

Learning reversible symplectic dynamics

Riccardo Valperga

Informatics Institute, University of Amsterdam, The Netherlands

R.VALPERGA@UVA.NL

Kevin Webster

Dmitry Turaev

Victoria Klein

Jeroen S. W. Lamb

Department of Mathematics, Imperial College London, United Kingdom

KEVIN.WEBSTER@IMPERIAL.AC.UK

D.TURAEV@IMPERIAL.AC.UK

VICTORIA.KLEIN18@IMPERIAL.AC.UK

JSW.LAMB@IMPERIAL.AC.UK

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Time-reversal symmetry arises naturally as a structural property in many dynamical systems of interest. While the importance of hard-wiring symmetry is increasingly recognized in machine learning, to date this has eluded time-reversibility. In this paper we propose a new neural network architecture for learning time-reversible dynamical systems from data. We focus in particular on an adaptation to symplectic systems, because of their importance in physics-informed learning.

Keywords: Physics-informed machine learning, time-reversal symmetry, symplectic neural networks, dynamical systems

1. Introduction

Neural networks are universal approximators: in principle, any sufficiently well-behaved function can be approximated, with arbitrary accuracy, using a neural network (Hornik et al., 1989). In practice, to achieve good approximations we need large datasets and for physical systems data acquisition can be costly. On the other hand, when dealing with physical systems, we very often use models which carry particular mathematical structures. Such structure can be recognised a priori, to guide - and improve the quality of - learning. In this paper, we address the problem of learning time-reversible and symplectic dynamics, motivated by their widespread occurrence in physics. Time-reversibility and symplecticity are fundamental properties that the corresponding learned dynamical system should have *exactly*, rather than approximately.¹

Time-reversing symmetry arises naturally as a structural property of many physical systems of interest (Lamb and Roberts, 1998). An elementary example is given by the Newtonian-type second order differential equations $\ddot{\mathbf{x}} = \mathbf{F}(\mathbf{x}, \dot{\mathbf{x}})$: when the forces \mathbf{F} do not depend on the velocities $\dot{\mathbf{x}}$ (or are even functions of the velocities), the equations do not change after the reversal of time $t \rightarrow -t$. This implies that if $(\mathbf{x}(t), \dot{\mathbf{x}}(t))$ is a solution, then $(\mathbf{x}(-t), -\dot{\mathbf{x}}(-t))$ is also a solution.

In general, a smooth system

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}), \tag{1}$$

1. In practical terms, this means that any deviation in the structure should be limited by the numerical accuracy of the computational device rather than the (larger) error of the optimisation algorithm.

of autonomous ordinary differential equations is said to possess a *reversing symmetry* if there exists an invertible smooth map R such that

$$\frac{dR(x)}{dt} = -\mathbf{F}(R(x)). \quad (2)$$

In this case, we call such system R -reversible. Similarly, a *discrete-time dynamical system* defined by an invertible map \mathcal{T} is called R -reversible if it possesses a reversing symmetry R :

$$R \circ \mathcal{T} = \mathcal{T}^{-1} \circ R. \quad (3)$$

The time- t flow map of an R -reversible system of differential equations is R -reversible for all $t \in \mathbb{R}$. *Symplectic dynamics* are generated by the *Hamiltonian* differential equations

$$\dot{\mathbf{q}} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}}, \quad \dot{\mathbf{p}} = -\frac{\partial \mathcal{H}}{\partial \mathbf{q}}, \quad (4)$$

where \mathcal{H} , the Hamiltonian or energy function, is a given function of the coordinates $\mathbf{q} \in \mathbb{R}^n$ and their conjugate momenta $\mathbf{p} \in \mathbb{R}^n$. This particular class of equations is important because of one of the most basic facts of physics: every *isolated* physical system has a Hamiltonian structure. Hamiltonian systems *preserve energy*, i.e., dynamics are constrained to the $(2n - 1)$ -dimensional manifold $\mathcal{H}(\mathbf{p}, \mathbf{q}) = \text{constant}$. The flow maps of a Hamiltonian system are symplectic in the following sense.²

A smooth map $\mathcal{T} : \mathcal{M} \rightarrow \mathbb{R}^{2n}$ (where \mathcal{M} is an open subset of \mathbb{R}^{2n}) is said to be symplectic if

$$(d_x \mathcal{T})^\top J (d_x \mathcal{T}) = J, \quad \forall x \in \mathcal{M}, \quad (5)$$

where $d_x \mathcal{T}$ is the Jacobian matrix of \mathcal{T} at the point x , and $J = \begin{bmatrix} 0 & \mathbb{I}_n \\ -\mathbb{I}_n & 0 \end{bmatrix}$.

When the Hamiltonian is an even function of the momenta, i.e., it is invariant with respect to the involution

$$R : (\mathbf{p}, \mathbf{q}) \rightarrow (-\mathbf{p}, \mathbf{q}), \quad (6)$$

system (4) is R -reversible. The classical example is given by mechanical systems where the Hamiltonian is the sum of potential and kinetic energies. However, not every Hamiltonian system is time-reversible (those involving an interaction with a magnetic field are, often, not). Also not every reversible system is Hamiltonian (non-holonomic mechanics provide many examples (Gonchenko et al., 2020)).

Note that the canonical reversing symmetry $R(p, q) = (-p, q)$ is an involution (i.e., $R^2 = id$) and is also *anti-symplectic*, i.e., it satisfies $R^\top J R = -J$. It is indeed natural to consider Hamiltonian systems with (anti-)symplectic time-reversal symmetries. In the examples discussed in this paper, reversing symmetries are always anti-symplectic, but our results also hold for symplectic reversing symmetries. The structure defined by a combination of symplecticity and an involutory reversing (anti-)symplectic symmetry is henceforth called *reversible symplectic structure*.³

2. For simplicity, we confine the discussion here to systems with a standard Darboux symplectic form.

3. Non-involutory reversing symmetries may also be considered, but they imply the existence of additional time-preserving symmetries, leading to reversible equivariant (symplectic) settings, which will not be considered in this paper.

The importance of geometric structure - such as symmetry or symplecticity - has begun to be recognized in the context of learning dynamical systems from observations (Burby et al., 2020; Jin et al., 2020), but - to date - this has eluded time-reversing symmetries. In this paper, we propose a new structure-preserving neural network for learning reversible dynamical systems. We focus in particular on learning reversible symplectic dynamics for their importance in physics-informed learning.

We implement our method for several examples of chaotic behavior including the paradigmatic Hénon-Heiles system and the periodically forced pendulum. Both in qualitative and quantitative aspects, we demonstrate improved performance and accuracy of our structure-preserving network, over the existing symplectic learning algorithms that are not time-reversible.

2. Setup

As usual in a supervised learning setting, the objective is to learn a map $\mathcal{T} : \mathcal{M} \rightarrow \mathcal{M}$ representing the evolution of a dynamical system, from observations: given a set $\{(x_i, y_i = \mathcal{T}(x_i))\}_{i=1}^N \subset \mathcal{M}$, the goal is to approximate \mathcal{T} with some map $\hat{\mathcal{T}} \in \mathbf{H}$, where \mathbf{H} a given function space which we refer to as the *hypothesis space*. We aim to find $\hat{\mathcal{T}}$ approximating \mathcal{T} with best fit, in the sense that

$$\hat{\mathcal{T}} = \arg \min_{\mathcal{T}^* \in \mathbf{H}} \frac{1}{N} \sum_{i=1}^N d(\mathcal{T}(x_i), \mathcal{T}^*(x_i)), \quad (7)$$

where d denotes the Euclidean metric $d(x, y) = \sqrt{\|x - y\|^2}$.

In the context of physics-informed learning, there may be a priori knowledge about the structure of the dynamical system that is to be learned. In this case, the hypothesis space should possess this structure intrinsically. Natural hypothesis spaces that approximate with arbitrary accuracy are infinite dimensional. The constructive approach is to build a countable family of hypothesis subspaces $\{\mathbf{H}_\ell\}_{\ell \in \mathbb{N}}$, with $\mathbf{H}_\ell \subset \mathbf{H}_{\ell+1}$ for all $\ell \in \mathbb{N}$, such that $\lim_{\ell \rightarrow \infty} \mathbf{H}_\ell = \mathbf{H}$ in an appropriate topology. In other words, $\cup_\ell \mathbf{H}_\ell$ universally approximates \mathbf{H} . When the learned function is represented by a neural network, the number ℓ is indicative of the number of degrees of freedom in the network. In practice, one settles for a certain \mathbf{H}_{ℓ^*} with ℓ^* large enough, so that targets are sufficiently well approximated. Key to a structure-preserving approach is that \mathbf{H}_ℓ has the desired structure *exactly* for all $\ell \in \mathbb{N}$. The objective of this paper is to propose choices of $\{\mathbf{H}_\ell\}_{\ell \in \mathbb{N}}$ that are reversible or reversible symplectic.

Neural networks architectures with hardwired symplectic structure have been proposed before. (Greydanus et al., 2019; Desai et al., 2021) learn a Hamiltonian function using a neural network. This requires numerical integration and training datasets with points that are sufficiently close to each other to estimate derivatives. Jin et al. (2020); Burby et al. (2020) develop algorithms for learning symplectic maps (including time-shifts by the flow of Hamiltonian systems). In either approach, time-reversibility was not yet taken into account.

3. Reversible symplectic neural networks

In this section we introduce reversible and reversible symplectic neural networks, which we subsequently show to be universal approximators for reversible and reversible symplectic dynamical systems.

3.1. Reversible neural networks

Consider the target space of R -reversible diffeomorphisms with an involutory symmetry R (i.e., $R \circ R = id_{\mathbb{R}^d}$):

$$\mathbf{H} = \left\{ \mathcal{T} \in \text{Diff}(\mathbb{R}^d) \mid \mathcal{T} = R \circ \mathcal{T}^{-1} \circ R \right\}. \quad (8)$$

We aim to find a family of hypothesis spaces $\{\mathbf{H}_\ell\}_{\ell \in \mathbb{N}}$ of R -reversible maps (i.e. $\mathbf{H}_\ell \subset \mathbf{H}$ for all $\ell \in \mathbb{N}$) such that $\lim_{\ell \rightarrow \infty} \mathbf{H}_\ell = \mathbf{H}$. Furthermore, the functions in \mathbf{H}_ℓ are to be parameterised in such a way that we can use common gradient-descent based optimization algorithms to find the best approximation to the unknown target function $\mathcal{T} \in \mathbf{H}$. We construct \mathbf{H}_ℓ using compositions of neural network-like maps.

Note that, in general, compositions of R -reversible maps are not R -reversible. However, given any set of invertible maps $\{f_n\}_{n=1}^\ell$ and letting $\hat{f}_i := R \circ f_i^{-1} \circ R$, it is readily verified that the following composition is R -reversible

$$F := \hat{f}_1 \circ \hat{f}_2 \circ \cdots \circ \hat{f}_\ell \circ f_\ell \circ f_{\ell-1} \circ \cdots \circ f_1. \quad (9)$$

We propose \mathbf{H}_ℓ to consist of functions of the form (9), where the functions $\{f_n\}_{n=1}^\ell$ are Real NVP bijective layers (Dinh et al., 2016) which have the important useful property that they are exactly (and easily) invertible. Real NPV bijective layers are defined as follows:

Definition 3.1 Given a d -dimensional input x , two functions $s, t : \mathbb{R}^{d'} \rightarrow \mathbb{R}^{d-d'}$, with $d' < d$, the output function $y_{1:d} = y(x_{1:d})$ of a Real NVP is

$$\begin{cases} y_{1:d'} = x_{1:d'} \\ y_{d'+1:d} = x_{d'+1:d} \odot \exp(s(x_{1:d'})) + t(x_{1:d'}), \end{cases} \quad (10)$$

where \odot is the element-wise product, and $y_{1:d} := (y_1, y_2, \dots, y_d)$.

3.2. Symplectic reversible neural networks

In analogy to Section 3.1, we consider R -reversible symplectic diffeomorphisms on a \mathbb{R}^{2n} as the target space:

$$\tilde{\mathbf{H}} = \left\{ \mathcal{T} \in \text{Diff}(\mathbb{R}^{2n}) \mid \mathcal{T} = R \circ \mathcal{T}^{-1} \circ R \text{ and } (D_x \mathcal{T})^T J (D_x \mathcal{T}) = J, \forall x \in \mathbb{R}^{2n} \right\}. \quad (11)$$

We propose to construct the hypothesis spaces $\tilde{\mathbf{H}}_\ell$ using compositions of the form (9), where the maps f_i are polynomial Hénon maps. Polynomial Hénon maps are symplectic transformations on $\mathbb{R}^n \times \mathbb{R}^n$, $(x, y) \mapsto (\bar{x}, \bar{y})$, with

$$\begin{cases} \bar{x} = y \\ \bar{y} = -x + \nabla V(y), \end{cases} \quad (12)$$

with $V : \mathbb{R}^n \rightarrow \mathbb{R}^n$ polynomial. We note here that, as required, the resulting composition (9) is symplectic if R is symplectic or anti-symplectic.

For practical implementation, it is useful that, like Real NVPs, the inverses of polynomial Hénon maps admit convenient and straightforward analytical expressions.

3.3. Universal approximation

We finally establish the universal approximation properties of the hypothesis spaces \mathbf{H}_ℓ and $\tilde{\mathbf{H}}_\ell$, proposed in Sections 3.1 and 3.2. Note that the composition (9) can be expressed as

$$\mathcal{T} = R \circ g^{-1} \circ R \circ g, \quad (13)$$

with $g := f_N \circ f_{N-1} \circ \cdots \circ f_1$. Let us first show that reversible maps indeed admit such a decomposition under mild assumptions. Let R be a linear involution. Recall that an orientation-preserving R -reversible C^2 -diffeomorphism \mathcal{T} of a d -dimensional ball into R^d is *smoothly isotopic to the identity*, i.e., there exists a smooth family of R -reversible diffeomorphisms $f : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ such that $f(x, 0) = x$, and $f(x, 1) = \mathcal{T}(x)$, for all x . Moreover, if \mathcal{T} is symplectic, that the diffeomorphism in the family f are also symplectic.

Theorem 3.1 *Let \mathcal{T} be an R -reversible diffeomorphism, with R being a linear involution. Let \mathcal{T} be smoothly isotopic to the identity. Then there exists a diffeomorphism $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$, such that $\mathcal{T} = R \circ g^{-1} \circ R \circ g$. If \mathcal{T} is symplectic, then g can be chosen symplectic.*

Proof Define the non-autonomous vector field h_s by the rule $h_s(f(x, s)) = \partial_s f(x, s)$, where f is the isotopy from the condition of the theorem. It is well-known that if the isotopy is symplectic, then h is Hamiltonian. Define the vector field

$$y_s = \frac{1}{2} R h_s \circ R.$$

We claim, that the inverse of the sought map g is the time-1 map defined along the orbits of the non-autonomous vector field y_s .

Indeed, it is enough to show that, for all $s \in [0, 1]$,

$$R f_s \circ g_s = g_s \circ R,$$

where g_s is the time- s map defined by the orbits of y_s (we denote $f_s = f(\cdot; s)$). Differentiating this with respect to s , we find that it is enough for the vector field y_s to satisfy the following identity:

$$R h_s \circ f_s + (R D_x f_s) y_s = y_s \circ R \circ f_s. \quad (14)$$

By the reversibility,

$$f_s \circ R \circ f_s = R.$$

Differentiating this equality with respect to s , we find

$$h_s \circ R + (D_{R f_s(x)} f_s) R h_s \circ f_s = 0,$$

and (14) indeed follows.

By construction, if h_s is Hamiltonian, then y_s is Hamiltonian too, so g is symplectic, as required. ■

Theorem 3.1 implies that the problem of approximating \mathcal{T} with arbitrary accuracy is reduced to the problem of approximating g with arbitrary accuracy. In the general (non-symplectic) reversible case, we can approximate any diffeomorphism g by Real NVP bijective layers in L^p according to the result of (Teshima et al., 2020). A stronger, C^∞ -approximation result, follows from (Turaev, 2015). When g is symplectic, it can be C^∞ -approximated by compositions of polynomial Hénon maps according to (Turaev, 2002).

4. Numerical experiments

Finally, we present some results concerning the approximation of Poincaré return maps of time-reversible Hamiltonian systems. We consider a periodically driven pendulum and the Hénon-Heiles system. In each of these examples, the dynamical systems are defined as solutions of explicitly given differential equations. However, the Poincaré maps that we aim to learn do not admit explicit expressions that can be derived directly from the differential equations. Hence, it is a natural objective to learn such maps using a neural network.

4.1. The Hénon-Heiles system

Consider a particle in \mathbb{R}^2 , with phase space coordinates $(x, y, p_x, p_y) \in \mathbb{R}^4$, whose equations of motion are generated by the Hamiltonian

$$\mathcal{H}(x, y, p_x, p_y) = \frac{1}{2}(p_x^2 + p_y^2 + x^2 + y^2) + \lambda(x^2y - \frac{y^3}{3}). \quad (15)$$

This dynamical system is known as the Hénon-Heiles System, after Michel Hénon and Carl Heiles who introduced it in their celebrated 1964 paper (Hénon and Heiles, 1964), as a model for the motion of a star in an axisymmetric potential. For a given value of the energy $E \in \mathbb{R}$ the orbits are constrained to 3-dimensional level sets of the Hamiltonian $\mathcal{E} = \mathcal{H}^{-1}(E)$.

The aim is to study the dynamics of the Hénon-Heiles system using a so-called Poincaré return map. A Poincaré section \mathcal{P} is a $(2n - 2)$ -dimensional submanifold of \mathcal{E} that is everywhere transversal to the flow. Let ϕ_t be the flow of the Hamiltonian system and \mathcal{P} a Poincaré section. Then the Poincaré return map $\mathcal{T} : \mathcal{P} \rightarrow \mathcal{T}(\mathcal{P}) \subseteq \mathcal{P}$ is defined as

$$\mathcal{T}(x) = \phi_{t_0}(x),$$

with $t_0 > 0$ least so that $\phi_{t_0}(x) \in \mathcal{P}$. If the section is smooth then $\mathcal{T} : \mathcal{P} \rightarrow \mathcal{T}(\mathcal{P})$ is a symplectic diffeomorphism. For more details on Poincaré return maps see (Broer and Takens, 2010).

It is readily verified that for the Hénon-Heiles system

$$\mathcal{P} = \{(x, y, p_x, p_y) \in \mathcal{E} \mid x = 0, p_x > 0\} \quad (16)$$

defines a Poincaré section on \mathcal{E} . In Fig. 1(a-c) we present some *Poincaré plots* (portraits containing selected orbits of initial conditions under the Poincaré map, illustrating the dynamics) for different values of the energy E . The Hénon-Heiles Hamiltonian (15) satisfies the property that $\mathcal{H}(x, y, p_x, p_y) = \mathcal{H}(x, y, -p_x, -p_y)$ which yields its dynamics reversible with respect to the reversing symmetry $R(x, y, p_x, p_y) = (x, y, -p_x, -p_y)$. As a consequence, it follows that the Poincaré return map \mathcal{T} on \mathcal{P} is also reversible, with respect to the involution $R_{\mathcal{P}} : (y, p_y) \rightarrow (y, -p_y)$. This time-reversal symmetry gives rise to the reflection symmetry $R_{\mathcal{P}}$ in the Poincaré plots in Fig. 1(a-c).

Poincaré maps provide important information about the dynamics. However, extensive studies of such maps obtained directly from the differential equations by numerical integration is often impractical due to the expense of the computations involved, tracking trajectories in the ambient space between hitting the surface of section. Training a neural network to learn a Poincaré map may be computationally expensive as well, but once trained it is fast to iterate.

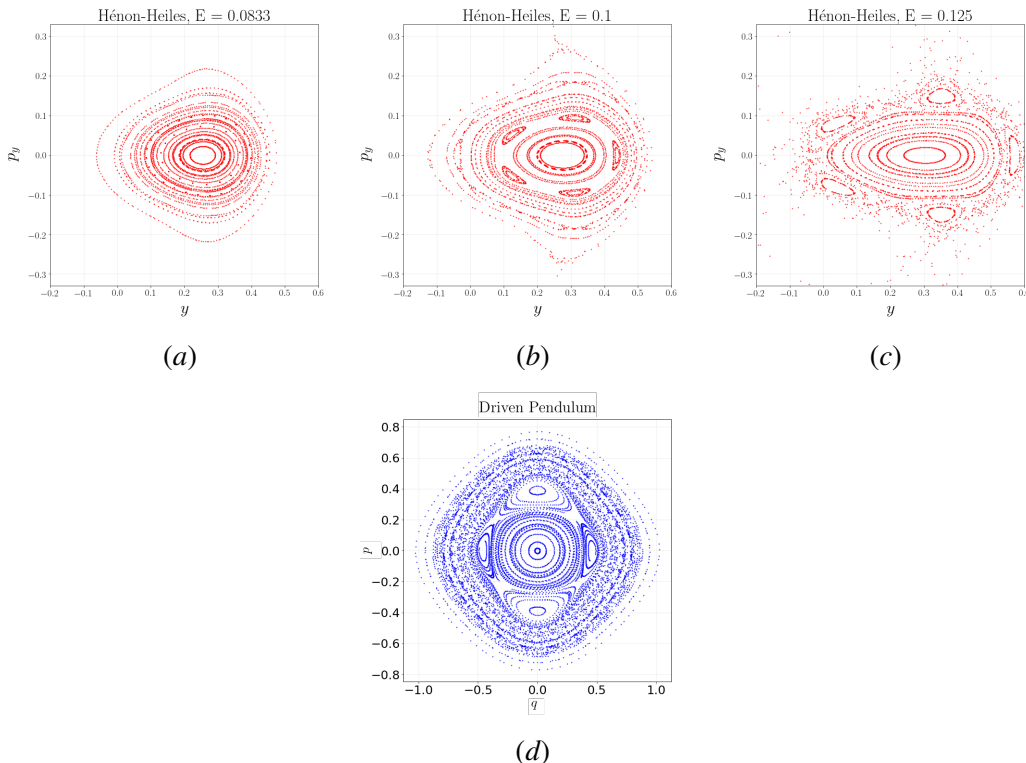


Figure 1: (a), (b), (c) Poincaré plots of the Hénon-Heiles system for different values of the energy E , and surface of section (16). (d) Poincaré plot of the periodically driven pendulum.

4.2. Periodically driven pendulum

Consider the time-periodic Hamiltonian

$$\mathcal{H}(p, q, t) = \frac{1}{2}p^2 - \nu^2 \cos(q) - \lambda [0.3pq \sin(2t) + 0.7pq \sin(3t)]. \quad (17)$$

This Hamiltonian represents a simple pendulum with natural frequency ν driven by a 2π -periodic force. The time- 2π map of this system defined by the flow $\Phi_{2\pi}$ is usually also referred to as a Poincaré map as one could view time as an augmented phase space coordinate and $t = 2\pi$ can be considered a surface of section in the extended phase space (p, q, t) . A Poincaré plot of this map is depicted in Fig. 1(d).

It can be shown that the time- 2π map $\Phi_{2\pi}$ is reversible symplectic with reversing symmetry $R(q, p) = (q, -p)$, which is also reflected by the observed symmetry in the Poincaré plot in Fig. 1(d).

4.3. Dataset and hypothesis spaces for experiments

For both systems the dataset is of form $\{(x_i, y_i), \mathcal{T}(x_i, y_i)\}_{i=1}^N$ where \mathcal{T} is the Poincaré map obtained via numerical integration. For both the pendulum and the Hénon-Heiles system, the optimizer has a small number of points available, i.e., $N = 100$ and $N = 300$ respectively.

We include the most significant experiments for three different physics-informed hypothesis spaces and the hypothesis space of plain MLPs:

- \mathbf{H}_{NN} : the set of MLPs, with depth 6, and a total of ≈ 5000 trainable parameters.
- \mathbf{H}_R : A function space \mathbf{H}_ℓ of $R_{\mathcal{P}}$ -reversible diffeomorphisms with 6 different, appropriately masked, Real NVP bijective layers of form (10) in which the functions s and t are represented by shallow MLPs.
- \mathbf{H}_{HR} : A function space $\tilde{\mathbf{H}}_\ell$ of R -reversible compositions of polynomial Hénon maps of degree 4, with depth 25.
- \mathbf{H}_{SN} : The set of symplectic neural networks (SympNets (Jin et al., 2020)) of depth 18, with 8 linear modules for each activation module.

The hyperparameters of these hypothesis spaces have been chosen so that they all have comparable numbers of trainable parameters and training times. All the experiments are implemented using Keras (Chollet et al., 2015) and have been run on a single Nvidia RTX 2080. The source code can be found [agithub.com/Ricvalp/SymplecticTimeReversibleNN](https://github.com/Ricvalp/SymplecticTimeReversibleNN). The models have been trained with Adam optimizer and a decaying learning rate; unless stated otherwise, the dataset has been divided into training/validation sets with a 9:1 ratio. Figure 2 reports the losses during training.

4.4. Results

To compare performances we have evaluated the trained evolution maps in the following way. First, we compare the resulting Poincaré plots. From these figures it is clear that evolution maps from intrinsically time-reversible hypothesis spaces result in Poincaré plots that look very similar to the ground truth obtained using numerical integration (Fig. 1). We then calculate the average Euclidean norm between a few ground truth trajectories and the predictions. Poincaré plots and errors are shown in Fig. 3 for the driven pendulum, and in Fig. 4 for the Hénon-Heiles system.

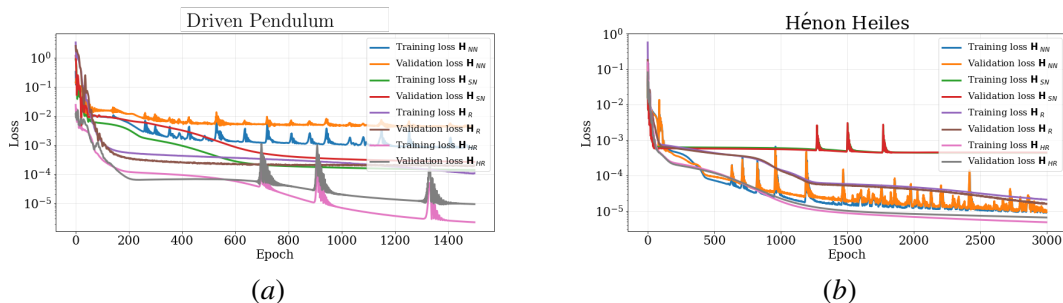


Figure 2: Training and validation loss versus epoch for the four hypothesis spaces and the two systems.

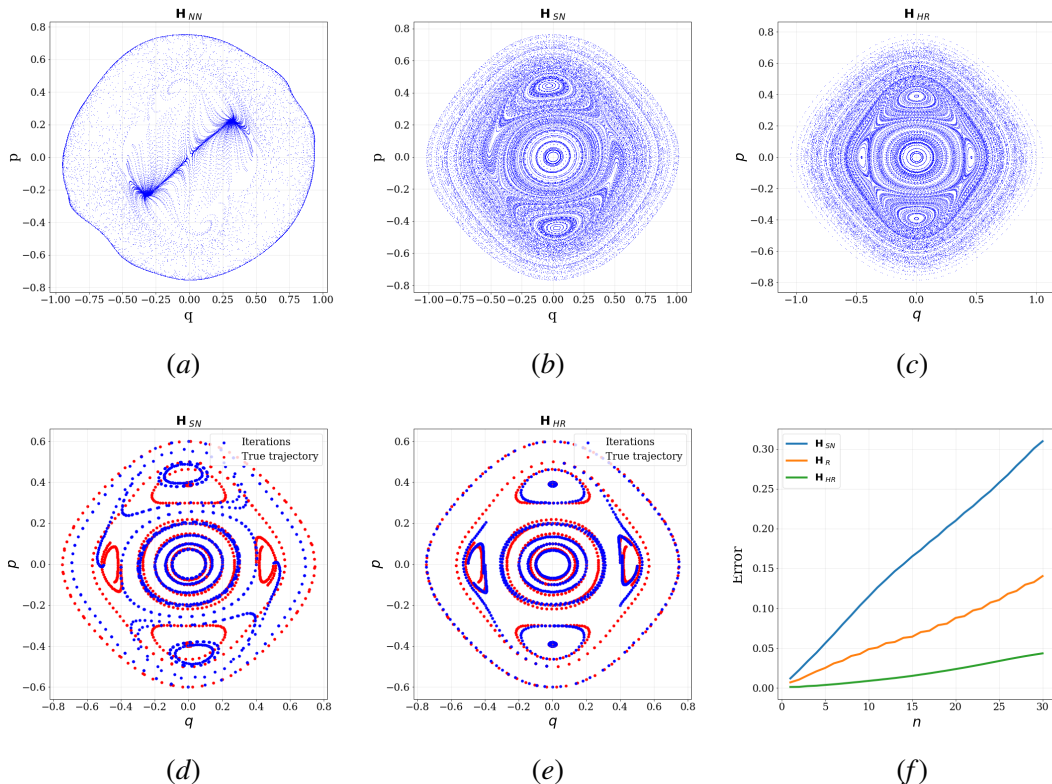


Figure 3: **Periodically driven pendulum.** Top: Poincaré plots obtained iterating the trained evolution maps from H_{NN} , H_{SN} , and H_{HR} . Bottom: a few predicted trajectories and the average prediction error, calculated as the average euclidean norm of the difference vs. the iteration n .

5. Conclusion

We have addressed the problem of learning the evolution map of a discrete-time dynamical system that is known to be either reversible or reversible symplectic. Many Hamiltonian systems (and their Poincaré maps) are reversible symplectic. We have proposed relevant hypothesis spaces with these structures in terms of trainable neural networks, extending the results of [Jin et al. \(2020\)](#) and [Burby et al. \(2020\)](#) to symplectic systems that are also reversible. We observed in numerical experiments that reversible symplectic networks perform better than existing symplectic ones, like SympNets. Also, even reversible networks without symplectic structure, involving non-symplectic Real NVP bijective layers, perform remarkably well on reversible symplectic systems.




Figure 4: **Hénon-Heiles system**. Top: Poincaré plots obtained iterating the trained evolution maps from H_{SN} , H_R , and H_{HR} . Bottom: a few predicted trajectories.

References

- H.W. Broer and F. Takens. Chapter 1 - preliminaries of dynamical systems theory. volume 3 of *Handbook of Dynamical Systems*, pages 1–42. Elsevier Science, 2010.
- J. W. Burby, Q. Tang, and R. Maulik. Fast neural Poincaré maps for toroidal magnetic fields. *Plasma Physics and Controlled Fusion*, 63(2):024001, 2020.
- François Chollet et al. Keras. <https://keras.io>, 2015.
- S. A. Desai, M. Mattheakis, D. Sondak, P. Protopapas, and S. J. Roberts. Port-Hamiltonian neural networks for learning explicit time-dependent dynamical systems. *Physical Review E*, 104(3):034312, 2021.
- L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- S. V. Gonchenko, A. S. Gonchenko, and A. O. Kazakov. Three types of attractors and mixed dynamics of nonholonomic models of rigid body motion. *Proceedings of the Steklov Institute of Mathematics*, 308(1):125–140, 2020.

- S. Greydanus, M. Dzamba, and J. Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- M. Hénon and C. Heiles. The applicability of the third integral of motion: some numerical experiments. *The astronomical journal*, 69:73–79, 1964.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- P. Jin, Z. Zhang, A. Zhu, Y. Tang, and G. Karniadakis. Sympnets: Intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems. *Neural Networks*, 132, 08 2020.
- J. S.W. Lamb and J. A.G. Roberts. Time-reversal symmetry in dynamical systems: a survey. *Physica D: Nonlinear Phenomena*, 112(1-2):1–39, 1998.
- T. Teshima, I. Ishikawa, K. Tojo, K. Oono, M. Ikeda, and M. Sugiyama. Coupling-based invertible neural networks are universal diffeomorphism approximators. *Advances in Neural Information Processing Systems*, 33:3362–3373, 2020.
- D. Turaev. Polynomial approximations of symplectic dynamics and richness of chaos in non-hyperbolic area-preserving maps. *Nonlinearity*, 16(1):123–135, 2002.
- D. Turaev. Maps close to identity and universal maps in the newhouse domain. *Communications in Mathematical Physics*, 335:1235–1277, 2015.