Incorporating Additive Separability into Hamiltonian Neural Networks for Regression and Interpretation

Khoo Zi-Yu
School of Computing
National University of Singapore
Singapore, Singapore
khoozy@comp.nus.edu.sg

Jonathan Sze Choong Low
Singapore Institute of Manufacturing Technology
Agency for Science, Technology and Research
Singapore, Singapore
sclow@simtech.a-star.edu.sg

Stéphane Bressan
School of Computing
National University of Singapore
Singapore, Singapore
steph@nus.edu.sg

Abstract

Hamiltonian neural networks are state-of-the-art models that regress the vector field of a dynamical system under the learning bias of Hamilton's equations. A recent observation is that incorporating a second-order bias regarding the additive separability of the Hamiltonian reduces the regression complexity and improves regression performance. We propose three *separable Hamiltonian neural networks* that incorporate additive separability within Hamiltonian neural networks using observational, learning and inductive biases. We show that the proposed models are more effective than a Hamiltonian neural network at regressing a vector field, and have the capability to interpret the kinetic and potential energy of the system.

1 Introduction

Hamiltonian neural networks [2, 4] are physics-informed neural networks with learning biases [6] given by Hamilton's equations and their corollaries. They unsupervised-ly regress the vector field [12, 11] and the Hamiltonian [8] of a dynamical system from discrete observations of its state space or vector field, and outperform their physics-uninformed counterparts in doing so [7].

A Hamiltonian system is additively separable if the Hamiltonian can be separated into additive terms, each dependent on disjoint subsets of the state variables [3]. Given two disjoint subsets of state variables $\vec{x}$ and $\vec{y}$, the Hamiltonian, $H(\vec{x}, \vec{y}) = T(\vec{x}) + V(\vec{y})$, has a mixed partial derivative of zero. T and V are arbitrary [3] but can represent the kinetic and potential energies of a mechanical Hamiltonian. Recently, Gruver et al. in their ICML 2022 spotlight paper, found that additive separability allowed the physics-informed neural network to "*avoid[...] artificial complexity from the coordinate system*" (the state variables) and improve its performance [5]. This motivates incorporating additive separability within Hamiltonian neural networks to regress additively separable Hamiltonians.

The main contribution of this work is three independent methods to incorporate additive separability into a Hamiltonian neural network. We call this family of Hamiltonian neural networks, *separable Hamiltonian neural networks*. The separable Hamiltonian neural networks follow the nomenclature of Karniadakis et al. [6] and incorporate additive separability using either an observational bias, learning bias, or inductive bias. The first model incorporates an observational bias by training on newly generated data that embody separability. The second model incorporates a learning bias through the

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

loss function of the Hamiltonian neural network. The third model incorporates an inductive bias through the architecture of the Hamiltonian neural network using conjoined multilayer perceptrons.

We empirically evaluate the performance of the proposed models against a baseline Hamiltonian neural network on a variety of representative additively separable Hamiltonian systems for the task of regressing an additively separable Hamiltonian vector field. We find that the proposed models outperform the baseline on the regression task. Evaluation of the proposed model with inductive bias also allows interpretation of the kinetic and potential energy of each Hamiltonian system.

2 Methodology

Four Hamiltonian neural networks, shown in Figure 1, are compared for the task of regressing the vector field of an additively separable Hamiltonian system from random samples of a vector field. One, the baseline, is uninformed of the additive separability of the Hamiltonian system. Three proposed models are informed, and incorporate observational, learning and inductive biases respectively.

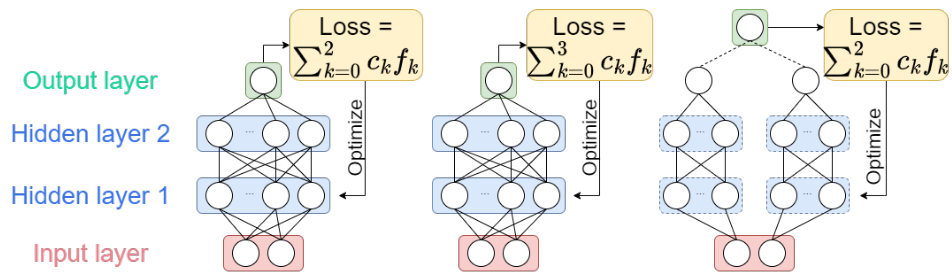

Figure 1: Architecture of the baseline and proposed Hamiltonian neural networks. Their loss functions, shown in yellow, comprise equations f_k with a coefficient c_k representing the real-valued weight of each f_k. (a) The Hamiltonian neural network (HNN) and proposed separable Hamiltonian neural network with observational bias (sHNN-O). They have the same architecture but differ in their input data. (b) The proposed separable Hamiltonian neural network with learning bias (sHNN-L). (c) The proposed separable Hamiltonian neural network with inductive bias (sHNN-I).

2.1 The baseline Hamiltonian neural network

We adapt Hamiltonian neural networks [2, 4] (Figure 1a) for the task of regressing the vector field. Hamiltonian neural networks incorporate a Hamiltonian learning bias into the model using the linear summation of three equations in its loss function. The first is an arbitrary pinning term that "pins" the regressed Hamiltonian to one among several solutions that are modulo and additive constant. The second and third ensure the model satisfies Hamilton's equations in Equation 1. In Equation 1, $\vec{x}$ and $\vec{y}$ are conjugate state variables of the Hamiltonian system, and typically referred to as the position and momentum of a mechanical Hamiltonian.

$$\frac{\mathrm{d}\vec{x}}{\mathrm{d}t} = \frac{\partial H(\vec{x},\vec{y})}{\partial \vec{y}} \quad \frac{\mathrm{d}\vec{y}}{\mathrm{d}t} = -\frac{\partial H(\vec{x},\vec{y})}{\partial \vec{x}} \tag{1}$$

To perform the task of regressing the vector field of a Hamiltonian system, the Hamiltonian neural network uses instantaneous, random samples of the state variables and vectors. The state variables are input to the model. The output is the regression surrogate, an estimator of the Hamiltonian. Training is supervised through the loss function, which uses the vectors corresponding to the state variables input to the model. The derivative of the surrogate Hamiltonian at all input state variables is the surrogate vector field, computed via automatic differentiation [9] of the Hamiltonian neural network.

The separable Hamiltonian neural networks adopt the Hamiltonian learning bias of the Hamiltonian neural network and further incorporate biases regarding additive separability. They perform the task of regressing the Hamiltonian vector field in the same way as the Hamiltonian neural network.

2.2 Incorporating a separability observational bias within a Hamiltonian neural network

The baseline Hamiltonian neural network can be informed of additive separability by incorporating an observational bias. Consider data of the vector field comprising instantaneous state variables, which are $(\vec{x}, \vec{y})$, and their instantaneous time derivatives. Additive separability of the Hamiltonian means that the time derivatives of the state variables are dependent solely on $\vec{x}$ or $\vec{y}$. Therefore, new samples of the vector field can be generated, by creating new combinations of $\vec{x}$ and $\vec{y}$ from the existing state variables and inferring their instantaneous time derivatives from existing samples of the vector field with the same $\vec{x}$ or $\vec{y}$. As each of the N samples can be combined with $N-1$ remaining samples, the quantity of existing data can be quadratically scaled. Training the Hamiltonian neural network on the new data incorporates the observational bias and allows the model to regress a surrogate Hamiltonian that reflects the additive separability of the data.

2.3 Incorporating a separability learning bias within a Hamiltonian neural network

The baseline Hamiltonian neural network can be informed of additive separability by incorporating a learning bias into the Hamiltonian neural network, using the linear summation of four equations in its loss function (Figure 1b). The first three equations follow from the Hamiltonian neural network. The fourth equation ensures the mixed partial derivative of the surrogate Hamiltonian $\hat{H}$ is zero. This incorporates a learning bias in the separable Hamiltonian neural network that favours convergence of the Hamiltonian neural network toward a surrogate Hamiltonian that is additively separable.

2.4 Incorporating a separability inductive bias within a Hamiltonian neural network

The baseline Hamiltonian neural network can be informed of additive separability by incorporating an inductive bias using a non-fully connected architecture (Figure 1c) comprising two smaller neural networks conjoined only at the output layer. Each smaller conjoined neural network has one output. Their sum (indicated by the dotted lines in Figure 1) is the surrogate Hamiltonian. This architecture ensures additive separability as each smaller neural network has an input of either $\vec{x}$ or $\vec{y}$, and is, therefore, a function of either $\vec{x}$ or $\vec{y}$. During forward propagation, the sum of the two conjoined neural networks ensures the surrogate Hamiltonian is always the sum of two independent functions of $\vec{x}$ and $\vec{y}$. The mixed partial derivative of the surrogate Hamiltonian is by design always zero. Additive separability is strictly satisfied. The separable Hamiltonian neural network with inductive bias regresses a surrogate Hamiltonian that implicitly and strictly satisfies additive separability.

3 Results

The four models are compared on the task of regressing the Hamiltonian vector field of an n-dimensional Hamiltonian. Training data comprising 512 samples of the state variables, $(\vec{x}, \vec{y})$, and their corresponding time derivatives or vector field. Training data is generated uniformly at random for each Hamiltonian system. In training, 20% of the training data is set aside as validation data for a dynamic stopping criterion using validation-based early stopping [10] and a batch size of 80 is used.

The models to be experimented on are designed with two hidden layers, an Adam optimizer and softplus activation. All models have an input layer with width $2 \times n$ corresponding to the $2 \times n$ state variables and an output layer with width one representing the surrogate Hamiltonian $H \in \mathbb{R}$. The models are trained until there is no improvement in the validation loss. Following Hamiltons' equations, the derivative of the surrogate Hamiltonian with respect to its inputs $\vec{x}$ and $\vec{y}$ is equivalent to the vector field. The derivative is computed via automatic differentiation. The performance of the models is measured by the error in regressing the vector field, using the vector error shown in Equation 2 [1]. $\hat{v}$ is the regressed vector from the derivative of the surrogate Hamiltonian and v is the true vector from test data. The test data set comprises $d = 10^{2n}$ vectors, evenly spaced in each of the n dimensions of the Hamiltonian system. Details regarding the Hamiltonian systems used for the experiments are in Appendix **??**. The code will be made available online.

$$E_V = \frac{1}{d} \sum_{k=1}^{d} \frac{||\hat{v}_k - v_k||_2}{||v_k||_2}. \tag{2}$$

For the setup of the Hamiltonian neural network and separable Hamiltonian neural network with learning bias, their loss functions comprise equations f_k, with a weight coefficient c_k. For both models, all equations are given the same weight of $c_k = 1$. For the model with learning bias, this balances the importance of additive separability and Hamilton's equations.

For the separable Hamiltonian neural network with observational bias, the training data comprises samples $(\vec{x}_i, \vec{y}_i)$ and time derivatives $(\frac{\mathrm{d}\vec{x}_i}{\mathrm{d}t}, \frac{\mathrm{d}\vec{y}_i}{\mathrm{d}t}) \forall i \in N = 512 \times 80\%$. New data is created comprising samples $(\vec{x}_i, \vec{y}_{i+m})$ with time derivatives $(\frac{\mathrm{d}\vec{x}_{i+m}}{\mathrm{d}t}, \frac{\mathrm{d}\vec{y}_i}{\mathrm{d}t})$ for $m = \{1, 2\}$, and appended to the original data. As m increases, the marginal improvements in regressing the vector field diminishes.

For the setup of the Hamiltonian neural network with inductive bias, forward propagation through the proposed model is computed as $x_z = \sigma(W_z \times x_{z-1} + B_z)$ where σ is the activation function, x_z is the output of layer z, W_z is the weight matrix of layer z of shape $2 \times L_z \times L_{z-1}$, and L_z and L_{z-1} are the widths of layers z and $z-1$ respectively, and B_z is the bias matrix of layer z of shape $2 \times 1 \times L_z$. The 2 in W_z and B_z correspond to the *two* conjoined neural networks, which are summed in the last layer. Further details regarding the setups are available in Appendix **??** and **??**.

Table 1: Vector error E_V and standard error (in brackets) of the baseline Hamiltonian neural network (HNN) and proposed separable Hamiltonian neural networks with observational (sHNN-O), learning (sHNN-L) and inductive (sHNN-I) biases. The lowest vector error is bolded.

Experiment	HNN	sHNN-O	sHNN-L	sHNN-I
Non-linear Pendulum	1.24E-02 (2.5E-03)	**6.80E-03** (1.7E-03)	9.97E-03 (1.7E-03)	8.54E-03 (1.3E-03)
Trigo	1.18E-02 (1.1E-03)	**8.17E-03** (7.2E-04)	1.22E-02 (1.2E-03)	8.19E-03 (9.1E-04)
Arctangent	5.09E-03 (2.9E-04)	**3.55E-03** (2.5E-04)	4.64E-03 (4.0E-04)	5.42E-03 (1.2E-03)
Logarithmic	4.32E-03 (2.2E-04)	**3.02E-03** (1.7E-04)	4.00E-03 (2.4E-04)	4.70E-03 (2.5E-04)
Anisotropic Oscillator	7.33E-03 (6.3E-04)	**4.53E-03** (4.5E-04)	5.67E-03 (5.7E-04)	5.43E-03 (4.8E-04)
Henon Heiles	1.11E-02 (8.9E-04)	**5.48E-03** (3.8E-04)	8.04E-03 (6.1E-04)	7.14E-03 (6.4E-04)
Toda Lattice	9.47E-03 (4.9E-04)	**4.65E-03** (3.6E-04)	5.39E-03 (3.2E-04)	7.28E-03 (8.9E-04)
Coupled Oscillator	1.10E-02 (1.4E-03)	**5.35E-03** (5.0E-04)	5.82E-03 (4.9E-04)	6.43E-03 (4.5E-04)

From Table 1, in general, all proposed separable Hamiltonian neural networks regress the Hamiltonian vector field with a lower vector error than the baseline Hamiltonian neural network. The proposed models leverage physics information regarding separability to penalise or prevent interaction between the state variables and this reduces the complexity of the Hamiltonian regression problem. The regression of the Hamiltonian vector field is therefore improved. The model with observational bias regresses the vector field best. The observational bias generates more samples of the data, which simultaneously emphasises additive separability and covers the input domain of the Hamiltonian system, which eases the interpolation task to improve the regression of the vector field. The models with learning and inductive biases only emphasise additive separability. An inductive bias generally outperforms a learning bias by restricting regressions to strictly satisfy separability, therefore forcing the model to simplify a complex regression problem into two smaller ones.

Incorporating knowledge regarding the additive separability of a mechanical Hamiltonian in the form of an inductive bias also improves the interpretability of the separable Hamiltonian neural network. The two conjoined neural networks become surrogates for the kinetic and potential energies of an otherwise black-box Hamiltonian system, and the instantaneous kinetic and potential energies can be evaluated. The absolute errors in regressing the kinetic and potential energy for five mechanical Hamiltonian systems are reported in Table 2. None of the other Hamiltonian neural networks allow such an interpretation of the kinetic and potential energies of the Hamiltonian system.

Table 2: Absolute error and standard error (in brackets) in regressing the kinetic and potential energy of the Hamiltonian system using the separable Hamiltonian neural network with inductive bias.

Experiment	Kinetic Energy	Potential Energy
Non-linear Pendulum	5.85E-04 (1.2E-04)	5.38E-03 (1.7E-03)
Anisotropic Oscillator	4.28E-04 (4.6E-05)	4.68E-04 (6.8E-05)
Henon Heiles	5.17E-04 (5.7E-05)	6.29E-04 (8.4E-05)
Toda Lattice	1.22E-03 (2.1E-04)	1.56E-03 (2.2E-04)
Coupled Oscillator	8.10E-04 (6.9E-05)	8.30E-04 (1.1E-04)

4 Conclusion

We propose three separable Hamiltonian neural networks that incorporate a second-order bias, or additive separability, within Hamiltonian neural networks. The resulting models are more effective than Hamiltonian neural networks in regressing Hamiltonian vector fields as they leverage additive separability to avoid artificial complexity between state variables. The best model for regression incorporates an observational bias, which additionally covers the input domain of the Hamiltonian system to ease the interpolation task. The best model for interpretability incorporates an inductive bias, to obtain insights regarding the kinetic and potential energies of the Hamiltonian system.

We are currently extending the approach to dynamically detect and incorporate separability as an inductive bias by rewiring the Hamiltonian neural network as it learns.

5 Acknowledgements

References

[1] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.

[2] 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.

[3] Zhijie Chen, Mingquan Feng, Junchi Yan, and Hongyuan Zha. Learning neural hamiltonian dynamics: A methodological overview, 2022.

[4] 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.

[5] Nate Gruver, Marc Anton Finzi, Samuel Don Stanton, and Andrew Gordon Wilson. Deconstructing the inductive biases of hamiltonian neural networks. In *International Conference on Learning Representations*, 2022.

[6] 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.

[7] Zi-Yu Khoo, Delong Zhang, and Stéphane Bressan. What's next? predicting hamiltonian dynamics from discrete observations of a vector field. In *Database and Expert Systems Applications: 33rd International Conference, DEXA 2022, Vienna, Austria, August 22–24, 2022, Proceedings, Part II*, page 297–302, Berlin, Heidelberg, 2022. Springer-Verlag.

[8] 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.

[9] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.

[10] Lutz Prechelt. Early stopping-but when? In *Neural Networks: Tricks of the Trade, This Book is an Outgrowth of a 1996 NIPS Workshop*, page 55–69, Berlin, Heidelberg, 1998. Springer-Verlag.

[11] Hiroki Sayama. *Introduction to the modeling and analysis of Complex Systems*. Open SUNY Textbooks, Milne Library, 2015.

[12] J. Wainwright and G. Tenti. Calculus 4 course notes for amath 231, September 2020.