
Physics-informed Discovery of State Variables in Second-Order and Hamiltonian Systems

Félix Chavelli*

National University of Singapore
Singapore, Singapore
chavelli@comp.nus.edu.sg

Khoo Zi-Yu*

National University of Singapore
Singapore, Singapore
khoozy@comp.nus.edu.sg

Dawen Wu

CNRS@CREATE LTD
Singapore, Singapore
dawen.wu@cnrsatcreate.sg

Jonathan Sze Choong Low

Agency for Science, Technology and Research
Singapore, Singapore
sclow@simtech.a-star.edu.sg

Stéphane Bressan

National University of Singapore
Singapore, Singapore
steph@nus.edu.sg

Abstract

The modelling of dynamical systems is a pervasive concern for not only describing but also predicting and controlling natural phenomena and engineered systems. Current data-driven approaches often assume prior knowledge of the relevant state variables or result in overparameterized state spaces. Boyuan Chen and his co-authors proposed a neural network model that estimates the degrees of freedom and attempts to discover the state variables of a dynamical system. Despite its innovative approach, this baseline model lacks a connection to the physical principles governing the systems it analyzes, leading to unreliable state variables. This research proposes a method that leverages the physical characteristics of second-order Hamiltonian systems to constrain the baseline model. The proposed model outperforms the baseline model in identifying a minimal set of non-redundant and interpretable state variables.

1 Introduction

Dynamical systems, systems that change over time, pervade the natural and engineered world, embodying the complex interactions and evolution observed across a multitude of fields, from physics and biology to economics and engineering. Their universality and applicability in modelling real-world phenomena underscore the imperative for their study, motivating works that offer insights into system behaviour, prediction, and control [17, 7, 18, 16, 6]. However, most data-driven methods for modelling dynamical systems assume that the relevant state variables are already known [5, 6] or use more parameters than necessary to represent possible configurations of the state variables [18, 7].

Recent work by Chen et al. [5] proposed a neural network model for dynamical system analysis that identifies the relevant state variables. The model estimated the degrees of freedom and discovered the state variables of a dynamical system from its images using an encoder-decoder architecture [5]. Yet, Chen et al.'s proposed model is detached from the physical underpinnings of the systems it models [5]; the proposed model uses an external intrinsic dimension estimator [2, 13, 4] which is ignorant of

physics principles to estimate a non-integer number of degrees of freedom. Rounding a non-integer to an integer can be influenced by confirmation bias. Furthermore, the obtained state variables are correlated, possibly redundant or entangled, with no guarantee of interpretability. Therefore, this research pivots towards harnessing the physical characteristics of dynamical systems, specifically their second-order [8] and Hamiltonian [14, 15] characteristics. By leveraging physics-informed machine learning, we propose to incorporate physics knowledge into the machine learning model: observational biases are introduced directly through data that embody the underlying physics [11], learning biases are introduced through appropriate loss functions which modulate convergence towards solutions that adhere to the underlying physics [11, 3], and inductive biases are incorporated by tailored interventions to a machine learning model architecture [11]. These embed physical constraints into Chen et al.’s baseline model, to identify a minimal set of non-redundant, interpretable state variables. Our central research question seeks to exploit these physical properties to achieve an optimal representation of the system’s dynamics, offering a balance between simplicity and accuracy that improves existing methodologies.

We introduce modifications to Chen et al.’s model using physics-informed machine learning in Section 2. In Section 3, we evaluate, compare, and find that the proposed model outperforms Chen et al.’s model on their original dataset by accurately identifying and interpreting the state variables.

2 Methodology

Chen et al.’s baseline model is a nested autoencoder [5]. In the first step, two consecutive video frames from a dynamical system are input to the outermost autoencoder. A compact representation of the input is extracted and the subsequent two consecutive frames are predicted. Then, an intrinsic dimension (ID) estimation [13] of the compact representation is performed. In the second step, the innermost autoencoder further compresses the compact representation into a latent space of dimension ID. These are the state variables of the dynamical system.

We leverage Karniadakis et al.’s framework to enhance Chen et al.’s model by introducing three proposed models, each incorporating different physics-informed machine learning biases. They are the Physics-Informed AutoEncoder (PI-AE), the Physics-Informed Variational AutoEncoder (PI-VAE), and the Hamiltonian Physics-Informed Variational AutoEncoder (HPI-VAE). The three proposed models make modifications to the innermost autoencoder. Further details regarding the architectures are supplemented in Appendix B.

The Physics-Informed AutoEncoder (PI-AE) builds on Chen et al.’s baseline. It incorporates an observational bias by enforcing the system’s second-order constraint to the latent space of the innermost autoencoder. It follows the nomenclature of second-order dynamical systems, and constrains pairs of latent variables of the autoencoder such that the first represents the position, and the second, the momentum of the dynamical system. The Physics-Informed Variational AutoEncoder (PI-VAE) builds on the PI-AE and incorporates a learning bias regarding the time-continuity of the dynamical system and the independence of the state variables. The constraint is incorporated by a variational autoencoder. The KL divergence term in the loss function enforces latent sparsity, as the learned distribution in the latent space has a Standard Multivariate Normal Distribution prior. By enforcing latent sparsity using the KL divergence term, the PI-VAE circumvents the need for an intrinsic dimension estimator. Furthermore, the intrinsic dimension estimator is used as an external tool to find the number of degrees of freedom of the dynamical system, and functions independently from the training and convergence of the encoder and decoder. Therefore, the PI-VAE is more parsimonious by design compared to the baseline and PI-AE as it bypasses the intrinsic dimension estimator. The Hamiltonian Physics-Informed Variational AutoEncoder (HPI-VAE) builds on the PI-VAE and incorporates an inductive bias. It modifies the PI-VAE architecture to include a Hamiltonian neural network which takes as input the latent variables. This model has three terms in its loss function. They are the reconstruction loss, the KL divergence, and Hamilton’s equations. Likewise, by enforcing latent sparsity using the KL divergence term, the HPI-VAE circumvents the need for an intrinsic dimension estimator. The HPI-VAE is more parsimonious by design compared to the baseline and PI-AE as it bypasses the intrinsic dimension estimator.

3 Experiments

The data comprises video frames of the systems which are either real and filmed with a camera, or numerically simulated from Chen et al. They comprise 100,000 frames per dynamical system, comprising 100 frames from 1,000 random trajectories. These datasets are divided into training, validation, and test sets, constituting approximately 80%, 10%, and 10% of the data, respectively.

The five dynamical systems are the reaction-diffusion system, a single pendulum system, a double pendulum system, a swing stick system, and an elastic pendulum system. Example frames of the five systems are shown in Appendix C.

The outermost autoencoder shared by all models takes as input two consecutive frames and outputs the next two consecutive frames. It has five consecutive convolutional layers and a latent space of size 64. The innermost autoencoder takes the latent space of the outermost autoencoder as input and reconstructs it as the output. It has four fully connected layers. The baseline and proposed PI-AE have a latent space of dimension ID . The proposed PI-VAE and HPI-VAE have a latent space of size 10 which exceeds the ground truth intrinsic dimension of the five dynamical systems. The β hyperparameter which adjusts the respective weights of the reconstruction loss and KL divergence is placed on the reconstruction loss term. Further details regarding the hyperparameters of each model are reported in Appendix D.

The number of degrees of freedom estimated by the baseline and proposed models for the five dynamical systems is shown in Table 1. For the baseline and PI-AE, this is the value given by Chen et al.’s intrinsic dimension estimator, rounded to a whole number. For the PI-VAE and HPI-VAE, the VAE architecture promotes sparsity of the latent space, hence no intrinsic dimension estimator is required. The number of degrees of freedom is the number of latent variables that are non-zero. A variance threshold of 0.01, corresponding to 1% of the targeted variance of the reduced prior, discriminates between the non-zero latent variables of interest and the uninformative latent variables. All models can successfully reconstruct the latent space of the outermost autoencoder with a reconstruction loss lower than 0.01. Table 1 shows that the PI-VAE and the HPI-VAE, with the VAE architecture, are able to directly find the correct number of degrees of freedom for the five dynamical systems without the use of the intrinsic dimension estimator. They hence avoid influence by or tendency toward confirmation bias.

Table 1: Latent space dimension of systems from different models

System	Ground truth	Baseline [5]	PI-AE	PI-VAE	HPI-VAE
Reaction-diffusion	2	2.16 \approx 2	2.16 \approx 2	2	2
Single pendulum	2	2.05 \approx 2	2.05 \approx 2	2	2
Double pendulum	4	4.71 \approx 4	4.71 \approx 4	4	4
Swingstick system	4	4.89 \approx 4	4.89 \approx 4	4	4
Elastic pendulum	6	5.34 \approx 6	5.34 \approx 6	6	6

Additionally, we comment on the non-zero latent variables of interest for the HPI-VAE and compare them to the latent variables of the baseline by visualizing the latent variables for a random trajectory of the dynamical system. We focus on the latent variables obtained for the three pendulum systems for which the ground truth is known. Further results can be found in Appendix E.

Single pendulum Figure 1 displays the latent variables obtained with each model, scaled to between -1 and +1. The baseline model (a) identifies two latent variables, `var0` and `var1`, which appear entangled or redundant. The HPI-VAE captures the system’s evolution using two unique latent variables. The first latent variable is shown. The second is its time derivative, omitted for visualization clarity but enforced with the physics-informed second-order constraint to its latent space. The plot of $\cos 2\theta(t)$ against time, t , is plotted as a black dotted line on all subplots of Figure 1. It can also be observed that the baseline exhibits discontinuous changes that align with the curve of $\cos 2\theta(t)$. However, in the HPI-VAE model, the latent variable is smoother, revealing a more continuous and coherent understanding of the pendulum’s dynamics with the use of the VAE architecture, which focuses on finding continuous latent variables.

Double pendulum Figure 2 displays the latent variables obtained with each model, scaled to between -1 and +1. The baseline model (a) identifies four latent variables, `var0`, `var1`, `var2`, and `var3`, which




Figure 1: Values of the latent variables obtained by the models (y-axis) against time (x-axis) for one trajectory of the simple pendulum. The subplots show the respective latent variables for the (a) baseline and (b) HPI-VAE model.

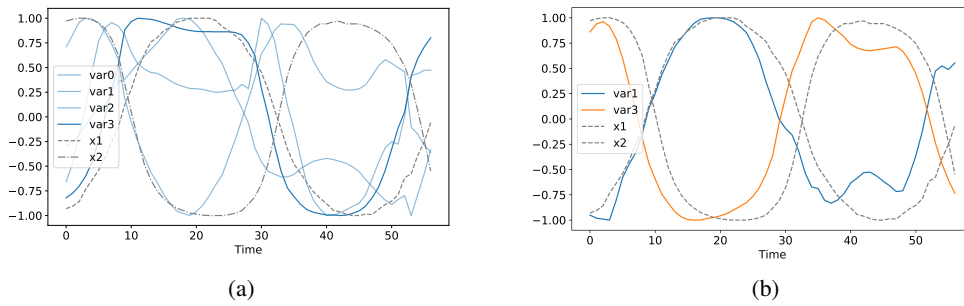


Figure 2: Values of the latent variables obtained by the models (y-axis) against time (x-axis) for one trajectory of the double pendulum. The subplots show the respective latent variables for the (a) baseline and (b) HPI-VAE model.

appear entangled or redundant. The HPI-VAE captures the system's evolution using four unique latent variables with the VAE architecture, which emphasizes disentanglement of the latent space. The first two latent variables are shown. The third and fourth are the time derivatives of the first two, omitted for visualization clarity but enforced with the physics-informed second-order constraint to its latent space. The horizontal positions of the two arms of the double pendulum, $x_1(t)$ and $x_2(t)$, are plotted in black dotted lines all subplots of Figure 2. For the baseline model, the third and fourth latent variables are correlated to $x_1(t)$ and $x_2(t)$, and the first two latent variables are correlated with the third. For the HPI-VAE, the first two latent variables are correlated with $x_1(t)$ and $x_2(t)$ which offers an even more refined and smooth portrayal of the dynamics of the two arms in the double pendulum system.

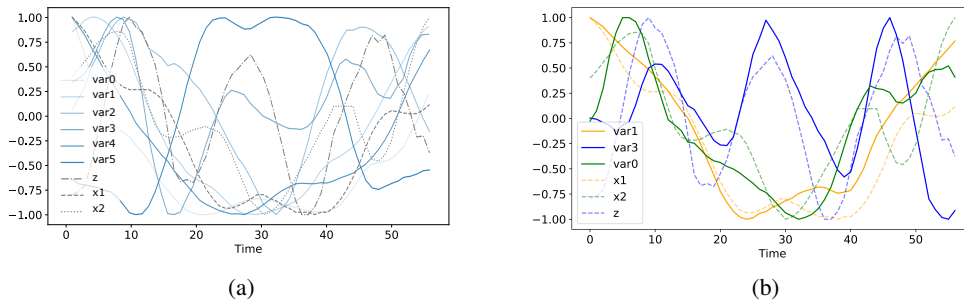


Figure 3: Values of the latent variables obtained by the models (y-axis) against time (x-axis) for one trajectory of the elastic pendulum. The subplots show the respective latent variables for the (a) baseline and (b) HPI-VAE model.

Elastic pendulum Figure 3 displays the latent variables obtained with each model, scaled to between -1 and +1. The baseline model (a) identifies six latent variables, var0, var1, var2, var3, var4, and var5. The HPI-VAE capture the system’s evolution using six unique latent variables. The first three latent variables are shown. The last three are the time derivatives of the first three, omitted for visualization clarity but enforced with the physics-informed second-order constraint to its latent space. The horizontal positions of the two arms of the elastic pendulum, $x_1(t)$ and $x_2(t)$, and the changing length of its main arm, z , are plotted in black dotted lines all subplots of Figure 2. For the baseline model, the fifth latent variable is correlated with $x_1(t)$. It is not possible to identify the correlation of the other variables with z and x_2 . However, the HPI-VAE can assign latent variables to model each arm’s position and the main arm’s fluctuating length. Although the match between the physical and model-derived variables in the HPI-VAE is not as precise as in simpler cases, it manages to trace the physical variables’ patterns more effectively, by encouraging the latent variables to converge to the canonical variables of a Hamiltonian system.

4 Conclusion

In conclusion, through the integration of physics-informed machine learning with variational autoencoders, we have allied physical knowledge and data-driven machine learning to enhance the interpretability and simplicity of modelling complex systems. Our approach marks an improvement over traditional methods by bypassing the need for an intrinsic dimension estimator and using a variational autoencoder to parsimoniously identify the degrees of freedom of dynamical systems during training. Furthermore, the latent variables are a minimal, non-redundant representation of dynamics that faithfully captures the system’s physical characteristics, especially the second-order constraint between pairs of latent variables. This advancement holds promise for a wide range of applications, from fundamental physics to engineering. We anticipate that the methodologies and insights gleaned from this work will help catalyze further research, fostering the development of more sophisticated, physics-informed models capable of tackling the complexities inherent in the natural and engineered world.

Future work should be directed toward the interpretation of latent variables. Different state variables have varying contributions to their respective dynamical systems, being able to differentiate the penalization of errors for different latent variables at large and small spatial or temporal scales will be beneficial to identifying more interpretable state variables.

5 Acknowledgements

References

- [1] Andrea Asperti. Sparsity in variational autoencoders, 2019.
- [2] Richard E. Bellman. *Adaptive Control Processes : A Guided Tour*. Princeton University Press, Princeton, 1961.
- [3] Tom Bertalan, Felix Dietrich, Igor Mezić , and Ioannis G. Kevrekidis. On learning hamiltonian systems from data. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(12):121107, dec 2019.
- [4] Félix Chavelli, Khoo Zi-Yu, Jonathan Sze Choong Low, and Stéphane Bressan. Assessing the effectiveness of intrinsic dimension estimators for uncovering the phase space dimensionality of dynamical systems from state observations. In Christine Strauss, Toshiyuki Amagasa, Gabriele Kotsis, A. Min Tjoa, and Ismail Khalil, editors, *Database and Expert Systems Applications*, pages 259–265, Cham, 2023. Springer Nature Switzerland.
- [5] Boyuan Chen, Kuang Huang, Sunand Raghupathi, Ishaan Chandratreya, Qiang Du, and Hod Lipson. Automated discovery of fundamental variables hidden in experimental data. *Nature Computational Science*, 2:433–442, 07 2022.
- [6] L. Mars Gao and J. Nathan Kutz. Bayesian autoencoders for data-driven discovery of coordinates, governing equations and fundamental constants, 2022.
- [7] Zhangyang Gao, Cheng Tan, Lirong Wu, and Stan Z. Li. Simvp: Simpler yet better video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3170–3180, June 2022.
- [8] Jerry Ginsberg. *Engineering dynamics*, volume 10, pages 397–400. Cambridge University Press, 2008.
- [9] Ian J. Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, Cambridge, MA, USA, 2016. <http://www.deeplearningbook.org>.
- [10] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [11] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, Jun 2021.
- [12] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [13] Elizaveta Levina and Peter Bickel. Maximum likelihood estimation of intrinsic dimension. *Advances in neural information processing systems*, 17, 2004.
- [14] Kenneth R. Meyer and Glen R. Hall. *Hamiltonian Differential Equations and the N-Body Problem*, pages 1–32. Springer New York, New York, NY, 1992.
- [15] David D. Nolte. The Tangled Tale of Phase Space. In *Galileo Unbound: A Path Across Life, the Universe and Everything*. Oxford University Press, 07 2018.
- [16] Riku Sasaki, Naoya Takeishi, Takehisa Yairi, and Koichi Hori. Neural gray-box identification of nonlinear partial differential equations. In *PRICAI 2019: Trends in Artificial Intelligence: 16th Pacific Rim International Conference on Artificial Intelligence, Cuvu, Yanuca Island, Fiji, August 26–30, 2019, Proceedings, Part II*, page 309–321, Berlin, Heidelberg, 2019. Springer-Verlag.
- [17] Yunbo Wang, Haixu Wu, Jianjin Zhang, Zhifeng Gao, Jianmin Wang, Philip S. Yu, and Mingsheng Long. Predrnn: A recurrent neural network for spatiotemporal predictive learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):2208–2225, 2023.

- [18] Hao Wu, Shilong Wang, Yuxuan Liang, Zhengyang Zhou, Wei Huang, Wei Xiong, and Kun Wang. Earthfarseeer: Versatile spatio-temporal dynamical systems modeling in one model. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, 2024*.

Appendices

A Background

Dynamical systems are defined by a set of time-dependent equations that characterize the evolution of a system’s state over time [8], and are represented as

$$\dot{\chi} = f(\chi, t) \quad (1)$$

where χ is a vector of the state variables of length equivalent to the number of degrees of freedom of the dynamical system, $\dot{\chi}$ is the time derivative of χ with respect to time t , and f is a function that defines the system dynamics [8].

Autonomous second-order dynamical systems are a class of dynamical systems characterized by an even number of state variables where half of them describe the position (τ), and the other half represent the associated momentum (ρ) [8]. Hamiltonian dynamical systems are a class of dynamical systems whose dynamics are governed by Hamilton’s equations

$$\dot{\tau} = \frac{\partial H}{\partial \rho}, \quad \dot{\rho} = -\frac{\partial H}{\partial \tau} \quad (2)$$

where H represents the total energy or Hamiltonian of the dynamical system [14, 15], and is a conserved quantity. Hamiltonian neural networks [3, 11, 10] are physics-informed machine learning models that incorporate learning biases based on Hamilton’s equations within neural networks. They regress the Hamiltonian of a dynamical system directly from its state variables using Hamilton’s equations and enforce the invariance of the total energy of the dynamical system.

An autoencoder is a neural network comprising an encoder function, which constructs an encoding of the input, and a decoder function, which produces a reconstruction of the input [9]. Sandwiched between them is a hidden layer that describes a code, or latent variables, used to represent the input [9]. Autoencoders are designed to be unable to learn to copy perfectly, usually by limiting the number of latent variables [9].

Variational Autoencoders (VAEs) extend the autoencoder framework by introducing a probabilistic approach to encoding inputs [12]. VAEs modify the loss function of an autoencoder by adding a Kullback-Leibler (KL) divergence term in the loss function,

$$\mathcal{L}_{VAE} = -\beta \cdot \mathbb{E}_{q(z|x)}[\log p(x|z)] + KL(q(z|x)||p(z)) \quad (3)$$

where x is the input and z its latent representation, $-\mathbb{E}_{q(z|x)}[\log p(x|z)]$ is the reconstruction loss in a variational context, $q(z|x)$ denotes the encoder’s distribution and $p(z)$ denotes a prior distribution $p(z)$. β adjusts the respective weights of the two terms. It can be placed on either the reconstruction or KL divergence term. The KL divergence term penalizes deviations of the learned distribution in the latent space from a chosen prior Gaussian distribution.

The KL divergence term enforces independence among the latent variables by pushing the encoded distributions to resemble the prior [1]. Therefore variables that do not contribute significantly to reducing the reconstruction loss become redundant and converge to the non-informative zero prior [1]. This process inherently minimizes the size of the latent space to the number of degrees of freedom of the dynamical system, by eliminating excess variables. Furthermore, the inclusion of the KL divergence encourages the VAE to find disentangled, semantically meaningful, statistically independent, and causal factors of variation in data [12]. The result is a more interpretable and minimal representation, which facilitates understanding of the underlying structure of the data.

B Baseline and proposed architectures

Chen et al.’s baseline model is a nested autoencoder [5] (Figure 4, A and B). In the first step, two consecutive video frames from a dynamical system are input to the outermost autoencoder. A compact representation of the input is extracted and the subsequent two consecutive frames are predicted (Figure 4, A). Then, an intrinsic dimension (ID) estimation [13] of the compact representation is performed. In the second step, the innermost autoencoder further compresses the compact representation into a latent space of dimension ID (Figure 4, B). These are the state variables of the dynamical system.

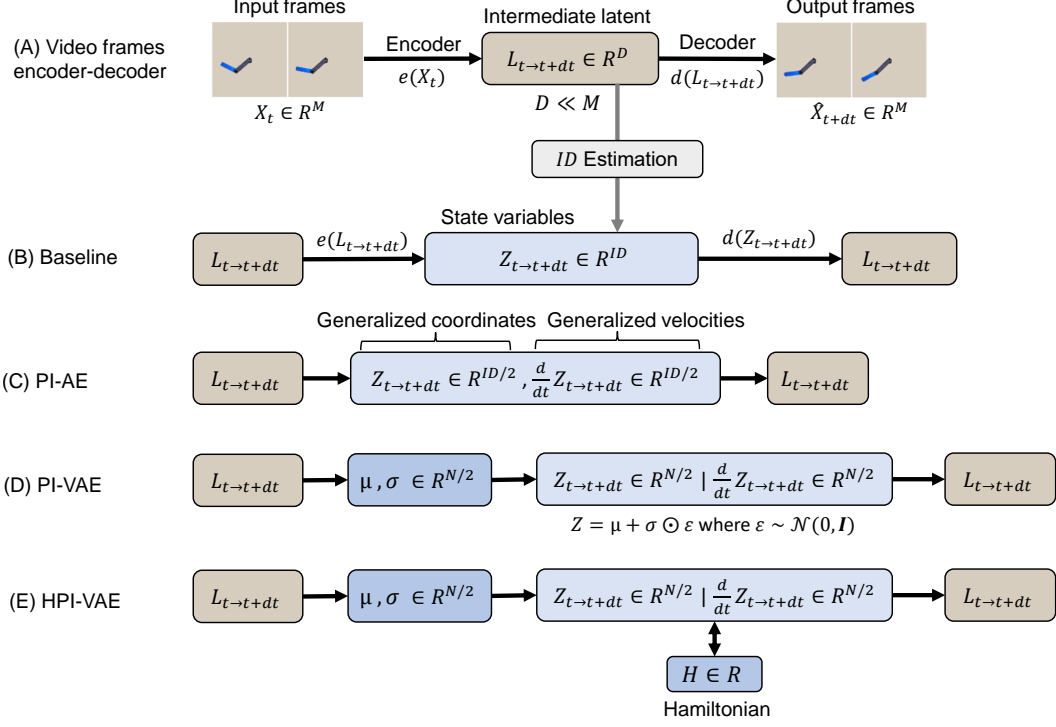


Figure 4: Chen et al.’s baseline (A and B) and our proposed models (C, D, E).

We leverage Karniadakis et al.’s framework to enhance Chen et al.’s model by introducing three proposed models, each incorporating different physics-informed machine learning biases. The three proposed models shown in Figure 1 C, D, and E utilize the outermost autoencoder from Chen et al. (Figure 4, A) and make modifications to the innermost autoencoder (Figure 4, B).

Table 2 sums up the biases introduced in each architecture. The three modifications are progressively incorporated into the baseline model using observational, learning, and inductive bias respectively. Through a systematic comparison in Section 3, we evaluate the capability of each model to capture the dynamics of various systems in a minimal, interpretable set of latent variables.

Table 2: Summary of baseline and proposed models

Model	Observational bias	Learning bias	Inductive bias
Baseline	-	-	-
PI-AE	2^{nd} order constraint	-	-
PI-VAE	2^{nd} order constraint	Latent sparsity	-
HPI-VAE	2^{nd} order constraint	Latent sparsity	Hamiltonian conservation

C Frames of the systems studied

The five dynamical systems¹ are the reaction-diffusion system (simulated dynamics of a planar spiral wave), a single pendulum system, a double pendulum system, a swing stick system (articulated arms fixed on a rigid base), and an elastic pendulum system (double pendulum which main arm length can vary).

¹These systems are used as is from the experiments of Chen et al. and are available for download here: <https://github.com/BoyuanChen/neural-state-variables>




Figure 5: Sequences of frames of the five dynamical systems: from the top in each row, the reaction-diffusion system, single pendulum system, double pendulum system, elastic pendulum system, and swing stick system from Chen et al.

D Hyperparameters for each architecture

Table 3 reports the value of β for the PI-VAE and HPI-VAE for the five dynamical systems. All other loss weights for all models are set to 1. All innermost autoencoders are trained for 1000 epochs and have reconstruction losses of less than 0.01.

Table 3: Reconstruction loss weights for PI-VAE and HPI-VAE

System	PI-VAE	HPI-VAE
Reaction-diffusion	7	7
Single pendulum	17	20
Double pendulum	30	40
Swingstick system	30	30
Elastic pendulum	50	80

E Results for the PI-AE and PI-VAE

Single pendulum Figure 6 displays the latent variables obtained with each model, scaled to between -1 and +1. The baseline model (a) identifies two latent variables, `var0` and `var1`, which appear entangled or redundant. The PI-AE, PI-VAE, and HPI-VAE capture the system’s evolution using two unique latent variables. The first latent variable is shown. The second is its time derivative, omitted for visualization clarity. For all models, the first of the two latent variables are correlated with $\cos 2\theta(t)$, where $\theta(t)$ is the angle of the pendulum arm with respect to its position of rest at time t given on the x-axis, although this correlation is more apparent with the PI-VAE and HPI-VAE. The plot of $\cos 2\theta(t)$ against time, t , is plotted as a black dotted line on all subplots of Figure 6. It can also be observed that both the baseline and PI-AE in (a) and (b) exhibit discontinuous changes that align with the curve of $\cos 2\theta(t)$. However, in the PI-VAE and HPI-VAE models, the latent variable is smoother, revealing a more continuous and coherent understanding of the pendulum’s dynamics. In particular, the latent variable of the HPI-VAE has sharper peaks, which are more similar to the peaks of the $\cos 2\theta(t)$ curve.

Double pendulum Figure 2 displays the latent variables obtained with each model, scaled to between -1 and +1. The baseline model (a) identifies four latent variables, `var0`, `var1`, `var2`, and `var3`, which appear entangled or redundant. The PI-AE, PI-VAE, and HPI-VAE capture the system’s evolution using four unique latent variables. The first two latent variables are shown. The third and fourth are

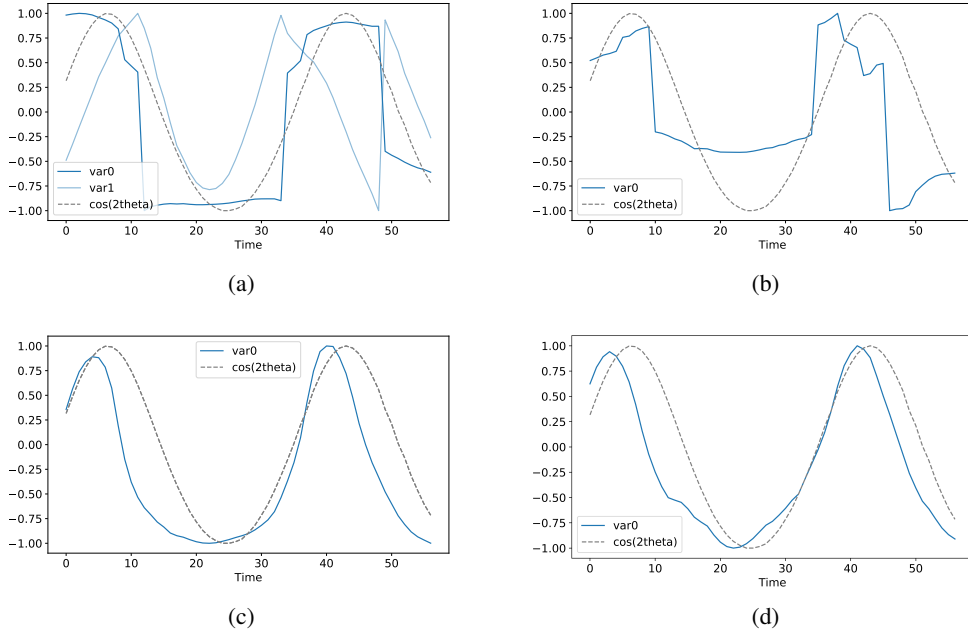


Figure 6: Values of the latent variables obtained by the models (y-axis) against time (x-axis) for one trajectory of the simple pendulum. The subplots show the respective latent variables for the (a) baseline, (b) PI-AE, (c) PI-VAE, and (d) HPI-VAE model.

the time derivatives of the first two, omitted for visualization clarity. The horizontal positions of the two arms of the double pendulum, $x_1(t)$ and $x_2(t)$, are plotted in black dotted lines all subplots of Figure 7. For the baseline model, the third and fourth latent variables are correlated with $x_1(t)$ and $x_2(t)$, and the first two latent variables are correlated with the third. For the PI-AE, PI-VAE, and HPI-VAE, the first two latent variables are correlated with $x_1(t)$ and $x_2(t)$. The PI-AE demonstrates the ability to distinguish and more accurately represent the joint evolution of both arms. This representation is further enhanced in the PI-VAE and HPI-VAE models, which offer an even more refined and smooth portrayal of the dynamics of the two arms in the double pendulum system.

Elastic pendulum Figure 8 displays the latent variables obtained with each model, scaled to between -1 and +1. The baseline model (a) identifies six latent variables, var0, var1, var2, var3, var4, and var5. The PI-AE, PI-VAE, and HPI-VAE capture the system's evolution using six unique latent variables. The first three latent variables are shown. The last three are the time derivatives of the first three, omitted for visualization clarity. The horizontal positions of the two arms of the elastic pendulum, $x_1(t)$ and $x_2(t)$, and the changing length of its main arm, z , are plotted in black dotted lines all subplots of Figure 7. For the baseline model, the fifth latent variable is correlated with $x_1(t)$. It is not possible to identify the correlation of the other variables with z and x_2 . Similarly, the latent variables of the PI-AE also do not correlate with z and x_2 , as seen in Figure 8 (b). However, the PI-VAE in Figure 8 (c) and HPI-VAE in Figure 8 (d) can assign latent variables to model each arm's position and the main arm's fluctuating length. Although the match between the physical and model-derived variables in PI-VAE and HPI-VAE models is not as precise as in simpler cases, these models manage to trace the physical variables' patterns more effectively.

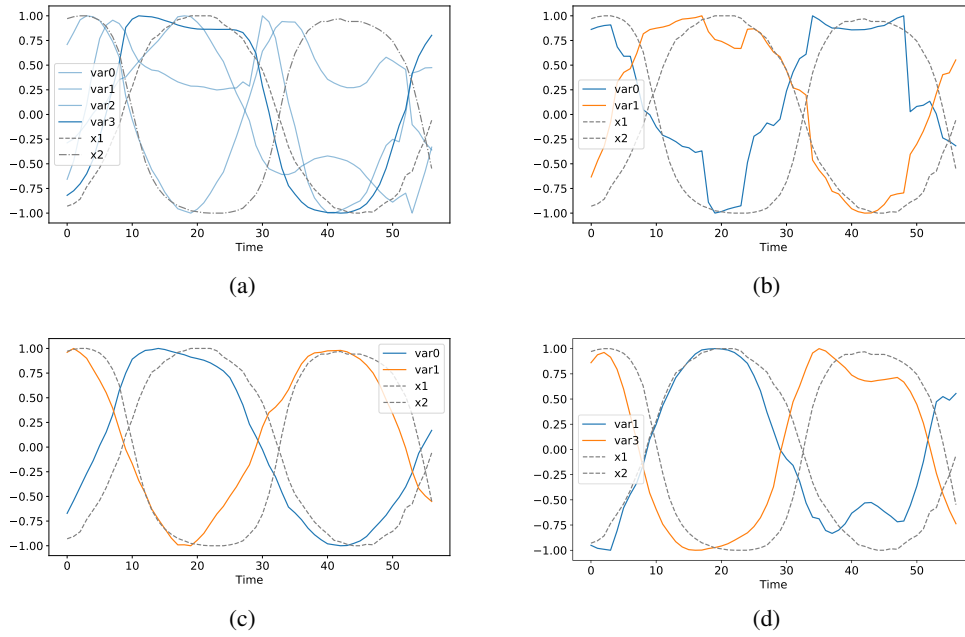


Figure 7: Values of the latent variables obtained by the models (y-axis) against time (x-axis) for one trajectory of the double pendulum. The subplots show the respective latent variables for the (a) baseline, (b) PI-AE, (c) PI-VAE, and (d) HPI-VAE model.

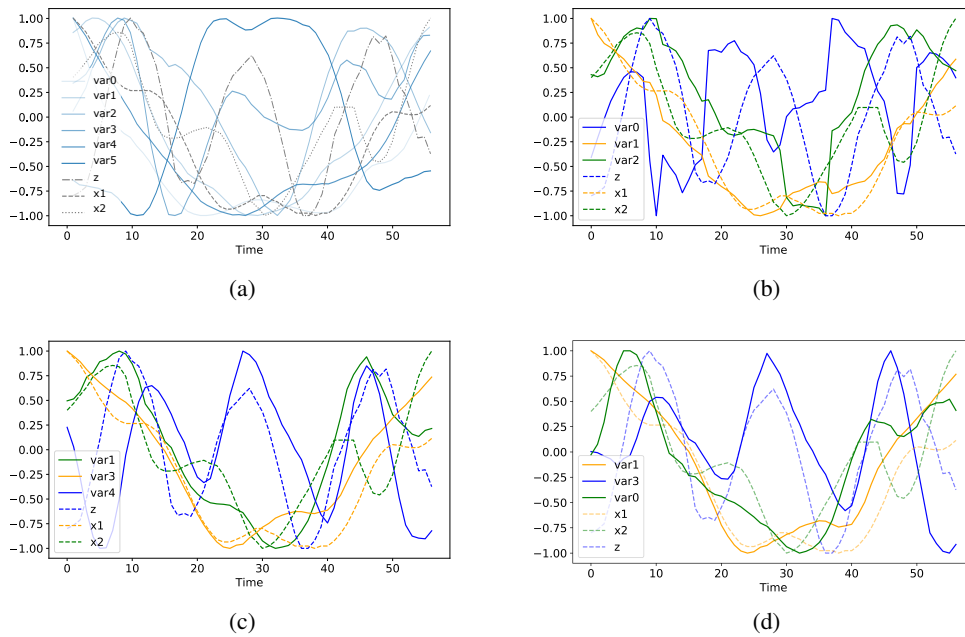


Figure 8: Values of the latent variables obtained by the models (y-axis) against time (x-axis) for one trajectory of the elastic pendulum. The subplots show the respective latent variables for the (a) baseline, (b) PI-AE, (c) PI-VAE, and (d) HPI-VAE model.