
Learning Dynamical Systems from Noisy Data with Inverse-Explicit Integrators

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We introduce the mean inverse integrator (MII), a novel approach to increase the
2 accuracy when training neural networks to approximate vector fields of dynamical
3 systems from noisy data. This method can be used to average multiple trajectories
4 obtained by numerical integrators such as Runge–Kutta methods. We show that the
5 class of mono-implicit Runge–Kutta methods (MIRK) has particular advantages
6 when used in connection with MII. When training vector field approximations,
7 explicit expressions for the loss functions are obtained when inserting the training
8 data in the MIRK formulae, unlocking symmetric and high order integrators that
9 would otherwise be implicit for initial value problems. The combined approach
10 of applying MIRK within MII yields a significantly lower error compared to the
11 plain use of the numerical integrator without averaging the trajectories. This is
12 demonstrated with experiments using data from several (chaotic) Hamiltonian
13 systems. Additionally, we perform a sensitivity analysis of the loss functions under
14 normally distributed perturbations, supporting the favourable performance of MII.

15 1 Introduction

16 Recently, many deep learning methodologies have been introduced to increase the efficiency and
17 quality of scientific computations [1, 2, 3, 4]. In physics-informed machine learning, deep neural
18 networks are purposely built so to enforce physical laws. As an example, Hamiltonian neural networks
19 (HNNs) [5] aim at learning the Hamiltonian function from temporal observations. The Hamiltonian
20 formalism was derived within classical mechanics for modelling a wide variety of physical systems.
21 The temporal evolution of such systems is fully determined when the Hamiltonian function is known,
22 and it is characterized by geometric properties such as the preservation of energy, the symplectic
23 structure and the time-reversal symmetry of the flow [6, 7].

24 Numerical integrators that compute solutions preserving such properties are studied in the field of
25 geometric numerical integration [7, 8]. Thus, deep learning, classical mechanics and geometric
26 numerical integration are all relevant to the development of HNNs. In this work, we try to identify
27 the optimal strategy for using numerical integrators when constructing loss functions for HNNs that
28 are trained on noisy and sparse data.

29 Generally, we aim at learning autonomous systems of first-order ordinary differential equations
30 (ODE)

$$\frac{d}{dt}y = f(y(t)), \quad y : [0, T] \rightarrow \mathbb{R}^n. \quad (1)$$

31 In the traditional setting, solving an initial value problem (IVP) means computing approximated
32 solutions $y_n \approx y(t_n)$ when the vector field $f(y)$ and an initial value $y(t_0) = y_0$ are known. The
33 focus of our study is the corresponding inverse problem; assuming knowledge of multiple noisy
34 samples of the solution, $S_N = \{\tilde{y}_n\}_{n=0}^N$, the aim is to approximate the vector field f with a neural

network model f_θ . We will assume that the observations originate from a (canonical) Hamiltonian system, with a Hamiltonian $H : \mathbb{R}^{2d} \rightarrow \mathbb{R}$, where the vector field is given by

$$f(y) = J \nabla H(y(t)), \quad J := \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix} \in \mathbb{R}^{2d \times 2d}. \quad (2)$$

This allows for learning the Hamiltonian function directly by setting $f_\theta(y) = J \nabla H_\theta(y)$, as proposed initially in [5].

Recently, many works highlight the benefit of using symplectic integrators when learning Hamiltonian neural networks [9, 10, 11, 12]. Here, we study what happens if, instead of using symplectic methods, efficient and higher-order MIRK methods are applied for inverse problems. We develop different approaches and apply them to learn highly oscillatory and chaotic dynamical systems from noisy data. The methods are general, they are not limited to separable Hamiltonian systems, and could indeed be used to learn any first-order ODE. However we focus our study on Hamiltonian systems, in order to build on the latest research on HNNs. Specifically, we compare our methods to the use of symplectic integrators to train Hamiltonian neural networks. Our contributions can be summarized as follows:

- We introduce the **mean inverse integrator** (MII), which efficiently averages trajectories of MIRK methods in order to increase accuracy when learning vector fields from noisy data (Definition 5.1).
- We present an **analysis of the sensitivity** of the loss function to perturbations giving insight into when the MII method yields improvement over a standard one-step scheme (Theorem 5.2).
- We show that **symplectic MIRK** methods have at most order $p = 2$ (Theorem 4.4). Particularly, the second-order implicit midpoint method is the symplectic MIRK method with minimal number of stages.

Finally, numerical experiments on several Hamiltonian systems benchmark MII against one-step training and symplectic recurrent neural networks (SRNN) [10], which rely on the Störmer–Verlet integrator. The structural difference between these three approaches is presented in Figure 2. Additionally, we demonstrate that substituting Störmer–Verlet with the classic Runge–Kutta method (RK4) in the SRNN framework yields significant reduction in error and allows accurate learning of non-separable Hamiltonian systems.

2 Related work

Hamiltonian neural networks was introduced in [5]. The numerical integration of Hamiltonian ODEs and the preservation of the symplectic structure of the ODE flow under numerical discretization have been widely studied over several decades [8, 7]. The symplecticity property is key and could inform the neural network architecture [13] or guide the choice of numerical integrator, yielding a theoretical guarantee that the learning target is actually a (modified) Hamiltonian vector field [14, 9], building on the backward error analysis framework [8]. Discrete gradients is an approach to numerical integration that guarantees exact preservation of the (learned) Hamiltonian, and an algorithm for training Hamiltonian neural networks using discrete gradient integrators is developed in [15] and extended to higher order in [16].

Since we for the inverse problem want to approximate the time-derivative of the solution, f , using only \tilde{y}_n , we need to use a numerical integrator when specifying the neural network loss function. For learning dynamical systems from data, explicit methods such as RK4 are much used [5, 17, 18]. However, explicit methods cannot in general preserve time-symmetry or symplecticity, and they often have worse stability properties compared to implicit methods [19]. Assuming that the underlying Hamiltonian is separable allows for explicit integration with the symplectic Störmer–Verlet method, which is exploited in [10, 20]. Symplecticity could be achieved without the limiting assumption of separability by training using the implicit midpoint method [12]. As pointed out in [12], this integrator could be turned into an explicit method in training by inserting sequential training data \tilde{y}_n and \tilde{y}_{n+1} . In fact, the MIRK class [21, 22] contains all Runge–Kutta (RK) methods (including the midpoint method) that could be turned into explicit schemes when inserting the training data. This is exploited in [23], where high-order MIRK methods are used to train HNNs, achieving accurate

84 interpolation and extrapolation of a single trajectory with large step size, few samples and assuming
85 zero noise.

86 The assumption of noise-free data limits the potential of learning from physical measurements
87 or applications on data sets from industry. This issue is addressed in [10], presenting symplectic
88 recurrent neural networks (SRNN). Here, Störmer–Verlet is used to integrate multiple steps and is
89 combined with initial state optimization (ISO) before computing the loss. ISO is applied after training
90 f_θ a given number of epochs and aims at finding the optimal initial value \hat{y}_0 , such that the distance
91 to the subsequent observed points $\tilde{y}_1, \dots, \tilde{y}_N$ is minimized when integrating over f_θ . While [10] is
92 limited by only considering separable systems, [24] aims at identifying the optimal combination of
93 third order polynomial basis functions to approximate a cubic non-separable Hamiltonian from noisy
94 data, using a Bayesian framework.

95 3 Background on numerical integration

96 Some necessary and fundamental concepts on numerical integration and the geometry of Hamiltonian
97 systems are presented below to inform the discussion on which integrators to use in inverse problems.
98 Further details could be found in Appendix C.

99 **Fundamental concepts:** An important subclass of the general first-order ODEs (1) is the class of
100 Hamiltonian systems, as given by (2). Often, the solution is partitioned into the coordinates $y(t) =$
101 $[q(t), p(t)]^T$, with $q(t), p(t) \in \mathbb{R}^d$. A separable Hamiltonian system is one where the Hamiltonian
102 could be written as the sum of two scalar functions, often representing the kinetic and potential
103 energy, that depend only on q and p respectively, this means we have $H(q, p) = H_1(q) + H_2(p)$.

104 The h flow of an ODE is a map $\varphi_{h,f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ sending an initial value $y(t_0)$ to the solution
105 of the ODE at time $t_0 + h$, given by $\varphi_{h,f}(y(t_0)) := y(t_0 + h)$. A numerical integration method
106 $\Phi_{h,f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a map approximating the exact flow of the ODE, so that

$$y(t_1) \approx y_1 = \Phi_{h,f}(y_0).$$

107 Here, $y(t_n)$ represents the exact solution and we denote with y_n the approximation at time $t_n =$
108 $t_0 + nh$. It should be noted that the flow map satisfies the following group property:

$$\varphi_{h_1,f} \circ \varphi_{h_2,f}(y(t_0)) = \varphi_{h_1,f}(y(t_0 + h_2)) = \varphi_{h_1+h_2,f}(y(t_0)). \quad (3)$$

109 In other words, a composition of two flows with step sizes h_1, h_2 is equivalent to the flow map over f
110 with step size $h_1 + h_2$. This property is not shared by numerical integrators for general vector fields.
111 The order of a numerical integrator $\Phi_{h,f}$ characterizes how the error after one step depends on the
112 step size h and is given by the integer p such that the following holds:

$$\|y_1 - y(t_0 + h)\| = \|\Phi_{h,f}(y_0) - \varphi_{h,f}(y(t_0))\| = \mathcal{O}(h^{p+1}).$$

113 **Mono-implicit Runge–Kutta methods:** Given vectors $b, v \in \mathbb{R}^s$ and a strictly lower triangular
114 matrix $D \in \mathbb{R}^{s \times s}$, a MIRK method is a Runge–Kutta method where $A = D + vb^T$ [25, 26] and we
115 assume that $[A]_{ij} = a_{ij}$ is the stage-coefficient matrix. This implies that the MIRK method can be
116 written on the form

$$y_{n+1} = y_n + h \sum_{i=1}^s b_i k_i, \quad (4)$$

$$k_i = f(y_n + v_i(y_{n+1} - y_n) + h \sum_{j=1}^s d_{ij} k_j).$$

117 Specific MIRK methods and further details on Runge–Kutta schemes is discussed in Appendix C.2.

118 **Symplectic methods:** The flow map of a Hamiltonian system is symplectic, meaning that its Jacobian
119 $\Upsilon_\varphi := \frac{\partial}{\partial y} \varphi_{h,f}(y)$ satisfies $\Upsilon_\varphi^T J \Upsilon_\varphi = J$, where J is the same matrix as in (2). As explained in [8, Ch.
120 VI.2], this is equivalent to the preservation of a projected area in the phase space of $[q, p]^T$. Similarly,
121 a numerical integrator is symplectic if its Jacobian $\Upsilon_\Phi := \frac{\partial}{\partial y_n} \Phi_{h,f}(y_n)$ satisfies $\Upsilon_\Phi^T J \Upsilon_\Phi = J$. It is
122 possible to prove [8, Ch. VI.4] that a Runge–Kutta method is symplectic if and only if the coefficients
123 satisfy

$$b_i a_{ij} + b_j a_{ji} - b_i b_j = 0, \quad i, j = 1, \dots, s. \quad (5)$$

124 4 Numerical integration schemes for solving inverse problems

125 We will now consider different ways to use numerical integrators when training Hamiltonian neural
 126 networks and present important properties of MIRK methods, a key component of the MII that is
 127 presented in Chapter 5.

128 **Inverse ODE problems in Hamiltonian form:** We assume to have potentially noisy samples
 129 $S_N = \{\tilde{y}\}_{n=0}^N$ of the solution of an ODE with vector field f . The inverse problem can be formulated
 130 as the following optimization problem:

$$\arg \min_{\theta} \sum_{n=0}^{N-1} \left\| \tilde{y}_{n+1} - \Phi_{h, f_{\theta}}(\tilde{y}_n) \right\|, \quad (6)$$

131 where $f_{\theta} = J\nabla H_{\theta}$ is a neural network approximation with parameters θ of a Hamil-
 132 tonian vector field f , and $\Phi_{h, f_{\theta}}$ is a one-step integration method with step length h .
 133 In the setting of inverse ODE problems, the availabil-
 134 ity of sequential points S_N could be exploited when
 135 a numerical method is used to form interpolation
 136 conditions, for $f_{\theta} \approx f$ for each n in the optimiza-
 137 tion problem (6). For example, \tilde{y}_n and \tilde{y}_{n+1} could
 138 be inserted in the implicit midpoint method, turning
 139 a method that is implicit for IVPs into an explicit
 140 method for inverse problems:

$$\Phi_{h, f_{\theta}}(\tilde{y}_n, \tilde{y}_{n+1}) = \tilde{y}_n + h f_{\theta}\left(\frac{\tilde{y}_n + \tilde{y}_{n+1}}{2}\right). \quad (7)$$

141 We denote this as the inverse injection, which defines
 142 an inverse explicit property for numerical integrators.

Definition 4.1 (Inverse injection). Assume that
 $\tilde{y}_n, \tilde{y}_{n+1} \in S_N$. Let the *inverse injection* for the
 integrator $\Phi_{h, f}(y_n, y_{n+1})$ be given by the substitu-
 tion $(\tilde{y}_n, \tilde{y}_{n+1}) \rightarrow (y_n, y_{n+1})$ such that

$$\hat{y}_{n+1} = \Phi_{h, f}(\tilde{y}_n, \tilde{y}_{n+1}).$$

143 **Definition 4.2** (Inverse explicit). A numerical one-step method Φ is called *inverse explicit* if it is
 144 explicit under the inverse injection.

145 This procedure is utilized successfully by several authors when learning dynamical systems from
 146 data, see e.g. [12, 27]. However, this work is the first attempt at systematically exploring numerical
 147 integrators under the inverse injection, by identifying the MIRK methods as the class consisting of
 148 inverse explicit Runge–Kutta methods.

149 **Proposition 4.3.** *MIRK-methods are inverse explicit.*

150 *Proof.* Since the matrix D in (4) is strictly lower triangular, the stages are given by

$$\begin{aligned} k_1 &= f(y_n + v_i(y_{n+1} - y_n)) \\ k_2 &= f(y_n + v_i(y_{n+1} - y_n) + h d_{21} k_1) \\ &\vdots \\ k_s &= f(y_n + v_i(y_{n+1} - y_n) + h \sum_{j=1}^{s-1} d_{sj} k_j) \end{aligned}$$

151 meaning that if y_n and y_{n+1} are known, all stages, and thus the next step $\hat{y}_{n+1} = y_n + h \sum_{i=1}^s b_i k_i$,
 152 could be computed explicitly. \square

153 Because of their explicit nature when applied to inverse ODE problems, MIRK methods are an
 154 attractive alternative to explicit Runge–Kutta methods; in contrast to explicit RK methods, they
 155 can be symplectic or symmetric, or both, without requiring the solution of systems of nonlinear




Figure 1: Venn diagram of Runge–Kutta (RK) subclasses: explicit RK (ERK), symplectic RK (SympRK), mono-implicit RK (MIRK) and symmetric RK (SymRK).




Figure 2: Differences of observation dependency, assuming $N = 2$ for explicit and mono-implicit one-step training, and explicit multi-step training with initial state optimization (green node \hat{y}_0).

equations, even when the Hamiltonian is non-separable. Figure 1 illustrates the relation between various subclasses and the specific methods are described in Table 1 in Appendix C. In addition, for s -stage MIRK methods, it is possible to construct methods of order $p = s + 1$ [22]. This is in general higher order than what is possible to obtain with s -stage explicit Runge–Kutta methods. Further computational gains could also be made by reusing evaluations of the vector field between multiple steps, which using MIRK methods allow for, as explained in Appendix I. The dependency structure on the data S_N of explicit RK (ERK) methods, MIRK methods and the SRNN method [10] is illustrated in Figure 2.

Maximal order of symplectic MIRK methods: From the preceding discussion, it is clear that symplectic MIRK methods are of interest when learning Hamiltonian systems from data, since they combine computational efficiency with the ability to preserve useful, geometric properties. Indeed, symplectic integrators in the training of HNNs have been considered in [9, 10, 11, 12, 13]. The subclass of symplectic MIRK methods is represented by the middle, dark blue field in the Venn diagram of Figure 1. The next result gives an order barrier for symplectic MIRK methods that was, to the best of our knowledge, not known up to this point.

Theorem 4.4. *The maximum order of a symplectic MIRK method is $p = 2$.*

Proof. This is a shortened version of the full proof, which can be found in Appendix F. A MIRK method is a Runge–Kutta method with coefficients $a_{ij} = d_{ij} + v_i b_j$. Requiring d_{ij}, b_i and v_i to satisfy the symplecticity conditions of (5) in addition to D being strictly lower triangular, yields the following restrictions

$$\begin{aligned} b_i d_{ij} + b_i b_j (v_j + v_i - 1) &= 0, & \text{if } i \neq j, \\ b_i &= 0 \text{ or } v_i = \frac{1}{2}, & \text{if } i = j, \\ d_{ij} &= 0, & \text{if } i > j. \end{aligned} \tag{8}$$

These restrictions result in an RK method that could be reduced to choosing a coefficient vector $b \in \mathbb{R}^s$ and choosing stages on the form $k_i = f(y_n + \frac{h}{2} \sum_j b_j k_j)$ for $i = 1, \dots, s$. It is then trivial to check that this method can only be of up to order $p = 2$. Note that for $s = 1$ and $b_1 = 1$ we get the midpoint method. \square

Numerical integrators outside the RK class: While this paper is mainly concerned with MIRK methods, several other types of numerical integrators could be of interest for inverse problems. *Partitioned Runge–Kutta methods* are an extension and not a subclass of RK methods, and can be symplectic and symmetric, while also being explicit for separable Hamiltonian systems. The Störmer–Verlet integrator of order $p = 2$ is one example. Higher order methods of this type are derived in [28] and used for learning Hamiltonian systems in [29, 30]. *Discrete gradient methods* [31, 32] are inverse explicit and well suited to train Hamiltonian neural networks using a modified automatic differentiation algorithm [15]. This method could be extended to higher order methods as shown in [16]. In contrast to symplectic methods, discrete gradient methods preserve the Hamiltonian exactly up to machine precision. A third option is *elementary differential Runge–Kutta methods* [33], where for instance [34] show how to use backward error analysis to construct higher order methods from modifications to the midpoint method. This topic is discussed further in Appendix H, where we also present a novel, symmetric discrete gradient method of order $p = 4$.

5 Mean inverse integrator for handling noisy data

Noisy ODE sample: It is often the case that the samples S_N are not exact measurements of the system, but perturbed by noise. In this paper, we model the noise as independent, normally distributed

196 perturbations

$$\tilde{y}_n = y(t_n) + \delta_n, \quad \delta_n \sim \mathcal{N}(0, \sigma^2 I), \quad (9)$$

197 where $\mathcal{N}(0, \sigma^2 I)$ represents the multivariate normal distribution. With this assumption, a standard
 198 result from statistics tells us that the variance of a sample-mean estimator with N samples converges
 199 to zero at the rate of $\frac{1}{N}$. That is, assuming that we have N samples $\tilde{y}_n^{(1)}, \dots, \tilde{y}_n^{(N)}$, then

$$\text{Var}[\bar{y}_n] = \text{Var}\left[\frac{1}{N} \sum_{j=1}^N \tilde{y}_n^{(j)}\right] = \frac{\sigma^2}{N}.$$

200 Using the inverse injection with the midpoint method, the vector field is evaluated in the average of
 201 \tilde{y}_n and \tilde{y}_{n+1} , reducing the variance of the perturbation by a factor of two, compared to evaluating the
 202 vector field in \tilde{y}_n , as is done in all explicit RK methods. Furthermore, considering the whole data
 203 trajectory S_N , multiple independent approximations to the same point $y(t_n)$ can enable an even more
 204 accurate estimate. This is demonstrated in the analysis presented in Theorem 5.2 and in Figure 4.

205 **Averaging multiple trajectories:** In the inverse ODE problem, we assume that there exists an *exact*
 206 vector field f whose flow interpolates the discrete trajectories S_N , and the flow of this vector field
 207 satisfies the group property (3). The numerical flow $\Phi_{h,f}$ for a method of order p satisfies this
 208 property only up to an error $\mathcal{O}(h^{p+1})$ over one step. In the presence of noisy data, compositions of
 209 one-step methods can be used to obtain multiple different approximations to the same point $y(t_n)$,
 210 by following the numerical flow from different nearby initial values $\tilde{y}_j, j \neq n$, and thus reduce the
 211 noise by averaging over these multiple approximations. Accumulation of the local truncation error is
 212 expected when relying on points further away from t_n . However, for sufficiently small step sizes h
 213 compared to the size of the noise σ , one can expect increased accuracy when averaging over multiple
 214 noisy samples.

215 As an example, assume that we know the points $\{\tilde{y}_0, \tilde{y}_1, \tilde{y}_2, \tilde{y}_3\}$. Then $y(t_2)$ can be approximated by
 216 computing the mean of the numerical flows $\Phi_{h,f}$ starting from different initial values:

$$\begin{aligned} \bar{y}_2 &= \frac{1}{3} (\Phi_{h,f}(\tilde{y}_1) + \Phi_{h,f} \circ \Phi_{h,f}(\tilde{y}_0) + \Phi_{-h,f}^*(\tilde{y}_3)) \\ &\approx \frac{1}{3} (\tilde{y}_0 + \tilde{y}_1 + \tilde{y}_3 + h(\Psi_{0,1} + 2\Psi_{1,2} - \Psi_{2,3})), \end{aligned} \quad (10)$$

217 where we by Φ^* mean the adjoint method of Φ , as defined in [8, Ch. V], and we let $\Psi_{n,n+1}$ be the
 218 increment of an inverse-explicit numerical integrator, so that

$$\Phi_{h,f}(\tilde{y}_n, \tilde{y}_{n+1}) = \tilde{y}_n + h\Psi_{n,n+1}.$$

219 For example, for the midpoint method, we have that $\Psi_{n,n+1} = f(\frac{\tilde{y}_n + \tilde{y}_{n+1}}{2})$. When stepping in
 220 negative time in (10), we use the adjoint method in order to minimize the number of vector field
 221 evaluations, also when non-symmetric methods are used (which implies that we always use e.g. $\Psi_{1,2}$
 222 and not $\Psi_{2,1}$). Note that in order to derive the approximation in (10), repeated use of the inverse
 223 injection allows the known points \tilde{y}_n to form an explicit integration procedure, where composition
 224 of integration steps are approximated by summation over increments $\Psi_{n,n+1}$. This approximation
 225 procedure is presented in greater detail in Appendix D.

226 **Mean inverse integrator:** The mean approximation over the whole trajectory \bar{y}_n , for $n = 0, \dots, N$,
 227 could be computed simultaneously, reusing multiple vector field evaluations in an efficient manner.
 228 This leads to what we call the mean inverse integrator. For example, when $N = 3$ we get

$$\begin{bmatrix} \bar{y}_0 \\ \bar{y}_1 \\ \bar{y}_2 \\ \bar{y}_3 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} \tilde{y}_0 \\ \tilde{y}_1 \\ \tilde{y}_2 \\ \tilde{y}_3 \end{bmatrix} + \frac{h}{3} \begin{bmatrix} -3 & -2 & -1 \\ 1 & -2 & -1 \\ 1 & 2 & -1 \\ 1 & 2 & 3 \end{bmatrix} \begin{bmatrix} \Psi_{0,1} \\ \Psi_{1,2} \\ \Psi_{2,3} \end{bmatrix},$$

229 and the same structure is illustrated in Figure 3.

230 **Definition 5.1** (Mean inverse integrator). For a sample S_N and an inverse-explicit integrator $\Psi_{n,n+1}$,
 231 the mean inverse integrator is given by

$$\bar{Y} = \frac{1}{N} \left(U\tilde{Y} + hW\Psi \right) \quad (11)$$

where $\tilde{Y} := [\tilde{y}_0, \dots, \tilde{y}_N]^T \in \mathbb{R}^{(N+1) \times m}$, $\Psi := [\Psi_{0,1}, \dots, \Psi_{N-1,N}]^T \in \mathbb{R}^{N \times m}$.

Finally, $U \in \mathbb{R}^{(N+1) \times (N+1)}$ and $W \in \mathbb{R}^{(N+1) \times N}$ are given by

$$[U]_{ij} := \begin{cases} 0 & \text{if } i = j \\ 1 & \text{else} \end{cases} \quad \text{and} \quad [W]_{ij} := \begin{cases} j-1-N & \text{if } j \geq i \\ j & \text{else} \end{cases}.$$

By substituting the known vector field f with a neural network f_θ and denoting the matrix containing vector field evaluations by Ψ_θ such that $\bar{Y}_\theta := \frac{1}{N}(U\tilde{Y} + hW\Psi_\theta)$, we can formulate an analogue to the inverse problem (6) by

$$\arg \min_{\theta} \|\tilde{Y} - \bar{Y}_\theta\|. \quad (12)$$

Analysis of sensitivity to noise: Consider the optimization problems using integrators either as one-step methods or MII by (6) resp. (12). We want to investigate how uncertainty in the data \tilde{y}_n introduces uncertainty in the optimization problem. Assume, for the purpose of analysis, that the underlying vector field $f(y)$ is known. Let

$$\begin{aligned} \mathcal{T}_n^{\text{OS}} &:= \tilde{y}_n - \Phi_{h,f}(\tilde{y}_{n-1}, \tilde{y}_n), \\ \mathcal{T}_n^{\text{MII}} &:= \tilde{y}_n - [\bar{Y}]_n \end{aligned}$$

be the *optimization target* or the expression one aims to minimize using a one-step method (OS) and the MII, where \bar{Y} is given by Definition 5.1. For a matrix A with eigenvalues $\lambda_i(A)$, the spectral radius is given by $\rho(A) := \max_i |\lambda_i(A)|$. An analytic expression that approximates $\rho(\mathcal{T}_n^{\text{OS}})$ and $\rho(\mathcal{T}_n^{\text{MII}})$ by linearization of f for a general MIRK method is provided below.

Theorem 5.2. Let $S_N = \{\tilde{y}_n\}_{n=0}^N$ be a set of noisy samples, equidistant in time with step size h , with Gaussian perturbations as defined by (9) with variance σ^2 . Assume that a MIRK integrator $\Phi_{h,f}$ is used as a one-step method. Then the spectral radius is approximated by

$$\rho_n^{\text{OS}} := \rho\left(\text{Var}[\mathcal{T}_n^{\text{OS}}]\right) \approx \sigma^2 \left\| 2I + hb^T(\mathbb{1} - 2v)(f' + f'^T) + h^2 Q^{\text{OS}} \right\|_2, \quad (13)$$

$$\rho_n^{\text{MII}} := \rho\left(\text{Var}[\mathcal{T}_n^{\text{MII}}]\right) \approx \frac{\sigma^2}{N} \left\| (1+N)I + hP_{nn} + \frac{h}{N} \sum_{\substack{j=0 \\ j \neq n}}^s P_{nj} + \frac{h^2}{N} Q^{\text{MII}} \right\|_2, \quad (14)$$

where $f' := f'(y_n)$ and P_{nj} , Q^{OS} and Q^{MII} (defined in (24) in Appendix G) are matrices independent of the step size h .

The proof is found in Appendix G. Let $\alpha := b^T(\mathbb{1} - 2v)$ denote the coefficients of the first order term in h of Equation (13). For any explicit RK method we have that $v = 0$ and since $b^T \mathbb{1} = 1$ (method of at least order one) we find that $\alpha_{\text{ERK}} = 1$. Considering the Butcher tableau of MIRK4 in Figure 9 we find that $\alpha_{\text{MIRK4}} = 0$. Thus, as $h \rightarrow 0$ we would expect quadratic convergence of MIRK4 and linear convergence of RK4 for ρ_n^{OS} to $2\sigma^2$. Considering MII (14) one would expect linear convergence for ρ_n^{MII} to σ^2 if N is large, as $h \rightarrow 0$.

A numerical approximation of ρ_n^{OS} and ρ_n^{MII} could be realized by a Monte-Carlo estimate. We compute the spectral radius $\hat{\rho}_n$ of the empirical covariance matrix of $\mathcal{T}_n^{\text{OS}}$ and $\mathcal{T}_n^{\text{MII}}$ by sampling $5 \cdot 10^3$ normally distributed perturbations δ_n with $\sigma^2 = 2.5 \cdot 10^{-3}$ to each point y_n in a trajectory of $N+1$ points and step size h . We then compute the




Figure 3: Illustration of the structure of the mean inverse integrator for $N = 3$.




Figure 4: Average of $\bar{\rho}$ over 10 trajectories. Shaded area represent one standard deviation.

trajectory average $\bar{\rho} = \frac{1}{N+1} \sum_{n=0}^N \hat{\rho}_n$, fix the end time $T = 2.4$, repeat the approximations for decreasing step sizes h and increasing N and compute the average of $\bar{\rho}$ for 10 randomly sampled trajectories S_N from the double pendulum system. The plot in Figure 4 corresponds well with what one would expect from Theorem 5.2 and confirms that first MIRK (with $v \neq 0$) and secondly MII reduces the sensitivity to noise in the optimization target.

6 Experiments

Methods and test problems: We train HNNs using different integrators and methods in the inverse problem (6). We use MIRK4 together with the MII method and compare to the implicit midpoint method, RK4 and MIRK4 applied as one-step methods, as well as ISO followed by Störmer–Verlet and RK4 integrated over multiple time-steps. The latter strategy, illustrated in Figure 2, was suggested in [10], where Störmer–Verlet is used. Separable networks $H_\theta(q, p) = H_{1,\theta}(q) + H_{2,\theta}(p)$ are trained on data from the Fermi–Pasta–Ulam–Tsingou (FPUT) problem and the Hénon–Heiles system. For the double pendulum, which is non-separable, a fully connected network is used for all methods except Störmer–Verlet, which requires separability in order to be explicit. The Hamiltonians are described in Appendix A and all systems have solutions $y(t) \in \mathbb{R}^4$.

After using the specified integrators in training, approximated solutions are computed for each learned vector field f_θ using the Scikit-learn implementation of DOP853 [35], which is also used to generate the training data. The error is averaged over $M = 10$ points and we find what we call the flow error by

$$e(f_\theta) = \frac{1}{M} \sum_{n=1}^M \|\hat{y}_n - y(t_n)\|_2, \quad y(t_n) \in S_M^{\text{test}}, \quad (15)$$

$$\hat{y}_{n+1} = \Phi_{h,f_\theta}(y_n).$$

Training data: Training data is generated by sampling $N_2 = 300$ random initial values y_0 requiring that $0.3 \leq \|y_0\|_2 \leq 0.6$. The data $S_{N_1, N_2} = \{\hat{y}_n^{(j)}\}_{n=0, j=0}^{N_1, N_2}$ is found by integrating the initial values with DOP853 with a tolerance of 10^{-15} for the following step sizes and number of steps: $(h, N_1) = (0.4, 4), (0.2, 8), (0.1, 16)$. The points in the flow are perturbed by noise where $\sigma \in \{0, 0.05\}$. Error is measured in $M = 10$ random points in the flow, within the same domain as the initial values. Furthermore, experiments are repeated with a new random seed for the generation of data and initialization of neural network parameters five times in order to compute the standard deviation of the flow error. The flow error is shown in Figure 6. Additional results are presented in Appendix B.

Neural network architecture and optimization: For all test problems, the neural networks have 3 layers with a width of 200 neurons and $\tanh(\cdot)$ as the activation function. The algorithms are implemented using PyTorch [36] and the code for performing ISO is a modification of the implementation by [10]¹. Training is done using the quasi-Newton L-BFGS algorithm [37] for 20 epochs without batching. This optimization algorithm is often used to train physics-informed neural networks [1] and in this setting it proved to yield superior results in comparison to the often used Adam optimizer. Further details are provided in Appendix E.

Results: As observed in Figure 6 and supported by the analytical result illustrated in Figure 4, the MII approach facilitates more accurate training from from noisy data than one-step methods. However, training with multiple integration steps in combination with ISO yields lower error when RK4 is used




Figure 5: Roll-out in time obtained by integrating over the learned vector fields when training on data from the double pendulum Hamiltonian.

¹<https://github.com/zhengdao-chen/SRNN> (CC-BY-NC 4.0 License)

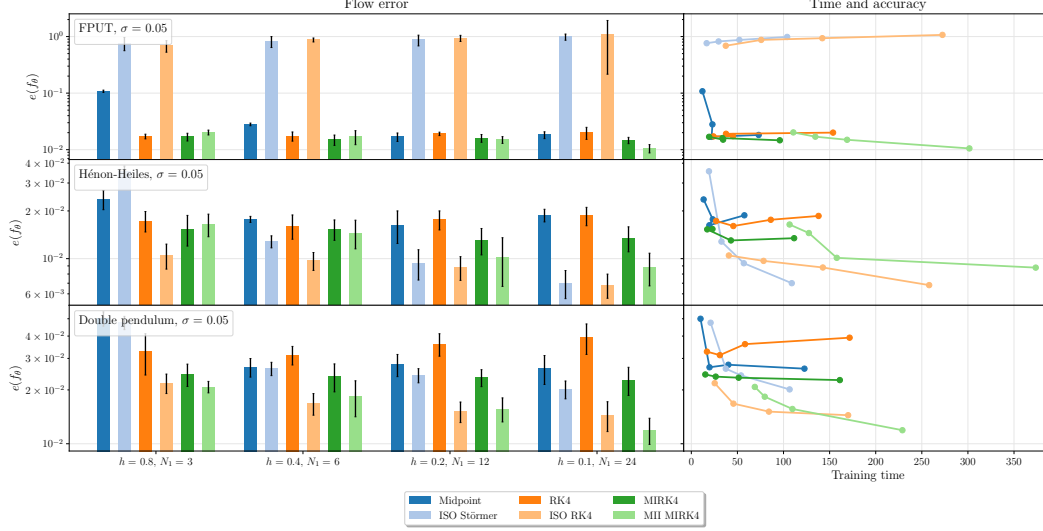


Figure 6: The flow error when learning vector fields using one-step methods directly (Midpoint, RK4 and MIRK4), ISO and multiple time-steps (ISO Störmer and ISO RK4) and MII (MII MIRK4). The error bars display the standard deviation after rerunning 5 experiments on data with $\sigma = 0.05$. The right subplot shows the computational time used in training against the flow error.

for the Hénon–Heiles problem and similar performance as MII on the double pendulum. We notice that the SRNN approach, i.e. ISO with Störmer–Verlet, is improved when switching to RK4, which means sacrificing symplecticity to achieve higher order. The results for FPUT stand out in Figure 6, since both ISO methods have large errors here. The roll-out in time of the learned vector fields is presented in Figure 8 in Appendix B, where the same can be observed. As also could be seen here, the FPUT Hamiltonian gives rise to highly oscillatory trajectories, and the errors observed in Figure 6 might indicate that ISO is ill-suited for this kind of dynamical systems.

Two observations could be made regarding the one-step methods without averaging or ISO. First, it is likely that the midpoint method has weaker performance for large step sizes due to its lower order, compared to both RK4 and MIRK4, despite the fact that it is a symplectic method. The same is clear from Figure 7 in Appendix B, which display the flow error when training on data without noise. Secondly, building on the sensitivity analysis, we observe that MIRK4 consistently attains higher accuracy than RK4, as expected from the Monte-Carlo simulation found in Figure 4.

7 Conclusion

In this work we present the mean inverse integrator, which allows both chaotic and oscillatory dynamical systems to be learned with high accuracy from noisy data. Within this method, integrators of the MIREK class are a key component. To analyse how noise is propagated when training with MII and MIREK, compared to much used explicit methods such as RK4, we developed a sensitivity analysis that is verified both by a Monte-Carlo approximation and reflected in the error of the learned vector fields. Finally, we build on the SRNN [10] by replacing Störmer–Verlet with RK4, and observer increased performance. When also considering the weak performance of the implicit midpoint method, this tells us that order might be of greater importance than preserving the symplectic structure when training HNNs. Both the MIREK methods, the mean inverse integrator and initial state optimization form building blocks that could be combined to form novel approaches for solving inverse problems and learning from noisy data.

Limitations: The experiments presented here assume that both the generalized coordinates q_n and the generalized momenta p_n could be observed. In a setting where HNNs are to model real and not simulated data, the observations might lack generalized momenta [38] or follow Cartesian coordinates, requiring the enforcement of constraints [17, 39]. Combining approaches that are suitable for data that is both noisy and follow less trivial coordinate systems is a subject for future research.

References

- [1] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. 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, 2019.
- [2] Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal differential equations for scientific machine learning. Aug 2020.
- [3] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [4] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- [5] Sam Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *CoRR*, abs/1906.01563, 2019.
- [6] Herbert Goldstein, Charles Poole, and John Safko. *Classical Mechanics*. Addison Wesley, 3 edition, 2001.
- [7] Benedict Leimkuhler and Sebastian Reich. *Simulating Hamiltonian Dynamics*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2005.
- [8] Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations; 2nd ed.* Springer, Dordrecht, 2006.
- [9] Christian Offen and Sina Ober-Blöbaum. Symplectic integration of learned Hamiltonian systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 32(1):013122, 2022.
- [10] Zhengdao Chen, Jianyu Zhang, Martin Arjovsky, and Léon Bottou. Symplectic recurrent neural networks. In *International Conference on Learning Representations*, 2020.
- [11] Aiqing Zhu, Pengzhan Jin, and Yifa Tang. Deep Hamiltonian networks based on symplectic integrators. *arXiv preprint arXiv:2004.13830*, 2020.
- [12] Marco David and Florian Méhats. Symplectic learning for Hamiltonian neural networks. *arXiv preprint arXiv:2106.11753*, 2021.
- [13] Pengzhan Jin, Zhen Zhang, Aiqing Zhu, Yifa Tang, and George Em Karniadakis. SympNets: Intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems. *Neural Networks*, 132:166–179, 2020.
- [14] Aiqing Zhu, Pengzhan Jin, Beibei Zhu, and Yifa Tang. Inverse modified differential equations for discovery of dynamics. *arXiv preprint arXiv:2009.01058*, 2020.
- [15] Takashi Matsubara, Ai Ishikawa, and Takaharu Yaguchi. Deep energy-based modeling of discrete-time physics. *Advances in Neural Information Processing Systems*, 33:13100–13111, 2020.
- [16] Sølve Eidnes. Order theory for discrete gradient methods. *BIT*, 62(4):1207–1255, 2022.
- [17] Elena Celledoni, Andrea Leone, Davide Murari, and Brynjulf Owren. Learning Hamiltonians of constrained mechanical systems. *J. Comput. Appl. Math.*, 417:Paper No. 114608, 12, 2023.
- [18] Alvaro Sanchez-Gonzalez, Victor Bapst, Kyle Cranmer, and Peter Battaglia. Hamiltonian Graph Networks with ODE Integrators.
- [19] Gerhard Wanner and Ernst Hairer. *Solving ordinary differential equations II*, volume 375. Springer Berlin Heidelberg, 1996.
- [20] Senwei Liang, Zhongzhan Huang, and Hong Zhang. Stiffness-aware neural network for learning Hamiltonian systems. 2022.

- [21] Jeff R Cash. A class of implicit Runge–Kutta methods for the numerical integration of stiff ordinary differential equations. *Journal of the ACM (JACM)*, 22(4):504–511, 1975.
- [22] K Burrage, FH Chipman, and Paul H Muir. Order results for mono-implicit Runge–Kutta methods. *SIAM journal on numerical analysis*, 31(3):876–891, 1994.
- [23] Håkon Noren. Learning Hamiltonian systems with mono-implicit Runge–Kutta methods. *arXiv preprint, arXiv:2303.03769*, 2023.
- [24] Harsh Sharma, Nicholas Galioto, Alex A Gorodetsky, and Boris Kramer. Bayesian identification of nonseparable Hamiltonian systems using stochastic dynamic models. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 6742–6749. IEEE, 2022.
- [25] W. M. G. van Bokhoven. Efficient higher order implicit one-step methods for integration of stiff differential equations. *BIT*, 20(1):34–43, 1980.
- [26] J. R. Cash and A. Singhal. Mono-implicit Runge–Kutta formulae for the numerical integration of stiff differential systems. *IMA J. Numer. Anal.*, 2(2):211–227, 1982.
- [27] Sølve Eidnes, Alexander J Stasik, Camilla Sterud, Eivind Bøhn, and Signe Riemer-Sørensen. Pseudo-Hamiltonian neural networks with state-dependent external forces. *arXiv preprint, arXiv:2206.02660*, 2022.
- [28] Haruo Yoshida. Construction of higher order symplectic integrators. *Physics letters A*, 150(5-7):262–268, 1990.
- [29] Shaan A Desai, Marios Mattheakis, and Stephen J Roberts. Variational integrator graph networks for learning energy-conserving dynamical systems. *Physical Review E*, 104(3):035310, 2021.
- [30] Daniel DiPietro, Shiyong Xiong, and Bo Zhu. Sparse symplectically integrated neural networks. *Advances in Neural Information Processing Systems*, 33:6074–6085, 2020.
- [31] GRW Quispel and Grant S Turner. Discrete gradient methods for solving ODEs numerically while preserving a first integral. *Journal of Physics A: Mathematical and General*, 29(13):L341, 1996.
- [32] Robert I McLachlan, G Reinout W Quispel, and Nicolas Robidoux. Geometric integration using discrete gradients. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 357(1754):1021–1045, 1999.
- [33] Ander Murua. *Métodos simplécticos desarrollables en P-series*. PhD thesis, PhD thesis. Valladolid: Universidad de Valladolid, 1995.
- [34] Philippe Chartier, Ernst Hairer, and Gilles Vilmart. Numerical integrators based on modified differential equations. *Mathematics of Computation*, 76(260):1941–1953, October 2007.
- [35] J.R. Dormand and P.J. Prince. A family of embedded Runge–Kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19–26, 1980.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- [37] Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999.
- [38] Yuhan Chen, Takashi Matsubara, and Takaharu Yaguchi. Neural symplectic form: learning hamiltonian equations on general coordinate systems. *Advances in Neural Information Processing Systems*, 34:16659–16670, 2021.
- [39] Marc Finzi, Ke Alexander Wang, and Andrew Gordon Wilson. Simplifying Hamiltonian and Lagrangian neural networks via explicit constraints. *arXiv preprint arXiv:2010.13581*, 2020.
- [40] Enrico Fermi, P Pasta, Stanislaw Ulam, and Mary Tsingou. Studies of the nonlinear problems. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 1955.

- [41] E. Hairer, S. P. Nørsett, and G. Wanner. *Solving ordinary differential equations. I*, volume 8 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, second edition, 1993. Nonstiff problems.
- [42] P. H. Muir. Optimal discrete and continuous mono-implicit Runge-Kutta schemes for BVODEs. *Adv. Comput. Math.*, 10(2):135–167, 1999.
- [43] J. R. Cash and D. R. Moore. A high order method for the numerical solution of two-point boundary value problems. *BIT*, 20(1):44–52, 1980.
- [44] Philippe Chartier. Symmetric Methods. In Björn Engquist, editor, *Encyclopedia of Applied and Computational Mathematics*, pages 1439–1448. Springer, Berlin, Heidelberg, 2015.
- [45] J. R. Cash and A. Singhal. High order methods for the numerical solution of two-point boundary value problems. *BIT*, 22(2):184–199, 1982.
- [46] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [47] Philippe Chartier, Ernst Hairer, and Gilles Vilmart. Numerical integrators based on modified differential equations. *Math. Comp.*, 76(260):1941–1953, 2007.
- [48] Ge Zhong and Jerrold E. Marsden. Lie-Poisson Hamilton-Jacobi theory and Lie-Poisson integrators. *Phys. Lett. A*, 133(3):134–139, 1988.
- [49] Takashi Matsubara and Takaharu Yaguchi. FINDE: Neural differential equations for finding and preserving invariant quantities. *arXiv preprint, arXiv:2210.00272*, 2022.

A Test problems

Fermi–Pasta–Ulam–Tsingou: This dynamical system is a model for a chain of $2m + 1$ alternating stiff and soft springs connecting $2m$ mass points. The chain is fixed in both ends [8, 40]. With the coordinate transformation suggested in [8, Ch. I.5.I] we have coordinates $[q, p]^T \in \mathbb{R}^{4m}$ where $q_i, i = 1, \dots, m$ represents a scaled displacement of the i -th stiff spring and $q_{i+m}, i = 1, \dots, m$ represents a scaled expansion of the i -th spring. q_i represents their velocities. Letting ω be the angular frequency of the stiff spring, in general the Hamiltonian is given by

$$H(q, p) = \frac{1}{2} \sum_{i=1}^m (p_i^2 + p_{i+m}^2) + \frac{\omega^2}{2} \sum_{i=1}^m q_{i+m}^2 + \frac{1}{4} \left(\sum_{i=1}^{m-1} (q_{i+1} - q_{i+m+1} - q_i - q_{i+m})^4 + (q_1 - q_{m+1})^4 + (q_m + q_{2m})^4 \right)$$

We consider the most trivial case of $m = 1$ and letting $\omega = 2$, yielding the quartic, separable Hamiltonian by

$$H(q_1, q_2, p_1, p_2) = \frac{1}{2} (p_1^2 + p_2^2) + 2q_2^2 + \frac{1}{4} \left((q_1 - q_2)^4 + (q_1 + q_2)^4 \right).$$

Double pendulum: Let q_i and p_i denote the angle and angular momentum of pendulum $i = 1, 2$. The double pendulum system has a Hamiltonian that is not separable, where $y = [q_1, q_2, p_1, p_2]^T \in \mathbb{R}^4$ and the Hamiltonian is given by

$$H(q_1, q_2, p_1, p_2) = \frac{\frac{1}{2} p_1^2 + p_2^2 - p_1 p_2 \cos(q_1 - q_2)}{1 + \sin^2(q_1 - q_2)} - 2 \cos(q_1) - \cos(q_2).$$

476 **Hénon–Heiles:** This model was introduced for describing stellar motion inside the gravitational
 477 potential of a galaxy, as described in [8]. This Hamiltonian is separable. However, it is a canonical
 478 example of a chaotic system and its properties are discussed more in detail in [6]. The Hamiltonian is
 479 given by

$$H(q_1, q_2, p_1, p_2) = \frac{1}{2}(p_1^2 + p_2^2) + \frac{1}{2}(q_1^2 + q_2^2) + q_1^2 q_2 - \frac{1}{3}q_2^3.$$

480 B Additional numerical results

481 Here we present additional numerical experiments. In Figure 7, the flow error when learning from
 482 data without noise, could be found. The roll-out in time of the learned Hamiltonian for the FPUT and
 483 Hénon–Heiles problem is presented in Figure 8.

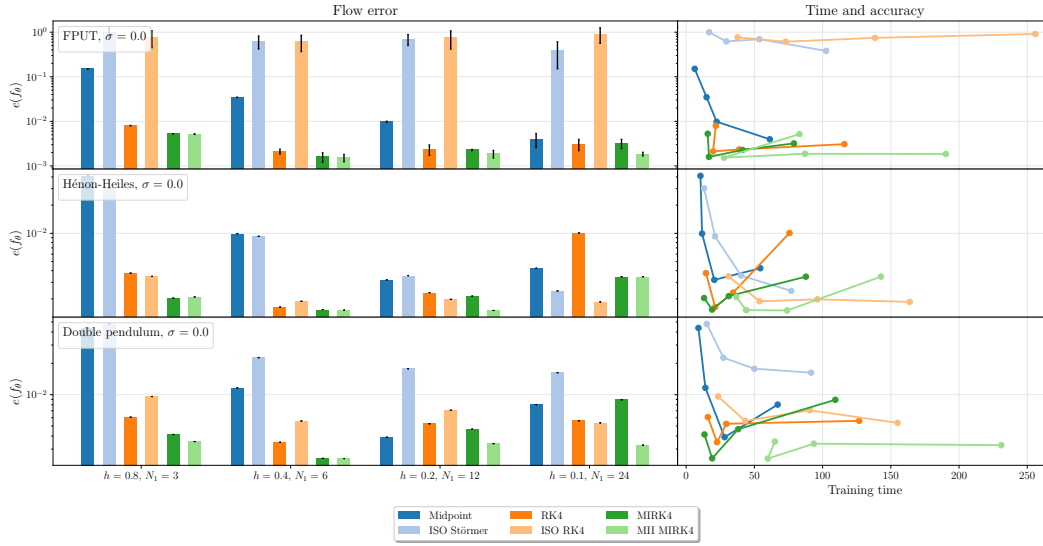


Figure 7: The flow error when learning vector fields using one-step methods directly (Midpoint, RK4 and MIRK4), ISO and multiple time-steps (ISO Störmer and ISO RK4) and MII (MII MIRK4). The error bars display the standard deviation after rerunning 5 experiments on data with $\sigma = 0$. The right subplot shows the computational time used in training against the flow error.

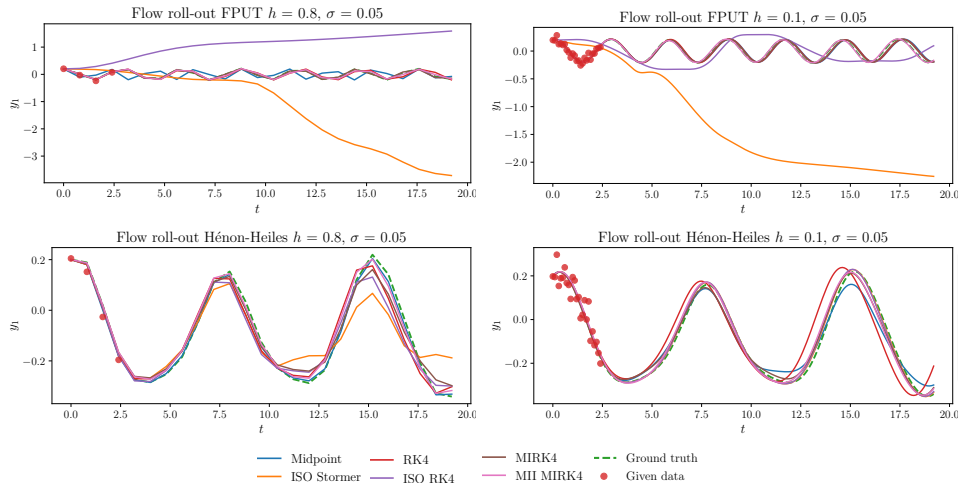


Figure 8: Roll-out in time obtained by integrating over the learned vector fields when training on data from the Fermi–Pasta–Ulam–Tsingou and Hénon–Heiles Hamiltonian.

484 C More on numerical integration

485 C.1 Runge–Kutta methods

486 A general Runge–Kutta method for an autonomus system with s stages is a one-step numerical
487 integrator given by

$$y_{n+1} = y_n + h \sum_{j=1}^s b_j k_j, \quad (16)$$

$$k_i = f\left(y_n + h \sum_{j=1}^s a_{ij} k_j\right), \quad i = 1, \dots, s.$$

488 A concrete method is determined by specifying the coefficient matrix $A \in \mathbb{R}^{s \times s}$ and the vector
489 $b \in \mathbb{R}^s$, and there are conditions for symplecticity and order associated with these [41]. The
490 conditions for order $p = 1$ require that the coefficient $c \in \mathbb{R}^s$ is determined by $c_i = \sum_{j=1}^s a_{ij}$. A
491 method could be compactly represented by a *Butcher tableau* which structures the coefficients the
492 following way

$$\begin{array}{c|c} c & A \\ \hline & b^T \end{array}$$

493 The two symplectic and symmetric Gauss-Legendre methods (found e.g. in [8]) with order $p = 4, 6$
494 and denoted as GL4 and GL6 in Table 1 are presented in below:

$$\begin{array}{c|cc} \frac{1}{2} - \frac{\sqrt{3}}{6} & \frac{1}{4} & \frac{1}{4} - \frac{\sqrt{3}}{6} \\ \frac{1}{2} + \frac{\sqrt{3}}{6} & \frac{1}{4} + \frac{\sqrt{3}}{6} & \frac{1}{4} \\ \hline & \frac{1}{2} & \frac{1}{2} \end{array} \quad \begin{array}{c|ccc} \frac{1}{2} - \frac{\sqrt{15}}{10} & \frac{5}{36} & \frac{2}{9} - \frac{\sqrt{15}}{15} & \frac{5}{36} - \frac{\sqrt{15}}{30} \\ \frac{1}{2} & \frac{5}{36} + \frac{\sqrt{15}}{24} & \frac{2}{9} & \frac{5}{36} - \frac{\sqrt{15}}{24} \\ \frac{1}{2} + \frac{\sqrt{15}}{10} & \frac{2}{9} + \frac{\sqrt{15}}{15} & \frac{5}{36} + \frac{\sqrt{15}}{30} & \frac{5}{36} \\ \hline & \frac{5}{18} & \frac{4}{9} & \frac{5}{18} \end{array}$$

Table 1: Properties of RK methods. *Symm.* is short for symmetric and *sympl.* for symplectic, and inv. for inverse.

Integration method	Name in figures	Order (p)	Stages (s)	Symm.	Sympl.	Inv. explicit	Explicit
Explicit Euler	E. Euler	1	1	no	no	yes	yes
Implicit Euler	I. Euler	1	1	no	no	yes	no
Runge–Kutta 4	RK4	4	4	no	no	yes	yes
Implicit midpoint	Midpoint	2	1	yes	yes	yes	no
MIRK3	MIRK3	3	2	no	no	yes	no
MIRK4	MIRK4	4	3	yes	no	yes	no
MIRK5	MIRK5	5	4	no	no	yes	no
MIRK6	MIRK6	6	5	yes	no	yes	no
Gauss Legendre 4	GL4	4	2	yes	yes	no	no
Gauss Legendre 6	GL6	6	4	yes	yes	no	no

495 C.2 Mono-Implicit Runge–Kutta methods

496 The MIRK methods are specified by a coefficient vector $b \in \mathbb{R}^s$, $v \in \mathbb{R}^s$ in addition to the strictly
497 lower triangular matrix $D \in \mathbb{R}^{s \times s}$ and could be represented by the an extended Butcher tableau in
498 the following manner

$$\begin{array}{c|c|c} c & v & D \\ \hline & & b^T \end{array}$$

499 In [22] it is proved that the maximum order of an s -stage MIRK method is $p = s + 1$ and several
500 methods with stages $s \leq 5$ are presented. Below, we specify the MIRK methods used in the numerical
501 experiments in addition to presenting their extended Butcher tableau in Figure 9.

0	0	0	0	0	0	0
1	1	0	0	0	0	0
$\frac{1}{2} - \frac{\sqrt{21}}{14}$	$\frac{1}{2} - \frac{7\sqrt{21}}{128}$	$\frac{1}{14} + \frac{\sqrt{21}}{98}$	$-\frac{1}{14} + \frac{\sqrt{21}}{98}$	0	0	0
$\frac{1}{2} + \frac{\sqrt{21}}{14}$	$\frac{1}{2} + \frac{7\sqrt{21}}{128}$	$\frac{1}{14} - \frac{\sqrt{21}}{98}$	$-\frac{1}{14} - \frac{\sqrt{21}}{98}$	0	0	0
$\frac{1}{2}$	$\frac{1}{2}$	$-\frac{5}{128}$	$\frac{5}{128}$	$\frac{7\sqrt{21}}{128}$	$-\frac{7\sqrt{21}}{128}$	0
		$\frac{1}{20}$	$\frac{1}{20}$	$\frac{49}{180}$	$\frac{49}{180}$	$\frac{16}{45}$

Figure 9: Extended Butcher tableau of MIRK methods with stage and order $(s, p) = (2, 3), (3, 4), (4, 5), (5, 6)$.

- Midpoint: The symmetric and symplectic MIRK method where $(s, p) = (1, 2)$ is equivalent to the midpoint method.
 - MIRK3: The method $(s, p) = (2, 3)$ found by choosing $c_1 = 1$ in [22].
 - MIRK4: The method $(s, p) = (3, 4)$ with $x_{31} = \frac{1}{8}$ in [42] and is first presented in [25, 43].
 - MIRK5: The method $(s, p) = (4, 5)$ presented in [22] choosing $c_2 = 0$ and $c_3 = \frac{3}{2}$. It should be noted that as long as $c_3 > 1$ the method is A-stable, however the particular choice of $c_3 = \frac{3}{2}$ is arbitrary.
 - MIRK6: The method $(s, p) = (5, 6)$ presented in [42], which is the $s = 5$ stage scheme in [22] choosing $c_3 = \frac{1}{2} - \frac{\sqrt{21}}{14}$. According to [42], this method is an improvement over earlier schemes on the same form which used $c_3 = \frac{1}{4}$.

513 C.3 Symmetric methods:

514 The exact flow of an ODE satisfies the following property known as (time) symmetry:

$$y(t_0) = \varphi_{h,f}^{-1}(y(t_0 + h)) = \varphi_{-h,f}(y(t_0 + h)),$$

where the superscript “ -1 ” denotes the inverse map. This is a desirable property also for the numerical approximation. A numerical integration method $\Phi_{h,f}$ is called symmetric if

$$\Phi_{h,f} = \Phi_{-h,f}^{-1}. \quad (17)$$

517 Symmetric numerical methods have the following properties [44]:

1. A symmetric integrator preserves the (time) symmetry of the exact flow.
2. The order p of a symmetric method is necessarily even.
3. Solutions of Hamiltonian systems satisfy the following reflection symmetry: if $(q(t), p(t))$ solves the Hamiltonian ODE, then $(q(-t), -p(-t))$ is also a solution, with $y(t) = [q(t), p(t)]^T$. Numerical solutions (q_n, p_n) obtained from a symmetric Runge–Kutta method satisfy the same reflection symmetry [7].

524 A Runge–Kutta method is symmetric if and only if

$$PA + AP - \mathbb{1}b^T = 0, \quad (18)$$

$$b = Pb, \quad (19)$$

where $\mathbb{1} := [1, \dots, 1]^T \in \mathbb{R}^s$ and $[p]_{ij} = \delta_{i,s+1-j}$ [44]. That is, P is the reflection of the identity matrix over the first axis. Inserting the definition of a MIRK method from (4), we get

$$PD + DP + (Pv + v - \mathbb{1})b^T = 0$$

$$b = Pb.$$

Symmetric MIRK methods of order $p = 2, 4, 6$ are presented in [45, 42] and specific examples are found in Figure 9.

D Details on the inverse injection in MII

Assume we are deriving the MII following the example in Equation (10) using the implicit midpoint method, where

$$y_{n+1} = y_n + hf\left(\frac{y_n + y_{n+1}}{2}\right) = y_n + h\Psi_{n,n+1}.$$

We thus find that the second term in (10), the composition of two steps starting in \tilde{y}_0 could be approximated by

$$\begin{aligned} \hat{y}_2 &= \Phi_{h,f} \circ \Phi_{h,f}(\tilde{y}_0) \\ &= \Phi_{h,f}(\tilde{y}_0) + hf\left(\frac{\Phi_{h,f}(\tilde{y}_0) + \hat{y}_2}{2}\right) \\ &\approx \Phi_{h,f}(\tilde{y}_0) + hf\left(\frac{\tilde{y}_1 + \tilde{y}_2}{2}\right) \\ &= \tilde{y}_0 + hf\left(\frac{\tilde{y}_0 + \Phi_{h,f}(\tilde{y}_0)}{2}\right) + h\Psi_{1,2} \\ &\approx \tilde{y}_0 + hf\left(\frac{\tilde{y}_0 + \tilde{y}_1}{2}\right) + h\Psi_{1,2} \\ &= \tilde{y}_0 + h\Psi_{0,1} + h\Psi_{1,2}. \end{aligned}$$

where the approximation \approx is obtained by the substitution $\tilde{y}_2 \rightarrow \hat{y}_2$ and $\tilde{y}_1 \rightarrow \Phi_{h,f}(\tilde{y}_0)$. The same procedure (repeatedly using the inverse injection) is generalized over longer trajectories and used to arrive at the MII method in Definition 5.1.

E Details on neural network training

The experiments were performed on a Apple M1 Pro chip with double precision. The PyTorch L-BFGS [36] algorithm is run with the following parameters:

- History size: 120.
- Gradient tolerance: 10^{-9} .
- Termination tolerance on parameter changes: 10^{-9} .
- Line search: Strong Wolfe.

Both MII and ISO works better when f_θ has been pre-trained to be a reasonable approximation of the underlying vector field f . Thus, for both MII and ISO training is run 10 epochs on the one-step method before training additional 10 epochs with MII (MII MIRK4) and ISO (ISO Störmer and ISO RK4). The ISO procedure (searching for the optimal initial value \hat{y}_0) utilizes the L-BFGS optimization algorithm from the SciPy library [46] with gradient tolerance of 10^{-6} and the maximum number of iterations limited to 10.

F Proof of Theorem 4.4

Proof. As stated in Equation (5) Runge–Kutta method as given by Equation (16) is symplectic if and only if

$$b_i a_{ij} + b_j a_{ji} - b_i b_j = 0.$$

553 Inserting the particular form for the MIRK coefficients, $a_{ij} = d_{ij} + v_i b_j$, we get

$$\begin{aligned} b_i(v_i b_j + d_{ij}) + b_j(v_j b_i + d_{ji}) - b_i b_j &= 0 \\ b_i d_{ij} + b_j d_{ji} + b_i b_j(v_j + v_i - 1) &= 0. \end{aligned} \quad (20)$$

554 As D is strictly lower triangular either $d_{ji} = 0$ or $d_{ij} = 0$, which for Equation (20) implies that

$$\begin{cases} b_i d_{ij} + b_i b_j(v_j + v_i - 1) = 0 & \text{if } i \neq j \\ b_i^2(2v_i - 1) = 0 & \text{if } i = j \end{cases}$$

555 Requiring d_{ij} , b_i and v_i to satisfy the symplecticity condition yields the following restriction

$$\begin{cases} b_i d_{ij} + b_i b_j(v_j + v_i - 1) = 0 & \text{if } i \neq j, \\ b_i = 0 \text{ or } v_i = \frac{1}{2} & \text{if } i = j. \end{cases} \quad (21)$$

556 Without loss of generality, we assume that the m first entries of $b \in \mathbb{R}^s$ are zero. Enforcing the
557 conditions of Equation (21) on $v \in \mathbb{R}^s$ we get for $1 \leq m \leq s$

$$\begin{aligned} b &= [0, \dots, 0, b_{m+1}, \dots, b_s]^T, \\ v &= [v_1, \dots, v_m, \frac{1}{2}, \dots, \frac{1}{2}]^T. \end{aligned}$$

558 In total, this gives the following constraints for v, b, D :

	$b_j = 0$	$b_j \neq 0$
$b_i = 0$	$d_{ij} \in \mathbb{R}$ $v_i, v_j \in \mathbb{R}$	$d_{ij} \in \mathbb{R}$ $v_i, v_j \in \mathbb{R}$
$b_i \neq 0$	$d_{ij} = 0$ $v_i, v_j \in \mathbb{R}$	$d_{ij} = 0$ $v_i, v_j = \frac{1}{2}$

559 Which for the Runge–Kutta method $A = D + vb^T$ gives a (RK) Butcher tableau of the form

0	0	\dots	0	$v_1 b_{m+1}$	\dots	$v_1 b_s$
d_{21}	0	\dots	0			
d_{31}	d_{32}		0	\vdots		\vdots
\vdots		\ddots				
$d_{m,1}$	\dots	$d_{m,m-1}$	0	$v_m b_{m+1}$	\dots	$v_m b_s$
0	\dots	\dots	0	$\frac{1}{2} b_{m+1}$	\dots	$\frac{1}{2} b_s$
\vdots			\vdots	\vdots		\vdots
0	\dots	\dots	0	$\frac{1}{2} b_{m+1}$	\dots	$\frac{1}{2} b_s$
0	\dots	\dots	0	b_{m+1}	\dots	b_s

560 Since the lower left submatrix is the zero matrix, this leaves the stages k_{m+1}, \dots, k_s unconnected
561 to the first m stages. In addition, as $b_i = 0$ for $i = 1, \dots, m$, these stages are not included in the
562 computation of the final integration step. The method is thus reducible to the lower right submatrix
563 of A and b_{m+1}, \dots, b_s . The reduced method is thus in general given by the following stage-values

$$k_i = f(y_n + \frac{h}{2} \sum_j b_j k_j).$$

564 It is trivial to check that if $\sum_i b_i = 1$ the method satisfies order conditions up to order $p = 2$, which
565 could be found in [8] to be

$$\sum_i b_i = 1, \quad \text{and} \quad \sum_{i,j} b_i a_{ij} = \frac{1}{2}.$$

566 However, the method fails to satisfy the first of the two conditions required for order $p = 3$, since

$$\sum_{i,j,k} b_i a_{ij} a_{ik} = \frac{1}{4} \sum_{i,j,k} b_i b_j b_k = \frac{1}{4} \neq \frac{1}{3}.$$

567 Hence, the maximum order of a symplectic MIRK method is $p = 2$. □

568 As a remark, it should be noted that the $s = 1$ stage, symplectic MIRK method found by setting
 569 $b_1 = 1$, $v_1 = \frac{1}{2}$ and $d_{11} = 0$ is simply the midpoint method $y_{n+1} = y_n + h k_1$ with $k_1 = f(\frac{y_n + y_{n+1}}{2})$.

570 G Proof of Theorem 5.2

571 *Proof.* Let $s_i(y_n, y_{n+1}) := y_n + v_i(y_{n+1} - y_n)$ and \tilde{y}_n be noisy data (9). Observe that we can obtain
 572 the following approximation to the MIRK stages (4) by

$$\begin{aligned} k_i &= f\left(\tilde{y}_n + v_i(\tilde{y}_{n+1} - \tilde{y}_n) + h \sum_{j=1}^s d_{ij} k_j\right) \\ &= f\left(s_i(y_n, y_{n+1}) + s_i(\delta_n, \delta_{n+1}) + \mathcal{O}(h)\right) \\ &= f(s_i(y_n, y_{n+1})) + f'(s_i(y_n, y_{n+1})) s_i(\delta_n, \delta_{n+1}) + \mathcal{O}(\|s(\delta_n, \delta_{n+1})\|^2) + \mathcal{O}(h) \\ &= f(y_n) + f'(y_n) s_i(\delta_n, \delta_{n+1}) + \mathcal{O}(\|s(\delta_n, \delta_{n+1})\|^2) + \mathcal{O}(h). \end{aligned} \quad (22)$$

573 Where in the final equality we expand $y_{n+1} = y_n + h f(y_n) + \mathcal{O}(h^2)$ to find

$$\begin{aligned} f(s_i(y_n, y_{n+1})) &= f(y_n + v_i(y_{n+1} - y_n)) \\ &= f(y_n + h v_i f(y_n) + \mathcal{O}(h^2)) \\ &= f(y_n) + \mathcal{O}(h). \end{aligned}$$

574 And similarly for $f'(s_i(y_n, y_{n+1}))$. In total, this means that the next MIRK-step could be approxi-
 575 mated by

$$\begin{aligned} y_{n+1} &= \tilde{y}_n + h \sum_{i=1}^s b_i k_i \\ &= \tilde{y}_n + h \sum_{i=1}^s b_i \left(f(y_n) + f'(y_n) s_i(\delta_n, \delta_{n+1}) \right) + \mathcal{O}(h \|s(\delta_n, \delta_{n+1})\|^2) + \mathcal{O}(h^2). \end{aligned} \quad (23)$$

576 First note, that if x is a multivariate normally distributed random variable $x \sim \mathcal{N}(0, \Sigma)$ then for a
 577 matrix $G \in \mathbb{R}^{n \times n}$ the variance of the linear transformation is given by $\text{Var}[Gx] := \text{Cov}[Gx, Gx] =$
 578 $G \Sigma G^T$. Now, using the approximation in Equation (23), we find the variance of the optimization

579 target $\mathcal{T}_{n+1}^{\text{OS}} = \tilde{y}_{n+1} - \Phi_{h,f}(\tilde{y}_n, \tilde{y}_{n+1})$ by

$$\begin{aligned}
\text{Var} \left[\tilde{y}_{n+1} - \tilde{y}_n - h \sum_{i=1}^s b_i k_i \right] &\approx \text{Var} \left[\delta_{n+1} - \delta_n - h \sum_{i=1}^s b_i \left(f'(y_n) s_i(\delta_n, \delta_{n+1}) \right) \right] \\
&= \text{Var} \left[\left(I - hb^T v f'(y_n) \right) \delta_{n+1} - \left(I + hb^T (\mathbb{1} - v) f'(y_n) \right) \delta_n \right] \\
&= \sigma^2 \left(I - hb^T v f'(y_n) \right) \left(I - hb^T v f'(y_n) \right)^T \\
&\quad + \sigma^2 \left(I + hb^T (\mathbb{1} - v) f'(y_n) \right) \left(I + hb^T (\mathbb{1} - v) f'(y_n) \right)^T \\
&= \sigma^2 \left[2I + hb^T (\mathbb{1} - 2v) (f'(y_n) + f'(y_n)^T) \right. \\
&\quad \left. + h^2 \underbrace{\left((b^T v)^2 + (b^T (\mathbb{1} - v))^2 \right) f'(y_n) f'(y_n)^T}_{:= Q^{\text{OS}}} \right] \\
&= \sigma^2 \left[2I + hb^T (\mathbb{1} - 2v) (f'(y_n) + f'(y_n)^T) + h^2 Q^{\text{OS}} \right].
\end{aligned}$$

580 Here, $\mathbb{1} := [1, \dots, 1]^T \in \mathbb{R}^s$. This is the variance estimate we wanted to find for MIRK methods
581 used as one-step integration schemes. Similarly, considering a point computed by the mean inverse
582 integrator \bar{y}_n , we find, using the stage approximation by Equation (22) that

$$\begin{aligned}
\bar{y}_n &= \frac{1}{N} \sum_{\substack{j=0 \\ j \neq n}}^N \tilde{y}_j + \frac{h}{N} \sum_{j=0}^{N-1} w_{n,j} \sum_{l=1}^s b_l k_l \\
&\approx \frac{1}{N} \sum_{\substack{j=0 \\ j \neq n}}^N \tilde{y}_j + \frac{h}{N} \sum_{j=0}^{N-1} w_{n,j} \sum_{l=1}^s b_l \left(f(y_j) + f'(y_j) s_l(\delta_j, \delta_{j+1}) \right)
\end{aligned}$$

583 Where we note that $w_{n,j} := [W]_{nj}$, from Definition 5.1 of the MII. Let $\bar{y}_n := [\bar{Y}]_n$. Computing the
584 variance of the optimization target $\mathcal{T}_i^{\text{MI}} = \tilde{y}_n - \bar{y}_n$ we find, by introducing \bar{P}_{nj} to simplify notation,
585 that

$$\begin{aligned}
\text{Var} [\tilde{y}_n - \bar{y}_n] &\approx \text{Var} \left[\delta_n - \frac{1}{N} \sum_{\substack{j=0 \\ j \neq n}}^N \delta_j - \frac{h}{N} \sum_{j=0}^{N-1} w_{n,j} f'(y_j) \left(b^T (\mathbb{1} - v) \delta_j + b^T v \delta_{j+1} \right) \right] \\
&= \text{Var} \left[\frac{1}{N} \sum_{\substack{j=0 \\ j \neq n}}^N \left(I + h f'(y_j) \underbrace{\left(\tilde{w}_{n,j} b^T (\mathbb{1} - v) + \tilde{w}_{n,j-1} b^T v \right)}_{:= \bar{P}_{nj}} \right) \delta_j \right] \\
&\quad + \text{Var} \left[\left(I - \frac{h}{N} f'(y_n) \underbrace{\left(\tilde{w}_{n,n} b^T (\mathbb{1} - v) + \tilde{w}_{n,n-1} b^T v \right)}_{:= \bar{P}_{nn}} \right) \delta_n \right] \\
&= \text{Var} \left[\frac{1}{N} \sum_{\substack{j=0 \\ j \neq n}}^N \left(I + h \bar{P}_{nj} \right) \delta_j \right] + \text{Var} \left[\left(I - \frac{h}{N} \bar{P}_{nn} \right) \delta_n \right] \\
&= \frac{\sigma^2}{N^2} \sum_{\substack{j=0 \\ j \neq n}}^N \left(I + h \bar{P}_{nj} \right) \left(I + h \bar{P}_{nj} \right)^T + \sigma^2 \left(I - \frac{h}{N} \bar{P}_{nn} \right) \left(I - \frac{h}{N} \bar{P}_{nn} \right)^T.
\end{aligned}$$

586 In the second line, $\tilde{w}_{n,j}$ is introduced which is elements of a matrix $\tilde{W} = [0|w_1|w_2|\dots|w_N|0] \in$
 587 $\mathbb{R}^{N \times N+1}$, or in other words the matrix you obtain by padding W right and left with a column of
 588 zeros. Expanding the terms and introducing matrices P_{nj} and Q^{MII} we finally find

$$\begin{aligned} \text{Var}[\tilde{y}_n - \bar{y}_n] &\approx \frac{\sigma^2}{N} \left[(1+N)I + h \underbrace{(\bar{P}_{nn} + \bar{P}_{nn}^T)}_{=P_{nn}} + \frac{h}{N} \sum_{\substack{j=0 \\ j \neq n}}^s \underbrace{(\bar{P}_{nj} + \bar{P}_{nj}^T)}_{:=P_{nj}} + \frac{h^2}{N} \sum_{j=0}^s \underbrace{\bar{P}_{nj} \bar{P}_{nj}^T}_{:=Q^{\text{MII}}} \right] \\ &= \frac{\sigma^2}{N} \left[(1+N)I + hP_{nn} + \frac{h}{N} \sum_{\substack{j=0 \\ j \neq n}}^s P_{nj} + \frac{h^2}{N} Q^{\text{MII}} \right]. \end{aligned}$$

589 Since for a symmetric matrix A we have that the spectral radii ρ (largest absolute value of eigenvalues)
 590 could be found by $\rho(A) = \|A\|_2$, we find for both variance approximations (covariance matrix is
 591 always symmetric) that

$$\begin{aligned} \rho\left(\text{Var}[\tilde{y}_n - \bar{y}_n]\right) &\approx \frac{\sigma^2}{N} \left\| (1+N)I + hP_{nn} + \frac{h}{N} \sum_{\substack{j=0 \\ j \neq n}}^s P_{nj} + \frac{h^2}{N} Q^{\text{MII}} \right\|_2 \\ \rho\left(\text{Var}[\tilde{y}_{n+1} - \Phi_{h,f}(\tilde{y}_n, \tilde{y}_{n+1})]\right) &\approx \sigma^2 \left\| 2I + hb^T(\mathbb{1} - 2v)(f'(y_n) + f'(y_n)^T) + h^2 Q^{\text{OS}} \right\|_2. \end{aligned}$$

592 Finally, we note that:

$$\begin{aligned} Q^{\text{OS}} &:= \left((b^T v)^2 + (b^T(\mathbb{1} - v))^2 \right) f'(y_n) f'(y_n)^T \\ \bar{P}_{nj} &:= f'(y_j) \left(\tilde{w}_{n,j} b^T(\mathbb{1} - v) + \tilde{w}_{n,j-1} b^T v \right) \\ P_{nj} &:= \bar{P}_{nj} + \bar{P}_{nj}^T \\ Q^{\text{MII}} &:= \sum_{j=0}^s \bar{P}_{nj} \bar{P}_{nj}^T. \end{aligned} \tag{24}$$

593 □

594 **H Higher-order inverse-explicit invariant-preserving symmetric** 595 **non-partitioned integrators**

596 We define invariant-preserving integrators as methods that preserve the Hamiltonian or other invariants
 597 of the exact solution, either exactly up to machine precision or within a bound, like symplectic
 598 methods. Although we argue in this paper that symplecticity is a less important property when
 599 learning Hamiltonian systems from data than for integration of a known system, we do not mean to
 600 suggest that invariant-preserving integrators may not be beneficial to some extent and have important
 601 qualities in the inverse problem also. However, we urge anyone who seeks to use invariant-preserving
 602 methods to also consider the order of the method and whether it is a symmetric inverse-explicit
 603 method. Although the maximum order of a symplectic inverse-explicit *Runge–Kutta* method is two,
 604 there exist higher-order inverse-explicit invariant-preserving integrators that are not Runge–Kutta
 605 methods.

606 Note that *partitioned* Runge–Kutta (PRK) methods is an extension that does not belong to the class of
 607 Runge–Kutta methods. This is important to clarify, since there exist PRK methods that are symmetric
 608 and explicit for separable systems. This marks a distinction from non-partitioned RK methods:
 609 these cannot be symmetric and explicit in general [8]. Several papers suggest using symplectic PRK
 610 methods for learning Hamiltonian systems [10, 30, 29], but these methods, although symmetric, only
 611 depend on one point to approximate the right hand side of each integration step, and thus do not
 612 average out any noise.

613 H.1 Symplectic elementary differential Runge–Kutta methods

614 Chartier et al. showed in [47] that an integrator can be applied to a modified vector field in such a way
 615 that it yields a higher order approximation of the original vector field while inheriting the geometric
 616 properties of the given integrator. As an example, they present the fourth-order modified implicit
 617 midpoint method

$$\frac{y_{n+1} - y_n}{h} = f(\bar{y}) + \frac{h}{12} \left(-Df(\bar{y})Df(\bar{y})f(\bar{y}) + \frac{1}{2}D^2f(\bar{y})f(\bar{y})f(\bar{y}) \right), \quad (25)$$

618 where $\bar{y} = (y_n + y_{n+1})/2$. This is an example of an elementary differential Runge–Kutta (EDRK)
 619 method [33], which relies on the calculation of (multi-order) derivatives of the vector field f , denoted
 620 here as $D^p f$ for order p . Automatic differentiation can be utilized also to get higher-order derivatives,
 621 and we note that f , Df and $D^2 f$ each only have to be evaluated once for each training step, since
 622 they are only evaluated at the one point \bar{y} . A sixth-order modification of the implicit midpoint method
 623 is also presented in [47], but that requires the calculation of up to fourth-order derivatives and might
 624 be considered prohibitively expensive.

625 H.2 Discrete gradient methods

626 Discrete gradient methods are a class of integrators that can preserve an invariant, e.g. the Hamiltonian,
 627 exactly [32]. This is in contrast to symplectic methods, which only preserve a perturbation of the
 628 invariant exactly and the exact invariant within some bound. We remark that no method can be both
 629 symplectic and exactly invariant-preserving in general [48]. Discrete gradient methods are defined
 630 strictly for invariant-preserving ODEs, which can be written on the form

$$\dot{y} = S(y)\nabla H(y), \quad (26)$$

631 for some skew-symmetric matrix $S(y)$ [32]. Then a discrete gradient is a function $\bar{\nabla}H : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$
 632 satisfying

$$\bar{\nabla}H(u, v)^T(u - v) = H(u) - H(v),$$

633 a discrete analogue to the invariant-preserving property $\dot{H}(y) = \nabla H(y)^T \dot{y} = 0$ of (26). A corre-
 634 sponding discrete gradient method is then given by

$$\frac{y_{n+1} - y_n}{h} = \bar{S}(y_n, y_{n+1}, h)\bar{\nabla}H(y_n, y_{n+1}), \quad (27)$$

635 for some approximation $\bar{S}(y_n, y_{n+1}, h)$ of $S(y)$ such that $\bar{S}(y, y, 0) = S(y)$, where h is the step
 636 size in time. A discrete gradient can at most be a second-order approximation of the gradient, but
 637 appropriate choices of \bar{S} can yield inverse-explicit integrators of arbitrarily high order [16]. Matsubara
 638 et al. have developed a discrete version of the automatic differentiation algorithm that makes it possible
 639 to efficiently calculate a discrete gradient of neural network functions, and demonstrated its use in
 640 training of HNNs [15] and for detecting invariants [49]. A fourth-order discrete gradient method is
 641 suggested for training HNNs in [16], given a constant S in (26). This is the scheme (27) with

$$\bar{S}(y_n, \cdot, h) = S + \frac{8}{9}hSQ(y_n, z_2)S - \frac{1}{12}h^2SD^2H(z_1)SD^2H(z_1)S,$$

642 with $z_1 = y_n + \frac{1}{2}hf(y_n)$, $z_2 = y_n + \frac{3}{4}hf(z_1)$ and $Q(u, v) := \frac{1}{2}(D_2\bar{\nabla}H(u, v)^T - D_2\bar{\nabla}H(u, v))$,
 643 where $D_2\bar{\nabla}H$ denotes the derivative of $\bar{\nabla}H$ with respect to the second argument, and $D^2H := D\nabla H$
 644 is the Hessian of H . This is not symmetric, so we propose here instead the *fourth-order symmetric*
 645 *invariant-preserving scheme* obtained by

$$\begin{aligned} \bar{S}(y_n, y_{n+1}, h) = & S + \frac{h}{2}S(Q(y_n, \frac{1}{3}y_n + \frac{2}{3}y_{n+1}) - Q(y_{n+1}, \frac{2}{3}y_n + \frac{1}{3}y_{n+1}))S \\ & - \frac{1}{12}(h)^2SD^2H(\bar{y})SD^2H(\bar{y})S. \end{aligned}$$

646 H.3 Numerical comparison of fourth-order integrators

647 We test four different fourth-order integrators on solving an initial value problem of the double
 648 pendulum described in Appendix A. We compute an approximation of the error of the solution at each

time by comparing to a solution obtained using RK4 with 10 times as many time steps. As seen in the left plot of Figure 10, the symmetric methods are clearly superior to the explicit RK4 method, when using the same step size. For integration, the advantage of RK4 is that it is more computationally efficient than the implicit methods, which facilitates taking smaller step sizes. However, as pointed out in Section 4, RK4 does not have this advantage over MIRK methods for the inverse problem.

Furthermore, although the higher-order MIRK methods we suggest to use in this paper are not symplectic and thus lack general energy preservation guarantees, we see from Figure 10 that they may still preserve the energy within a bound for specific problems. In fact, for the double pendulum problem considered here, the non-symplectic MIRK4 method preserves the energy slightly better than the symplectic MIMP4 scheme up to time $T = 500$. The invariant-preserving discrete gradient method preserves the Hamiltonian to machine precision.

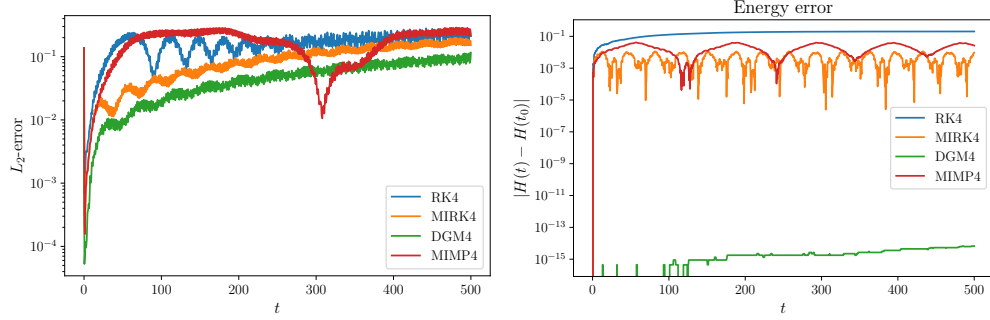


Figure 10: Global error (*left*) and energy error (*right*) of the solution of the double pendulum problem obtained using four different integrators. The initial condition is $y_0 = [0.1, 0.3, -0.4, 0.2]^T$, and the step size for all integrators is $h = \frac{1}{2}$.

659

660 I Computational cost

The fourth-order MIRK method from Table 9 (MIRK4) is between twice and thrice as expensive as the implicit midpoint method, depending on the training strategy. That is, if no batching is performed and f is evaluated at all points in the training set at each iteration of the optimization, then the number of function evaluations for a trajectory with n points is $n - 1$ for the implicit midpoint method and $2n - 1$ for MIRK4. However, if batching is done and function evaluations cannot generally be reused for successive points, the total number of function evaluations at each epoch may increase to $3n - 3$ for MIRK4.

In general, the cost of an s -stage MIRK method depends on both the training strategy and whether the end points y and \hat{y} are two of the stages. If batching is not done and y and \hat{y} are two of the stages, then computational cost at each epoch is $\mathcal{O}(m(n + (s - 2)(n - 1)))$, where m is the number of trajectories of n points in each. The maximum cost with batching is the same as the cost if y and \hat{y} are not two of the stages: $\mathcal{O}(ms(n - 1))$. This cost is equivalent to that of an explicit s -stage RK method.

673