


# Symplectic integration of learned Hamiltonian systems

C. Offen <sup>1</sup> and S. Ober-Blöbaum<sup>1</sup>

Paderborn University, Department of Mathematics, Warburger Str. 100, 33098 Paderborn, Germany

(\*Electronic mail: christian.offen@uni-paderborn.de)

(Dated: 13 December 2021)

Hamiltonian systems are differential equations which describe systems in classical mechanics, plasma physics, and sampling problems. They exhibit many structural properties, such as a lack of attractors and the presence of conservation laws. To predict Hamiltonian dynamics based on discrete trajectory observations, incorporation of prior knowledge about Hamiltonian structure greatly improves predictions. This is typically done by learning the system's Hamiltonian and then integrating the Hamiltonian vector field with a symplectic integrator. For this, however, Hamiltonian data needs to be approximated based on the trajectory observations. Moreover, the numerical integrator introduces an additional discretisation error. In this paper, we show that an inverse modified Hamiltonian structure adapted to the geometric integrator can be learned directly from observations. A separate approximation step for the Hamiltonian data is avoided. The inverse modified data compensates for the discretisation error such that the discretisation error is eliminated. The technique is developed for Gaussian Processes.

**Combining trajectory data with prior knowledge about structural properties of a dynamical system is known to greatly improve predictions of the system's motions. The article introduces the new technique *Symplectic Shadow Integration*, which incorporates *modified* structures into learned models. The modified structures compensate discretisation errors which limit the accuracy of existing approaches.**

## I. INTRODUCTION

### A. Hamiltonian systems and symplectic integrators

A Hamiltonian system on the phase space  $M = \mathbb{R}^{2n}$  is a differential equation of the form

$$\dot{z} = J^{-1} \nabla H(z), \quad J = \begin{pmatrix} 0 & -I_n \\ I_n & 0 \end{pmatrix}, \quad (1)$$

where  $H: M \rightarrow \mathbb{R}$  and  $I_n$  is an  $n$ -dimensional identity matrix. Hamiltonian systems arise in classical mechanics, plasma physics, electrodynamics, sampling problems (Hamiltonian Monte Carlo methods) and many other applications. Trajectories of Hamiltonian systems conserve  $H$ . In classical mechanics this corresponds to energy conservation. The dynamical system has no attractors. Moreover, its flow map  $\phi_t: M \rightarrow M$  is symplectic, i.e. it fulfils

$$\phi_t'(z)^\top J \phi_t'(z) = J \quad \forall z \in M, \quad (2)$$

where  $\phi_t'(z)$  denotes the Jacobian matrix of  $\phi_t$  at  $z$ . The symplectic structure has the remarkable effect that (symplectic) symmetries of  $H$  yield conserved quantities of the flow by Noether's theorem. It is, therefore, rewarding to preserve symplecticity and symmetries when discretising (1) such that the numerical flow inherits qualitative features such as conserved quantities, complete integrability, and regular and chaotic regions from the exact flow. Integrators which preserve symplectic structure are called *symplectic integrators*. A key feature is that the numerical trajectories ob-

tained by a symplectic integrator conserve a modified Hamiltonian or shadow Hamiltonian, which can be calculated explicitly using backward error analysis techniques.<sup>1</sup> Moreover, symplectic maps preserve phase space volume. This makes symplectic integration schemes relevant for sampling techniques such as Hamiltonian Monte Carlo methods.<sup>2</sup> Next to these effects concerning initial value problems, preservation of symplectic structure is crucial to capture the bifurcation behaviour of solutions to boundary value problems in Hamiltonian systems,<sup>3-5</sup> variational PDEs,<sup>6</sup> and in optimal control problems.<sup>7</sup>

### B. Symplectic Shadow Integration (SSI)

Techniques to identify Hamiltonian functions  $H$  from data were developed by Bertalan et. al.<sup>8</sup> For training it relies on the availability of derivatives of the flow map, so it cannot learn  $H$  from observed trajectories directly. We will modify the idea and identify an *inverse modified Hamiltonian*  $\bar{H}$  directly from the data without any numerical approximations of derivatives. Inverse modified equations and inverse modified Hamiltonians were introduced by Zhu et. al.,<sup>9</sup> where they are used as an analysis tool for neural networks. An inverse modified Hamiltonian  $\bar{H}$  is adapted to a symplectic integrator such that if the integrator is applied to  $\dot{z} = J^{-1} \nabla \bar{H}(z)$  the discretisation error of the numerical scheme gets compensated for.

We introduce the following procedure, which we coin *Symplectic Shadow Integration*:

1. *Preparation.* Choose a symplectic integrator and a step size  $h$  compatible with the discrete trajectory observations.
2. *Inverse system identification.* Learn the inverse modified Hamiltonian  $\bar{H}$  from data.
3. *Integration.* Apply the symplectic integrator to the inverse modified Hamiltonian system  $\dot{z} = J^{-1} \nabla \bar{H}(z)$  to obtain a numerical flow map.

4. *System identification.* If required, compute  $H$  from  $\bar{H}$  using backward error analysis techniques to obtain physical insight or for verification.

### C. Relation to other approaches

Before we review the idea of backward error analysis and provide details on the steps of the Symplectic Shadow Integration (SSI) procedure, let us outline the advantages of the integration technique over other more direct approaches to predict Hamiltonian dynamics from data and contrast SSI to techniques in the literature.

#### *Comparison to learning the flow map directly (Strategy 1).*

One could use established learning techniques, such as artificial neural networks, Gaussian processes, or kernel methods to learn the flow map of the system directly from trajectory data. SSI has the following advantages over this approach.

- Hamiltonian structure is incorporated into the learned system. This guarantees important qualitative aspects of the prediction such as energy conservation, preservation of phase space volume and topological properties of the phase portrait.
- Only a real valued map  $\bar{H}$  needs to be learned rather than the flow map, reducing the dimension of the learning problem and data requirements.
- Hamiltonian structure can be identified. It provides physical insight into the dynamics and can be used for verification. More precisely, the predicted motions using the SSI technique are the exact motions of an identified Hamiltonian, which can be computed from  $\bar{H}$ . The performance of SSI can, therefore, be evaluated using backward error analysis.
- SSI provides a framework to incorporate further prior knowledge about conservation laws, such as (angular) momentum conservation, through a combination with symmetric learning, for instance using symmetric kernels for Gaussian processes.<sup>10</sup> The conservation laws are then guaranteed by a discrete Noether theorem.

In this context, we mention Symplectic Neural Networks (SympNets).<sup>11</sup> SympNets can be used to learn the flow map of Hamiltonian systems and the learned map is guaranteed to be symplectic. In contrast, SSI learns a scalar valued map  $\bar{H}$  related to the Hamiltonian of the system. Next to neural networks, SSI can be used with Gaussian Processes and kernel methods, which will be the focus of this work. A technique analogous to SSI has been developed for artificial neural networks in the recent preprint.<sup>12</sup>

Another technique to incorporate geometric structure is to learn a generating function of the symplectic flow map.<sup>13,14</sup> The learned flow map is then guaranteed to be symplectic. In

contrast, SSI identifies Hamiltonian structure, albeit the approaches can yield the same learning problem for some simple numerical schemes. However, SSI uses higher order terms to correct predictions of the Hamiltonian function in a post processing step.

#### *Comparison to learning the exact Hamiltonian and then using a symplectic integrator (Strategy 2).*

Techniques have been developed to learn the Hamiltonian rather than the flow map from data using Gaussian process regression or artificial neural networks.<sup>8</sup> This approach requires a subsequent application of a classical numerical integrator to predict motions. SSI has the following advantages over learning the exact Hamiltonian  $H$  and applying a symplectic integrator to (1).

- The numerical integrator introduces a discretisation error in addition to uncertainty in the Hamiltonian due to limited training data. SSI compensates this discretisation error such that high accuracy and excellent energy behaviour can be achieved despite large step sizes.
- Step size selection is decoupled from accuracy requirements. This is beneficial if the learned Hamiltonian and its gradient are expensive to evaluate.
- SSI uses the trajectory data directly. There is no need to approximate data of the underlying vector field, which would include additional discretisation errors.

Hamiltonian neural networks<sup>15</sup> constitute an example for a technique covered by strategy 2: the Hamiltonian of the dynamical system is parametrised as an artificial neural network. For training, observations of velocities and derivatives of conjugate momenta are required. If a special form of the Hamiltonian is assumed, such as the mechanical form  $H(q, p) = p^\top M(q)p + V(q)$ , the components  $M$  and  $V$  can be parametrised separately by neural networks.<sup>16</sup> Exploiting the mechanical form of the Hamiltonian, only position and velocity observations are required. Moreover, control terms can be added. In contrast, SSI applies to all canonical Hamiltonian systems, does not assume a particular form of the Hamiltonian, and requires observations of position and momentum data only. Observations of derivatives of these quantities are not needed.

Hamiltonian neural networks can suffer from two types of discretisation errors: the training data needs to contain velocity data and information on the derivative of conjugate momenta. Typically, these need to be approximated from observed position and momentum data, which introduces a discretisation error. Another discretisation error occurs when the learned Hamiltonian is integrated using a numerical method. These errors have been analysed by Zhu et. al.<sup>9</sup> SSI can be trained directly on position and momentum data and all discretisation errors are compensated, which is important when high accuracy is required.

A different approach to compensate discretisation errors can be found in the work by Park et. al.<sup>17</sup> It applies to spin




FIG. 1: Illustration of the relation of the vector field  $f$ , the modified vector field  $\tilde{f}$ , and the inverse modified vector field  $\tilde{f}$ . Backward error analysis (BEA) provides a tool to calculate the effect of an integrator on a vector field, i.e. to compute  $\tilde{f}$  from  $f$  as formal power series in the integrator's step size.

systems: a neural network is trained and used to compensate the discretisation error of a numerical time stepping scheme. Therefore, large time steps can be used and computations can be accelerated. In contrast, SSI uses analytical methods to correct discretisation errors and applies to all canonical Hamiltonian systems. Again, large time steps can be used without losing accuracy. Moreover, SSI does not require explicit knowledge of the Hamiltonian but identifies Hamiltonian structure from data.

The paper proceeds as follows: Section II contains a review of (inverse) modified Hamiltonians and of backward error analysis techniques. After a brief review of Gaussian process regression, Section III explains how to learn inverse modified Hamiltonians using Gaussian processes for SSI. The paper proceeds with numerical experiments in Section IV and concludes with remarks on possible extensions of SSI in Section V.

## II. (INVERSE-) BACKWARD ERROR ANALYSIS (BEA)

Backward error analysis (BEA) is a well-established tool of traditional numerical analysis. Its role in this work is threefold: it provides a theoretical justification of the SSI technique, is integral when SSI is used to perform system identification, and is employed to evaluate the quality of numerical experiments.

When a differential equation  $\dot{z} = f(z)$  is discretised using a numerical method, one obtains a flow map  $\tilde{\phi}$  that approximates the exact flow  $\phi$  of the system. The forward error measures the difference between  $\phi$  and  $\tilde{\phi}$ . In contrast, backward error analysis seeks a modified differential equation  $\dot{z} = \tilde{f}(z)$  whose exact flow map coincides with  $\tilde{\phi}$ . One can then compare the exact vector field  $f$  with the modified (or shadow) vector field  $\tilde{f}$  and understand the properties of the numerical method through the properties of  $\tilde{f}$ . The idea has been successfully applied to analyse long-term behaviour of numerical methods, especially symplectic integrators.<sup>1,18</sup> For a given integration method, an *inverse modified differential equation*<sup>9</sup>  $\dot{z} = \tilde{f}(z)$  is a differential equation such that the integration scheme applied to  $\tilde{f}$  yields the exact flow  $\phi$ , in other words  $\tilde{f} = f$ . The relations of  $f$ ,  $\tilde{f}$ , and  $\tilde{f}$  are illustrated in Figure 1.

Symplectic integrators have the remarkable property that applied to a Hamiltonian vector field  $f = J^{-1}\nabla H$  the modified vector field  $\tilde{f}$  is of the form<sup>1</sup>  $\tilde{f} = J^{-1}\nabla \tilde{H}$  and the inverse modified vector field  $\tilde{f}$  is of the form<sup>9</sup>  $\tilde{f} = J^{-1}\nabla \tilde{H}$ . See Fig-




FIG. 2: Illustration of modified and inverse modified Hamiltonian equations. Formal power series  $\tilde{H}$  and  $\tilde{H}$  in the step size of the integrator exist such that the diagram commutes if the integration scheme is symplectic.




FIG. 3: Figures 1 and 2 only show a formal analysis. The power series  $\tilde{H}$  typically does not converge. For a rigorous analysis,  $\tilde{H}$  is optimally truncated at an index  $N(h)$ . The exact flow  $\phi^{[N(h)]}$  of the truncated Hamiltonian  $\tilde{H}^{[N(h)]}$  agrees with the numerical flow  $\psi_h$  up to an exponentially small error.

ure 2 for an illustration. The Hamiltonian structure of  $\tilde{f}$  and  $\tilde{f}$  provides an explanation for the excellent preservation properties of symplectic integrators: the energy error  $H(z_j) - H(z_0)$  oscillates within a band of width  $\mathcal{O}(h^p)$  on exponentially long time intervals if the integrator is of order  $p$  and the trajectory stays within a compact set of the phase space. This needs to be contrasted with the generic energy error behaviour of non-symplectic methods which is of order  $\mathcal{O}(th^p)$ . Modified and inverse modified vector fields can be computed as formal power series in the step size  $h$ . Although the formal power series  $\tilde{f}$  and  $\tilde{H}$  typically do not converge, they still govern the dynamics as optimal truncation results are available.<sup>1</sup> The situation is illustrated in Figure 3.

Let us review the Symplectic Euler method and the implicit midpoint method and their traditional backward error analysis formulas. These will not only be employed to analyse the quality of our numerical schemes but will also play a role when we identify the energy of dynamical systems using SSI. Moreover, we will relate the inverse modified Hamiltonians of the two schemes with generating functions of appropriate types.

The Symplectic Euler method is a 1st order symplectic integrator. Applied to a partitioned system of differential equations

$$\begin{pmatrix} \dot{q} \\ \dot{p} \end{pmatrix} = f(q, p), \quad q \in \mathbb{R}^n, p \in \mathbb{R}^n \quad (3)$$

with step size  $h$  it reads

$$\begin{pmatrix} \bar{q} \\ \bar{p} \end{pmatrix} = \begin{pmatrix} q \\ p \end{pmatrix} + hf(\bar{q}, p). \quad (4)$$

Of the  $2n$  equations, the first  $n$  equations define  $\bar{q}$  implicitly, while  $\bar{p}$  can be computed explicitly from the last  $n$  equations once  $\bar{q}$  is known. If  $f(q, p) = J^{-1}\nabla H(q, p)$ , then the first terms of  $\tilde{H}$  are given by<sup>1</sup>

$$\begin{aligned} \tilde{H} &= H - \frac{h}{2} H_q^\top H_p \\ &+ \frac{h^2}{12} \left( H_q^\top H_{pp} H_q + H_p^\top H_{qq} H_p + 4(H_p^\top H_{qp} H_q) \right) + \dots \end{aligned} \quad (5)$$

Here  $H_q$  denotes the gradient  $H_q = \left( \frac{\partial H}{\partial q^j} \right)_{j=1}^n$  and  $H_{qq} = \left( \frac{\partial^2 H}{\partial q^i \partial q^j} \right)_{i,j=1}^n$ .  $H_p$ ,  $H_{pp}$  and  $H_{qp}$  are defined analogously. The inverse modified Hamiltonian is the Hamiltonian  $\bar{H}$  such that  $\bar{H} = H$ . Its power series is easily found by plugging the ansatz  $\bar{H} = H + h\bar{H}^1 + h^2\bar{H}^2 + \mathcal{O}(h^3)$  into  $\tilde{H} = H$  and comparing powers of  $h$ . For the Symplectic Euler method we obtain

$$\begin{aligned} \bar{H} &= H + \frac{h}{2} H_q^\top H_p \\ &+ \frac{h^2}{6} \left( H_q^\top H_{pp} H_q + H_p^\top H_{qq} H_p + H_p^\top H_{qp} H_q \right) + \mathcal{O}(h^3). \end{aligned}$$

The power series  $\bar{H}$ ,  $H$ ,  $\tilde{H}$  are related as shown in Figure 2.

The Symplectic Euler method with step size  $h$  applied to  $\bar{f}(q, p) = J^{-1}\nabla \bar{H}(q, p)$  reads

$$\begin{aligned} \bar{q} &= q + h\bar{H}_p(q, \bar{p}) \\ \bar{p} &= p - h\bar{H}_q(q, \bar{p}). \end{aligned}$$

We observe that  $S_2(q, P) = q^\top P + h\bar{H}(q, P)$  is a generating function of type 2 for the exact flow map  $(\bar{q}, \bar{p}) = \phi_h(q, p)$ . Since  $\phi_h$  is close to the identity, the existence of  $\bar{H}$  is guaranteed<sup>1</sup> for sufficiently small  $h$ .

Another example of a symplectic integration method is the implicit midpoint rule, which is second order accurate. Applied to the differential equation  $\dot{z} = f(z)$  it reads

$$\bar{z} = z + hf\left(\frac{\bar{z} + z}{2}\right). \quad (6)$$

If  $f(z) = J^{-1}\nabla H(z)$  then the modified Hamiltonian is given as<sup>1</sup>

$$\tilde{H} = H - \frac{h^2}{24} f^\top \text{Hess}(H) f + \mathcal{O}(h^4) \quad (7)$$

and the inverse modified Hamiltonian as

$$\bar{H} = H + \frac{h^2}{24} f^\top \text{Hess}(H) f + \mathcal{O}(h^4). \quad (8)$$

Here  $\text{Hess}(H)$  denotes the Hessian matrix of  $H$ . Only even powers of  $h$  occur in  $\tilde{H}$  and  $\bar{H}$ .

As for the Symplectic Euler method, the inverse modified Hamiltonian to the implicit midpoint rule can be interpreted as a generating function expressed in the coordinates  $\frac{\bar{q}+q}{2}, \frac{\bar{p}+p}{2}$  for the exact flow map  $\phi_h$ . Its existence is guaranteed for sufficiently small  $h$ . Details can be found in Appendix A.

*Remark II.1.* B-series and P-series methods constitute a large class of numerical schemes which include Runge–Kutta methods and partitioned Runge–Kutta methods. Explicit formulas for inverse modified Hamiltonians for symplectic B- and P-series methods have been calculated and their expressions can be given using the theory of rooted trees (in the sense of graph theory)<sup>1</sup>. Using the relation  $\tilde{H} = H$ , we derive a recursion for the terms  $\bar{H} = \bar{H}_1 + h\bar{H}_2 + h^2\bar{H}_3 + \dots$  for consistent symplectic B- and P-series methods: in the notation of Theorem IX.9.8 in the book by Hairer et. al.<sup>1</sup> for a B series method we obtain

$$\bar{H}_1 = H, \quad \bar{H}_{k+1} = \sum_{j=2}^{k+1} \sum_{\tau \in T^*, |\tau|=j} \frac{b(\tau)}{\sigma(\tau)} \bar{H}_{k-j+2}(\tau)$$

for  $k = 0, 1, \dots$ . A corresponding formula for P series methods can be obtained analogously by inverting the traditional backward error analysis formula given in Theorem IX.10.9.<sup>1</sup>  $\triangle$

Applying the argumentation of traditional backward error analysis<sup>1</sup> to inverse modified equations, the existence of  $\bar{H}$  as a formal power series for any symplectic method was proved in the work by Zhu et. al.<sup>9</sup> For non-symplectic methods,  $\bar{H}$  does not exist.

In this paper, first the inverse modified Hamiltonian  $\bar{H}$  will be learned from data and then the vector field  $J^{-1}\nabla \bar{H}$  will be integrated using the corresponding symplectic integrator. System identification is then performed by computing  $\tilde{H}$  to obtain  $H$ , i.e. by applying the traditional backward error analysis formulas (5) or (7) to  $\bar{H}$ .

*Remark II.2.* Backward error analysis is known to describe the behaviour of numerical solutions well, not only for small but also for moderate to large time steps.<sup>1</sup> We can expect Symplectic Shadow Integration (SSI) to work well for those time step sizes for which backward error analysis techniques apply. In particular, we will be able to use discretisation parameters  $h$  of moderate size in the following numerical experiments.  $\triangle$

### III. LEARNING INVERSE MODIFIED HAMILTONIANS FROM TRAJECTORY DATA

With its existence established (as formal power series in the general case and explicitly proved for the Symplectic Euler method and the implicit midpoint rule by relating their inverse modified Hamiltonians to certain generating functions), we proceed to learning an inverse modified Hamiltonian  $\bar{H}$  from data using Gaussian process regression.

For an introduction to Gaussian process regression for machine learning see the book by Rasmussen and Williams<sup>19</sup>. Let us briefly recall some relevant aspects: in Gaussian process regression a function  $S: M \rightarrow \mathbb{R}$  is modelled as a sequence of random variables  $(\hat{S}(z))_{z \in M}$  over an index



set  $M$ . Each random vector  $(\hat{S}(z))_{z \in \hat{M}}$ , with  $\hat{M}$  a finite subset of  $M$ , is multivariate normally distributed with covariance matrix  $K(Z, Z) := (k(z, w))_{z, w \in \hat{M}}$  for a covariance or kernel function  $k: M \times M \rightarrow \mathbb{R}$  and mean  $m(Z) := (m(z))_{z \in \hat{M}}$  for a function  $m: M \rightarrow \mathbb{R}$ . For given finite data  $(Z', S(Z')) := \{(z', S(z'))\}_{z' \in M'}$ , mean function  $m$ , and kernel  $k$ , the posterior distribution of the random vector  $\hat{S}(Z)|(Z', S(Z')) := (\hat{S}(z))_{z \in \hat{M}}|(Z', S(Z'))$  is again multivariate normally distributed. Its mean is given as

$$\begin{aligned} \mathbb{E}[\hat{S}(Z)|(Z', S(Z'))] \\ = m(Z) + K(Z, Z')K(Z', Z')^{-1}(S(Z') - m(Z')), \end{aligned}$$

where  $K(Z, Z') = (k(z, z'))_{z \in \hat{M}, z' \in M'}$ . The covariance matrix is given as

$$\begin{aligned} \text{Cov}[\hat{S}(Z)|(Z', S(Z'))] \\ = K(Z, Z) - K(Z, Z')K(Z', Z')^{-1}K(Z', Z). \end{aligned}$$

The mean  $\mathbb{E}[\hat{S}(Z)|(Z', S(Z'))]$  can be used as a prediction of  $(S(z))_{z \in \hat{M}}$ , while the variance at a given point  $z$  can be interpreted as a measure of the model uncertainty in the prediction of  $S(z)$ .

A technique to learn the Hamiltonian  $H$  of a system from values of the Hamiltonian vector field using Gaussian processes was introduced by Bertalan et. al.<sup>8</sup> We modify the idea such that we can use data points  $(y, \phi_h(y))$  of the Hamiltonian flow  $\phi$  rather than of the Hamiltonian vector field. Therefore, our method applies when there are trajectory observations available but the underlying vector field is unknown. Moreover, instead of learning  $H$  we learn the inverse modified Hamiltonian  $\bar{H}$ . The exact Hamiltonian  $H$  is then computed in a post-processing step. The following technique directly extends to all kernel methods with sufficiently smooth kernels.

Let  $Z = (z_1, \dots, z_N)$  be  $N$  points in the phase space  $M$  for which we would like to predict the values  $\bar{H}(Z) = (\bar{H}(z_1), \dots, \bar{H}(z_N))$  of the inverse modified Hamiltonian corresponding to the Symplectic Euler method with step size  $h$ . The prediction is based on observed data of the flow which maps a collection of  $\tilde{N}$  points of the phase space  $Y = (y_1, \dots, y_{\tilde{N}}) = ((q_1, p_1), \dots, (q_{\tilde{N}}, p_{\tilde{N}}))$  to the collection of points  $\bar{Y} = (\bar{y}_1, \dots, \bar{y}_{\tilde{N}}) = ((\bar{q}_1, \bar{p}_1), \dots, (\bar{q}_{\tilde{N}}, \bar{p}_{\tilde{N}}))$  after time  $h$ . To apply Gaussian Process regression, the  $2n$ -dimensional phase space  $M$  is interpreted as an index set. To derive the method, first we assume that we already have the corresponding values  $\bar{H}(Z)$ . This will help us to derive a linear system of equations for  $\bar{H}(Z)$ .

Let  $\hat{H}$  be a Gaussian process with index set  $M$ , a continuously differentiable kernel  $k: M \times M \rightarrow \mathbb{R}$ , and a constant zero mean function  $m \equiv 0$ . Now  $\bar{H}$  can be predicted at a new point  $y \in M$  as the conditional expectation

$$\mathbb{E}[\hat{H}(y)|(Z, \bar{H}(Z))] = k(y, Z)^\top k(Z, Z)^{-1} \bar{H}(Z). \quad (9)$$

Here  $k(Z, Z) \in \mathbb{R}^{N \times N}$  is the covariance matrix  $k(Z, Z)_{i,j} = k(z_i, z_j)$  and  $k(y, Z) \in \mathbb{R}^{N \times 1}$  is given as  $k(y, Z) = k(y, z_j)$ . Differentiation of (9) with respect to  $y$  yields

$$\mathbb{E}[\hat{H}(y)|(Z, \bar{H}(Z))] = \nabla_1 k(y, Z)^\top k(Z, Z)^{-1} \bar{H}(Z) =: J \bar{f}(y) \quad (10)$$

with  $\nabla_1 k(y, Z) \in \mathbb{R}^{N \times 2n}$  given as  $(\nabla_1 k(y, Z))_{i,j} = \frac{\partial k}{\partial y_j} k(y, z_i)$ . Here  $J$  is the symplectic structure matrix from (1).

Let  $Y = (y_1, \dots, y_{\tilde{N}}) = ((q_1, p_1), \dots, (q_{\tilde{N}}, p_{\tilde{N}}))$  be a collection of  $\tilde{N}$  points of the phase space and let  $\bar{Y} = (\bar{y}_1, \dots, \bar{y}_{\tilde{N}}) = ((\bar{q}_1, \bar{p}_1), \dots, (\bar{q}_{\tilde{N}}, \bar{p}_{\tilde{N}}))$  denote the corresponding values of the Hamiltonian flow after time  $h$  (training data). Imposing that  $\bar{Y}$  was obtained from  $Y$  using the Symplectic Euler method on  $\bar{f}$  yields the relations

$$\begin{aligned} \begin{pmatrix} \bar{q}_j \\ \bar{p}_j \end{pmatrix} &= \begin{pmatrix} q_j \\ p_j \end{pmatrix} + h \bar{f}(\bar{q}_j, p_j) \\ &= \begin{pmatrix} q_j \\ p_j \end{pmatrix} + h J^{-1} \nabla_1 k((\bar{q}_j, p_j), Z)^\top k(Z, Z)^{-1} \bar{H}(Z) \end{aligned} \quad (11)$$

for  $j = 1, \dots, \tilde{N}$ . The Hamiltonian of a Hamiltonian system is defined up to an additive constant. We can, therefore, impose  $\mathbb{E}[\hat{H}(y_0)|(Z, \bar{H}(Z))] = \bar{H}_0$  for any  $\bar{H}_0 \in \mathbb{R}$  (normalisation). Together with (11) we obtain a linear system

$$\begin{pmatrix} \nabla_1 k((\bar{q}_1, p_1), Z)^\top k(Z, Z)^{-1} \\ \vdots \\ \nabla_1 k((\bar{q}_{\tilde{N}}, p_{\tilde{N}}), Z)^\top k(Z, Z)^{-1} \\ k(y_0, Z)^\top k(Z, Z)^{-1} \end{pmatrix} \begin{pmatrix} \bar{H}(z_1) \\ \vdots \\ \bar{H}(z_N) \end{pmatrix} = \frac{1}{h} \begin{pmatrix} J(\bar{y}_1 - y_1) \\ \vdots \\ J(\bar{y}_{\tilde{N}} - y_{\tilde{N}}) \\ \bar{H}_0 \end{pmatrix} \quad (12)$$

for  $\bar{H}(Z) = (\bar{H}(z_1), \dots, \bar{H}(z_N))^\top$ , which consists of  $n\tilde{N} + 1$  equations for  $N$  unknowns. The last equation corresponds to the normalisation.

We conclude that given data  $Y, \bar{Y}$  of the flow map of a dynamical system, e.g. obtained from observations of trajectories at times  $\tau$  and  $\tau + h$ , values of the inverse modified Hamiltonian  $\bar{H}(Z)$  can be predicted over points  $Z$  of the phase space by solving 12 in the least square sense. The solution has the following interpretation: if (12) is solved exactly, the collection  $(Z, \bar{H}(Z))$  has the property that if the mean of the Gaussian process  $\hat{H}|Z, \bar{H}(Z)$  is used to predict  $\bar{H}(Y)$  then an application of the Symplectic Euler method to  $J^{-1} \nabla \bar{H}(y_j)$  at  $y_j$  recovers  $\bar{y}_j$  for all  $j = 1, \dots, \tilde{N}$ .

#### IV. NUMERICAL EXPERIMENTS

We apply Symplectic Shadow Integration (SSI) with the Symplectic Euler method (SE) and midpoint rule (MP) to the mathematical pendulum and the Hénon–Heiles system. To learn the inverse modified Hamiltonian, we employ Gaussian Process regression with radial basis functions

$$k(x, y) = k_c \exp\left(-\frac{1}{\epsilon^2} \|x - y\|^2\right) \quad (13)$$

as kernels. When analysing the quality of our numerical results, we particularly focus on whether the phase portrait topology has been captured correctly such that long-term predictions yield qualitatively correct results. This is done by applying backward error analysis to compute the system whose exact flow coincides (up to a truncation error) with the SSI prediction, which is then compared to the exact system.

For comparison to other approaches, we fit a Gaussian process (GP) directly to the training data  $(Y, \bar{Y})$  using again radial basis functions as kernels, where the parameters  $k_c, e$  are fitted using marginal likelihood estimation. For this, we employ the Python package `scikit-learn`.<sup>20</sup> This corresponds to strategy 1 of Section IC (learning the flow map directly without using Hamiltonian structure). For comparison with strategy 2 type approaches, we apply the Symplectic Euler method or midpoint rule to the exact system with the same step size as used in the SSI scheme. This corresponds to strategy 2 of Section IC in the infinite data limit, where the Hamiltonian has been learned up to machine precision.

Source code can be found in our GitHub repository.<sup>21</sup>

### A. Mathematical Pendulum

We consider the Hamiltonian system (1) of the mathematical pendulum with Hamiltonian

$$H(q, p) = \frac{1}{2}p^2 + (1 - \cos(q)).$$

To generate training data, we use a Halton sequence  $Y = (y_1, \dots, y_N)$  on the phase space  $M = [-2\pi, 2\pi] \times [-1.2, 1.2]$  of length  $N = 160$ . Values  $\bar{Y} = (\bar{y}_1, \dots, \bar{y}_N)$  are obtained by integrating (1) up to time  $h = 0.3$  with high precision. For this we use  $n_{LP} = 800$  steps of the Störmer-Verlet scheme.<sup>1</sup> As a kernel for the GP we use radial basis functions (13) with parameters  $k_c = 1$  and  $e = 2$ . Setting  $Z = Y$  in (12) and the step size  $h = 0.3$  we obtain values  $\bar{H}(Z)$ . We can now compute values for  $\bar{H}(z)$  or  $\nabla \bar{H}(z)$  for any  $z \in M$  using (9) or (10), respectively. For numerical stability, a Tikhonov regularisation  $k(Z, Z) \mapsto k(Z, Z) + \sigma I_{nN}$  with  $\sigma = 10^{-13}$  is applied and SciPy's<sup>22</sup> Cholesky solver is used, wherever multiplication with the inverse covariance matrix  $k(Z, Z)^{-1}$  is required. Higher derivatives of  $\bar{H}$  are obtained using derivatives of (10).

Figure 4 compares the behaviour of a trajectory obtained using the Symplectic Euler method (SE), when SE is used with SSI and when SE is applied to the exact system  $\dot{z} = J^{-1}\nabla H(z)$  directly, which corresponds to strategy 2 in the infinite data limit. We see that the SSI prediction conserves the exact energy  $H$  much better than the strategy 2 prediction. Indeed, the SSI trajectory is visually indistinguishable from a reference solution. Moreover, if the experiment is repeated with an increased dataset of  $N = 700$  training points obtained from a Halton sequence, the amplitude of the energy error oscillation further decreases and looks like a random walk which does not leave a band of width  $\approx 4 \cdot 10^{-7}$  during the simulation time. For strategy 2 to have such a small energy error on the given time-interval, we need to decrease the step-sizes by a factor of 10000 since SE is just of first order. This demonstrates that SSI successfully compensates the discretisation error introduced by SE and that rather large step size  $h = 0.3$  can be used while maintaining high accuracy.

In the following, we continue with the smaller data set  $N = 160$  and compare to strategy 1, i.e. learning the Hamiltonian flow map directly by fitting a Gaussian process to the training data  $(Y, \bar{Y})$ . Figure 5 shows that the learned flow is

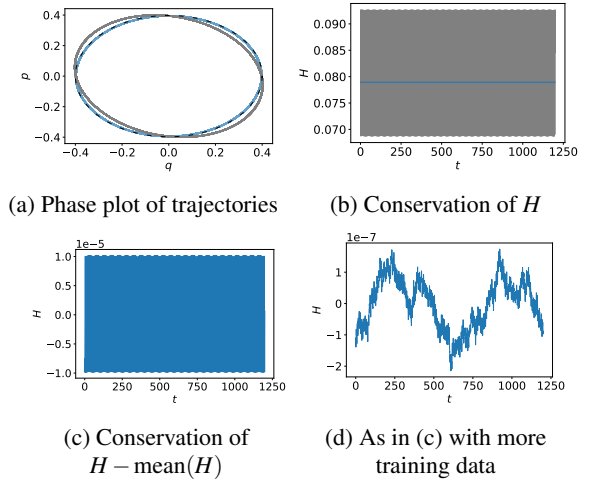


FIG. 4: Pendulum experiment with SE, step size  $h = 0.3$ , initialisation at  $z_0 = (0.4, 0)$ . The SSI approach is shown in blue, a direct application of SE in grey. Figure (a) shows a phase plot of both trajectories as well as a reference solution (dashed) which is exactly covered by the SSI trajectory. The Hamiltonian  $H$  is plotted along the SSI and SE trajectory in (b). The amplitude of the blue oscillation is significantly smaller than of the grey oscillation. (c) depicts the energy of the SSI motion, where the arithmetic mean of  $H$  along the trajectory was subtracted. When the experiment is repeated with an abundance of training data ( $N = 700$ ) the energy error looks like a random walk (d).




FIG. 5: Pendulum experiment. The energy conservation properties of SSI are compared to a GP fitted to the flow map data without incorporation of the Hamiltonian structure. Its trajectory initialised at  $z_0 = (0.4, 0)$  (green) spirals outwards and its energy erroneously grows approximately linearly, while the trajectory from SSI (blue) is periodic and has excellent energy behaviour even on long time intervals.

not energy conserving, which is not surprising as we have not explicitly incorporated Hamiltonian structure. The trajectory starting from  $z_0 = (0.4, 0)$  has a steady energy growth and does not capture the periodicity. This demonstrates that structure preservation is crucial in this example.

To analyse qualitative aspects of the SSI scheme, we compute the systems whose exact dynamics coincides with SSI predictions. For this, we apply the backward error analysis formula (5) to the learned inverse modified Hamiltonian  $\bar{H}$  and obtain a power series  $\tilde{\bar{H}}$ . Several truncations are evalu-




FIG. 6: Pendulum experiment (SE). (a) Level sets of  $\tilde{H}^{[2]}$  almost coincides with level sets of  $H$  (dashed). (b) and (c) show that the truncations of  $\tilde{H}$  to 0th (blue), 1st (orange), and 2nd (green) order evaluated along the blue SSI trajectory of Figure 4 (a) are increasingly better preserved. This justifies that a Hamiltonian system for the truncation  $\tilde{H}^{[2]}$  describes the numerical motions of the SSI.

ated along the trajectory in Figure 6. As the oscillations of the second truncation  $\tilde{H}^{[2]}$  occur within a band of width  $4 \cdot 10^{-4}$ , a truncation to 2nd order describes the Hamiltonian that governs the numerical dynamics sufficiently well for our purposes. Indeed, a contour plot of  $\tilde{H}^{[2]}$  on  $M$  is visually indistinguishable from a contour plot of the exact  $H$ . Here we have used a uniform mesh with  $120 \times 120$  points. A Hamiltonian is only defined up to a constant by the system's motion. Rather than measuring the  $L^2$  distance of  $H$  and  $\tilde{H}^{[2]}$  over  $M$ , we compute the standard deviation of  $H_{\text{diff}} = H - \tilde{H}^{[2]}$  for a uniform distribution on the training domain  $M$ , i.e.

$$\sigma(H_{\text{diff}}) = \sqrt{\mathbb{E}[(H_{\text{diff}} - \mathbb{E}[H_{\text{diff}}])^2]}, \text{ with } \mathbb{E}[f] = \frac{\int_M f dv}{dv(M)}$$

and with  $dv$  denoting the Lebesgue measure on  $M$ . Approximating  $M$  using a  $120 \times 120$  uniform mesh, we obtain  $\sigma(H_{\text{diff}}) \leq 5.2 \cdot 10^{-4}$ . This explains the excellent energy conservation of SSI motions and shows that qualitative aspects such as periodic motions and even the behaviour close to the separatrix are guaranteed, since the phase portrait of the system governing the SSI motions is close to the exact phase portrait. This is in particular important for long-term simulations. Furthermore, this demonstrates that SSI can be used for system identification.

The experiments are repeated with the second order accurate Implicit Midpoint rule (MP) using the same parameters but  $N = 400$  data points in the training process. Phase plots of the trajectory computed with SSI and with a direct application of MP to the exact system  $\dot{z} = J^{-1} \nabla H(z)$  (strategy 2 in the infinite data limit) both visually coincide with an exact




FIG. 7: Pendulum experiment (MP). The energy behaviour of MP (gray) applied to the exact system is oscillatory while the SSI approach shows energy conservation up to a random walk (blue). In all plots, the mean value of the energy trajectory has been subtracted before plotting.




FIG. 8: Pendulum experiment (MP). Repetition of the experiment of Figure 6 with MP instead of SE. Notice that the 0th and 1st as well as the 2nd and 3rd truncation of  $\tilde{H}$  coincide due to the symmetry of MP.

trajectory (not shown). However, while the energy error of the strategy 2 approach is oscillatory with amplitudes bounded by  $3 \cdot 10^{-5}$ , the energy error of the SSI solution shows a random walk within a band of width  $4 \cdot 10^{-7}$  on long time scales (Figure 7). For a similar energy error on the given time-interval we would need to decrease the step size in the strategy 2 approach by a factor of approximately 9. Moreover, the error shows a linear trend such that the step size needs to be decreased further when simulation time is increased (not shown).

Similarly as in the experiments with SE, the truncation  $\tilde{H}^{[2]}$  computed with the backward error analysis formula (8) describes the numerical dynamics accurately as it is up to small errors a conserved quantity of the numerical flow (Figure 8). Notice that the 0th and 1st order truncation and the 2nd and 3rd order truncations of  $\tilde{H}$  coincide since the midpoint rule is symmetric. The standard deviation of  $H_{\text{diff}} = H - \tilde{H}^{[2]}$  (computed as before) is smaller than  $\sigma(H_{\text{diff}}) \leq 9.4 \cdot 10^{-4}$ . Again, this shows that SSI replicates qualitative aspects of the exact system even in long term simulations.

We see that using the Symplectic Euler scheme or the Implicit midpoint rule in the SSI technique yields similar behaviour and there appears to be no gain from using a higher order method for SSI.




FIG. 9: Contour plot of the exact Hénon–Heiles potential  $V$  (dashed) and the potential  $\tilde{V}^{[2]}(q) := \tilde{H}^{[2]}(q, (0,0))$  computed from the learned inverse modified Hamiltonian  $\tilde{H}$ . The level sets of  $\tilde{V}^{[2]}$  and  $V$  are not distinguishable in this plot.

## B. Hénon–Heiles system

The Hénon–Heiles system is a Hamiltonian system with

$$H(q, p) = \frac{1}{2} \|p\|^2 + V(q), \quad V(q) = \frac{1}{2} \|q\|^2 + \mu \left( q_1^2 q_2 - \frac{q_2^3}{3} \right).$$

We set the parameter to  $\mu = 0.8$ . A contour plot of  $V$  is shown in Figure 9. The system has bounded as well as unbounded motions. More precisely, bounded connected components of level sets of  $V$  and  $H$  correspond to values of  $V$  or  $H$  within the interval  $I = [0, \frac{1}{6\mu^2})$ . All connected components of level sets to values in  $\mathbb{R} \setminus I$  are unbounded. Accurate preservation of energy is crucial, when integrating trajectories, since trajectories on close-by energy level sets can show very different long-term behaviour. Classically, this forces one to use very small time steps in numerical integration schemes. However, as SSI compensates discretisation errors, we can use a moderate time-step. We repeat the experiments from Section IV A with  $h = 0.1$ ,  $N = 800$ ,  $n_{LP} = 600$ , kernel parameter  $e = 5$ , and the four dimensional tesseract  $M = [-1, 1]^4$  as domain for training data.

The exact motion initialised at  $z_0 = (q_0, p_0) = ((0.675499, 0.08), (0, 0))$  lies on a bounded connected component of its energy level set  $H(z_0)$ . While the trajectory predicted by SSI based on SE correctly captures this behaviour (see Figure 10), a direct application of SE to the exact system  $\dot{z} = J^{-1} \nabla H(z)$  with the same step size gives a trajectory which is unbounded. The reason is that  $H(z_0)$  is very close to the critical value  $\frac{1}{6\mu^2}$ . Indeed  $|H(z_0) - \frac{1}{6\mu^2}| \approx 6.2 \cdot 10^{-7}$ . SE conserves energy only up to an oscillation within a band of width  $\approx 0.03$  such that the numerical trajectory repeatedly jumps to unbounded energy level sets and eventually escapes (Figure 11). In contrast, the trajectory obtained using SSI has a significantly better energy behaviour looking like a random walk, which is bounded within a band of width  $\approx 2 \cdot 10^{-5}$  until time  $T = 5 \cdot 10^4$ . While this does not exclude that the trajectory escapes from an energy perspective, it makes the escape less likely. Indeed, no escape occurs during the simulation time. The step size would need to be decreased by a factor of at least 10000 to obtain comparable energy conservation properties with the strategy 2 approach.

This demonstrates that compensating discretisation errors




FIG. 10: Hénon–Heiles experiment. Phase plot of trajectory from SSI initialised at  $z_0$  computed up to time  $T = 800$  and  $T = 40000$ . The motion (blue) densely fills the area bounded by the level set  $V^{-1}(q_0)$  (black).

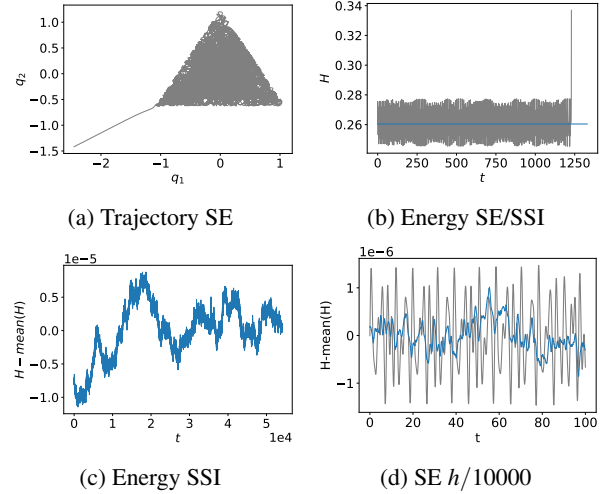


FIG. 11: Hénon–Heiles experiment. The trajectory obtained from an application of SE to the exact differential equation  $\dot{z} = J^{-1} \nabla H(z)$  (grey) erroneously leaves its energy level set and diverges. (a) shows the trajectory up to time  $t = 1229.5$ . (b) compares energy conservation of SSI (blue) with SE (grey). (c) shows the random walk like long-term energy behaviour of SSI. (d) compares energy conservation of SSI and SE on a smaller time interval, where SE's time-step was decreased to  $h/10000$ .

using the SSI framework can be highly beneficial, when accurate energy preservation is important.

For comparison with strategy 1, we learn the flow map of the system directly rather than the inverse modified Hamiltonian  $\tilde{H}$ . This is done by fitting a GP directly to the training data  $(Y, \tilde{Y})$ , again using radial basis functions as kernels (13) and fitting the hyperparameters  $k_e, e$  using marginal likelihood estimation. Figure 12 shows a trajectory initialised at  $z_0 = (q_0, (0, 0))$ . It fails to densely fill the area bounded by  $V^{-1}(q_0)$ . Instead, it monotonously loses energy and artificially converges to the origin. Moreover, access to the learned flow map does not directly reveal information on the Hamiltonian structure of the system, which could be used for system identification.

To analyse the behaviour of numerical motions predicted with SSI and to demonstrate SSI's potential for system identification, we apply the backward error analysis formula (5) to



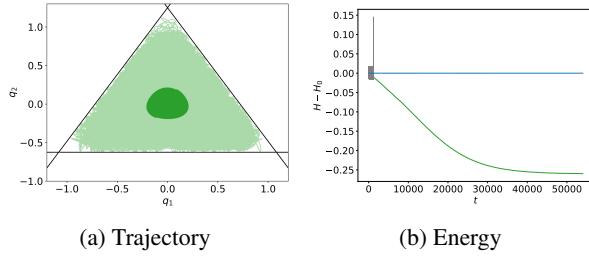


FIG. 12: A motion predicted from a GP fitted to the flow map data  $(Y, \bar{Y})$  converges to an artificial attractor and fails to densely fill the area bounded by  $V^{-1}(q_0)$ . (a) shows a plot of the trajectory in light green up to time  $T_1 = 30000$  and in darker green from time  $T_1$  to  $T = 54000$ . (b) compares the energy behaviour of the trajectory from (a) (green), from SSI (blue), and from the direct SE integration (Figure 11 a) (grey). Only SSI shows the correct long-term energy behaviour.

the learned Hamiltonian  $\bar{H}$  to obtain the formal power series  $\tilde{H}$ . Truncations of  $\tilde{H}$  now describe up to the truncation error the Hamiltonian of a dynamical system which governs the numerical motion. Since the truncation to second order  $\tilde{H}^{[2]}$  is sufficiently well preserved for our purposes (Figure 13), we use  $\tilde{H}^{[2]}$  to recover the potential  $V$  as  $\tilde{V}^{[2]}(q) := \tilde{H}^{[2]}(q, (0, 0))$ . Contour plots of the exact and recovered potential are visually indistinguishable (Figure 9). Notice that also the separatrix is reproduced correctly, which is related to the favourable energy conservation properties of SSI and explains the excellent long term behaviour for motions initialised close to the separatrix. To quantify the difference between the numerical and exact phase portrait, we consider a uniform distribution on the training domain  $M$  and compute the standard deviation  $\sigma = \sigma(H - \tilde{H}^{[2]})$  approximated using a uniform  $20 \times 20 \times 20$  grid. We obtain  $\sigma < 7 \cdot 10^{-4}$ . This confirms that the numerical and exact phase portrait are close on  $M$ , which guarantees qualitatively correct long-term behaviour.

## V. FUTURE WORK

Symplectic Shadow Integration is compatible with techniques to incorporate symmetries into learning processes.<sup>10</sup> If symmetry groups of the system are known, equivariant kernels adapted to the system's symmetries can be used to model  $\bar{H}$ . If the symmetries are symplectic and preserved by the integrator, the numerical flow will share the symmetry group of the exact flow. This is especially valuable when the action of the symmetry group is symplectic. In such cases symmetries relate to conserved quantities by Noether's theorem,<sup>23</sup> which are then picked up by the numerical flow. This leads to excellent preservation of topological properties of the phase portrait and can be used, for instance, to conserve complete integrability.<sup>1</sup>

In Section III the inverse modified Hamiltonian  $\bar{H}$  is

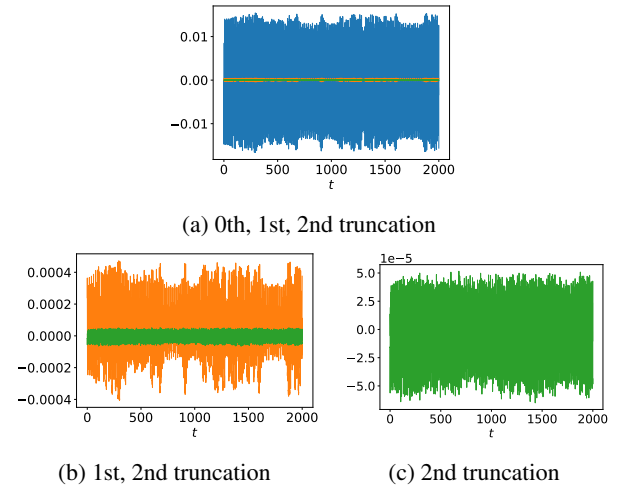


FIG. 13: Hénon–Heiles experiment. Truncations of the formal power series  $\tilde{H}$  are increasingly better conserved along the motion predicted by SSI. This justifies using a system with Hamiltonian  $\tilde{H}^{[2]}$  to analyse the dynamical system that describes the numerical motions obtained with SSI. (Arithmetic means are subtracted before plotting the energy behaviour.)

obtained as the mean function of a Gaussian process  $\hat{H}$ . The technique can easily be extended to general kernel based method with continuously differentiable kernels. For the system identification part, higher differentiability may be required. When Gaussian Processes are used, information about uncertainty of  $\hat{H}$  could be used to predict uncertainty of trajectories computed with SSI. Moreover, SSI can be extended to non-canonical Hamiltonian systems, if a symplectic integration method is known for the system type, which is the case for spin systems,<sup>24</sup> for instance. Furthermore, building on variational backward error analysis techniques,<sup>25</sup> the authors plan to develop a variational version of SSI which learns inverse modified Lagrangian functions that compensate discretisation errors of variational integrators.

For a summary of the findings, also see our poster.<sup>26</sup>

## DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available from the repository Christian-Offen/symplectic-shadow-integration<sup>21</sup> (<https://github.com/Christian-Offen/symplectic-shadow-integration>, <https://dx.doi.org/10.5281/zenodo.5159766>).

## Appendix A: Relation of inverse modified Hamiltonians to generating functions for the implicit midpoint rule

*Lemma.* Consider the application of the implicit midpoint rule to a Hamiltonian system with step size  $h$ . Denote its inverse modified Hamiltonian by  $\bar{H}$ . Around any point in the phase space, there exists a generating function  $S$  of the exact flow map  $\phi_h$  such that the scaled inverse modified Hamiltonian  $h\bar{H}$  fulfils the same partial differential equations as  $S$ .

*Proof.* The graph  $\Gamma = \{(q, p, \bar{q}, \bar{p}) | (\bar{q}, \bar{p}) = \phi_h(q, p)\}$  is an embedded Lagrangian submanifold in a symplectic manifold with symplectic structure  $\Omega = \sum_{j=1}^n (d\bar{q}^j \wedge d\bar{p}^j - dq^j \wedge dp^j)$ . The following 1-form  $\alpha$  is a primitive of  $\Omega$ , i.e.  $d\alpha = \Omega$ :

$$\alpha = \sum_{j=1}^n \left( (\bar{q}^j - q^j) d\left(\frac{\bar{p}^j + p^j}{2}\right) - (\bar{p}^j - p^j) d\left(\frac{\bar{q}^j + q^j}{2}\right) \right).$$

The (pullback of the) 1-form  $\alpha$  is closed on  $\Gamma$ , since  $\Gamma$  is a Lagrangian submanifold. Therefore, primitives  $S$  with  $dS = \alpha$  exist locally on  $\Gamma$ . For small step sizes,  $\frac{\bar{q}+q}{2}, \frac{\bar{p}+p}{2}$  constitutes a local coordinate system on  $\Gamma$ . Writing

$$\begin{aligned} dS &= \sum_{j=1}^n \frac{\partial S}{\partial \left(\frac{\bar{q}^j + q^j}{2}\right)} \left(\frac{\bar{q}^j + q^j}{2}, \frac{\bar{p}^j + p^j}{2}\right) d\left(\frac{\bar{q}^j + q^j}{2}\right) \\ &\quad + \sum_{j=1}^n \frac{\partial S}{\partial \left(\frac{\bar{p}^j + p^j}{2}\right)} \left(\frac{\bar{q}^j + q^j}{2}, \frac{\bar{p}^j + p^j}{2}\right) d\left(\frac{\bar{p}^j + p^j}{2}\right) \end{aligned}$$

and comparing the coefficients of the total differentials of  $\alpha$  and  $dS$  yield equations, which after substituting  $h\bar{H}$  for  $S$  read

$$\begin{aligned} \bar{q} &= q + h\nabla_p \bar{H} \left( \frac{\bar{q} + q}{2}, \frac{\bar{p} + p}{2} \right) \\ \bar{p} &= p - h\nabla_q \bar{H} \left( \frac{\bar{q} + q}{2}, \frac{\bar{p} + p}{2} \right). \end{aligned} \tag{A1}$$

Here,  $\nabla_q \bar{H}$  and  $\nabla_p \bar{H}$  denote derivatives of  $\bar{H}$  with respect to the first input argument or second input argument, respectively. The system (A1) coincides with the implicit midpoint rule applied to  $\bar{f} = J^{-1} \nabla \bar{H}$ .  $\square$

<sup>1</sup>E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, Springer Series in Computational Mathematics (Springer Berlin Heidelberg, 2013).

<sup>2</sup>M. Betancourt, “A conceptual introduction to Hamiltonian Monte Carlo,” (2018), arXiv:1701.02434 [stat.ME].

<sup>3</sup>R. I. McLachlan and C. Offen, “Preservation of bifurcations of Hamiltonian boundary value problems under discretisation,” *Foundations of Computational Mathematics (FoCM)* **20**, 1363–1400 (2020).

<sup>4</sup>R. I. McLachlan and C. Offen, “Bifurcations of solutions to Hamiltonian boundary value problems,” *Nonlinearity* **31**, 2895–2927 (2018).

<sup>5</sup>C. Offen, *Analysis of Hamiltonian boundary value problems and symplectic integration (Doctoral Thesis)* (Massey University, 2020).

<sup>6</sup>L. M. Kreusser, R. I. McLachlan, and C. Offen, “Detection of high codimensional bifurcations in variational PDEs,” *Nonlinearity* **33**, 2335–2363 (2020).

<sup>7</sup>C. Offen and S. Ober-Blöbaum, “Bifurcation preserving discretisations of optimal control problems,” (2021), arXiv:2107.13853 [math.OC].

<sup>8</sup>T. Bertalan, F. Dietrich, I. Mezić, and I. G. Kevrekidis, “On learning Hamiltonian systems from data,” *Chaos: An Interdisciplinary Journal of Nonlinear Science* **29**, 121107 (2019).

<sup>9</sup>A. Zhu, P. Jin, and Y. Tang, “Inverse modified differential equations for discovery of dynamics,” (2020), arXiv:2009.01058 [math.NA].

<sup>10</sup>S. Ridderbusch, C. Offen, S. Ober-Blöbaum, and P. Goulart, “Learning ODE models with qualitative structure using Gaussian processes,” (2021), accepted contribution to 60th IEEE Conference on Decision and Control 2021, arXiv:2011.05364 [cs.LG].

<sup>11</sup>P. Jin, Z. Zhang, A. Zhu, Y. Tang, and G. E. Karniadakis, “Sympnets: Intrinsic structure-preserving symplectic networks for identifying hamiltonian systems,” *Neural Networks* **132**, 166–179 (2020).

<sup>12</sup>M. David and F. Méhats, “Symplectic learning for Hamiltonian Neural Networks,” (2021), arXiv:2106.11753 [cs.LG].

<sup>13</sup>K. Rath, C. G. Albert, B. Bischl, and U. von Toussaint, “Symplectic Gaussian process regression of maps in Hamiltonian systems,” *Chaos: An Interdisciplinary Journal of Nonlinear Science* **31**, 053121 (2021).

<sup>14</sup>R. Chen and M. Tao, “Data-driven prediction of general Hamiltonian dynamics via learning exactly-symplectic maps,” in *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, Proceedings of Machine Learning Research, Vol. 139, edited by M. Meila and T. Zhang (PMLR, 2021) pp. 1717–1727.

<sup>15</sup>S. Greidanus, M. Dzamba, and J. Yosinski, “Hamiltonian Neural Networks,” in *Advances in Neural Information Processing Systems*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019).

<sup>16</sup>Y. D. Zhong, B. Dey, and A. Chakraborty, “Symplectic ODE-Net: Learning Hamiltonian dynamics with control,” (2020), arXiv:1909.12077 [cs.LG].

<sup>17</sup>S. Park, W. Kwak, and H. K. Lee, “Accelerated spin dynamics using deep learning corrections,” *Scientific Reports* **10** (2020), 10.1038/s41598-020-70558-1.

<sup>18</sup>B. Leimkuhler and S. Reich, *Simulating Hamiltonian Dynamics* (Cambridge University Press, 2005).

<sup>19</sup>C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 2005).

<sup>20</sup>F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research* **12**, 2825–2830 (2011).

<sup>21</sup>C. Offen, “Software: GitHub repository @ Christian-Offen/symplectic-shadow-integration,” (2021).

<sup>22</sup>P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods* **17**, 261–272 (2020).

<sup>23</sup>E. L. Mansfield, “Variational problems with symmetry,” in *A Practical Guide to the Invariant Calculus*, Cambridge Monographs on Applied and Computational Mathematics (Cambridge University Press, 2010) pp. 206–240.

<sup>24</sup>R. I. McLachlan, K. Modin, and O. Verdier, “Symplectic integrators for spin systems,” *Physical Review E* **89** (2014), 10.1103/physreve.89.061301.

<sup>25</sup>M. Vermeeren, “Modified equations for variational integrators,” *Numerische Mathematik* **137**, 1001–1037 (2017).

<sup>26</sup>C. Offen and S. Ober-Blöbaum, “Symplectic integration of learned Hamiltonian systems, poster presented at the workshop Theory of deep learning (9 August 2021 to 13 August 2021), Isaac Newton Institute, University of Cambridge, Cambridge, UK,” (2021).