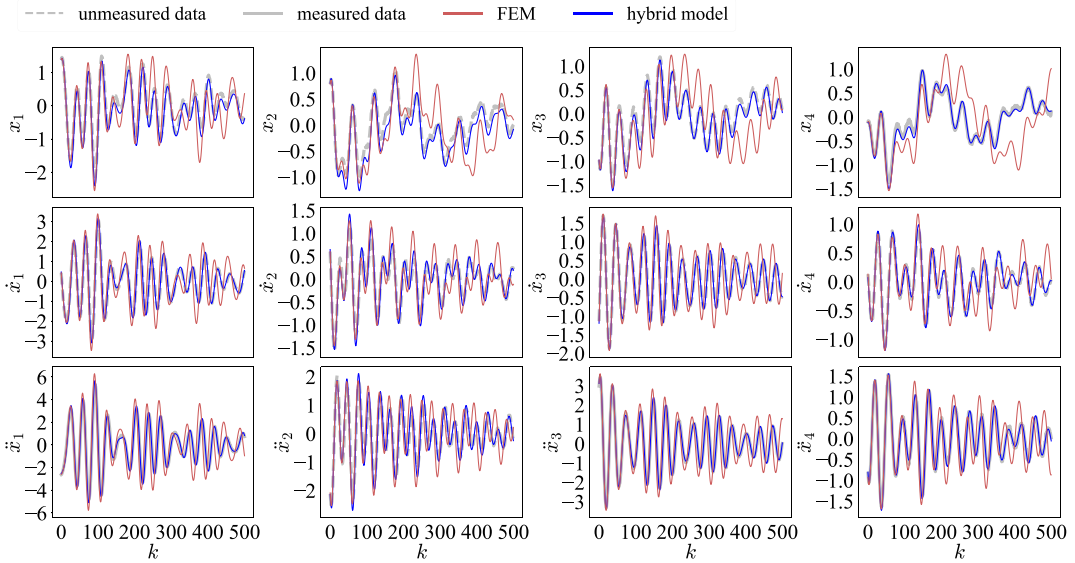


Figure 3. Recovered full-order response for the testing data set (only $\ddot{x}_1, \ddot{x}_3, \ddot{x}_4$, and x_4 are measured). (a) Linear case; (b) Mildly nonlinear case.

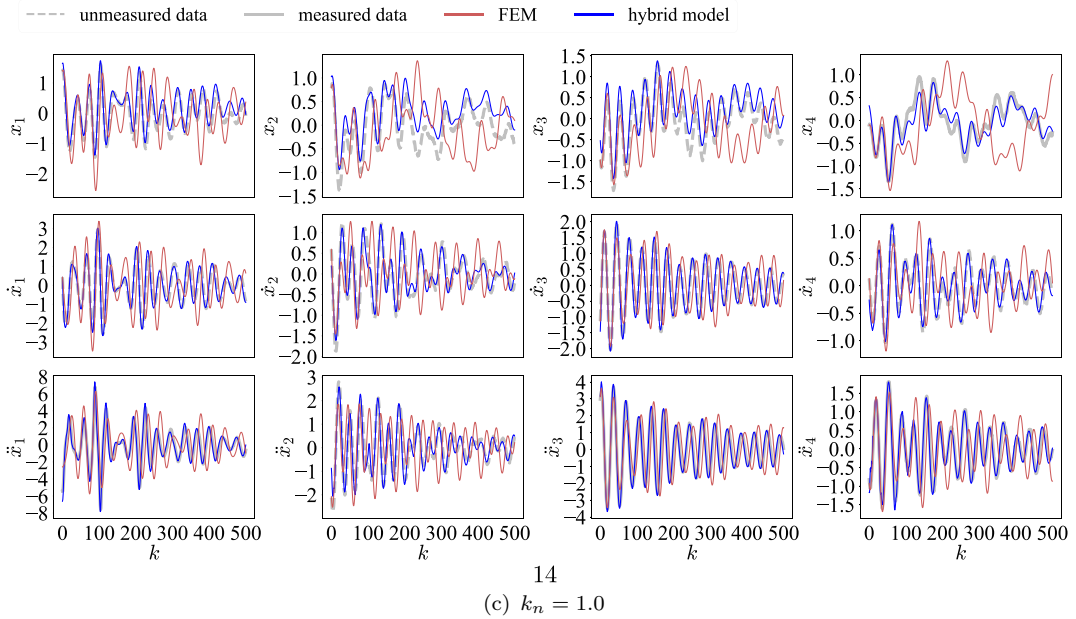


Figure 3. (Continued) (c) Nonlinear case.

Table 2. Performance metrics for the numerical study.

	Neural modal ODE		FEM	
	NRMSE	R^2	NRMSE	R^2
$k_n = 0.0$	0.0342	0.9760	0.1549	0.5635
$k_n = 0.5$	0.0496	0.9407	0.1823	0.4240
$k_n = 1.0$	0.0584	0.9431	0.2399	0.1244

To measure the dynamic response of the bridge model, as highlighted in Figure 4, eight MEMS (Micro-Electro-Mechanical System) accelerometers—labeled as A1–A8—are deployed on the structure, and a wired connection is used to collect the acceleration data to a digital data acquisition system. Acceleration measurements are collected at a sampling rate of 100 Hz, while the collected raw data is low-pass filtered at 30 Hz, as the dominant power in the spectrum of the raw signal lies below 30 Hz.

A “pull-and-release” action was used to excite the bridge model. A 1 kg iron weight was hung on node 19 with a wire. When the bridge model and the weight were both stationary, the wire was abruptly cut, inducing a damped free vibration of the bridge model, and those bridge responses were recorded with the accelerometers. It is worth mentioning that the weight was hanged at the exact lateral center of the beam, so the out-of-plane vibration such as torsion was supposed to be negligible.

Five repeated tests were performed. Four of these tests were used for training, with the remaining test serving as a testing dataset.

4.2. Finite element modeling

A two-dimensional (2-D) finite element model (FEM) of the scaled bridge has been developed, in a MATLAB environment (MATLAB, 2019), which serves as the physics-based model to be adopted within the proposed deep learning framework. We consider a 2-D model as the expected motion and the deployed

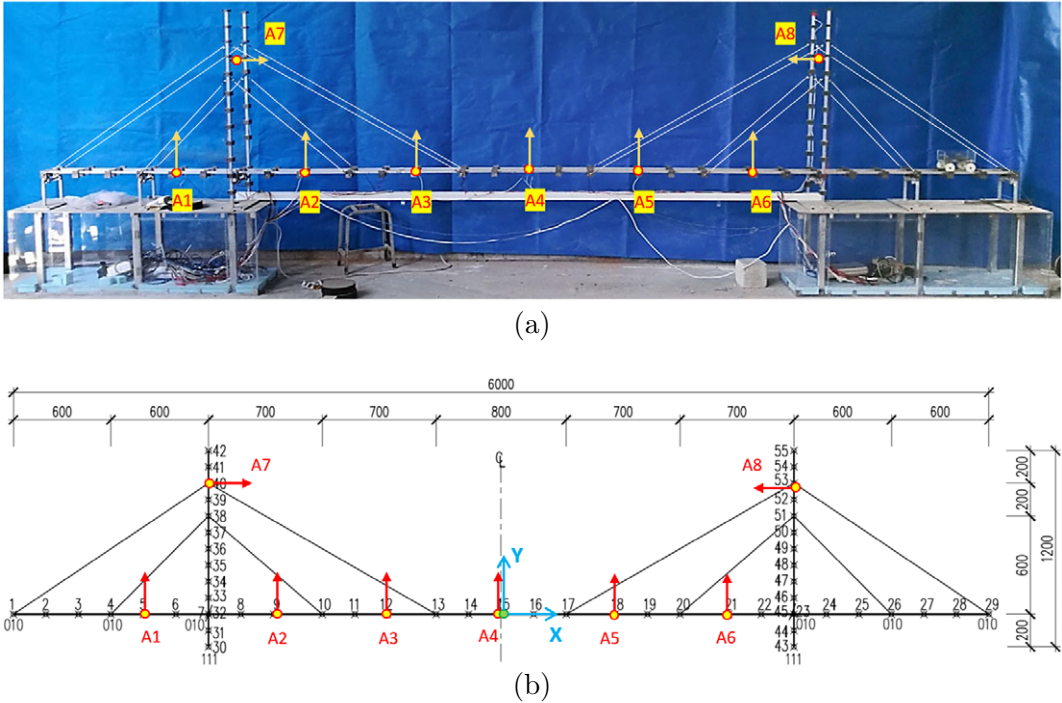


Figure 4. Scale model cable-stayed bridge: (a) in situ photo; (b) diagram of the finite element model (unit: mm). The eight deployed accelerometers are labeled as A1, A2, ..., A8, with arrows indicating the sensing directions; the coordinate system is defined by the direction of “X–Y” in the diagram.

sensors lie within a plane. The dimension, boundary conditions, node number, coordinate system, and sensor position of the FEM are displayed in Figure 4b. Each node corresponds to three DOFs: horizontal (x), vertical (y), and rotation. The notation “010” in Figure 4b signifies that vertical movement is restricted, while the horizontal and rotational movement are free (nodes 1, 4, 7, 23, 26, 29 are of this case); “111” signifies that all the three possible DOFs are restricted (nodes 30 and 43 are of this case). The beam and towers are simulated using the Euler–Bernoulli beam element, and the cables are modeled with the tension-only truss element. The total number of DOFs of the FEM model is 153, after applying the boundary conditions.

Eigenanalysis is performed on the FEM model of this bridge, and the first four mode shapes (ϕ_1 to ϕ_4) and corresponding frequencies ($\frac{\omega_1}{2\pi}$ to $\frac{\omega_4}{2\pi}$) are shown in Figure 5, which are a horizontal drifting mode (1.6387 Hz), followed by three vertical bending modes (3.4529, 6.3667, and 11.2516 Hz).

4.3. Model implementation

In this example, for modeling the latent dynamics via equation (4), we adopt the first 10 modes to construct the latent dynamics, that is, $p = 10$; $\Lambda = \text{diag}(\omega_1^2, \omega_2^2, \dots, \omega_{10}^2)$ and $\Gamma = \text{diag}(2\zeta_1\omega_1, 2\zeta_2\omega_2, \dots, 2\zeta_{10}\omega_{10})$. The decoder $\Phi_p = [\phi_1, \phi_2, \dots, \phi_{10}] \in \mathbb{R}^{153 \times 10}$, mapping the lower dimensional latent variables back to the full order of 153; ϕ_1 to ϕ_{10} are the first 10 mode shapes.

To train the model, the channels A1, and A3–A8 are used, while it is noted that the channel A2 is left out (considered as “unmeasured”) to be used for evaluating the performance of reconstruction, that is, the model uses the sensor data at a few DOFs to reconstruct a full-order response.

The data set includes multiple repeated free-vibration cases of the bridge, introduced by cutting a string that hangs a 1 kg mass on node 19. The whole data set is divided into batches for training the model, and the number of time steps for each batch is equally 500. Thus, for each batch, the initial conditions are

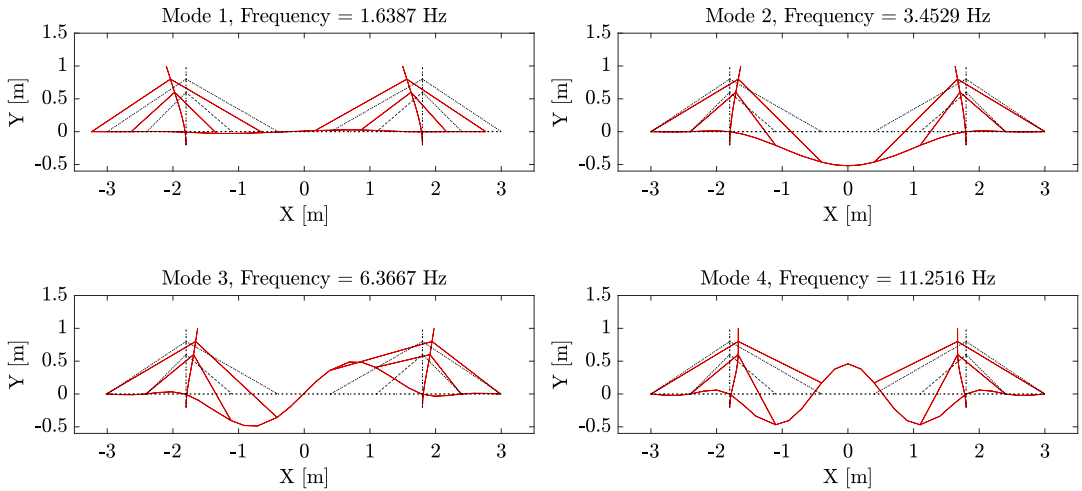


Figure 5. The first four mode shapes (denoted by red lines) derived from the eigenvalue analysis of the FEM model.

Table 3. Implementation details for the experimental study.

	Encoder		Modeling latent dynamics	Decoder
	RNN ($\mathbf{x}_{0:n_t}$) $n_t = 10$	MLP (\mathbf{x}_0)	NN(\mathbf{z}) in equation (4)	Φ_p
No. of hidden layers	1	2	2	Invariant, $\Phi_p = [\phi_1, \phi_2, \dots, \phi_{10}]$
No. of neurons in each hidden layer	128	128	128	

different, which is beneficial for training the encoder of the model. In addition, we normalize the measured acceleration across from A1 to A8, so that the maximum amplitude is 1.0, which is unitless. The details of the involved NNs are listed in Table 3.

4.4. Results

Once the model has been trained, the trained model is used for predicting the structural responses. The corresponding predictions of acceleration A1–A8 are shown in Figure 6, denoted by the blue lines. This prediction is compared with the actual measurements in gray color and predictions by the FEM models in red color. One can see that the FEM model offers satisfactory results, while some channel predictions are out of phase and fail to accurately follow the actual measurement, most possibly due to the inaccurate modeling of damping (this can be clearly observed in the A4 channel). The prediction from the proposed hybrid model is evidently more accurate than the FEM model, almost aligning with the actual measurements.

It is noted that the data of the A2 channel is unmeasured and not used for training the hybrid model, denoted by dashed gray lines. The prediction shown in the A2 plot comes from the full-order reconstructed responses. One can see that the reconstruction of A2 still highly agrees with the actual data, even though it is not used for the training.

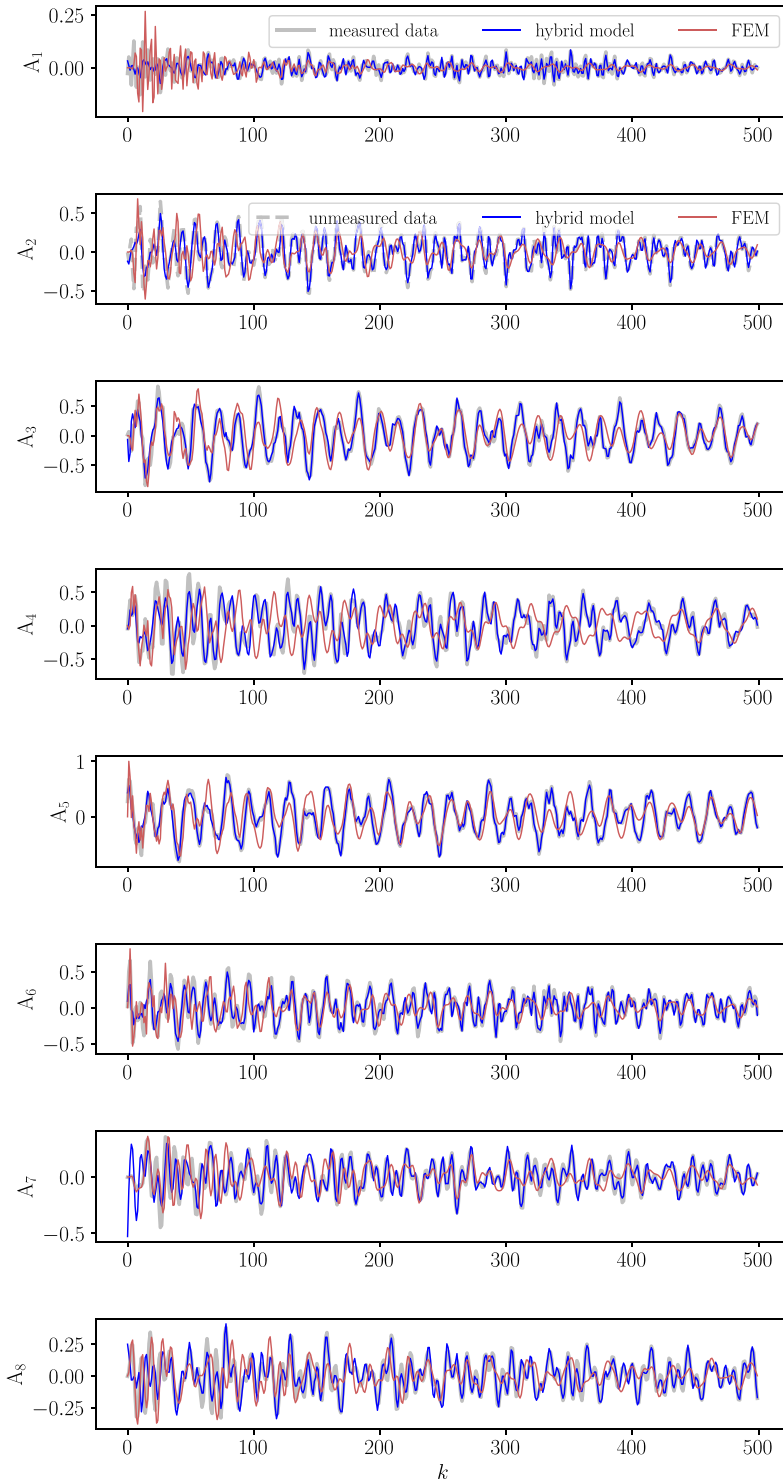


Figure 6. Comparisons of acceleration responses prediction between actual measurements, the proposed hybrid model (neural modal ODEs), and FEM model (A_1 – A_8 are normalized unitless data with maximum value of 1; the horizontal axis k denotes the time step).

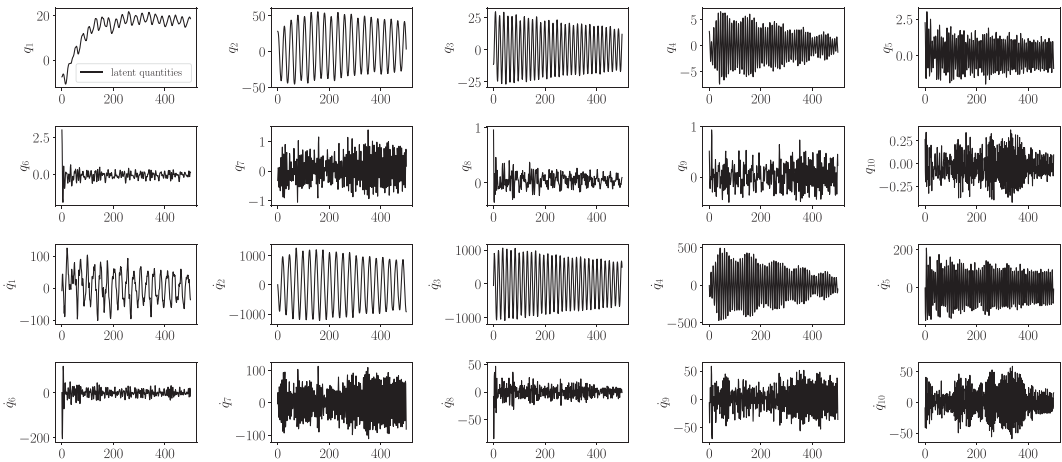


Figure 7. The learned latent variables $\mathbf{q} = [q_1, q_2, \dots, q_{10}]^T$ and $\dot{\mathbf{q}} = [\dot{q}_1, \dot{q}_2, \dots, \dot{q}_{10}]^T$ (the x -axis in each subplot is time step).

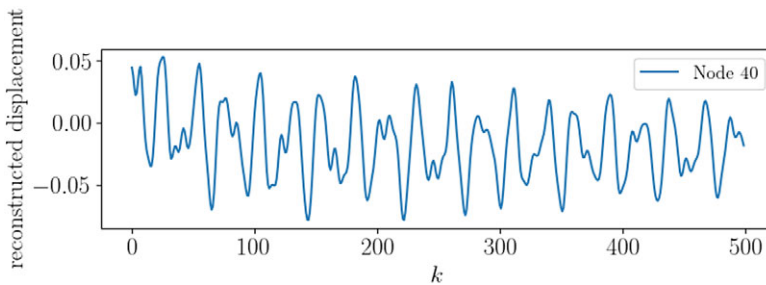


Figure 8. Reconstructed displacement at Node 40 (noted that in this example, as the acceleration is normalized, the reconstructed displacement is only a scaled version).

Figure 7 shows the corresponding learned time history of latent variables $\mathbf{q} = [q_1, q_2, \dots, q_{10}]^T$ and $\dot{\mathbf{q}} = [\dot{q}_1, \dot{q}_2, \dots, \dot{q}_{10}]^T$, related to displacement and velocity in modal coordinates, respectively. It is observed that: (a) q_1 to q_{10} retains the order from low-frequency to high-frequency, that we impose in the physics-informed term. In addition, these “modes” are near mono-frequent, almost preserving the decoupled structure; (b) by examining the amplitude of the latent variables, we are able to tell the contribution level of each mode. q_1, q_2, q_3 , and q_4 (hence, $\dot{q}_1, \dot{q}_2, \dot{q}_3$, and \dot{q}_4) have the highest amplitudes, dominating the vibration, while the amplitudes of other higher modes are much smaller (close to residuals). This is well understandable since for this free vibration, only the first several modes are fully excited while others are weakly present; (c) it is interesting to see that q_1 initiates from a value and then oscillates around an equilibrium which is not close to zero. q_1 is the modal displacement corresponding to the first horizontal drifting mode, which can only be picked up by A7 and A8 in the horizontal direction, at Nodes 40 and 53. We show that in Figure 8 as an example, after the decoder, the reconstructed displacement at Node 40 retains a reasonable vibration: initiating from a value and then oscillating around zero.

As stated in equation (6a), in this trained model, one has the flexibility of reconstructing different types of responses. For example, we reconstruct the full-order displacement responses via $\mathbf{x}_t^{\text{full}} = \Phi_p(\mathbf{q}_t)$.

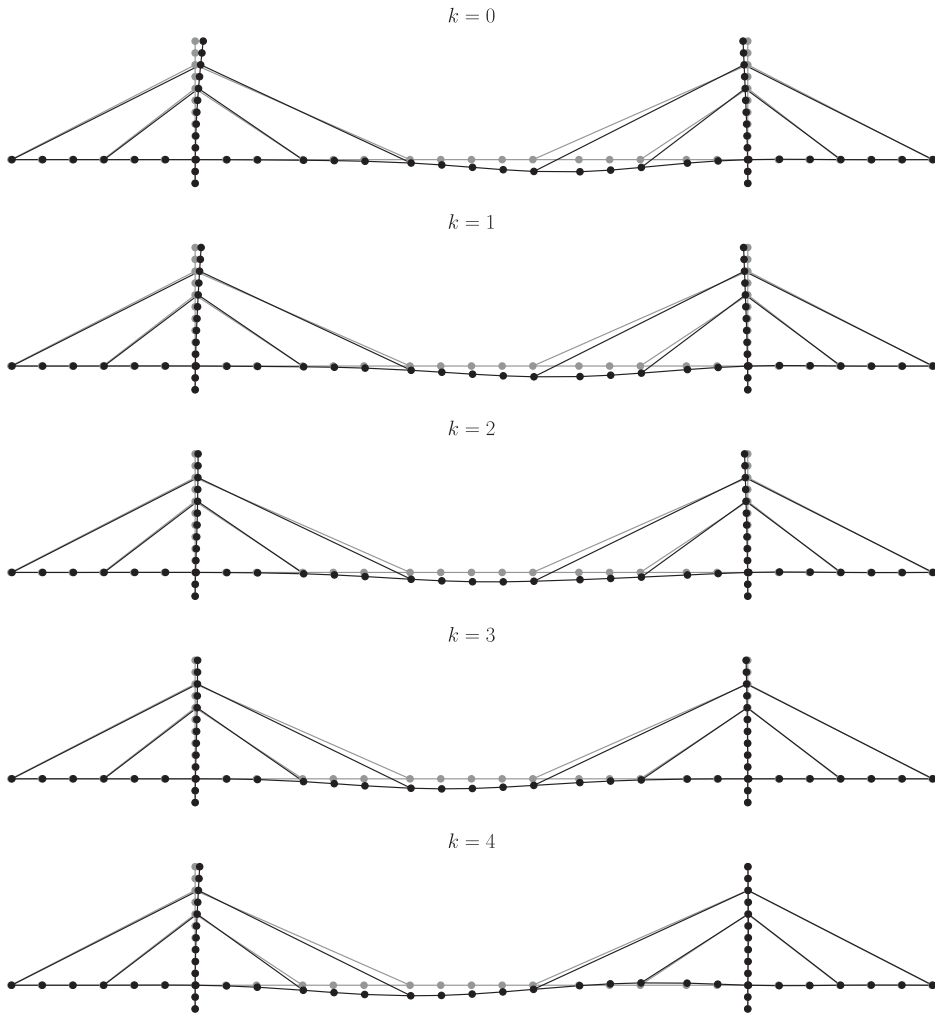


Figure 9. The reconstruction of the full-order displacement responses (using the first five snapshots as an example).

Figure 9 shows five consecutive snapshots of the full-order reconstructed displacements away from the equilibrium position, and a more intuitive video is provided in the auxiliary files. It is observed that the reconstruction preserves the legitimate spatial relationships between each node, due to the reason that the decoder is imposed by invariant normal modes. We also did an experiment using $\Phi_p + \text{NN}$ (normal modes added by a trainable NN to consider the imperfection of the normal modes) as a decoder. However, we find and conclude that this is not an appropriate decoder since the learned NN breaks the inherent spatial relationship between each node.

Since in this data set no reliable displacement measurements were achieved, in order to validate the accuracy of reconstructed full-order displacement from limited acceleration data, we compare the initial deformation ($k = 0$) with the one derived by the FEM model. The comparison result shown in Figure 10 indicates that the reconstructed displacement from the measured acceleration data highly agrees with the computed one by the FEM model. Thus, it is valid to see that the proposed hybrid model is capable of spatially extrapolating the dynamics and also of reconstructing other types of responses from a certain type of measurement (e.g., in this study case, the displacement and velocity are successfully reconstructed from the acceleration).




Figure 10. The comparisons of the initial deformation of the bridge between the reconstruction from the proposed hybrid model (neural modal ODEs) and FEM model.

5. Conclusions

In this article, we propose a framework for integrating physics-based modeling with deep learning for modeling large civil/mechanical dynamical systems. The framework couples a dynamical VAE with a Physics-informed Neural ODE scheme. The autoencoder encodes a limited amount of sensed data into an estimate of the initial conditions of the latent space. This allows for the construction of a generative model which aims at predicting the latent system dynamics via a learned Physics-informed Neural ODE. The predicted dynamic response is then mapped back onto the measured physical space via an invariant decoder, which is effectuated on the basis of the eigenmodes derived from a physics-based model. The framework assimilates physics-related features from a physics-based model into a deep learning model, to yield a learned generative model, which is not eventually data-dependent and leads to an interpretable architecture. The delivered models are able to reconstruct the full field structural response, meaning response in unmeasured locations, given limited sensing locations. Future work will investigate boosting the decoder via assimilation of a Bayesian NN.

6. Discussions

We want to further clarify that the extrapolation capability cannot be guaranteed if the dynamic regime differs significantly from the training data we used to train the model, which is also typically the limitation of most deep learning methods. From the numerical study, it can be observed that the proposed framework is able to capture unseen scenarios, when these do not excite a significantly higher level of nonlinearity. This is why, we comment on the framework being applicable for mildly nonlinear systems, implying that in presence of severe nonlinearity the extrapolation potential is limited.

Supplementary Materials. To view supplementary material for this article, please visit <http://doi.org/10.1017/dce.2022.35>.

Author Contributions. Z.L.: Conceptualization (Lead), Data curation (Supporting), Formal analysis (Lead), Investigation (Lead), Methodology (Lead), Software (Lead), Validation (Lead), Visualization (Lead), Writing—original draft (Lead), and Writing—review and editing (Lead); W.L.: Formal analysis (Supporting), Investigation (Supporting), Software (Supporting), Validation (Supporting), and Writing—original draft (Supporting); X.J.: Data curation (Supporting); K.B.: Software (Supporting); L.S.: Data curation (Lead); E.C.: Conceptualization (Supporting), Investigation (Supporting), Methodology (Supporting), Project administration (Lead), Resources (Lead), Supervision (Lead), Validation (Equal), Writing—original draft (Supporting), and Writing—review and editing (Equal).

Competing Interests. The authors declare no competing interests exist.

Data Availability Statement. A demonstrative code (Section 3) that implements the proposed method is openly available at <https://github.com/zlaidyn/Neural-Modal-ODE-Demo>.

Funding Statement. The research was conducted at the Singapore-ETH Centre, which was established collaboratively between ETH Zurich and the National Research Foundation Singapore. This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) program.

References

- Abgrall R, Amsallem D and Crisovan, R** (2016) Robust model reduction by L^1 -norm minimization and approximation via dictionaries: Application to nonlinear hyperbolic problems. *Advanced Modeling and Simulation in Engineering Sciences* 3(1) doi 10.1186/s40323-015-0055-3, <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84997090816&doi=10.1186/s40323-015-0055-3&partnerID=40&md5=111e4c6afabc3ea7450515a4dc4dc1d>.
- Amsallem D, Zahr MJ and Washabaugh K** (2015) Fast local reduced basis updates for the efficient reduction of nonlinear systems with hyper-reduction. *Advances in Computational Mathematics* 41(5), 1187–1230. <http://doi.org/10.1007/s10444-015-9409-0>
- Bae HJ and Koumoutsakos P** (2022) Scientific multi-agent reinforcement learning for wall-models of turbulent flows. *Nature Communications* 13(1), 1–9.
- Balajewicz M, Amsallem D and Farhat C** (2016) Projection-based model reduction for contact problems. *International Journal for Numerical Methods in Engineering* 106(8), 644–663.
- Carlberg K, Farhat C, Cortial J and Amsallem D** (2013) The gnat method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows. *Journal of Computational Physics* 242, 623–647.
- Chen RT, Rubanova Y, Bettencourt J and Duvenaud D** (2018) Neural ordinary differential equations. arXiv preprint, arXiv:1806.07366.
- Craig RR and Kurdila AJ** (2006) *Fundamentals of Structural Dynamics*. Hoboken, NJ: John Wiley & Sons.
- Cranmer M, Greydanus S, Hoyer S, Battaglia P, Spergel D and Ho S** (2020) Lagrangian neural networks. arXiv preprint, arXiv:2003.04630.
- Farrar CR and Worden K** (2012) *Structural Health Monitoring: A Machine Learning Perspective*. Hoboken, NJ: John Wiley & Sons.
- Girin L, Leglaive S, Bie X, Diard J, Hueber T and Alameda-Pineda X** (2020) Dynamical variational autoencoders: A comprehensive review. arXiv preprint, arXiv:2008.12595.
- Hinton GE and Zemel RS** (1994) Autoencoders, minimum description length, and Helmholtz free energy. *Advances in Neural Information Processing Systems* 6, 3–10.
- Kalman RE** (1960) On the general theory of control systems. In *Proceedings First International Conference on Automatic Control*. Moscow: USSR, pp. 481–492.
- Kamariotis A, Chatzi E and Straub D** (2022) Value of information from vibration-based structural health monitoring extracted via Bayesian model updating. *Mechanical Systems and Signal Processing* 166, 108465.
- Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S and Yang L** (2021) Physics-informed machine learning. *Nature Reviews Physics* 3(6), 422–440.
- Karpatne A, Watkins W, Read J and Kumar V** (2017) Physics-guided neural networks (PGNN): An application in lake temperature modeling. arXiv preprint, arXiv:1710.11431.
- Kashinath K, Mustafa M, Albert A, Wu JL, Jiang C, Esmailzadeh S, Azizzadenesheli K, Wang R, Chattopadhyay A, Singh A, Manepalli A, Chirila D, Yu R, Walters R, White B, Xiao H, Tchelepi HA, Marcus P, Anandkumar A, Hassanzadeh P and Prabhat** (2021) Physics-informed machine learning: Case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A* 379(2194), 20200093.
- Kingma DP and Welling M** (2013) Auto-encoding variational bayes. arXiv preprint, arXiv:1312.6114.
- Krishnan RG, Shalit U and Sontag D** (2015) Deep kalman filters. arXiv preprint, arXiv:1511.05121.
- Krishnan R, Shalit U and Sontag D** (2017) Structured inference networks for nonlinear state space models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31. San Francisco, California USA: AAAI PRESS.
- Lai Z, Alzugaray I, Chli M and Chatzi E** (2020) Full-field structural monitoring using event cameras and physics-informed sparse identification. *Mechanical Systems and Signal Processing* 145, 106905.
- Lai Z, Mylonas C, Nagarajaiah S and Chatzi E** (2021) Structural identification with physics-informed neural ordinary differential equations. *Journal of Sound and Vibration* 508, 116196.
- Lai Z and Nagarajaiah S** (2019) Sparse structural system identification method for nonlinear dynamic systems with hysteresis/inelastic behavior. *Mechanical Systems and Signal Processing* 117, 813–842.
- Liang Y, Lee H, Lim S, Lin W, Lee K and Wu C** (2002) Proper orthogonal decomposition and its applications—Part I: Theory. *Journal of Sound and Vibration* 252(3), 527–544.
- Liu W, Lai Z, Bacsá K and Chatzi E** (2022) Physics-guided deep markov models for learning nonlinear dynamical systems with uncertainty. *Mechanical Systems and Signal Processing* 178, 109276.
- Lusch B, Kutz JN and Brunton SL** (2018) Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications* 9(1), 1–10.
- Marconia J, Tisob P, Quadrellia DE and Braghina F** (2021) An enhanced parametric nonlinear reduced order model for imperfect structures using neumann expansion. arXiv preprint, arXiv:2102.01739.
- MATLAB** (2019) *R2019b*. Natick, MA: The MathWorks Inc.
- Ou Y, Tatsis KE, Dertimanis VK, Spiridonakos MD and Chatzi EN** (2021) Vibration-based monitoring of a small-scale wind turbine blade under varying climate conditions. Part I: An experimental benchmark. *Structural Control and Health Monitoring* 28(6), e2660.
- Peherstorfer B and Willcox K** (2016) Dynamic data-driven model reduction: Adapting reduced models from incomplete data. *Advanced Modeling and Simulation in Engineering Sciences* 3(1), 1–22.

- Qian E, Kramer B, Peherstorfer B and Willcox K** (2020) Lift & learn: Physics-informed machine learning for large-scale nonlinear dynamical systems. *Physica D: Nonlinear Phenomena* 406, 132401.
- Raissi M, Perdikaris P and Karniadakis GE** (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378, 686–707.
- Roehrl MA, Runkler TA, Brandstetter V, Tokic M and Obermayer S** (2020) Modeling system dynamics with physics-informed neural networks based on Lagrangian mechanics. *IFAC-PapersOnLine* 53(2), 9195–9200.
- Sankararaman S and Mahadevan S** (2013) Bayesian methodology for diagnosis uncertainty quantification and health monitoring. *Structural Control and Health Monitoring* 20(1), 88–106.
- Satake N, Suda K-I, Arakawa T, Sasaki A and Tamura Y** (2003) Damping evaluation using full-scale data of buildings in Japan. *Journal of Structural Engineering* 129(4), 470–477.
- Schmid PJ** (2010) Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics* 656, 5–28.
- Simpson T, Dervilis N and Chatzi E** (2021) Machine learning approach to model order reduction of nonlinear systems via autoencoder and lstm networks. *Journal of Engineering Mechanics* 147(10), 04021061.
- Strppmen EN** (2014) The finite element method in dynamics. In *Structural Dynamics*. Cham: Springer International, pp. 161–204. https://doi.org/10.1007/978-3-319-01802-7_4
- Sun L, Li Y, Zhu W and Zhang W** (2020) Structural response reconstruction in physical coordinate from deficient measurements. *Engineering Structures* 212, 110484.
- Swischuk R, Kramer B, Huang C and Willcox K** (2020) Learning physics-based reduced-order models for a single-injector combustion process. *AIAA Journal* 58(6), 2658–2672.
- Tatsis K, Agathos K, Chatzi E and Dertimanis V** (2022) A hierarchical output-only Bayesian approach for online vibration-based crack detection using parametric reduced-order models. *Mechanical Systems and Signal Processing* 167, 108558. <https://doi.org/10.1016/j.ymsp.2021.108558>. Available at <https://www.sciencedirect.com/science/article/pii/S0888327021008967>.
- Vettori S, DiLorenzo E, Peeters B and Chatzi E** (2022) Virtual sensing for wind turbine blade full field response estimation in operational modal analysis. In *Model Validation and Uncertainty Quantification*, vol. 3. Cham: Springer, pp. 49–52.
- Vlachas PR, Arampatzis G, Uhler C and Koumoutsakos P** (2022) Multiscale simulations of complex systems by learning their effective dynamics. *Nature Machine Intelligence* 4(4), 359–366. <http://doi.org/10.1038/s42256-022-00464-w>
- Vlachas K, Tatsis K, Agathos K, Brink AR and Chatzi E** (2021) A local basis approximation approach for nonlinear parametric model order reduction. *Journal of Sound and Vibration* 502, 116055.
- Vlachas K, Tatsis K, Agathos K, Brink AR, Quinn D and Chatzi E** (2022) On the coupling of reduced order modeling with substructuring of structural systems with component nonlinearities. In *Dynamic Substructures*, vol. 4. Cham: Springer, pp. 35–43.
- Wagg D, Worden K, Barthorpe R and Gardner P** (2020) Digital twins: State-of-the-art and future directions for modeling and simulation in engineering dynamics applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering* 6(3), 030901.
- Waisman H, Chatzi E and Smyth AW** (2010) Detection and quantification of flaws in structures by the extended finite element method and genetic algorithms. *International Journal for Numerical Methods in Engineering* 82(3), 303–328.
- Willard J, Jia X, Xu S, Steinbach M and Kumar V** (2020) Integrating physics-based modeling with machine learning: A survey. *arXiv preprint arXiv:2003.04919* 1(1), 1–34.
- Wu JL, Xiao H and Paterson E** (2018) Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Physical Review Fluids* 3(7), 074602.
- Yildiz C, Heinson M and Lähdesmäki H** (2019) Ode2vae: Deep generative second order odes with bayesian neural networks.
- Zhang R, Liu Y and Sun H** (2020) Physics-guided convolutional neural network (phycnn) for data-driven seismic response modeling. *Engineering Structures* 215, 110704.
- Zhu Y, Zabarar N, Koutsourelakis P-S and Perdikaris P** (2019) Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics* 394, 56–81.

Cite this article: Lai Z, Liu W, Jian X, Bacsá K, Sun L and Chatzi E (2022). Neural modal ordinary differential equations: Integrating physics-based modeling with neural ordinary differential equations for modeling high-dimensional monitored structures. *Data-Centric Engineering*, 3, e34. doi:10.1017/dce.2022.35