



Physics guided neural networks for modelling of non-linear dynamics

Haakon Robinson^a, Suraj Pawar^b, Adil Rasheed^{a,c,*}, Omer San^{b,1}

^a Department of Engineering Cybernetics, Norwegian University of Science and Technology, O. S. Bragstads plass 2, Trondheim, NO-7034, Norway

^b School of Mechanical and Aerospace Engineering, Oklahoma State University, 201 General Academic Building Stillwater, OK 74078, United States of America

^c Mathematics and Cybernetics, SINTEF Digital, Klæbuveien 153, Trondheim, NO-7037, Norway

ARTICLE INFO

Article history:

Received 23 September 2021

Received in revised form 16 May 2022

Accepted 19 July 2022

Available online 26 July 2022

Keywords:

Physics guided neural networks

Non-linear dynamics

Ordinary differential equations

ABSTRACT

The success of the current wave of artificial intelligence can be partly attributed to deep neural networks, which have proven to be very effective in learning complex patterns from large datasets with minimal human intervention. However, it is difficult to train these models on complex dynamical systems from data alone due to their low data efficiency and sensitivity to hyperparameters and initialisation. This work demonstrates that injection of partially known information at an intermediate layer in a DNN can improve model accuracy, reduce model uncertainty, and yield improved convergence during the training. The value of these *physics-guided neural networks* has been demonstrated by learning the dynamics of a wide variety of nonlinear dynamical systems represented by five well-known equations in nonlinear systems theory: the Lotka–Volterra, Duffing, Van der Pol, Lorenz, and Henon–Heiles systems.

© 2022 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

A dynamical system is a system whose state varies over time and obeys differential equations that involve time derivatives. The equations can either be analytically or numerically solved to predict the future state of the system. The evolution of the weather, progression of chemical reactions, the spread of diseases, and dynamics of vehicles can all be modelled as dynamical systems.

Dynamical equations are typically derived from first principles using well understood physical laws. In this work, we call this approach physics-based modelling (PBM), and also use PBM to refer to the model itself (Fig. 1(a)). In developing a PBM, the observed physics are related to theory to produce equations, assumptions are made, and simplifications are sometimes imposed to make the solutions computationally tractable. PBMs typically have sound foundations from first principles, are interpretable, generalise well, and there exist robust theories for analysing properties such

as stability and uncertainty. However, by relying on assumptions and simplifications, as well as our limited understanding of complex physical processes, we run the risk of not describing the desired phenomena with sufficient accuracy and certainty. Many PBMs do not account for unknown/unresolved physics, can be computationally expensive, do not adapt to new scenarios automatically and can be susceptible to numerical instabilities.

Data-driven modelling (DDM) (see Fig. 1(b)) is rapidly emerging as a tool that can address problems that resist traditional modelling methods, and some even controversially regard it as a full replacement for PBM (Karpathy, 2017). These models can learn both known and unknown physics directly from data without prior knowledge of any physical laws, achieving good performance while maintaining computational efficiency for inference. Deep neural networks (DNNs) in particular have enabled superhuman performance in tasks long considered impossible to solve for computers, such as the game of Go (Silver et al., 2016). Works such as the protein-structure model proposed by Senior et al. (2020) show that the use of DNNs has also begun to penetrate into scientific applications. Within the realm of dynamical systems, DDMs have been demonstrated to learn dynamics directly from the data. For example, Saha et al. (2021) use physics-incorporated convolutional recurrent neural networks for dynamical systems forecasting and source identification.

Despite these advantages, there remain some challenges before these models can find their way into high stake or safety-critical applications. They typically require vast amounts of data to train and can struggle to generalise to situations not represented well by the data. In contrast to established PBM methods,

* Corresponding author at: Department of Engineering Cybernetics, Norwegian University of Science and Technology, O. S. Bragstads plass 2, Trondheim, NO-7034, Norway.

E-mail addresses: haakon.robins@ntnu.no (H. Robinson), supawar@okstate.edu (S. Pawar), adil.rasheed@ntnu.no (A. Rasheed), osan@okstate.edu (O. San).

¹ The third and fourth authors are grateful for the support received by the Research Council of Norway and the industrial partners of the following projects: EXAIGON—Explainable AI systems for gradual industry adoption (Grant No. 304843), the U.S. DOE Early Career Research Program support (DE-SC0019290), and the National Science Foundation support (DMS-2012255).

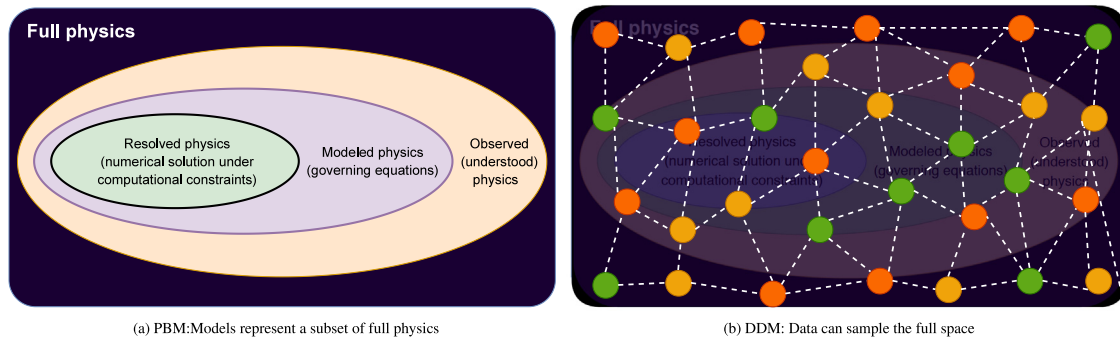


Fig. 1. Building a model naturally requires assumptions and simplifications. The result is that of the physics we can observe and understand, only a small part of the physics can be described using models, and even less can be numerically simulated. In contrast, large datasets can cover the full space, enabling general ML models to provide predictions in the absence of understanding or models.

there is a lack of robust theory for the analysis of properties such as stability and robustness, and practitioners often have to fall back on empirical testing to assure the safety of their models.

To counter the issues associated with both the PBM and DDM, a new paradigm is emerging, which we call Hybrid Analysis and modelling (HAM). The HAM approach combines the generalisability, interpretability, robust foundation, and understanding of PBM with the accuracy, computational efficiency, and automatic pattern-identification capabilities of DDM. In their recent surveys, Willard et al. (2020) and San et al. (2021) provide comprehensive overviews of techniques for integrating DDM with PBM. Many of the hybridisation techniques fall into the following categories: (i) Embedding PBMs inside NNs, (ii) Model order reduction, (iii) Physics-based regularisation terms, (iv) Data-driven equation discovery, (v) Error correction approaches, and (vi) Sanity check mechanisms using PBMs.

In the following sections, we present related work and discuss the advantages and disadvantages of the approaches.

1.1. Methods for embedding PBMs directly into NNs

This is perhaps the most straightforward approach to hybridisation. More advanced methods typically create a differentiable PBM that can be used as a layer in a NN. An example of this is OptNet proposed by Amos and Kolter (2017), a differentiable convex optimisation solver that can be used as a layer in a network. In related work, de Avila Belbute-Peres et al. (2018) propose the differentiable physics engine, a rigid body simulator that can be embedded into a NN. They demonstrate that it is possible to learn a mapping from visual data to the positions and velocities of objects, which are then updated using the simulator. Similar ideas are used by Yu et al. (2020) to simulate a structural dynamics problem by designing a hybrid recurrent NN (RNN) that contains an implicit numerical integrator. The advantage of these approaches is that they are usually quite data-efficient. A challenge is that these embedded PBMs are often iterative methods, making both inference and training more expensive.

1.2. Model order reduction methods

Reduced-order modelling (ROM) is a successful and widely adopted methodology (Quarteroni & Rozza, 2014). A ROM method typically projects complex partial differential equations onto a lower dimensional space based on the singular value decomposition of the offline high fidelity simulation data. This yields a set of ordinary differential equations (ODEs) that can be efficiently solved (Ahmed et al., 2021). ROMs have been used to accelerate high-fidelity numerical solvers by several orders of magnitude (Fonn et al., 2019). However, ROMs tend to become unstable

in the presence of unknown/unresolved complex physics. To alleviate these problems, recent research has shown how unknown and hidden physics within a ROM framework can be accounted for using DNNs (Pawar et al., 2020; Pawar et al., 2020). Despite these benefits, ROMs require full knowledge of the original equation before they can be applied.

1.3. Physics-based regularisation terms

By incorporating a PBM in the objective function, DDMs can be biased towards known physical laws during training. A recent work (Raissi et al., 2019) that has seen a lot of interest is the physics-informed neural network (PINN), where a NN is used to represent the solution to a PDE and deviations from the equation at a sample of points are penalised by an additional loss term. PINNs can be used to solve problems such as heat transfer, as was done by Zobeiry and Humfeld (2021) for parts in a manufacturing process. The PINN approach has also been extended by Arnold and King (2021) to allow for control in a state-space setting. In related work, Shen et al. (2021) create a model for classifying bearing health by training a NN on physics-based features and regularising the model using the output from a physics-based threshold model. These approaches require precise knowledge of the loss term. Complex regularisation terms may impact the training process due to the increased cost of computing the loss function.

1.4. Data-driven equation discovery

Sparse regression based on l_1 regularisation and symbolic regression based on gene expression programming have been shown to be very effective in discovering hidden or partially known physics directly from data. Notable work using this approach can be found in Champion et al. (2019) and Vaddireddy et al. (2020). One of the limitations of this class of method is that, in the case of sparse regression, additional features are required to be handcrafted, while in the case of symbolic regression, the resulting models can be often unstable and prone to overfitting.

1.5. Error correction and sanity check mechanisms

Corrective Source Term Approach (CoSTA) is a method proposed by Blakseth et al. (2022) that explicitly addresses the problem of unknown physics. This is done by augmenting the governing equations of a PBM with a DNN-generated corrective source term that takes into account the remaining unknown/neglected physics. One added benefit of the CoSTA approach is that the physical laws can be used to keep a sanity check on the predictions of the DNN used, i.e. checking conservation laws. A similar

approach has also been used to model unresolved physics in turbulent flows (Maulik et al., 2019; Pawar et al., 2020). However, even these approaches assume a specific structure for at least the known part of the equation.

1.6. Proposal: Physics-guided neural networks

From the previous discussion, it is clear that almost all the HAM approaches discussed above require information about the structure of the equation representing the physics, which is not always available. We often have a very simplistic understanding of the physics. For example, we can have some understanding of the diurnal variation of solar radiation but not about its influences on atmospheric flow. To exploit even such a small amount of knowledge, Pawar et al. (2021a) proposed a physics guided machine learning (PGML) approach. The basic idea behind the PGML approach is to inject partial knowledge into one of the layers within a DNN to guide the training process. The partial knowledge can, for example, come from a simplistic model or an empirical law (Pawar et al., 2021b, 2022).

This paper extends the PGML concept to modelling nonlinear dynamical systems. Since we limit the model space to neural networks, we call the approach physics-guided neural network (PGNN). Through a series of experiments involving a variety of equations representing nonlinear dynamical systems like Lotka–Volterra, Duffing, Van der Pol, Lorenz, and Henon–Heiles equations, we attempt to answer the following questions:

- What are the effects of knowledge injection on the training convergence?
- How does the accuracy/performance change with the choice of injection layer?
- Is there any correlation between model uncertainty and knowledge injection?

A brief background and rationale for the proposed method is given in Section 2. Section 3 details the selected dynamical systems that are considered in our study. Finally, the results are discussed in Section 4, and conclusions and recommendation for future work made in Section 5.

2. Physics-guided neural networks

The basic idea behind PGNN is to generalise the Principal Component Regression (PCR). In PCR, instead of regressing the dependent variable on the explanatory variables directly, the latent variables derived from the explanatory variables after the application of PCA are used as regressors. By replacing the high dimensional explanatory variables (containing redundancy) with much lower dimensional latent variables as the input to the regression model, one can significantly reduce the complexity of the regressors and make them more robust. However, there are two major problems associated with PCR. Firstly, the latent variables computed using PCA can only be a linear combination of the explanatory variables. Secondly, the regression task is decoupled from the latent variable computation.

In the PGNN approach, both the computation of the latent variables and the regression are combined within a neural network framework with bottleneck layers representing the latent variable layers. Additionally, the latent variables can be supplemented with additional features (partial knowledge) to improve the accuracy and reduce the uncertainty of the trained PGNN model. If the additional features were used in combination with the explanatory features as input to the DNN, chances are high that they would get corrupted during the training process.

We now present the rationale behind the PGNN approach to modelling nonlinear dynamics. Deep learning has been recently

used in many studies to model the spatio-temporal dynamics of high-dimensional systems (Pathak et al., 2018; Pawar et al., 2019; Vlachas et al., 2018). Given some dynamical system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, a NN can be trained directly on the mapping $\mathbf{f}(\cdot)$ by sampling repeatedly from the system. After training, the network can then be numerically integrated in order to perform predictions on the future states of the system, e.g. by computing the forward Euler step $\mathbf{x}_{k+1} = \mathbf{x}_k + h \mathbf{f}(\mathbf{x}_k)$. Consider a dataset generated by a more general dynamical system:

$$\mathcal{L}\mathbf{x} = f(g(\mathbf{x}), h(\mathbf{x})) \tag{1}$$

where \mathcal{L} is a linear differential operator, and $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$ are functions of the state. Now, the following scenarios can arise:

1. Eq. (1) is fully known meaning that the operator \mathcal{L} , and the functions $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$ are precisely known
2. The operator \mathcal{L} is known but one or two of the functions $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$ are unknown
3. The operator \mathcal{L} is known, but the functions $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$ are all unknown
4. The operator \mathcal{L} and the functions $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$ are all unknown.

In the first scenario, one can employ a purely PBM approach based on the known equations. The only advantage of DDM over PBM is possibly superior computational performance, which may enable real-time applications. In the second and third scenarios, while the problem can in theory be solved entirely using DDM, it would be unwise to ignore the known part completely. Incorporating them into the DDM may simplify the learning task as well as improve generalisation. In the fourth scenario, PBM is impossible, and it is necessary to model the process entirely from data.

Assume now that only $h(x)$ is known in Eq. (1). Then (1) can be learned using a PGNN with an $h(x)$ injected at hidden layer of the neural network, as shown in Fig. 2. By stacking known features into an intermediate layer, they can be utilised more effectively. The significance of which layer is used for injection is unknown.

Here, we briefly explain the NN and PGNN architecture. A neural network is designed using several layers consisting of a predefined number of neurons. Each neuron has a weighted connection to all neurons in the previous layer, and a bias term. This is represented as an affine transformation as shown below

$$\mathbf{z}^l = \mathbf{W}^l \boldsymbol{\chi}^{l-1} + \mathbf{b}^l, \tag{2}$$

where $\boldsymbol{\chi}^{l-1}$ is the output of the $(l - 1)$ th layer, \mathbf{W}^l is the matrix of weights representing the incoming connection strengths the l th layer, and \mathbf{b}^l is the bias vector. For notational simplicity we define $\boldsymbol{\chi}^0 = \mathbf{x}$. The transformed input is then passed through a node's activation function ζ , which is some nonlinear function. The introduction of this nonlinearity prevents the chain of affine transformations from simplifying and allows the neural network to learn highly complex relations between the input and output. The output of the l th layer can be written as

$$\boldsymbol{\chi}^l = \zeta(\mathbf{z}^l), \quad \boldsymbol{\chi}^0 = \mathbf{x} \tag{3}$$

where ζ is the activation function. Some possible choices are the ReLU, tanh, and sigmoid activation functions. We refer the reader to Goodfellow et al. (2016) for a more complete overview. If there are L layers between the input and the output in a neural network, then the output of the neural network can be represented mathematically as follows

$$\dot{\mathbf{x}} = \zeta_L(\mathbf{W}^L, \mathbf{b}^L, \dots, \zeta_2(\mathbf{W}^2, \mathbf{b}^2, \zeta_1(\mathbf{W}^1, \mathbf{b}^1, \mathbf{x}))) \tag{4}$$

where \mathbf{x} and $\dot{\mathbf{x}}$ are the independent and dependent variables of the system, respectively. The above equation can also be written as

$$\dot{\mathbf{x}} = \zeta_L(\cdot; \Theta_L) \circ \dots \circ \zeta_2(\cdot; \Theta_2) \circ \zeta_1(\cdot; \Theta_1) \tag{5}$$

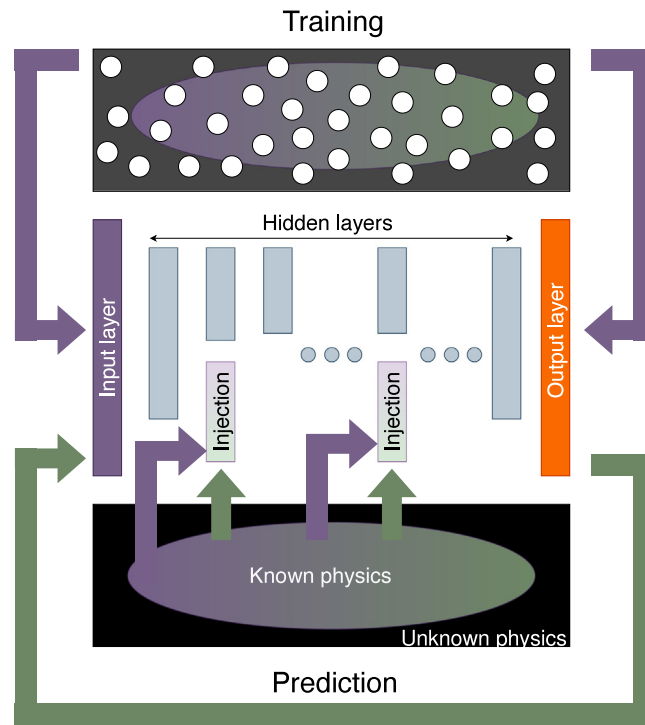


Fig. 2. PGNN framework: The purple arrows correspond to the training phase while the green arrows correspond to the prediction phase. (Assume that the data is represented by the white circles.) (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

where Θ represents the weights and biases of the corresponding layer of the neural network. For the PGNN framework, the information from the known part of the system is injected into an intermediate layer of the neural network as follows

$$\dot{\mathbf{x}} = \zeta_L(\cdot; \Theta_L) \circ \dots \circ \underbrace{\mathcal{C}(\zeta_i(\cdot; \Theta_i), h(\mathbf{x}))}_{\text{Known function injection}} \circ \dots \circ \zeta_1(\mathbf{x}; \Theta_1), \quad (6)$$

where $\mathcal{C}(\cdot, \cdot)$ represents the concatenation operation and the known information about the system, i.e., $h(\mathbf{x})$ is injected at i th layer. However, the choice of this layer is significant, and there is currently no way to know a priori which layer will yield the best results. In this paper, we investigate this by providing knowledge injections at each layer.

3. Methodology

To test the applicability of the PGNN approach, experiments were performed on five nonlinear dynamical systems. For each system, suitable injection terms were identified. The same NN architecture (3 hidden layers) was used in all cases to reduce the number of experiments. The functions were then injected with the following configurations: no injection, injection in the first layer, second layer, and third layer. Then, for each injection configuration, an ensemble of 10 models was trained on the data. This was done in order to estimate the model uncertainty.

3.1. Choice of equations

The systems were selected to cover a wide range of nonlinear phenomena including periodic and aperiodic solutions, limit cycles, and chaos. These properties are shown in Table 1.

3.1.1. Lotka–Volterra

The Lotka–Volterra equations are often used to describe the interactions of a population of predators x and a population of

Table 1

Possible types of solutions of the chosen nonlinear systems for the chosen parameters.

System	Periodic	Limit cycle	Chaotic
Lotka–Volterra	✓		
Duffing			✓
Van der Pol	✓	✓	
Lorenz			✓
Henon–Heiles	✓		✓

prey y :

$$\begin{aligned} \dot{x} &= \alpha x - \beta xy, \\ \dot{y} &= \delta xy - \gamma y, \end{aligned} \quad (7)$$

where \dot{y} and \dot{x} represent the instantaneous growth rates of the two populations due to predation, overpopulation, and starvation. The solutions are periodic; as the prey population x grows, the predators y can eat more and reproduce. This leads to a decline in x , causing y to drop as the predators starve. The two variables thus appear as similar waves, with y lagging behind x . In this paper, the values $\alpha = 0.1$, $\beta = 0.05$, $\delta = 0.1$, $\gamma = 1.1$ were used.

The Lotka–Volterra equations, and predator–prey models in general, remain of theoretical and practical interest today. Such systems have been successfully used to model ecological communities (Bunin, 2017), infections (Ghanbari & Djilali, 2020), and economic growth cycles (Goodwin, 1982; Harvie et al., 2007; Veneziani & Mohun, 2006). We will attempt to inject the nonlinear term xy , as it appears in both equations.

3.1.2. Duffing

The Duffing equation is a non-linear second-order differential equation that describes an oscillator with complex, sometimes chaotic behaviour. The Duffing equation was originally the result of Georg Duffing’s systematic study of nonlinear oscillations (Hamel, 1921). Interest in the equation was later revived with the advent of chaos theory. Since then, the system has come

to be regarded as one of the prototype systems in chaos theory (Strogatz, 2015), and related equations continue to find applications today, e.g. to describe the rolling of ships (Wawrzynski, 2018). The equation is

$$\ddot{x} = \gamma \cos(\omega t) - \delta \dot{x} - \alpha x - \beta x^3 \tag{8}$$

where $x(t)$ is the displacement at time t and the term $\gamma \cos(\omega t)$ represents a sinusoidal driving force. The cubic term describes an asymmetry in the restoring force of a spring that softens or stiffens as it is stretched. The parameters used in this work are $\delta = 1, \alpha = 0.5, \beta = 1, \gamma = 3$ and $\omega = 0.4$. Note that this is a time varying system depending on t . This is challenging to model using a neural network as-is, as the input t is unbounded. Instead, we reparametrise the system as follows:

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= \gamma \psi - \delta y - \alpha x - \beta x^3 \\ \dot{\psi} &= -\omega \theta \\ \dot{\theta} &= \omega \psi \\ \psi(0) &= 1, \quad \theta(0) = 0 \end{aligned} \tag{9}$$

This enables us to treat the system as if it were time-invariant. Note that from the ML perspective this is equivalent to feature engineering, as we provide the features $\cos(\omega t)$ and $\sin(\omega t)$ as additional inputs to the model. For knowledge injection, we use the x^3 term, and we also reuse the $\cos(\omega t)$ term to see if providing redundant features has any effect.

3.1.3. Van der Pol oscillator

The Van der Pol equation was discovered by Van der Pol (1960) while studying triode vibrations. It describes a nonlinear oscillator that approaches a limit cycle over time. Systems like this are immensely useful in a variety of fields. For example, coupled Van der Pol systems have been used by Rompala et al. (2007) to model biological circadian rhythms, and by Lucero and Schoentgen (2013) to model the asymmetries in vocal folds. Kuiuete et al. (2018) have even applied a variant of the system to encrypt images in real-time. The Van der Pol oscillator can be written as:

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= \mu(1 - x^2)y + x. \end{aligned} \tag{10}$$

where $x(t)$ is the displacement, and $\mu = 3$ is a scalar that controls the effects of the nonlinear damping term. In this work, we set $\mu = 3$.

The system tends to a stable limit cycle for all initial conditions. As x approaches the maximum amplitude of the oscillation, \dot{x} increases. When reaching the maximum, \dot{x} rapidly switches sign and x begins to decrease slowly, building up speed in the same way as it approaches the minimum. When the equation is forced with an additional sinusoidal term it can exhibit chaos, however this is not done in this work. The x^2y is fairly complex, and we select this for knowledge injection.

3.1.4. Lorenz system

The Lorenz system was originally developed to describe atmospheric convection by Lorenz (1963), but later became one of the most well-studied systems in chaos theory, and is often credited with the explosion of interest in the subject (Strogatz, 2015). The Lorenz equations have since been studied in connection with real physical phenomena, such as unstable spiking in lasers (Haken, 1975) and turbulence (Ruelle, 1976).

The system has the following form:

$$\begin{aligned} \dot{x} &= \sigma(y - x) \\ \dot{y} &= x(\rho - z) - y \\ \dot{z} &= xy - \beta z \end{aligned} \tag{11}$$

For $\rho < 1$, the origin is globally stable. When $\rho > 1$, the system has three fixed points: $(0, 0, 0)$, and $(\pm \sqrt{\beta(\rho - 1)}, \pm \sqrt{\beta(\rho - 1)}, \rho - 1)$, the latter of which we call C_+ and C_- . For $\rho > \frac{\sigma(\sigma + \beta + 3)}{\sigma - \beta - 1}$ the solutions of the system become non-periodic and chaotic where almost all initial states will converge to an invariant fractal set called the Lorenz attractor (Viswanath, 2004). Here we use $\sigma = 10, \rho = 28$, and $\beta = 8/3$, the values originally used by Lorenz. The terms xy and xz are natural candidates for injection. In this work we only test xy .

3.1.5. Henon–Heiles

Henon and Heiles (1964) originally developed these equations to study the movement of a star around a galactic centre while restricted to a plane. The system is still used to study the escape dynamics of orbits (Zotos, 2015). It is governed by the following Hamiltonian:

$$H = \frac{1}{2}(\dot{x}^2 + \dot{y}^2) + \frac{1}{2}(x^2 + y^2) + x^2y - \frac{y^3}{3} \tag{12}$$

This can be reformulated as a set of ODEs:

$$\begin{aligned} \ddot{x} &= -x - 2\lambda xy \\ \ddot{y} &= -y - 2\lambda(x^2 - y^2) \end{aligned} \tag{13}$$

In this work we set $\lambda = 1$. The solution set features a large number of periodic orbits, chaotic orbits, and escape trajectories when the energy of the system is sufficiently high (Zotos, 2014). The escape sets exhibit a rich fractal structure which adds additional complexity to the system behaviour (Zotos, 2017). This system has several features that can be injected. In this work, we try xy and y^2 , with x^2 omitted because it appears in the equation similarly to y^2 .

3.2. Data generation and pre-processing

For each of the five dynamical systems, we generated training data for the neural networks, and a test set to judge if the trained model can generalise to previously unseen states. This was done by manually choosing a set of initial conditions x_0 and generating the corresponding trajectories using the RKF45 solver with adaptive timestepping until a final time T . We used the implementation in the SciPy software stack (Virtanen et al., 2020), which is based on the Dormand–Prince pair of formulas (Dormand & Prince, 1980). The resulting data was then interpolated to generate a regular timeseries with timestep h . The time derivative at each data point was estimated as the forward difference $(f(x^+) - f(x))/h$. The datasets can then be described as a list of pairs $\mathcal{D} = \{(\mathbf{x}_k, \mathbf{y}_k)\}$, where \mathbf{x}_k and \mathbf{y}_k are the k th state and time derivative respectively. A validation set was constructed by reserving 20% of the data. The models are never trained on the validation set, but we monitor their performance on this data in order to measure their performance on unseen data. The initial conditions and other parameters that were used for each system are provided in Table 2. The test trajectory was generated from the last initial condition for each system, as discussed in Section 4.

3.3. Neural network architectures and training

The same network architecture was used in all cases to allow for a better comparison. The networks were given 3 hidden layers with 32, 64, and 32 neurons respectively. The injection was performing by concatenating the injection term to the selected hidden layer. Each model ensemble consisted of 10 neural networks. This was found to yield decent uncertainty estimates. The models were implemented in Tensorflow (Martín Abadi et al., 2015) and trained using the ADAM optimiser (Kingma & Ba, 2014) with default parameters. The models were trained on batches of

Table 2

Parameters and initial conditions used to generate the datasets. The training set was constructed by simulating each system with the initial conditions shown below using RKF45 with adaptive timestep until time T , and then estimating the pairs $(\mathbf{x}(t), \dot{\mathbf{x}}(t))$ at regular time intervals of length h . 20% of these pairs were reserved for the validation set. The test set was generated from a different initial condition, as shown below.

Model	Timestep h (s)	Final time T (s)	Initial conditions (IC) for training and validation	IC for testing
Lotka–Volterra	0.05	200	(2, 1), (10, 1), (12, 1), (15, 1), (20, 1), (22, 1), (25, 1)	(5, 1)
Duffing	0.05	200	(1, 1), (0, 1), (-1, 1), (1, -1), (0, -1), (-1, -1)	(1, 0.5)
Van der Pol	0.005	20	(0, 6), (0, -2), (-1, 2), (1, -4), (0, 0.1), (1, 3), (-2, 5)	(2, -5)
Lorenz	0.005	25	(1, 1, 1), (5, 1, 1), (1, 5, 1), (1, 1, 5), (-5, 1, 1), (1, -5, 1)	(1, 1, -5)
Henon–Heiles	0.05	100	(0.1, 0.5, 0, 0), (0.3, 0.4, 0, 0), (-0.35, 0.4, 0, 0), (0.3, -0.1, 0, 0)	(-0.325, 0.4, 0, 0)

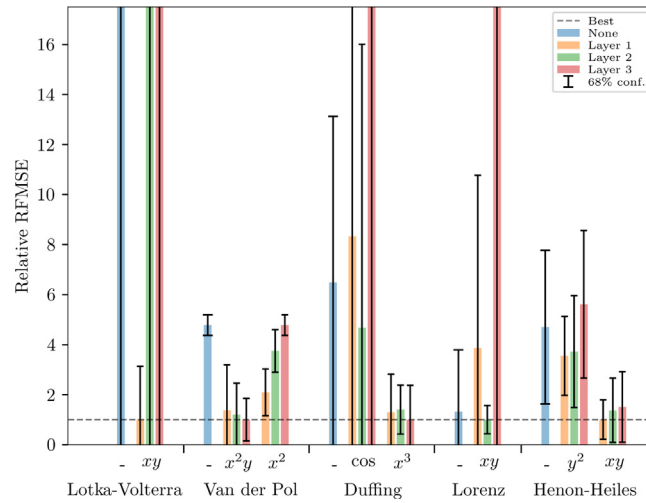


Fig. 3. The relative RFMSE for all model ensembles across different systems. The values for each system have been divided by the minimum RFMSE in each group for better comparison, such that the top performers for each system have a value of 1. Note that in some cases the rolling forecasts have diverged. Because of this, we do not compute the RFMSE on the full trajectory. Instead we use the final times [25 s, 70 s, 2.5 s, 2.5 s, 15 s] for each system respectively. The 68% confidence intervals shown here were chosen to improve clarity while still allowing for a comparison between models.

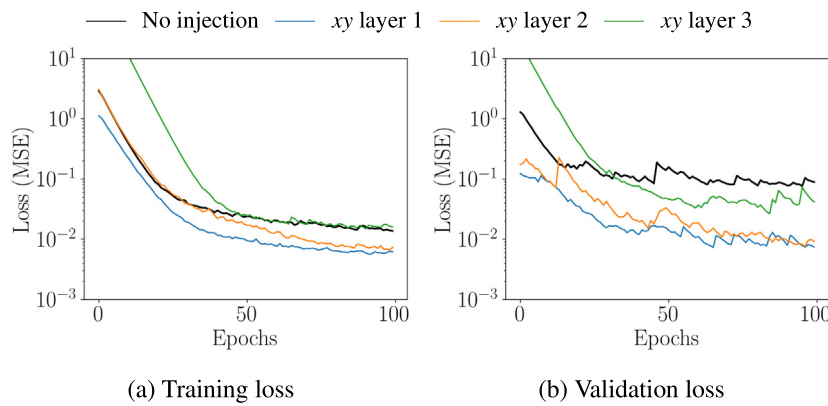


Fig. 4. Comparison of the training and validation loss for the Lotka–Volterra system for different injection configurations.

32 samples at a time (this number is known as the batch size) for a total of 100 epochs. An epoch is defined as the number of batch iterations after which the model will have trained on all data within the training set. We describe each batch as a set of indices $\mathcal{B} \subset \mathcal{N}$ that correspond to data in \mathcal{D} .

Since this is a regression problem, the mean-squared error (MSE) was utilised as a loss function. The loss for the training batch \mathcal{B} is then

$$L_{MSE}(\mathcal{B}; \theta) = \sum_{k \in \mathcal{B}} \|\mathbf{y}_k - \hat{f}(\mathbf{x}_k; \theta)\| \tag{14}$$

where $(\mathbf{x}_k, \mathbf{y}_k)$ is the k th pair in the dataset \mathcal{D} , as described in Section 3.2. Regularisation methods such as weight decay are usually used during training to prevent overfitting. We did not

encounter any overfitting issues, and therefore we do not apply any regularisation to reduce the number of comparisons.

3.4. Model evaluation

It was found that simply reporting the MSE on the test set did not clearly show how the models performed. Therefore, we chose to report the model performance as the MSE between a rolling forecast and the test trajectory, which we refer to as the rolling forecast MSE (RFMSE). During the rolling forecast stage, the initial condition for the first time step is provided. This information is used to predict the forecast state at the next time step using a forward Euler step, which is then used to predict the next

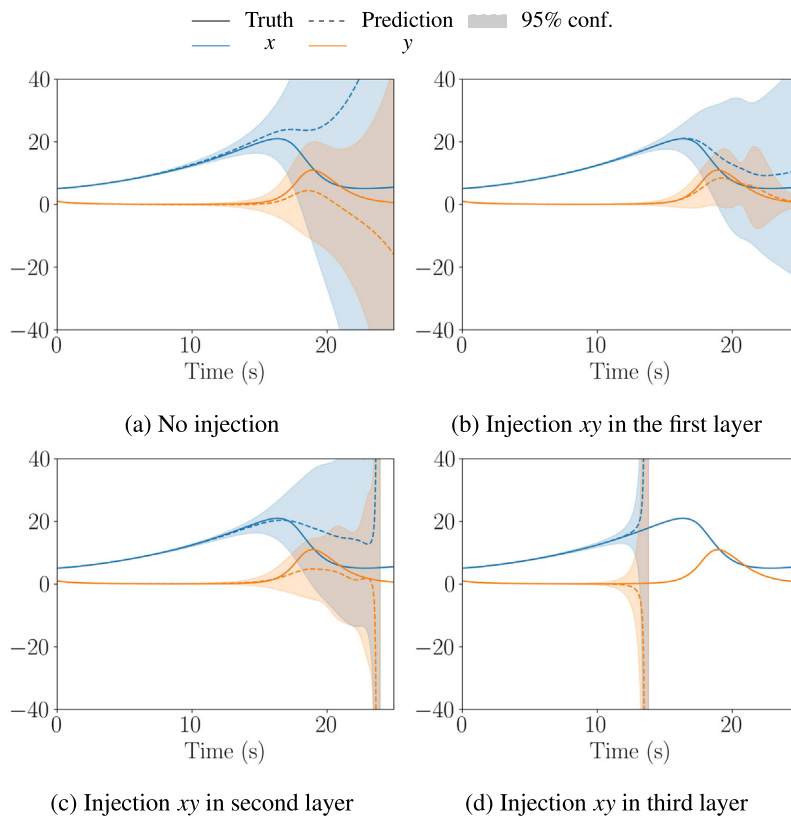


Fig. 5. Rolling forecast for the Lotka–Volterra system with and without injection at different layers. The best results are achieved through knowledge injection in the first layer. Injecting into the second and third layers appears to cause blowup.

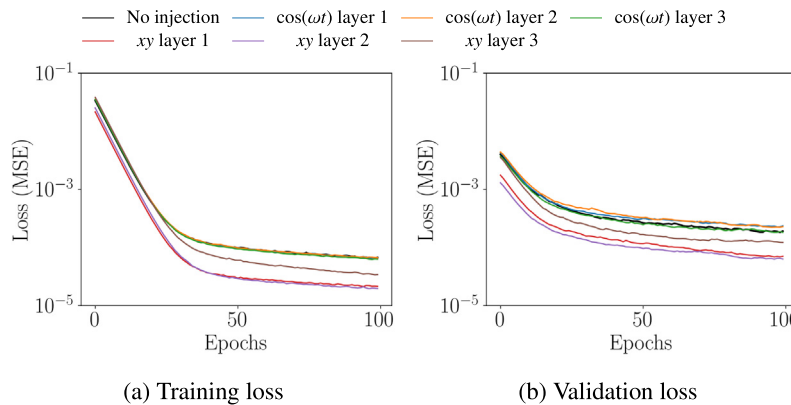


Fig. 6. Comparison of the training loss for the Duffing oscillator for different injection configurations.

time step until the final time step is reached. The timesteps shown in Table 2 were used. The resulting trajectories of the model ensemble were then compared to the true trajectory of the system. We believe that reporting the RFMSE more accurately reflects the actual use case of these models and makes it easier to qualitatively see how knowledge injection can affect the predictive accuracy and model uncertainty within each model class.

4. Results and discussion

In this section, we report the performance of the ensembles in terms of their training/validation loss, as well as the RFMSE on the test trajectory (see Section 3.4). The model uncertainty within

each model class is shown using 95% confidence bounds around the average predictions. We also report the mean training and validation loss for each ensemble. The loss signals of all models were smoothed using an exponential moving average filter using a weight of 0.2 before being averaged. This was done to improve the clarity of the plots, and allows us to compare overall trends between ensembles as well as the stability of the training. First, an overview of the results is presented, and then we provide a more detailed look at the best performing injection term within each model class.

4.1. Overview of results

The RFMSE for each model class is visualised in Fig. 3. Note that the data has been normalised due to the different scales of

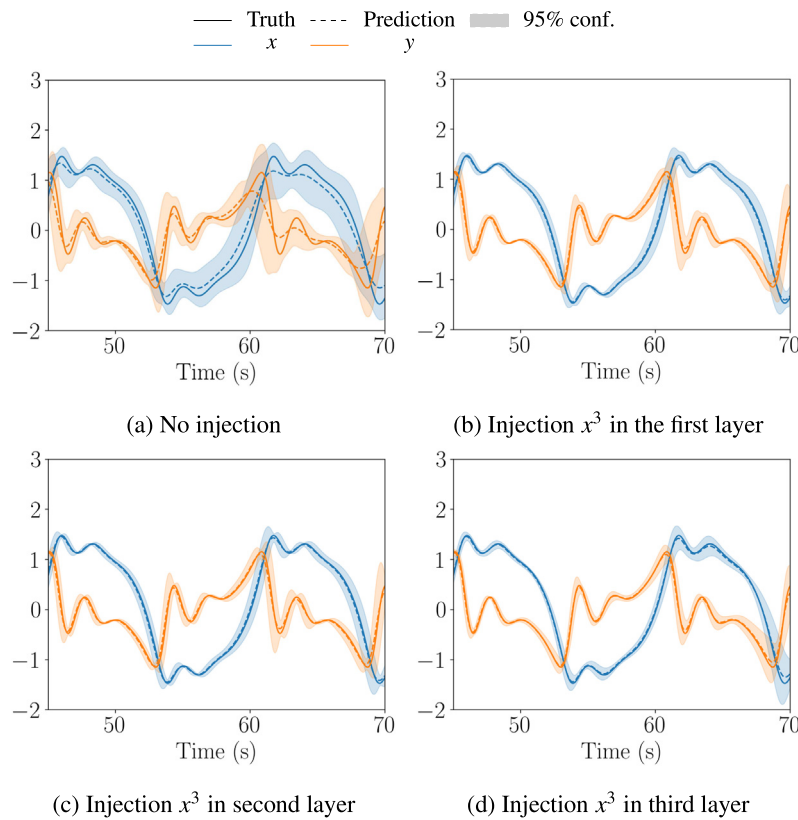


Fig. 7. Rolling forecast for the Duffing oscillator with and without injection at different layers. A small time segment from 45s–70s from the forecast is shown here to highlight the differences between the models.

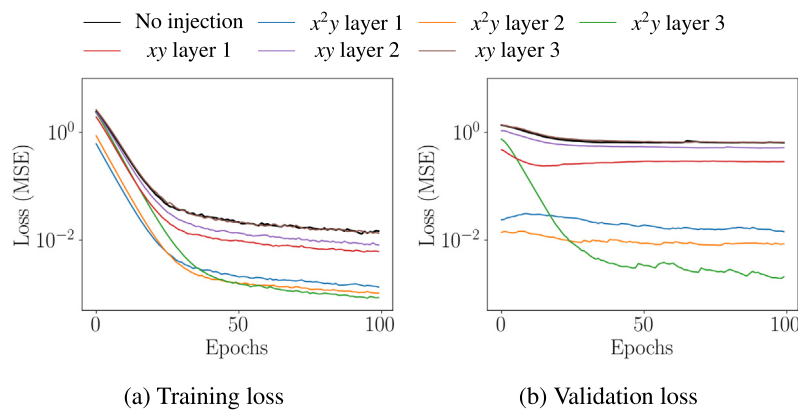


Fig. 8. Comparison of the training and validation loss for the Van der Pol oscillator for different injection configurations.

each test set, such that a value of 1 represents the top performer for each system. Additionally, because the predictions of the models blow up in some cases, we compute the RFMSE on a shorter time interval: [25 s, 70 s, 2.5 s, 2.5 s, 15 s] for the Lotka–Volterra, Duffing, Lorenz, Van der Pol, and Henon–Heiles systems respectively. For all systems, the best performing ensemble on average was a PGNN, often by a significant margin. The choice of injection layer appears to be a significant factor, although at this stage the data shows no conclusive pattern. This is surprising, as all of the nonlinear terms show up as additive terms in the equations, and there does not appear to be a good reason for the difference. Methods such as layer-wise relevance propagation (Montavon et al., 2019) could be adopted to interpret the impact

of the choice of layer for knowledge injection and we consider it as part of our future work.

4.2. Lotka Volterra system with xy injection

Fig. 4 shows that injecting the xy term in any layer caused the networks to reach a lower validation loss more quickly. This improvement was greatest when the injection was placed in the first hidden layer. Fig. 5 shows the mean rolling forecast of the ensembles with a 95% confidence interval. Injection in the first layer significantly improves the accuracy and the uncertainty of the forecast. However, the forecast quickly blows up for the models with second and third layer injections. This is especially pronounced for the third layer injection.

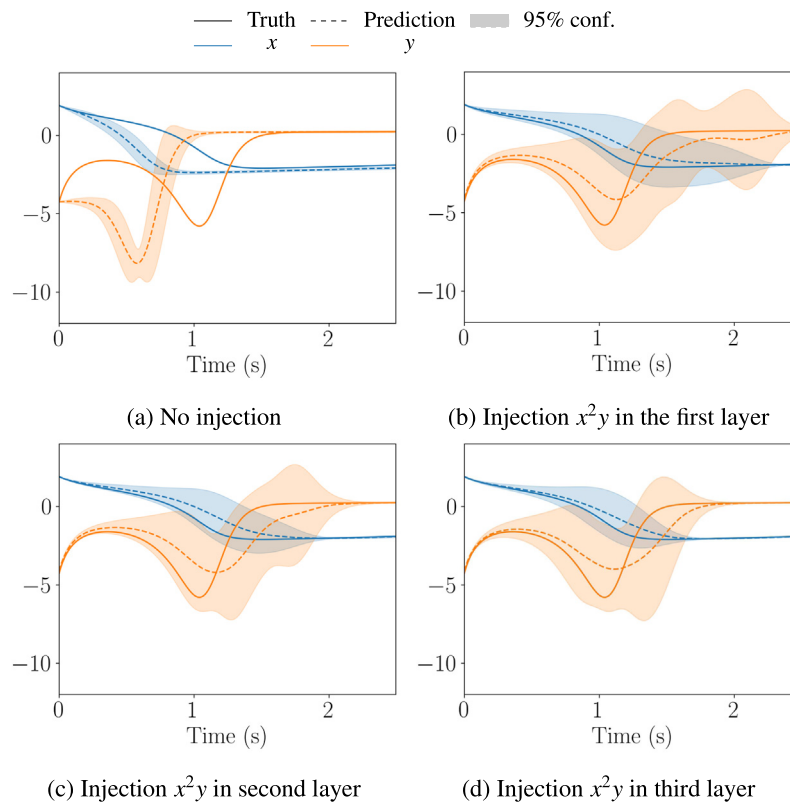


Fig. 9. Rolling forecast for the Van der Pol oscillator with and without injection at different layers.

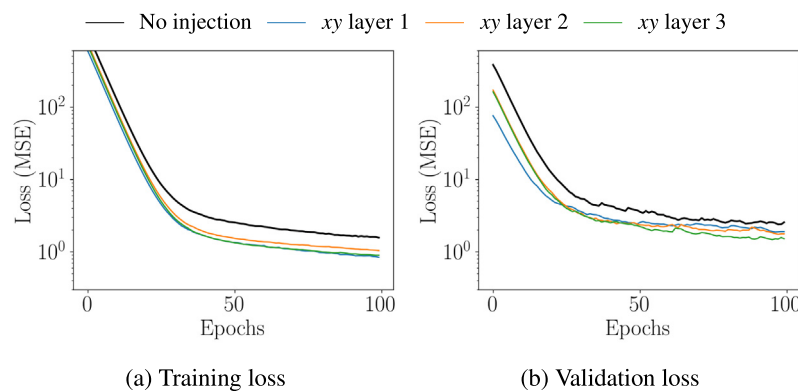


Fig. 10. Comparison of the training and validation loss for the Lorenz system for different injection configurations.

4.3. Duffing system with x^3 injection

Fig. 6 compares the training and validation loss of the models injected with x^3 and $\cos(\omega t)$ terms. Both training and validation loss are significantly improved with the x^3 injection, while $\cos(\omega t)$ appears to have little effect.

The predicted trajectories for x^3 models can be seen in Fig. 7, while the $\cos(\omega t)$ models have been omitted for brevity. Note that the figure shows a time segment from 75 s–100 s in order to highlight the differences between the models. We observe that knowledge injection improves the accuracy and model uncertainty in all cases, and all ensembles perform similarly.

It is interesting that although $\cos(\omega t)$ is available as an input to the network (due to the parameterisation described in Section 3.1.2), injecting the same term changes the RFMSE significantly, despite being redundant information. However, this is not reflected in the training and validation loss, where the baseline model and the models injected with $\cos(\omega t)$ appear

to have nearly identical training characteristics. We found that when forecasting over longer periods, the $\cos(\omega t)$ models yielded unstable predictions with a higher rate of blowup, while the other models tended to decay instead.

4.4. Van der Pol system with x^2y injection

For this system, the functions x^2y and x^2 were injected in all three layers. Fig. 8 shows how the x^2y injected models improve validation loss by 1–2 orders of magnitude better than the model without injection.

This large improvement in loss can be understood by inspecting Fig. 9, which shows a rolling forecast on the test trajectory near a fast transient. The figure shows that the models without injection fail to properly capture the transient, and also fail to converge to the slow dynamics afterwards. The long duration of the slow dynamics likely leads to a large build-up of error over the course of the full trajectory.

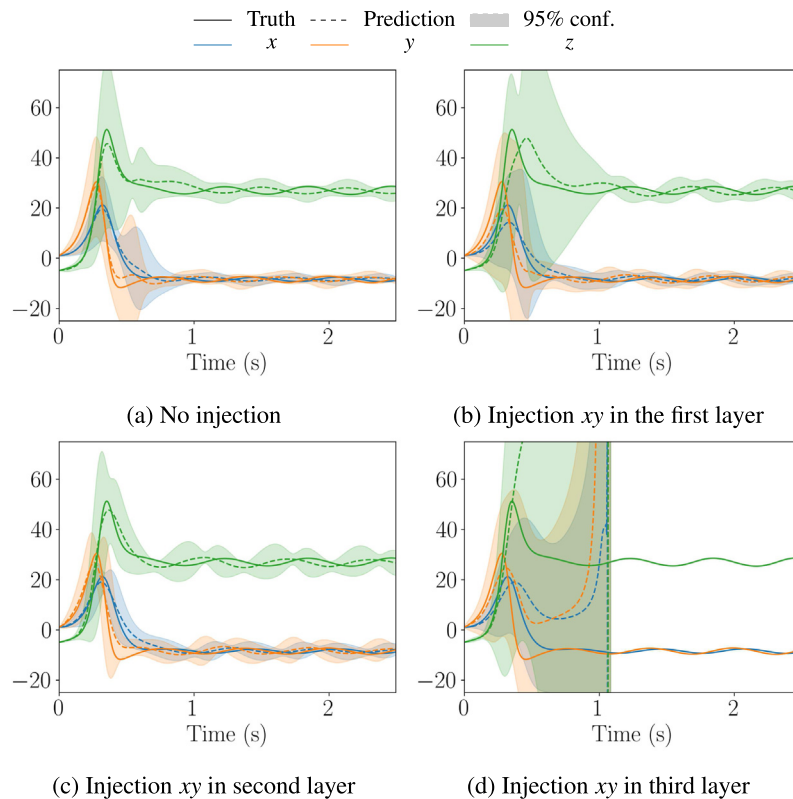


Fig. 11. Rolling forecast for the Lorenz system without and with injection at different layers. Here the injection is most effective at the second layer, while third layer injection causes the predictions to blow up.

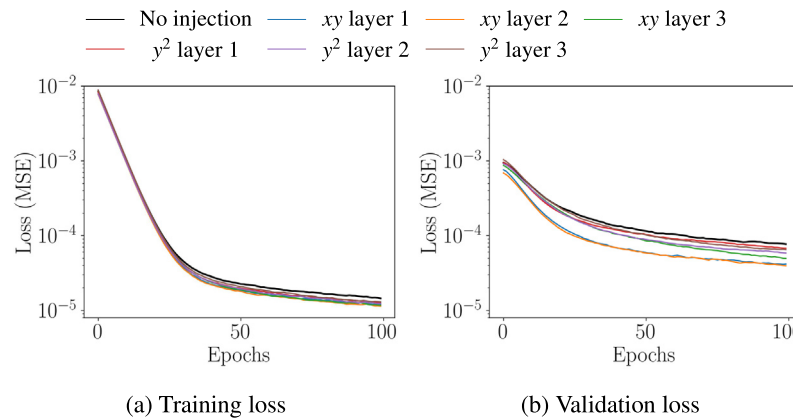


Fig. 12. Comparison of the training and validation loss for the Henon Heiles system with different injection configurations.

Surprisingly, the ensemble with no knowledge injection also exhibits very low model uncertainty. This might again be explained by the fast and slow dynamics of the Van der Pol oscillator. The dataset is likely unbalanced, dominated by slower varying states due to the longer duration of the slow dynamics. This could cause poor performance on the fast transients, which can be seen as relatively rare events.

Fig. 9 shows that this is greatly improved through knowledge injection. All injected models track the transient more closely, reach the correct value for the slow dynamics, and the true trajectory is within the 95% confidence intervals for all model classes.

The model uncertainty is naturally increased at the transient, and appears to shrink to zero when the slow dynamics set in.

4.5. Lorenz system with xy injection

The function xy was used for knowledge injection. Fig. 10 clearly shows that knowledge injection improved the training convergence and final validation loss over 100 epochs. Fig. 11 shows that the ensemble without injection stays close to the true trajectory, but oscillates out of phase with the ground truth. Injection in the first and second layer seems to reduce this lag,

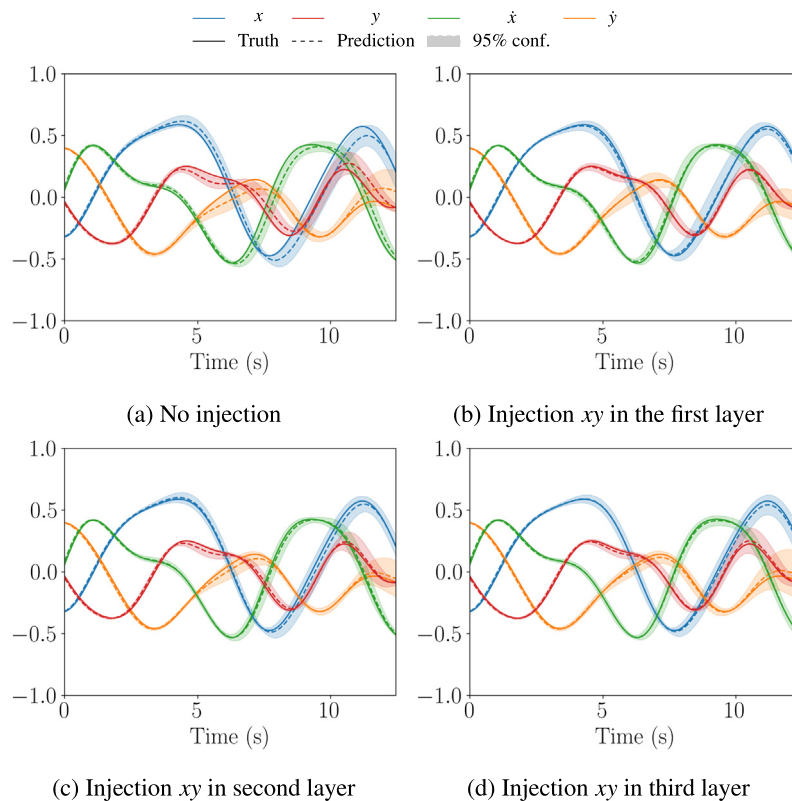


Fig. 13. Rolling forecast for the Henon–Heiles system without and with xy injection at different layers. Injecting functions appears to significantly aid the learning of this system. The best results are achieved by into the last hidden layer.

and the second layer injection appears to perform marginally better with lower model uncertainty. The predictions from the ensemble with third layer injections quickly diverge.

4.6. Henon–Heiles system with xy injection

For this system, all models performed very similarly, as can be seen in Fig. 13 for the xy injection. Fig. 12 shows that the injected models converge slightly faster and reach an overall lower validation loss.

5. Conclusion and future work

The physics-guided neural network (PGNN) framework was applied to a set of five distinct dynamical systems represented by first and second order non-linear ordinary differential equations. Three of the systems had two suitable injection terms which were also investigated, for a total of 8 combinations of system and injection terms. All possible injection layers were evaluated and compared. The main conclusions from the work are as follows:

- Knowledge injection can accelerate training and lead to better convergence. However, knowledge injection does not guarantee an improvement in performance.
- Accuracy of the models can in general be improved through knowledge injection. For the Van der Pol system, knowledge injection helped the models capture the fast transients, which were relatively underrepresented in the dataset. However, for some systems like Henon–Heiles system the improvements were not significant.
- The study remains inconclusive regarding the impact of knowledge injection on model uncertainty. For the Duffing oscillator we see a shrinkage in the uncertainty but for Van der Pol, we see an increase. However, we should stress

that lower uncertainty with poor prediction would not be desirable.

The first limitation of this work is that the choice of injection layer was found to have a significant impact on performance, and the results do not show how such a choice could be made a priori. A second limitation is that we do not address how to identify the correct information to inject into the hidden layers of the network. We simply tried all combinations of injection layers and terms, but it is clear that this will scale poorly for network architectures with more hidden layers or multiple injection terms. Hyperparameter search algorithms (for example genetic algorithms due to the discrete optimisation variable) may be helpful tools to choose an effective injection layer. However, this still involves training multiple models. A simpler solution to both problems is to make all injection features available to all layers via skip connections, although preliminary results in this direction have shown that this does not reach the same level of performance as the best single injection layer. Combining this approach with sparsifying regularisation may bias the network towards selecting the best injection layer and term. Testing the PGNN approach with deeper architectures and recurrent NNs could also elucidate the role of the injection layer, and should be investigated. Another method that could help identify suitable injection terms is symbolic regression based on gene expression programming (GEP). By running a large SR ensemble on the data and selecting the most frequently appearing terms, it might be possible to collect useful injection terms.

We see extensions of PGNNs as being useful for modelling more complex systems with rich dynamics and interactions with the environment. Significant assumptions are typically made about the nature of environmental forces and practitioners often defer to the data, e.g. when modelling the average wind force on marine vessels (Blendermann, 1994; Fossen, 2021; Isherwood,

1972). There is already much work where more advanced ML techniques such as reinforcement learning (RL) are applied to these systems (Meyer et al., 2020). The improved training characteristics and low overhead of PGNNs may prove useful in these RL contexts, where data efficiency is typically quite poor. Furthermore, the learned weighting of explicit features arguably makes the models more interpretable, which is often seen as desirable in safety-critical contexts, although more work is needed in this direction. PGNNs may also require fewer parameters than conventional DNNs to model the same data, which could enable the use of existing robustness verification algorithms (Liu et al., 2019), vastly improving confidence in these systems during deployment.

The attractiveness of prior knowledge injection is that it generalises two of the most common hybridisation methods: input feature engineering and output error correction. By injecting arbitrary features into the intermediate layers of a neural network, we begin to open the proverbial black box and recognise its potential as a general-purpose feature-learner and feature-selector, powered by stochastic optimisation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

A.R. is grateful for the support received by the Research Council of Norway and the industrial partners of the following projects: EXAIGON–*Explainable AI systems for gradual industry adoption* (Grant No. 304843). O.S. gratefully acknowledges the U.S. DOE Early Career Research Program support (DE-SC0019290), and the National Science Foundation support (DMS-2012255).

References

Ahmed, S. E., Pawar, S., San, O., Rasheed, A., Iliescu, T., & Noack, B. R. (2021). On closures for reduced order models—A spectrum of first-principle to machine-learned avenues. *Physics of Fluids*, 33(9), Article 091301.

Amos, B., & Kolter, J. Z. (2017). OptNet: Differentiable optimization as a layer in neural networks. In *International conference on machine learning*, (pp. 136–145). PMLR, URL: <https://proceedings.mlr.press/v70/amos17a.html>.

Arnold, F., & King, R. (2021). State-space modeling for control based on physics-informed neural networks. *Engineering Applications of Artificial Intelligence*, 101, Article 104195, <https://doi.org/10.1016/j.engappai.2021.104195>. URL: <https://www.sciencedirect.com/science/article/pii/S0952197621000427>.

Blakseth, S. S., Rasheed, A., Kvamsdal, T., & San, O. (2022). Deep neural network enabled corrective source term approach to hybrid analysis and modeling. *Neural Networks*, 146, 181–199.

Blendermann, W. (1994). Parameter identification of wind loads on ships. *Journal of Wind Engineering and Industrial Aerodynamics*, 51(3), 339–351, [https://doi.org/10.1016/0167-6105\(94\)90067-1](https://doi.org/10.1016/0167-6105(94)90067-1). URL: <https://www.sciencedirect.com/science/article/pii/0167610594900671>.

Bunin, G. (2017). Ecological communities with Lotka-Volterra dynamics. *Physical Review E*, 95(4), Article 042414. <http://dx.doi.org/10.1103/PhysRevE.95.042414>, URL: <https://link.aps.org/doi/10.1103/PhysRevE.95.042414>.

Champion, K., Lusch, B., Kutz, J. N., & Brunton, S. L. (2019). Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45), 22445–22451. <http://dx.doi.org/10.1073/pnas.1906995116>, URL: <https://www.pnas.org/content/116/45/22445>. arXiv: <https://www.pnas.org/content/116/45/22445.full.pdf>.

de Avila Belbute-Peres, F., Smith, K., Allen, K., Tenenbaum, J., & Kolter, J. Z. (2018). End-to-end differentiable physics for learning and control. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 31) (p. 1). Curran Associates, Inc., URL: <https://proceedings.neurips.cc/paper/2018/file/842424a1d0595b76ec4fa03c46e8d755-Paper.pdf>.

Dormand, J. R., & Prince, P. J. (1980). A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1), 19–26.

Fonn, E., Brummelen, H. v., Kvamsdal, T., & Rasheed, A. (2019). Fast divergence-conforming reduced basis methods for steady Navier–Stokes flow. *Computer Methods in Applied Mechanics and Engineering*, 346, 486–512, <https://doi.org/10.1016/j.cma.2018.11.038>.

Fossen, T. I. (2021). *Marine craft hydrodynamics and motion control* (2nd ed). Wiley.

Ghanbari, B., & Djilali, S. (2020). Mathematical analysis of a fractional-order predator-prey model with prey social behavior and infection developed in predator population. *Chaos, Solitons & Fractals*, 138, Article 109960. <http://dx.doi.org/10.1016/j.chaos.2020.109960>, URL: <https://www.sciencedirect.com/science/article/pii/S0960077920303593>.

Goodfellow, I., Courville, A., & Bengio, Y. (2016). *Deep learning*. MIT Press, URL: <https://www.deeplearningbook.org/>.

Goodwin, R. M. (1982). A growth cycle. In R. M. Goodwin (Ed.), *Essays in economic dynamics* (pp. 165–170). London: Palgrave Macmillan UK, http://dx.doi.org/10.1007/978-1-349-05504-3_12, https://doi.org/10.1007/978-1-349-05504-3_12.

Haken, H. (1975). Analogy between higher instabilities in fluids and lasers. *Physics Letters. A*, 53(1), 77–78. [http://dx.doi.org/10.1016/0375-9601\(75\)90353-9](http://dx.doi.org/10.1016/0375-9601(75)90353-9), URL: <https://www.sciencedirect.com/science/article/pii/0375960175903539>.

Hamel (1921). Georg Duffing, Ingenieur: Erzwungene Schwingungen bei veränderlicher Eigenfrequenz und ihre technische Bedeutung. Sammlung Vieweg. Heft 41/42, Braunschweig 1918. VI+134 S. ZAMM - *Journal of Applied Mathematics and Mechanics / Zeitschrift Für Angewandte Mathematik Und Mechanik*, 1(1), 72–73. <http://dx.doi.org/10.1002/zamm.19210010109>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/zamm.19210010109>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/zamm.19210010109>.

Harvie, D., Kelmanson, M., & Knapp, D. (2007). A dynamical model of business-cycle asymmetries: Extending goodwin. *Economic Issues*, 12.

Henon, M., & Heiles, C. (1964). The applicability of the third integral of motion: Some numerical experiments. *Astronomical Journal*, 69, 73. <http://dx.doi.org/10.1086/109234>, URL: <http://adsabs.harvard.edu/abs/1964AJ....69...73H>.

Isherwood, R. (1972). Wind resistance of merchant ships. *The Royal Institution of Naval Architects*, 115, 327–338.

Karpathy, A. (2017). Software 2.0. URL: <https://karpathy.medium.com/software-2-0-a64152b37c35>.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).

Kuiate, G. F., Rajagopal, K., Kingni, S. T., Tamba, V. K., & Jafari, S. (2018). Autonomous van der pol–duffing snap oscillator: analysis, synchronization and applications to real-time image encryption. *International Journal of Dynamics and Control*, 6(3), 1008–1022. <http://dx.doi.org/10.1007/s40435-017-0373-z>, <https://doi.org/10.1007/s40435-017-0373-z>.

Liu, C., Arnon, T., Lazarus, C., Barrett, C. W., & Kochenderfer, M. J. (2019). Algorithms for verifying deep neural networks. CoRR abs/1903.06758. arXiv: [1903.06758](https://arxiv.org/abs/1903.06758). eprint: 1903.06758.

Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130–141. [http://dx.doi.org/10.1175/1520-0469\(1963\)020<0130:dnf>2.0.co;2](http://dx.doi.org/10.1175/1520-0469(1963)020<0130:dnf>2.0.co;2).

Lucero, J. C., & Schoentgen, J. (2013). Modeling vocal fold asymmetries with coupled van der Pol oscillators. *Proceedings of Meetings on Acoustics*, 19(1), Article 060165. <http://dx.doi.org/10.1121/1.4798467>, URL: <https://asa.scitation.org/doi/abs/10.1121/1.4798467>.

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Jia, Y., Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Xiaoqiang Zheng (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*. URL: <https://www.tensorflow.org/>.

Maulik, R., San, O., Rasheed, A., & Vedula, P. (2019). Sub-grid modelling for two-dimensional turbulence using neural networks. *Journal of Fluid Mechanics*, 858, 122–144, <https://doi.org/10.1017/jfm.2018.770>.

Meyer, E., Robinson, H., Rasheed, A., & San, O. (2020). Taming an autonomous surface vehicle for path following and collision avoidance using deep reinforcement learning. *IEEE Access*, 8, 41466–41481.

Montavon, G., Binder, A., Lapuschkin, S., Samek, W., & Müller, K.-R. (2019). Layer-wise relevance propagation: an overview. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, 193–209.

Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018). Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical Review Letters*, 120(2), Article 024102.

Pawar, S., Ahmed, S. E., San, O., & Rasheed, A. (2020). Data-driven recovery of hidden physics in reduced order modeling of fluid flows. *Physics of Fluids*, 32, Article 036602.

Pawar, S., Ahmed, S. E., San, O., & Rasheed, A. (2020). An evolve-then-correct reduced order model for hidden fluid dynamics. *Mathematics*, 8(4), 570.

Pawar, S., Rahman, S., Vaddirreddy, H., San, O., Rasheed, A., & Vedula, P. (2019). A deep learning enabler for nonintrusive reduced order modeling of fluid flows. *Physics of Fluids*, 31(8), Article 085101.

- Pawar, S., San, O., Aksoylu, B., Rasheed, A., & Kvamsdal, T. (2021). Physics guided machine learning using simplified theories. *Physics of Fluids*, 33(1), Article 011701.
- Pawar, S., San, O., Nair, A., Rasheed, A., & Kvamsdal, T. (2021). Model fusion with physics-guided machine learning: Projection-based reduced-order modeling. *Physics of Fluids*, 33(6), Article 067123.
- Pawar, S., San, O., Rasheed, A., & Vedula, P. (2020). A priori analysis on deep learning of subgrid-scale parameterizations for Kraichnan turbulence. *Theoretical and Computational Fluid Dynamics*, 34, 429–455.
- Pawar, S., San, O., Vedula, P., Rasheed, A., & Kvamsdal, T. (2022). Multi-fidelity information fusion with concatenated neural networks. *Scientific Reports*, 12(1), 1–13.
- Quarteroni, A., & Rozza, G. (2014). *Reduced order methods for modeling and computational reduction* (Vol. 9). New York: Springer.
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (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.
- Rompala, K., Rand, R., & Howland, H. (2007). Dynamics of three coupled van der Pol oscillators with application to circadian rhythms. *Communications in Nonlinear Science and Numerical Simulation*, 12(5), 794–803. <http://dx.doi.org/10.1016/j.cnsns.2005.08.002>, URL: <https://www.sciencedirect.com/science/article/pii/S1007570405001206>.
- Ruelle, D. (1976). The Lorenz attractor and the problem of turbulence. In L. Streit (Ed.), *Acta physica austriaca, Quantum dynamics: Models and mathematics* (pp. 221–239). Vienna: Springer. http://dx.doi.org/10.1007/978-3-7091-8473-8_14.
- Saha, P., Dash, S., & Mukhopadhyay, S. (2021). Physics-incorporated convolutional recurrent neural networks for source identification and forecasting of dynamical systems. *Neural Networks*, <https://doi.org/10.1016/j.neunet.2021.08.033>. URL: <https://www.sciencedirect.com/science/article/pii/S0893608021003464>.
- San, O., Rasheed, A., & Kvamsdal, T. (2021). Hybrid analysis and modeling, eclecticism, and multifidelity computing toward digital twin revolution. *GAMM-Mitteilungen*, 44(2), Article e202100007.
- Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Židek, A., Nelson, A. W. R., Bridgland, A., Penedones, H., Petersen, S., Simonyan, K., Crossan, S., Kohli, P., Jones, D. T., Silver, D., Kavukcuoglu, K., & Hassabis, D. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792), 706–710. <http://dx.doi.org/10.1038/s41586-019-1923-7>, URL: <https://www.nature.com/articles/s41586-019-1923-7>. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7792 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Machine learning;Protein structure predictions Subject_term_id: machine-learning;protein-structure-predictions.
- Shen, S., Lu, H., Sadoughi, M., Hu, C., Nemani, V., Thelen, A., Webster, K., Darr, M., Sidon, J., & Kenny, S. (2021). A physics-informed deep learning approach for bearing fault detection. *Engineering Applications of Artificial Intelligence*, 103, Article 104295. <http://dx.doi.org/10.1016/j.engappai.2021.104295>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197621001421>.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489. <http://dx.doi.org/10.1038/nature16961>, URL: <https://www.nature.com/articles/nature16961>.
- Strogatz, S. H. (2015). *Nonlinear dynamics and chaos: With applications to Physics, Biology, Chemistry, and Engineering* (2nd ed). CRC Press.
- Vaddireddy, H., Rasheed, A., Staples, A. E., & San, O. (2020). Feature engineering and symbolic regression methods for detecting hidden physics from sparse sensors. *Physics of Fluids, Editor's Pick*, 32, Article 015113, <https://doi.org/10.1063/1.5136351>.
- Van der Pol, B. (1960). A theory of the amplitude of free and forced triode vibrations. I, 701–710, 754–762; *Selected scientific papers (Vol. I)*. North Holland.
- Veneziani, R., & Mohun, S. (2006). Structural stability and Goodwin's growth cycle. *Structural Change and Economic Dynamics*, 17(4), 437–451. <http://dx.doi.org/10.1016/j.strueco.2006.08.003>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0954349X06000300>.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., .. van Mulbregt, P., SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, 17, 261–272. <http://dx.doi.org/10.1038/s41592-019-0686-2>.
- Viswanath, D. (2004). The fractal property of the Lorenz attractor. *Physica D: Nonlinear Phenomena*, 190(1), 115–128. <http://dx.doi.org/10.1016/j.physd.2003.10.006>, URL: <https://www.sciencedirect.com/science/article/pii/S0167278903004093>.
- Vlachas, P. R., Byeon, W., Wan, Z. Y., Sapsis, T. P., & Koumoutsakos, P. (2018). Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks. *Proceedings of the Royal Society of London, Series A (Mathematical and Physical Sciences)*, 474(2213), Article 20170844.
- Wawrzynski, W. (2018). Bistability and accompanying phenomena in the 1-DOF mathematical model of rolling. *Ocean Engineering*, 147, 565–579. <http://dx.doi.org/10.1016/j.oceaneng.2017.11.013>, URL: <https://www.sciencedirect.com/science/article/pii/S0029801817306807>.
- Willard, J., Jia, X., Xu, S., Steinbach, M., & Kumar, V. (2020). Integrating physics-based modeling with machine learning: A survey. arXiv preprint arXiv:2003.04919.
- Yu, Y., Yao, H., & Liu, Y. (2020). Structural dynamics simulation using a novel physics-guided machine learning method. *Engineering Applications of Artificial Intelligence*, 96, Article 103947. <http://dx.doi.org/10.1016/j.engappai.2020.103947>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197620302670>.
- Zobeiry, N., & Humfeld, K. D. (2021). A physics-informed machine learning approach for solving heat transfer equation in advanced manufacturing and engineering applications. *Engineering Applications of Artificial Intelligence*, 101, Article 104232. <http://dx.doi.org/10.1016/j.engappai.2021.104232>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197621000798>.
- Zotos, E. E. (2014). Classifying orbits in the classical Hénon–Heiles Hamiltonian system. *Nonlinear Dynamics*, 79(3), 1665–1677. <http://dx.doi.org/10.1007/s11071-014-1766-6>, URL: <http://dx.doi.org/10.1007/s11071-014-1766-6>.
- Zotos, E. E. (2015). Comparing the escape dynamics in tidally limited star cluster models. *Monthly Notices of the Royal Astronomical Society*, 452, 193–209, ADS Bibcode: 2015MNRAS.452..193Z. <http://dx.doi.org/10.1093/mnras/stv1307>. URL: <https://ui.adsabs.harvard.edu/abs/2015MNRAS.452..193Z>.
- Zotos, E. E. (2017). An overview of the escape dynamics in the Hénon–Heiles Hamiltonian system. *Meccanica*, 52(11–12), 2615–2630. <http://dx.doi.org/10.1007/s11012-017-0647-8>, URL: <http://dx.doi.org/10.1007/s11012-017-0647-8>.