Variational Learning of Euler–Lagrange Dynamics from Data

Sina Ober-Blöbaum, Christian Offen*

ªPaderborn University, Warburger Str.100, Paderborn, 33098, Germany

Abstract

The principle of least action is one of the most fundamental physical principle. It says that among all possible motions connecting two points in a phase space, the system will exhibit those motions which extremise an action functional. Many qualitative features of dynamical systems, such as the presence of conservation laws and energy balance equations, are related to the existence of an action functional. Incorporating variational structure into learning algorithms for dynamical systems is, therefore, crucial in order to make sure that the learned model shares important features with the exact physical system. In this paper we show how to incorporate variational principles into trajectory predictions of learned dynamical systems. The novelty of this work is that (1) our technique relies only on discrete position data of observed trajectories. Velocities or conjugate momenta do *not* need to be observed or approximated and *no* prior knowledge about the form of the variational principle is assumed. Instead, they are recovered using backward error analysis. (2) Moreover, our technique compensates discretisation errors when trajectories are computed from the learned system. This is important when moderate to large step-sizes are used and high accuracy is required. For this, we introduce and rigorously analyse the concept of inverse modified Lagrangians by developing an inverse version of variational backward error analysis. (3) Finally, we introduce a method to perform system identification from position observations only, based on variational backward error analysis.

Keywords: Lagrangian learning, variational backward error analysis, modified Lagrangian, variational integrators, physics informed learning
2020 MSC: 65P10
2020 MSC: 93B30
2020 MSC: 68T05

1. Introduction

Physics informed learning has sparked extensive interest since embedding prior knowledge about physical laws can greatly improve the quality of predictions of dynamics based on observation data. One of the most fundamental principles is Hamilton's principle or the principle of least action. Among all curves connecting two points in a configuration space only those correspond to physical motions which extremise an action functional. The existence of an action functional induces some subtle structure on the differential equations governing the motions, which we will refer to as *variational structure*. When trying to identify a system from data, it, therefore, appears to be advantageous to learn the action rather than the differential equations which govern the motions. This does not only reduce data requirements since only a scalar valued map rather than the right hand side of a differential equation needs to be learned, but, more importantly, it guarantees that the learned system shares qualitative aspects with physical systems. For instance, autonomous variational systems conserve energy. When an autonomous physical system is learned and the learned system is variational, then the learned system is conservative as well.

Lagrangian Shadow Integration (LSI), introduced in this paper, provides a novel framework to predict motions of dynamical systems from data while incorporating variational structure. The data consists of discrete observations

---

*corresponding author [iD]

Email address: christian.offen@uni-paderborn.de (Christian Offen)
URL: https://www.uni-paderborn.de/en/person/85279 (Christian Offen), ORCiD: 0000-0002-5940-8057 (Christian Offen)

of trajectories (snapshot data). We assume that only position data is available, in contrast to other approaches such as Lagrangian Neural Networks (LNN) [1], which assume the availability of velocity and acceleration information. We incorporate variational structure as prior knowledge without making any assumptions on the specific form of the variational principle, in contrast to [2], for instance, which assumes a mechanical form. The novelty of our approach lies in the fact that we do not learn the exact variational structure but an inverse modified structure. The inverse modified variational principle has the property that when integrated using a variational integrator *no* discretisation error occurs. We prove the existence of inverse modified variational structure (up to truncation error) and explain the compensation of discretisation errors. Our analysis heavily makes use of *variational backward error analysis*, which was developed in [3]. Moreover, we show how to use Lagrangian Shadow Integration to identify structure and conserved quantities of dynamical systems based on observations of position data only, without any requirements of further knowledge about the system's structure.

In the proceeding of this section we will briefly review classical theory on Lagrangian dynamics and variational integrators, introduce the idea behind Lagrangian Shadow Integration, and contrast the technique to other approaches. We will restrict ourselves to autonomous systems in this work. Section 2 contains a more detailed explanation of the Lagrangian Shadow Integration framework. While Lagrangian Shadow Integration can be combined with many machine learning techniques, we show how to realise the technique using kernel based methods. Section 3 contains numerical examples. Finally, Section 4 provides rigorous theoretical underpinning of our technique and connects the framework with highly developed analysis techniques from Geometric Integration.

1.1. Lagrangian dynamics and variational integrators

In many dynamical systems motions $q \colon [t_0, t_N] \to Q$ connecting two points q_0, q_N of a state space $Q = \mathbb{R}^n$ are necessarily stationary points of an action functional

$$S(q) = \int_{t_0}^{t_N} L(q(t), \dot{q}(t)) \mathrm{d}t, \tag{1}$$

i.e. for any $v \colon [t_0, t_N] \to \mathbb{R}^n$ with $v(t_0) = 0 = v(t_N)$ we have

$$\lim_{\epsilon \to 0} \frac{S(q + \epsilon v) - S(q)}{\epsilon} = 0.$$

Equivalently, motions q fulfil the Euler–Lagrange equations

$$0 = \mathcal{E}(L) = \frac{\partial L}{\partial q}(q, \dot{q}) - \frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \tag{2}$$

subject to the boundary conditions $q(t_0) = q_0$, $q(t_N) = q_N$. L is called the Lagrangian of the dynamical system. In line with the examples presented in this article, we have restricted ourselves to the autonomous case. The variational structure governing the motions is related to structural properties of the dynamical system. For instance, motions preserve the Hamiltonian

$$H(q, \dot{q}) = \sum_{j=0}^{n} \dot{q}^j \frac{\partial L}{\partial \dot{q}^j} - L(q, \dot{q}), \tag{3}$$

which in classical mechanics corresponds to energy preservation in many cases. Moreover, Noether's theorem relates variational symmetries of L to further conserved quantities: for instance, rotational invariance of L yields conservation of angular momenta.

To guarantee that numerical solutions of (2) inherit physical properties of the exact system related to variational structure, rather than applying a numerical integrator to (2) directly, one discretises the action functional and looks for its stationary points. The discrete dynamics are governed by a discrete variational principle and a discrete version of Noether's theorem exists to derive discrete conservation laws [4]. Integrators derived in such a way are called *variational integrators*. More precisely, if $t_0, t_1, \ldots, t_N$ is a uniform grid with step-size h of the interval $[t_0, t_N]$ then S in (1) is replaced by a discrete action of the form

$$S_\Delta((q_j)_{j=0}^N) = \sum_{j=0}^{N-1} L_\Delta(q_j, q_{j+1}), \tag{4}$$

2

where the discrete Lagrangian $L_\Delta(a, b)$ is an approximation of the *exact discrete Lagrangian* $\int_0^h L(q(t), \dot{q}(t))$, where q solves[1] (2) subject to the boundary conditions $q(0) = a$, $q(h) = b$. Critical points $\nabla S_\Delta((q_j)_{j=0}^N) = 0$ of the discrete action (4) fulfil the discrete Euler–Lagrange equations

$$\nabla_2 L_\Delta(q_{j-1}, q_j) + \nabla_1 L_\Delta(q_j, q_{j+1}) = 0, \qquad j = 1, \ldots, N - 1 \tag{5}$$

on inner grid points. Here $\nabla_1 L_\Delta$ and $\nabla_2 L_\Delta$ denote the derivative with respect to the first or second input argument of L_Δ, respectively.

Example 1.1. *Popular discretisation methods are the midpoint rule*

$$L_\Delta(q_0, q_1) = hL\left(\frac{q_0 + q_1}{2}, \frac{q_1 - q_0}{h}\right)$$

or the trapezoidal rule

$$L_\Delta(q_0, q_1) = \frac{h}{2}L\left(q_0, \frac{q_1 - q_0}{h}\right) + \frac{h}{2}L\left(q_1, \frac{q_1 - q_0}{h}\right).$$

Both yield second order schemes.

To compute an approximation to a solution to an initial value problem $q(t_0) = q_0$, $\dot{q}(t_0) = \dot{q}_0$ for (2), variational integration proceeds as follows.

- A conjugate momentum $p_0 = \frac{\partial L}{\partial \dot{q}}(q_0, \dot{q}_0)$ is computed (Legendre transform).

- q_1 is approximated by solving the discrete conjugate momenta equation $-p_0 = \nabla_1 L_\Delta(q_0, q_1)$ for q_1.

- The discrete Euler–Lagrange equations (5) is used to compute $q_2, q_3, \ldots$ successively.

- If of interest, discrete conjugate momenta can be computed using

$$p_j = -\nabla_1 L_\Delta(q_j, q_{j+1}) \quad \text{or} \quad p_{j+1} = \nabla_2 L_\Delta(q_j, q_{j+1})$$

and velocities can be computed by solving $p_j = \frac{\partial L}{\partial \dot{q}}(q_j, \dot{q}_j)$ for $\dot{q}_j$ in a post-processing step.

1.2. Lagrangian Shadow Integration

Lagrangian Shadow Integration provides a framework to predict motions of unknown dynamical systems with variational structure form trajectory observations. It consists of a step in which the underlying continuous dynamical system is identified using machine learning techniques and another step in which the system is integrated numerically. Both steps make use of the system's variational structure. We assume that discrete snapshot data of trajectories to a time step h are available. Lagrangian Shadow Integration proceeds as follows.

1. Choose a variational integrator.
2. Learn an inverse modified Lagrangian L_{invmod} of the system. This is not the exact Lagrangian L of the dynamical system but an adapted (shadow) version of L which is tailored to the variational integration scheme.
3. Apply the variational integration technique to L_{invmod} to compute trajectories and, if needed, conjugate momenta.
4. Bonus. Identify the exact Lagrangian L and energy H using backward error analysis techniques from the learned L_{invmod}. Compute velocities to the trajectory data of step (3).

---

[1]A unique solution exists if $h > 0$ is sufficiently small and a and b sufficiently close to one another.

The inverse modified Lagrangian L_{invmod} can be learned directly from the observed snapshot data. Moreover, it has the property that in step (3) the discretisation error of the variational integrator is compensated, which yields excellent conservation properties. The concept of inverse modified Lagrangians L_{invmod}, introduced in this article, can be thought of as an inverse version of the concept of modified Lagrangians L_{mod} [3], which will be reviewed in the next section. In short, a *modified Lagrangian* governs the numerical solutions obtained by a variational integration scheme applied to a system defined by a Lagrangian L (the exact system). An *inverse modified Lagrangian* is a continuous Lagrangian L_{invmod} such that a variational integration scheme applied to L_{invmod} will produce motions of the system defined by L (the exact system).

Lagrangian Shadow Integration can be contrasted to other approaches of predicting the motions of dynamical systems from data.

- An incorporation of variational structure in prediction model guarantees that structural properties of the dynamics are preserved such that predicted motions show an excellent energy behaviour, for instance. Moreover, Lagrangian Shadow Integration could be combined with learning techniques which guarantee known symmetry properties of L_{invmod} such as GIM kernels [5] or group invariant neural networks [6, 7, 8, 9]. In this way, if symmetries of the system are known, the discrete Noether theorem guarantees the existence of corresponding conservation laws. Moreover, we only learn the scalar valued map L_{invmod} rather than the flow map $(x, v) \mapsto \Phi(x, v)$ of the underlying second order differential equation. This reduces the dimensionality of the problem.

- The observed data only contains position data without velocity information. Any technique that aims to learn the flow map $(x, v) \mapsto \Phi(x, v)$ of the underlying second order differential equation but does not make use of prior knowledge about variational structure would have to approximate velocities or conjugate momenta before teaching a model for Φ. The approximation of unknown initial velocities of the observed trajectories introduces an additional approximation error in the data. In contrast, Lagrangian Shadow Integration can directly use the observed data without additional approximations or without requiring further prior knowledge. This contrasts Lagrangian Shadow Integration to other techniques which

  - assume that velocity data or even higher derivatives of trajectories is available [1]
  - or assume more prior knowledge about the system such as [2, 10]. Prior knowledge about the form of the Lagrangian or kinetic term fix much of the phase space structure. Lagrangian Shadow Integration can discover the full Lagrangian from data without restrictions on its form.

- Moreover, approaches such as Lagrangian Neural Networks [1] and related approaches [11, 2, 10] aim at learning the exact Lagrangian. When such a system is integrated, an additional approximation error occurs which Shadow Lagrangian Integration compensates by construction. Moreover, since a Lagrangian is not uniquely determined by the motions of a dynamical system, techniques which aim to identify the exact Lagrangian might learn a Lagrangian which is not well suited for numerical integration, as we will demonstrate in Section 3.

- Shadow Lagrangian Integration allows the employment of established analysis tools from geometric integration technique. These can be actively employed for verification and system identification.

- Another approach is to learn the Hamiltonian of a dynamical system instead of its Lagrangian. In contrast to Lagrangian Shadow Integration, these approaches typically assume that the symplectic structure of the system is known and assume that observations of momentum data is available in addition to position data: Hamiltonian Neural Networks [12] learn the Hamiltonian of a system from position and momentum data of trajectories. Bertalan et al. [13] learn Hamiltonians using Gaussian Processes and introduce a technique to learn the symplectic structure of a system from position and momentum observations of trajectories. SympNets [14] learn the symplectic flow map of a Hamiltonian systems using a neural network architecture which is symplectic by construction. Other approaches learn the generating function of the symplectic flow map [15]. In a previous article [16] the authors introduced *Symplectic Shadow Integration* which learns an inverse modified Hamiltonian using kernel based methods. The present article can be seen as a Lagrangian version of [16].

4

- Backward error analysis has been used to improve the accuracy of classical numerical methods for solving differential equations in [17], for instance. For Hamiltonian systems it has been used in a data-driven context in [16, 18]. While traditional backward error analysis of structure preserving methods [4] applies to the Hamiltonian description of dynamical systems, here we apply *variational backward error analysis*, which stays fully on the variational side and was developed in [3]. In this work, we develop an inverse version of variational backward error analysis and apply it in a data-driven context. (Inverse) variational backward error analysis is used as a tool to analyse and verify the learning methods, for system identification, and for the computation of velocities to trajectories for which only position data has been observed.

- Up to a coordinate transform, learning an inverse modified Lagrangian L_{invmod} from data, as proposed in Lagrangian Shadow Integration, corresponds to learning a discrete Lagrangian L_Δ, as done in [19]. This is sufficient for solving boundary value problems or to extend discrete trajectories. However, to solve initial value problems with initial data of the form $(q_0, \dot{q}_0)$ or to perform system identification, variational backward error analysis as performed in Lagrangian Shadow Integration, is required.

2. The Lagrangian Shadow Integration Framework

After reviewing the concept of modified Lagrangians, we introduce inverse modified Lagrangians and show how to predict motions once an inverse modified Lagrangian is learned. We then discuss the non-uniqueness of Lagrangians and develop regularisation terms that are needed in the learning process. Finally, we show how to learn inverse modified Lagrangians using neural networks as well as Gaussian processes.

2.1. Modified Lagrangians

Let us review the concept of modified Lagrangians that were introduced in [3]. Related computational tools are reviewed in more detail in Section 4.

When a dynamical system defined by a Lagrangian $L(q, \dot{q})$ is integrated using a variational integrator with step-size h, then numerical solutions fulfil discrete Euler–Lagrange equations

$$\nabla_2 L_\Delta(q_{j-1}, q_j) + \nabla_1 L_\Delta(q_j, q_{j+1}) = 0, \qquad j = 1, \ldots, N - 1,$$

where L_Δ is the discrete Lagrangian corresponding to the integrator. Variational backward error analysis seeks a continuous Lagrangian $L_{\text{mod}}(q, \dot{q})$ such that any solution q to the modified Euler–Lagrange equations $\mathcal{E}(L_{\text{mod}}) = 0$ fulfils the discrete Euler–Lagrange equations, i.e.

$$\nabla_2 L_\Delta(q(t - h), q(t)) + \nabla_1 L_\Delta(q(t), q(t + h)) = 0. \tag{6}$$

Modified Lagrangians L_{mod} are sought as formal power series in the step size h such that when truncated to any order in h then (6) is fulfilled up to higher order terms. Except in special cases, the formal power series L_{mod} will not converge. However, optimal truncation results have been proved for an analogous theory on the Hamiltonian side [4] and a dynamical meaning of truncations of L_{mod} can be observed in numerical experiments [20, 21]. The relation between L, L_Δ, and L_{mod} is illustrated in Figure 1.

2.2. The concept of inverse modified Lagrangians

In the following, we consider (consistent) variational discretisation schemes $L \mapsto L_\Delta$, which map a Lagrangian L to a discrete Lagrangian L_Δ. We further require that the discretisation scheme is linear in the Lagrangian, i.e. $sL^1 + L^2 \mapsto sL_\Delta^1 + L_\Delta^2$.

The *exact discrete Lagrangian* L_Δ^{ex} to a step size h and a system defined by a Lagrangian L is defined as

$$L_\Delta^{\text{ex}}(q_0, q_1) = \int_{t_0}^{t_0 + h} L(q(t), \dot{q}(t)) \mathrm{d}t, \tag{7}$$
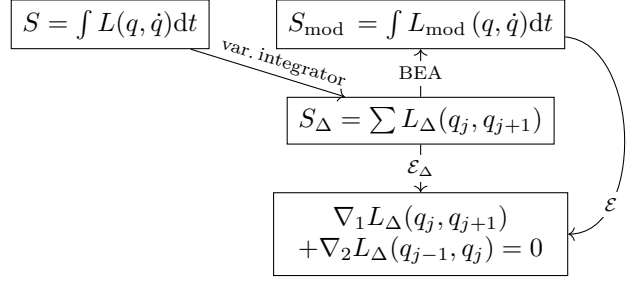
Figure 1: Variational backward error analysis (BEA) seeks a modified Lagrangian L_{mod} such that solutions to the modified Euler–Lagrange equations $\mathcal{E}(L_{\mathrm{mod}}) = 0$ fulfil the discrete Euler–Lagrange equations $\mathcal{E}_\Delta(L_\Delta) = 0$ with $q_j = q(t + jh)$.

where q solves the Euler–Lagrange equations (2) subject to $q(t_0) = q_0$, $q(t_0 + h) = q_1$. The inverse modified Lagrangian to a Lagrangian L and an integration scheme is a continuous Lagrangian L_{invmod} such that $L_{\mathrm{invmod}\,\Delta}$ coincides with the exact discrete Lagrangian L_Δ^{ex} (7), i.e.

$$L_{\mathrm{invmod}\,\Delta}(q_0, q_1) = L_\Delta^{\mathrm{ex}}(q_0, q_1). \tag{8}$$

As the discussion in Section 4 will show, L_{invmod} exists only as a formal power series in the step size h. Equality in (8) means that the left and right hand side coincide up to any order in h. Again, in the sense of formal power series, the discrete conjugate momenta of $L_{\mathrm{invmod}\,\Delta}$ coincide with the exact conjugate momenta of L, i.e.

$$\nabla_1 L_{\mathrm{invmod}\,\Delta}(q_0, q_1) = \frac{\partial L_{\mathrm{invmod}\,\Delta}}{\partial q_0}(q_0, q_1) = -\frac{\partial L}{\partial \dot{q}}(q_0, \dot{q}_0) = -p(t_0) \tag{9}$$

$$\nabla_2 L_{\mathrm{invmod}\,\Delta}(q_0, q_1) = \frac{\partial L_{\mathrm{invmod}\,\Delta}}{\partial q_1}(q_0, q_1) = \frac{\partial L}{\partial \dot{q}}(q_1, \dot{q}_1) = p(t_1). \tag{10}$$

The corresponding calculation for exact discrete Lagrangian can be found in [4, 22], for instance. The exact Lagrangian L can be recovered from L_{invmod} to any order in h using variational backward error analysis [3]. The necessary computational tools will be reviewed in Theorem 4.2. For instance, for the midpoint rule in the one-dimensional case we have

$$L = L_{\mathrm{invmod}} + \frac{1}{24}h^2 \left(\frac{\left(L_{\mathrm{invmod}}^{(1,0)} - L_{\mathrm{invmod}}^{(1,1)}\dot{q} \right)^2}{L_{\mathrm{invmod}}^{(0,2)}} - L_{\mathrm{invmod}}^{(2,0)}\dot{q}^2 \right) + O\left(h^4 \right). \tag{11}$$

For the trapezoidal rule we obtain

$$L = L_{\mathrm{invmod}} + \frac{1}{24}h^2 \left(2L_{\mathrm{invmod}}^{(2,0)}(\dot{q})^2 + \frac{\left(L_{\mathrm{invmod}}^{(1,0)} - L_{\mathrm{invmod}}^{(1,1)}\dot{q} \right)^2}{L_{\mathrm{invmod}}^{(0,2)}} \right) + O\left(h^4 \right).$$

Here $L^{(i,j)} = \frac{\partial^{i+j}}{(\partial q)^i (\partial \dot{q})^j}(q, \dot{q})$ denote derivatives which can be computed using automatic differentiation. Higher order terms and higher dimensional examples can be found or generated using the accompanying Mathematica scripts [23].

The relation between a Lagrangian, a modified Lagrangian, and an inverse modified Lagrangian is illustrated in Figure 2.

2.3. Computation of motions and idea of Lagrangian Shadow Integration

The idea of Lagrangian Shadow Integration is to learn L_{invmod} such that the Euler–Lagrange equations to $L_{\mathrm{invmod}\,\Delta}$ are consistent with the observed trajectory data. For this, we will ignore the fact that L_{invmod} only exists in the sense of formal power series and simply model L_{invmod} as a neural network or Gaussian Process which maps $(q, \dot{q})$ to a scalar output. Once L_{invmod} is learned, an approximation L to the (unknown) exact Lagrangian is computed using a variational backward error analysis formula (such as (11) in case of the midpoint rule in one dimension). A trajectory to initial position and velocity $(q_0, \dot{q}_0)$ can then be computed as follows.
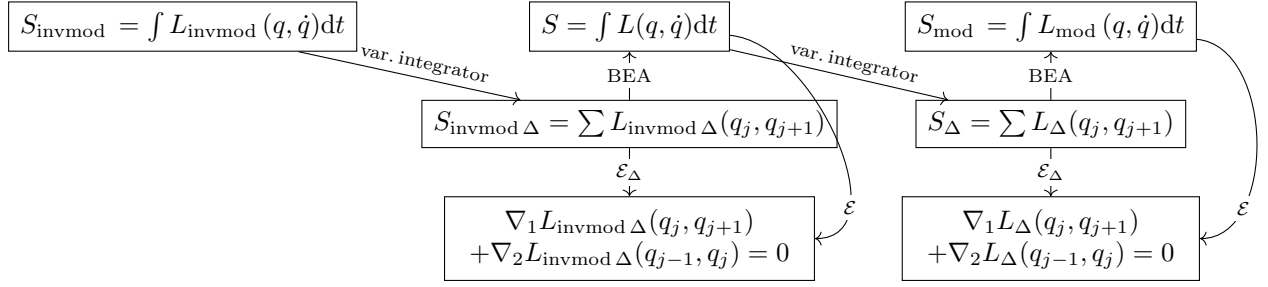
6

Figure 2: This figure extends Figure 1 by a column to the left and two boxes in the centre column. To a given variational integration scheme, the inverse modified Lagrangian L_{invmod} is a continuous Lagrangian such that the discrete Lagrangian $L_{\text{invmod}\,\Delta}$ to L_{invmod} coincides with the exact Lagrangian L up to any order in the step size h. The exact Lagrangian L can be recovered from L_{invmod} using variational backward error analysis (BEA). Up to any order in the step size h, solutions to the exact Euler–Lagrange equations $\mathcal{E}(L) = 0$ fulfil the discrete Euler–Lagrange equations $\mathcal{E}_\Delta(L_{\text{invmod}}) = 0$ to the Lagrangian L_{invmod} with $q_j = q(t + jh)$.

- Set $p_0 = \frac{\partial L}{\partial \dot{q}}(q_0, \dot{q}_0)$.

- Solve $\nabla_1 L_{\text{invmod}\,\Delta}(q_0, q_1) + p_0 = 0$ for q_1.

- For $j = 1, \ldots, N-1$ solve the discrete Euler–Lagrange equations for $L_{\text{invmod}\,\Delta}$

$$\nabla_2 L_{\text{invmod}\,\Delta}(q_{j-1}, q_j) + \nabla_1 L_{\text{invmod}\,\Delta}(q_j, q_{j+1}) = 0, \qquad j = 1, \ldots \tag{12}$$

  for q_{j+1}.

- If a velocity value $\dot{q}_j$ is needed, it can be computed by solving

$$\nabla_2 L_{\text{invmod}\,\Delta}(q_{j-1}, q_j) = \frac{\partial L}{\partial \dot{q}}(q_j, \dot{q}_j)$$

  for q_j, where a formula for L has been calculated from L_{invmod} to the desired order in h using variational backward error analysis.

Remark 2.1. *Notice that values for the conjugate momenta*

$$\frac{\partial L}{\partial \dot{q}}(q_j, \dot{q}_j) = p_j = -\nabla_1 L_{\text{invmod}}(q_j, q_{j+1}) = \nabla_2 L_{\text{invmod}}(q_{j-1}, q_j)$$

do not necessarily carry a physical meaning since they depend on the choice of a Lagrangian, which is not uniquely determined by a system's dynamics.

2.4. Non-uniqueness of Lagrangians

Lagrangians are not uniquely determined by the system's dynamics. In particular, constant Lagrangians have trivial Euler–Lagrange equations $\mathcal{E}(L) \equiv 0$ and are, therefore, consistent with any dynamical data. To avoid learning singular Lagrangians, we develop a normalisation condition for L and L_{invmod}. For this, let us review the non-uniqueness of the Lagrangian L with its conjugate momenta p and also discuss the ambiguity of symplectic structures and Hamiltonians.

Lagrangian dynamics can be translated to Hamiltonian dynamics. For this, conjugate momenta are introduced as $p = \frac{\partial L}{\partial \dot{q}}(q, \dot{q})$. Provided that L is regular, i.e. $\left(\frac{\partial^2 L}{\partial \dot{q}^j \partial \dot{q}^i}\right)_{i,j=1}^{n}$, $(q, \dot{q})$ and (q, p) provide local coordinate systems on the tangent bundle TQ and the cotangent bundle T^*Q, respectively, which are both locally identified with $\mathbb{R}^{2n}$ in the following. The Hamiltonian is given as

$$H(q, p) = \sum_{j=1}^{n} \dot{q}^j((q, p))p^j - L(q, \dot{q}(q, p)),$$

7

where q^j is expressed in terms of q, p. The Hamiltonian H is conserved along motions. The corresponding symplectic structure ω and Liouville volume μ are given as

$$\omega = \sum_{j=1}^{n} dq^j \wedge dp^j, \quad \mu = \frac{1}{n!} \underbrace{\omega \wedge \ldots \wedge \omega}_{n \text{ times}}.$$

In coordinates $(q, \dot{q})$ the Liouville volume form can be expressed as

$$\mu = \frac{\partial L}{\partial \dot{q}^1} \cdot \ldots \cdot \frac{\partial L}{\partial \dot{q}^n} dq^1 \wedge d\dot{q}^1 \wedge \ldots \wedge dq^n \wedge d\dot{q}^n.$$

The Lagrangian L is defined by the dynamics up to scalar multiples and up to a total derivative of a scalar valued map ξ in q, i.e. the Lagrangian $\tilde{L} = sL + \frac{d}{dt}\xi(q)$ defines the same dynamics as L but its conjugate momenta is $\tilde{p} = sp + \nabla\xi$. The symplectic structures $\omega = dq \wedge dp$ and $\tilde{\omega} = dq \wedge d\tilde{p}$ relate as $\tilde{\omega} = s\omega$. Once the symplectic structure is fixed, the Hamiltonian H is defined up to a constant by the dynamics.

In the following, we will consider the following non-triviality condition for L_{invmod}, where $E^{2n} = ([0, 1] \times [0, 1])^n$ is the unit hypercube in $\mathbb{R}^{2n}$ with coordinates $q_1, \dot{q}_1, \ldots, q_n, \dot{q}_n$:

$$c = \int_{E^{2n}} \mu = \int_{E^{2n}} \frac{\partial L}{\partial \dot{q}^1} \cdot \ldots \cdot \frac{\partial L}{\partial \dot{q}^n} dq^1 d\dot{q}^1 \ldots dq^n d\dot{q}^n, \tag{13}$$

where $c \neq 0$ is a constant controlling the scaling of L. Notice that while (13) removes the scaling ambiguity, it does not remove the ambiguity of summands which are total derivatives, even if a further normalisation such as $L(0, 0) = 0$ is added. However, it forces L to be non-trivial and regular. In the following, we use the approximation

$$c = \int_{E^{2n}} \mu = \int_{E^{2n}} \frac{\partial L_{\text{invmod}}}{\partial \dot{q}^1} \cdot \ldots \cdot \frac{\partial L_{\text{invmod}}}{\partial \dot{q}^n} dq^1 d\dot{q}^1 \ldots dq^n d\dot{q}^n \tag{14}$$

of (13), which coincides with (13) up to $O(h)$-terms (in the infinite data limit).

Remark 2.2. *An alternative (more ad hoc) non-triviality condition is*

$$\frac{\partial L_{\text{invmod}}}{\partial \dot{q}}(q_{\text{base}}, \dot{q}_{\text{base}}) - p_{\text{base}} = 0, \qquad p_{\text{base}} \neq 0,$$

which imposes that the value of the conjugate momentum of L_{invmod} at an arbitrary point $(q_{\text{base}}, \dot{q}_{\text{base}})$ is p_{base}. Again, this does not impose any restrictions on the dynamical system. However, in the following we will use (14) as it balances several points of the phase space and is more geometric.

2.5. Learning the inverse modified Lagrangian

To teach a model for L_{invmod}, the idea is to use the discrete Euler–Lagrange equations (12) for all available data triples (q_{j-1}, q_j, q_{j+1}) together with a discretisation of the non-triviality condition (14). In addition, we will use a simple quadrature on (14) and obtain the condition $\ell_{\text{nontrivial}} = 0$ with

$$\ell_{\text{nontrivial}} = c - \frac{1}{2^{2n}} \sum_{\text{corners of } E^{2n}} \frac{\partial L_{\text{invmod}}}{\partial \dot{q}^1} \cdot \ldots \cdot \frac{\partial L_{\text{invmod}}}{\partial \dot{q}^n}. \tag{15}$$

Without the condition $\ell_{\text{nontrivial}} = 0$ any constant inverse modified Lagrangian $L_{\text{invmod}} \equiv \text{const}$ would be consistent with the data because the discrete Euler–Lagrange equations are fulfilled on the data for any $L_{\text{invmod}\Delta} \equiv \text{const}$.

2.5.1. ANN

If L_{invmod} is modelled as a differentiable artificial neural network, we can define

$$\ell_{\text{data}} = \sum_{q \in \text{trj}} \sum_{j=1}^{\text{length}(q)-1} \|\nabla_2 L_{\text{invmod}\Delta}(q_{j-1}, q_j) + \nabla_1 L_{\text{invmod}\Delta}(q_j, q_{j+1})\|^2.$$

Here trj denotes the set of all trajectories of the training data. Minimisation of ℓ_{data} enforces the discrete Euler–Lagrange equations to hold on the training data. More concretely, for the midpoint rule ℓ_{data} reads

$$
\ell_{\text{data}} = \sum_{q \in \text{trj}} \sum_{j=1}^{\text{length}(q)-1} \left\| \frac{1}{2} \frac{\partial L_{\text{invmod}}}{\partial q} \left(\frac{q_{j-1}+q_j}{2}, \frac{q_j - q_{j-1}}{h} \right) + \frac{1}{h} \frac{\partial L_{\text{invmod}}}{\partial \dot q} \left(\frac{q_{j-1}+q_j}{2}, \frac{q_j - q_{j-1}}{h} \right) \right.
$$
$$
\left. + \frac{1}{2} \frac{\partial L_{\text{invmod}}}{\partial q} \left(\frac{q_j + q_{j+1}}{2}, \frac{q_{j+1} - q_j}{h} \right) - \frac{1}{h} \frac{\partial L_{\text{invmod}}}{\partial \dot q} \left(\frac{q_j + q_{j+1}}{2}, \frac{q_{j+1} - q_j}{h} \right) \right\|^2 .
\tag{16}
$$

The network can then be trained using a weighted sum of the non-triviality term $\|\ell_{\text{nontrivial}}\|^2$ and the term ℓ_{data}. The term $\|\ell_{\text{nontrivial}}\|^2$ can be interpreted as a regularising term.

2.5.2. Gaussian Process Regression

In the numerical examples of this paper, we will use Gaussian Process Regression. For an introduction we refer to [24]. The method can be applied to all kernel based methods with sufficiently differentiable kernel.

In the following, denote the tangent bundle over the state space $Q = \mathbb{R}^n$ by TQ which is identified with $\mathbb{R}^{2n}$. If $\hat L_{\text{invmod}}$ is a Gaussian process with zero mean and kernel $k \colon TQ \times TQ \to \mathbb{R}$ and (a priori unknown) values $L_{\text{invmod}}(Z) = \left(L_{\text{invmod}}(z_1), \ \ldots, \ L_{\text{invmod}}(z_M) \right)^\top$ for $Z = \left(z_1, \ \ldots, \ z_M \right) \subset TQ$, then the value $L_{\text{invmod}}(y)$ for some $(q, \dot q) \in TQ$ can be predicted as the expectation

$$
\mathbb{E}[\hat L_{\text{invmod}}((q, \dot q))|(Z, L_{\text{invmod}}(Z))] = k((q, \dot q), Z)^\top k(Z, Z)^{-1} L_{\text{invmod}}(Z)
\tag{17}
$$

with $k((q, \dot q), Z)^\top = \left(k((q, \dot q), z_1), \ \ldots, \ k((q, \dot q), z_M) \right)$ and $k(Z, Z) = (k(z_i, z_j))_{i,j=1}^M$.

Since the discretisation scheme I is linear in the Lagrangian, we have

$$
L_{\text{invmod}\,\Delta}((q_a, q_b)) = \mathbb{E}[\hat L_{\text{invmod}\,\Delta}((q_a, q_b))|(Z, L_{\text{invmod}}(Z))] = I[k]((q_a, q_b), Z)^\top k(Z, Z)^{-1} L_{\text{invmod}}(Z)
\tag{18}
$$

In case of the midpoint rule

$$
I[k]((q_a, q_b), Z) = \left(k\left(\left(\tfrac{q_a + q_b}{2}, \tfrac{q_b - q_a}{h} \right), z_1 \right), \ \ldots, \ k\left(\left(\tfrac{q_a + q_b}{2}, \tfrac{q_b - q_a}{h} \right), z_M \right) \right).
$$

In case of the trapezoidal rule

$$
I[k]((q_a, q_b), Z) = \left(\frac{1}{2} k\left(\left(q_a, \frac{q_b - q_a}{h} \right), z_j \right) + \frac{1}{2} k\left(\left(q_b, \frac{q_b - q_a}{h} \right), z_j \right) \right)_{j=1}^M .
$$

In general, $I[k]((q_a, q_b), Z) = \left(I(y \mapsto k(y, z_j))(q_a, q_b) \right)_{j=1}^M$. The right hand side of (18) can now be differentiated with respect to q_a or q_b to obtain expressions for $\nabla_1 L_{\text{invmod}\,\Delta}((q_a, q_b))$ and $\nabla_2 L_{\text{invmod}\,\Delta}((q_a, q_b))$. These can be substituted into the discrete Euler–Lagrange equations (12) for each data triple (q_{j-1}, q_j, q_{j+1}). For the midpoint rule this yields a linear system consisting of the following equation for each data triple (q_{j-1}, q_j, q_{j+1}) in the training set.

$$
\left[\frac{1}{h} \left(\nabla_{\dot q} k\left(\left(\frac{q_{j-1}+q_j}{2}, \frac{q_j - q_{j-1}}{h} \right), Z \right)^\top - \nabla_{\dot q} k\left(\left(\frac{q_{j+1}+q_j}{2}, \frac{q_{j+1} - q_j}{h} \right), Z \right)^\top \right) \right.
$$
$$
\left. + \frac{1}{2} \left(\nabla_q k\left(\left(\frac{q_{j-1}+q_j}{2}, \frac{q_j - q_{j-1}}{h} \right), Z \right)^\top - \nabla_q k\left(\left(\frac{q_{j+1}+q_j}{2}, \frac{q_{j+1} - q_j}{h} \right), Z \right)^\top \right) \right] \cdot k(Z, Z)^{-1} L_{\text{invmod}}(Z) = 0
$$

Together with the non-triviality condition $\ell_{\text{nontrivial}} = 0$

$$
\left(\frac{1}{2^{2n}} \sum_{v \in \text{corners}(E^{2n})} \mathbf{1}_n^\top \nabla_{\dot q} k(v, Z)^\top \right) k(Z, Z)^{-1} L_{\text{invmod}}(Z) = 0
$$

9

where $\mathbf{1}_n = \begin{pmatrix} 1, & \ldots, & 1 \end{pmatrix}$ and the normalising condition

$$k((q^*, \dot{q}^*), Z)^\top k(Z, Z)^{-1} L_{\text{invmod}}(Z) = 0$$

for some $(q^*, \dot{q}^*) \in TQ$, the system consists of $n \cdot K + 2$ equations for M unknowns. Here K is the number of triples (q_{j-1}, q_j, q_{j+1}) from the training data. Now a least-square or minimal norm solution of the linear system is computed. This means we have computed predictions for $k(Z, Z)^{-1} L_{\text{invmod}}(Z)$ at points $Z = \begin{pmatrix} z_1, & \ldots, & z_M \end{pmatrix} \subset TQ$ based on the available training data triples on Q.

In computations, we chose Z to consists of the points in TQ for which the training data provides maximal information and then predict any further values of L_{invmod} or its derivatives using (17) or derivatives thereof. In case of the midpoint rule, we chose $Z = \left(\frac{q_j + q_{j+1}}{2}, \frac{q_{j+1} - q_j}{h} \right)_{(q_j, q_{j+1})}$ where (q_j, q_{j+1}) runs over all data pairs in the training data set.

Remark 2.3. *Notice that it suffices to solve the linear system $(n \cdot K + 2) \times M$ dimensional system for the product $B = k(Z, Z)^{-1} L_{\text{invmod}}(Z)$ and use this expression in (17) or derivatives thereof. A computation of $k(Z, Z)^{-1}$ is not needed.*

Remark 2.4. *As there are more pairs (q_j, q_{j+1}) than triples (q_{j-1}, q_j, q_{j+1}) in the training data set, the (non-homogeneous) linear system is underdetermined.[2] Moreover, the non-triviality condition and normalisation does not fully remove the ambiguity of Lagrangians explained in Section 2.4. Assuming the existence[3] of L_{invmod}, we expect the system to be of rank smaller than $n \cdot K + 2$ and compute a minimal norm solution.*

2.5.3. Comparison - Lagrangian Gaussian Process (LGP)

For comparison, we also develop a method which can learn the Lagrangian when not only position but also velocity and acceleration data has been observed. The following can be seen as a Gaussian process version of Lagrangian neural networks (LNN) [1]. However, in contrast to the approach taken in [1], our algorithm avoids matrix inversion. In Figure 2 LGP corresponds to learning L with motions computed using the discrete Euler–Lagrange equation $\mathcal{E}_\Delta(L_\Delta) = 0$.

1. We model a Lagrangian as a Gaussian Process and impose the (continuous) Euler–Lagrange equations at the training data

$$\mathcal{E}(L)^k = \frac{\partial L}{\partial \dot{q}^k} - \frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial L}{\partial \dot{q}^k} = \frac{\partial L}{\partial \dot{q}^k} - \sum_{j=0}^{n} \left(\frac{\partial^2 L}{\partial q^j \partial \dot{q}^k} \dot{q}^j - \frac{\partial^2 L}{\partial \dot{q}^j \partial \dot{q}^k} \ddot{q}^j \right) \tag{19}$$

together with the non-triviality condition and normalisation that we employed for learning the inverse modified Lagrangian. More precisely, if the training data is presented as $Z = (z_1, \ldots, z_K) = ((q_1, \dot{q}_1), \ldots, (q_K, \dot{q}_K))$ then

$$r_{j,i} = \nabla_{\dot{q}} k((q_j, \dot{q}_j), z_i) - \mathrm{D}^2_{q,\dot{q}} k((q_j, \dot{q}_j), z_i) \dot{q}_j - \mathrm{D}^2_{\dot{q},\dot{q}} k((q_j, \dot{q}_j), z_i) \ddot{q}_j.$$

is a column vector of length n. Here $\mathrm{D}^2_{q,\dot{q}} k((q_j, \dot{q}_j), z_i)$ denotes the $n \times n$ matrix $\left(\frac{\partial^2}{\partial q^{l_1} \partial \dot{q}^{l_2}} k((q_j, \dot{q}_j), z_i) \right)_{l_1, l_2 = 1}^n$ with a corresponding interpretation of $\mathrm{D}^2_{\dot{q},\dot{q}}$. We then build the $n \times K$ matrix

$$R_j = \begin{pmatrix} r_{j,1}, \ldots, r_{j,K} \end{pmatrix}$$

which are stacked above each other in the vertical direction to form an $nK \times K$-dimensional matrix. Now the system

$$R k(Z, Z)^{-1} L(Z) = 0 \tag{20}$$

corresponds to the condition that (19) is fulfilled for all data points.

---

[2]Leaving aside some obscure cases.

[3]In Section 4 we will prove that L_{invmod} exists as a formal power series in h. Motivated by classical results in backward error analysis (BEA), we conjecture optimal truncation results such that L_{invmod} exists up to exponentially small terms in the step-size h.

2. Finally (20) together with the same non-triviality and normalisation condition as employed for learning the inverse modified Lagrangian is solved in the least square sense for $B = k(Z,Z)^{-1}L(Z) \in \mathbb{R}^{nK}$. Notice that a computation of $k(Z,Z)^{-1}$ is not required.

3. The learned Lagrangian L can be evaluated at a point $(q, \dot{q})$ as $L(q, \dot{q}) = k((q, \dot{q}), Z)B$.

For comparison, a neural network version of the above idea would employ a loss function of the form $\ell = \|\ell_{\text{nontrivial}}\|^2 + \ell_{\text{data}}$ with $\ell_{\text{nontrivial}}$ as in (15) with L instead of L_{invmod} and

$$\ell_{\text{data}} = \sum_{i=1}^{K} \sum_{k=1}^{n} \left\| \frac{\partial L}{\partial \dot{q}^k}(q_i, \dot{q}_i) - \sum_{j=0}^{n} \left(\frac{\partial^2 L}{\partial q^j \partial \dot{q}^k}(q_i, \dot{q}_i)\dot{q}_i{}^j - \frac{\partial^2 L}{\partial \dot{q}^j \partial \dot{q}^k}(q_i, \dot{q}_i)\ddot{q}_i{}^j \right) \right\|^2,$$

which corresponds to enforcing the Euler–Lagrange equations to be consistent with the training data.

Remark 2.5 (Comparison with LNN). *In contrast, Lagrangian neural networks (LNN), as introduced in [1], use a loss function based on comparing the acceleration data $\ddot{q}_i$ with*

$$\left(\left(\frac{\partial^2 L}{\partial \dot{q}^j \partial \dot{q}^k}(q_i, \dot{q}_i) \right)_{j,k=1}^{n} \right)^{-1} \left(\nabla_q L(q_i, \dot{q}_i) - \left(\frac{\partial^2 L}{\partial q^j \partial \dot{q}^k}(q_i, \dot{q}_i) \right)_{j,k=1}^{n} \dot{q}_i \right).$$

In LNN the implicit assumption that $\left(\frac{\partial^2 L}{\partial \dot{q}^j \partial \dot{q}^k}(q_i, \dot{q}_i) \right)_{j,k=1}^{n}$ is invertible, informs the network that the Lagrangian is regular. The regularity assumption prohibits the LNN to learn the constant solution $L \equiv$ const which would be consistent with the data but does not generate any dynamics. In contrast, our approach avoids the matrix inversion and imposes non-triviality by a term which normalises the symplectic volume of a unit hypercube in the space with coordinates $(q, \dot{q})$. We can, therefore, control the scaling of the learned Lagrangian and the corresponding symplectic structure and Hamiltonian.

3. Numerical Experiments

In the following, we will demonstrate the performance of Lagrangian Shadow Integration (LSI) on position data of the mathematical pendulum and of the Hénon–Heiles system. Within LSI we will use the second order accurate variational midpoint discretisation introduced in Example 1.1 as a base method. All integrations required to compute the training data set are performed by applying the Störmer-Verlet scheme with a very small step-size to the exact equations of motion. After integration, we retain only position data, since LSI uses position data only and does not require observations of velocities, accelerations, or prior knowledge about conjugate momenta. Based on the training data, the inverse modified Lagrangian is modelled as a Gaussian process and learned as described in Section 2.5.2.

For comparison, we employ Lagrangian Gaussian Processes (LGPs) (see Section 2.5.3) whose training targets are the exact Lagrangian. The learned Lagrangians are then integrated using the variational midpoint rule with the same step-size as we use for LSI. Velocity and acceleration data for training LGP are interpolated using central finite differences. The approximations are second order accurate in the discretisation parameter h. Notice that this reduces the available training data since no second order accurate acceleration data is available at the first two and last two snapshots of each trajectory.

All Gaussian Processes (GP) employed in this section use scaled radial basis functions

$$k(x, y) = c_k \exp\left(-\frac{\|x - y\|^2}{\epsilon^2} \right) \tag{21}$$

with parameters ϵ related to the typical length scale and c_k acting as a scaling factor.

3.1. Pendulum

The mathematical pendulum is the dynamical system defined by the second order ordinary differential equations $\ddot{q} = -\sin(q)$, which is the Euler–Lagrange equation to $L_{\text{ref}}(q, \dot{q}) = \frac{1}{2}\dot{q}^2 + \cos(q)$. Motions preserve the energy $H_{\text{ref}}(q, \dot{q}) = \frac{1}{2}\dot{q}^2 - \cos(q)$. We consider the step-size $h = 0.5$. To generate training data, we compute $N = 400$
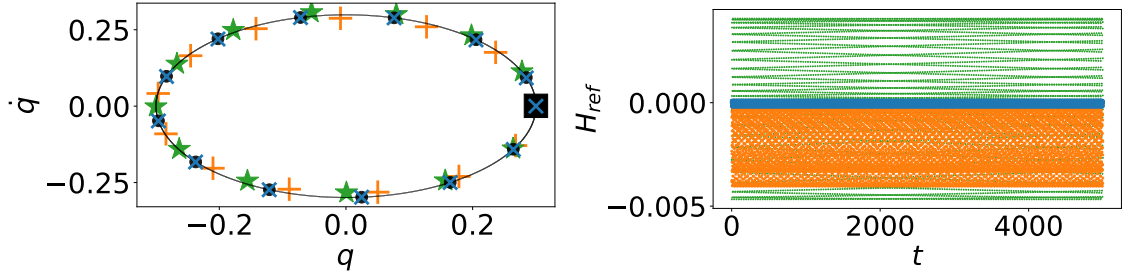
Figure 3: Left: The first 13 snapshots to the step-size $h = 0.5$ of a trajectory initialised at $(q, \dot{q}) = (0.3, 0)$ (■) computed with LSI (×), with LGP (+), and with LGPExact (∗) are shown. A reference solution is given in solid black, with position and velocity at the snapshot times marked by •. Right: Energy behaviour of LSI (blue), LGP (orange), and LGPExact (green) (subsampled).

trajectories of length $N_l = 6$ with initial values $(q, \dot{q})$ obtained from a Halton sequence on the domain $\Omega = [-\pi, \pi] \times [-1.2, 1.2]$. Only the position data in $[-\pi, \pi]$ is retained.

In the following, the parameters of the radial basis functions (21) of the GPs within LSI and LGP are set to $\epsilon = 5$ and $c_k = 1$. Figure 3 shows snapshot data of a trajectory computed with LSI as well as with LGP, whereas for LGP all required velocity and acceleration data was approximated using central finite differences. For additional comparison, we plot snapshot data of a trajectory computed with LGPExact. LGPExact is a Lagrangian Gaussian Process as well. However, it is trained with exact velocity and acceleration data. Although this violates our assumption that only position data is available, the experiment is included as it helps us to analyse the effects of discretisation errors in LGP. For training of LGPExact we used the same position data as for LSI. Acceleration data was computed from the differential equation $\ddot{q} = -\sin(q)$ of the pendulum. Since the right hand side of the differential equation does not explicitly depend on $\dot{q}$, velocity data was obtained as a Halton sequence of length $N \cdot N_l = 2400$ on the interval $(-1.2, 1.2)$.

While all methods preserve Lagrangian structure, predictions by LGPExact contain the discretisation error induced by the variational midpoint rule, which is considerable at $h = 0.5$ even though the scheme is second order accurate. Additionally, LGP contains another discretisation error, since its training data was obtained using central finite differences. LSI is designed to compensate both discretisation errors and successfully matches the snapshot data of the reference solution.

Further, we plot in Figure 3 the long-term energy behaviour of the numerical trajectories, i.e. we evaluate H_{ref} on the computed snapshot data. All methods show oscillatory energy error behaviour, which is expected, since all methods preserve the variational structure. We observe that the oscillations of LSI are smaller than those of LGP and LGPExact. This is expected since only LSI contains no inherent discretisation error. Interestingly, a comparison of the energy errors of LGPExact and LGP suggests that in LGP discretisation errors from training and the application of the variational midpoint rule compensate one another partially.

To analyse the quality of predictions obtained by LSI, it is convenient to pass to the Hamiltonian setting due to the non-uniqueness of Lagrangians. We will, therefore, focus on the modified energy $H^{[[k]]}$ preserved by the scheme. Up to errors due to limited training data and truncation errors, we expect the modified energy to have the same level sets as H_{ref}. To compute $H^{[[k]]}$, first $L^{[[k]]}$ is computed from the learned inverse modified Lagrangian L_{invmod} using the backward error analysis formula (11), where k denotes the truncation index (also see Theorem 4.2). The modified Hamiltonian $H^{[[k]]}$ is then computed using (3) applied to $L^{[[k]]}$ (instead of L). Figure 4 shows that $H^{[[2]]}$ governs the numerical motions well as it is preserved up to an oscillation within a band of width $\approx 10^{-6}$, which is small enough for our purposes. We can, therefore, use $H^{[[2]]}$ to analyse qualitative aspects of numerical motions predicted by LSI: up to truncation error, for a modified symplectic structure $\omega^{[[2]]} = \frac{\partial L^{[[2]]}}{\partial \dot{q}} \wedge q$ the function $H^{[[2]]}$ is the Hamiltonian which governs the numerical motion. The contour plot Figure 5 shows that the level sets of H_{ref} closely match the level sets of $H^{[[2]]}$. Analogous plots for LGP and LGPExact are given for comparison. We see that even in long-term simulations the numerical motions obtained by LSI show qualitatively the correct behaviour. Motions of LGP and LGPExact, however, will have a pronounced systematic bias.

Remark 3.1. *It might appear surprising that LGPExact performs quite poorly: the Lagrangian L_{LGPExact} learned*
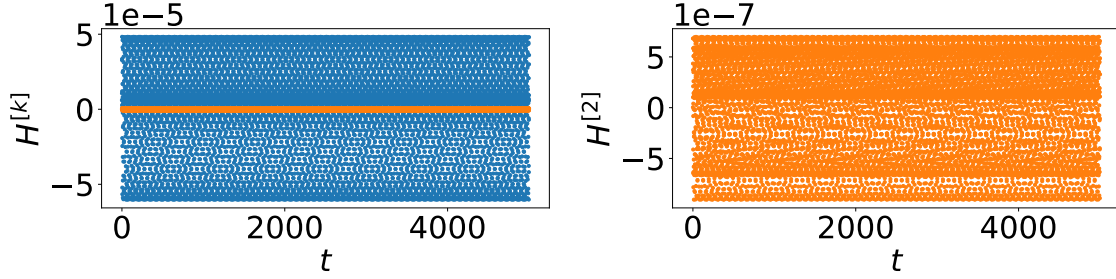
Figure 4: $H^{[[0]]}$ (blue) and $H^{[[2]]}$ (orange) are up to an oscillation within a band of width $\approx 10^{-4}$ and $\approx 10^{-6}$, respectively, conserved quantities of the numerical motion predicted by LSI initialised at $(q, \dot{q}) = (0.3, 0)$. $H^{[2]}$ can, therefore, be used to perform system identification or to verify LSI.
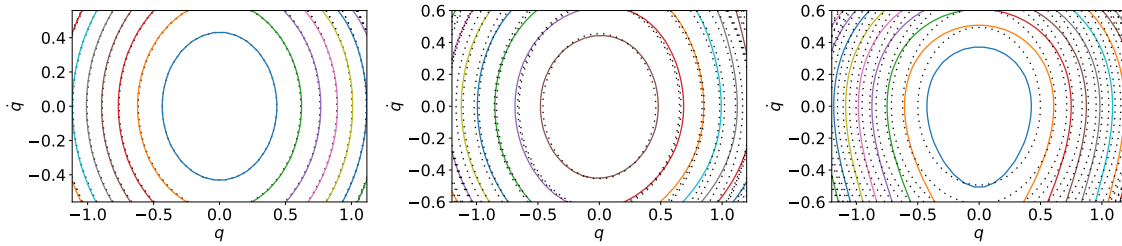

Figure 5: Contour plot of the Hamiltonian which governs the numerical motion for LSI (left), LGP (centre), and LGPExact (right). Contours of H_{ref}, which intersect the solid contour lines, are plotted as dotted lines for reference.
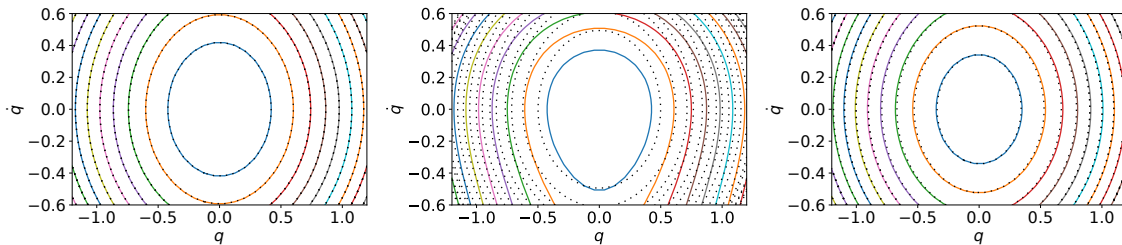

Figure 6: Contour plot of the Hamiltonian identified using LGPExact (left), the Hamiltonian governing the numerical motions when L_{LGPExact} is integrated (centre), and when the reference Lagrangian L_{ref} is integrated (right). This shows that although L_{LGPExact} is a highly accurate Lagrangian for the dynamical system, it is not well suited for numerical integration.

by LGPExact is consistent with the dynamical system (Figure 6 – left). However, L_{LGPExact} is not very suitable for numerical integration: the numerical motions follow the contour lines of the centre plot of Figure 6. For comparison, for the reference Lagrangian L_{ref}, the numerical motions follow energy level sets that closely match the level sets of H_{ref}. This shows that even if highly accurate velocity data and acceleration data is available, Shadow Integration can be worthwhile as it compensates discretisation errors in the integration step also for Lagrangians that would otherwise not be well suited for numerical purposes.

Recall that in this context Hamiltonians are defined only up to additive constants and scaling factors by the dynamics since the symplectic structure is not fixed. To quantify the difference between the level sets of H_{ref} and $H^{[[2]]}$ we measure how parallel their gradients are. More precisely, we measure the area $\mu(q, \dot{q})$ of the parallelograms spanned by the normalised gradients $\nabla H_{\mathrm{ref}}/\|\nabla H_{\mathrm{ref}}\|$ and $\nabla H^{[[2]]}/\|H^{[[2]]}\|$ at $(q, \dot{q})$. In this 2-dimensional example, we have

$$\mu(q, \dot{q}) = \left| \det \left(\frac{\nabla H_{\mathrm{ref}}(q, \dot{q})}{\|\nabla H_{\mathrm{ref}}(q, \dot{q})\|}, \frac{\nabla H^{[[2]]}(q, \dot{q})}{\|\nabla H^{[[2]]}(q, \dot{q})\|} \right) \right|.$$

We sum μ over an equidistant mesh with 30×30 data points on the domain $\Omega_0 = [-1.2, 1.2] \times [-0.6, 0.6]$ and obtain after division by the number of data points $\nu_{\mathrm{LSI}}^{\mathrm{BEA}} \approx 0.01$. For LGP we obtain $\nu_{\mathrm{LGP}}^{\mathrm{BEA}} \approx 0.05$ and $\nu_{\mathrm{LGPExact}}^{\mathrm{BEA}} \approx 0.1$ for LGPExact. It is interesting to observe that LGPExact performs significantly worse than LGP. Again, this suggests that some of the inherent discretisation errors of LGP compensate each other but to a lesser extend than in the systematic approach of the LSI framework.

When system identification is the goal rather than an analyis of the numerical motions obtained by applying the variational midpoint rule to a Lagrangian, for LGP and LGPExact we can compute H_{LGP} and H_{LGPExact} by applying (3) to the learned Lagrangians L_{LGP} and L_{LGPExact} directly rather than applying the backward error analysis formula (11) first (which is necessary when LSI is used). In this context LGP only suffers from a discretisation error in the training data, while in the experiment of Figure 3 an additional discretisation error of the variational midpoint rule is present. We get $\nu_{\mathrm{LGP}} \approx 0.03$. Compared to $\nu_{\mathrm{LSI}} = \nu_{\mathrm{LSI}}^{\mathrm{BEA}} \approx 0.01$ this shows that LSI has an edge over LGP when it comes to system identification. Moreover, we can conclude that even if L_{LGP} was integrated with tiny time-steps in the motion prediction test, solutions would improve but cannot reach the quality of LSI predictions.

Besides, LGPExact does not suffer from any discretisation error in this test and we obtain $\nu_{\mathrm{LGPExact}} \approx 3.4 \cdot 10^{-5}$, which is by magnitudes better than the value for LGP and LSI. However, LGPExact was trained with exact velocity and acceleration data, which might not be available.

3.2. Hénon-Heiles

The conservative motion of a particle of unit mass subject to the potential

$$V(q) = \frac{1}{2}\|q\|^2 + \alpha \left(q_1^2 q_2 - \frac{q_2^3}{3} \right)$$

is governed by the Euler–Lagrange equations (2) for the Lagrangian $L(q, \dot{q}) = \frac{1}{2}\|\dot{q}\|^2 - V(q)$. The energy is given by $H(q, \dot{q}) = \frac{1}{2}\|\dot{q}\|^2 + V(q)$ and is conserved along motions. The potential V has two critical values $c_0 = 0$ and $c_1 = \frac{1}{6\alpha^2}$. The set $V^{-1}(c_1)$ separates the phase space into six unbounded components and one bounded component A containing the origin [25]. (For a contour plot of V see Figure 9.)

In the following numerical experiments we used $\alpha = 0.8$ and the step-size $h = 0.1$. Training was performed on $N = 200$ trajectories of length $N_l = 5$. The initial values of the trajectories were obtained from a Halton sequence on the domain $\Omega = [-0.8, 0.8]^4$. Again, velocity data was discarded and only the position data of the training set is used in the following. Figure 7 shows predictions of trajectories computed by LSI and by LGP. For the underlying GPs we used radial basis functions (21) with parameters $c_k = 1$ and $\epsilon = 10$. For LGP missing velocity and acceleration data was approximated using second order accurate central finite differences. For further comparison, we compute flow map data by adding velocity data to the position data from the training data set with central finite difference approximations. Then a GP is fitted to the flow map data using scaled radial basis functions in combination with hyperparameter fitting employing the Python package `scikit-learn`. This method does not incorporate geometric structure and will be referred to by GPFlow. The trajectories in Figure 7 are initialised at $(q_0, \dot{q}_0) = (0.675499, 0.08, 0, 0)$. Their energy
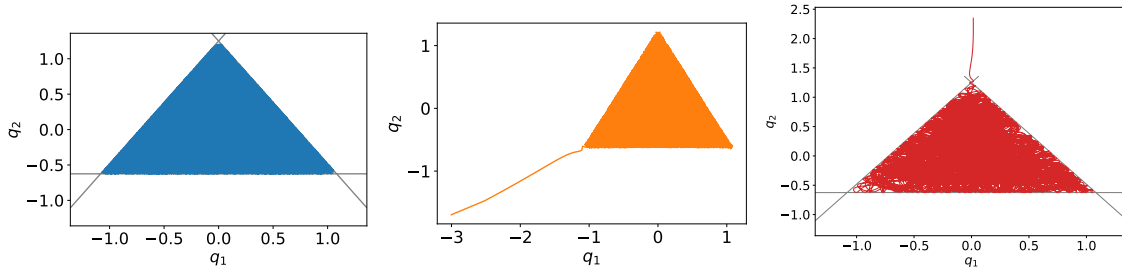
14

Figure 7: Trajectory initialised at $(q_0, \dot{q}_0) = (0.675499, 0.08, 0, 0)$ predicted by LSI until $T = 1.774 \cdot 10^4$ (left), by LGP until $T = 7.069 \cdot 10^3$ (centre), and by a GP directly fitted to approximated flow data until $T = 1.4574 \cdot 10^3$ (GPFlow). All but the LSI prediction stay in the correct region of the phase space until time $T = 1.774 \cdot 10^4$.
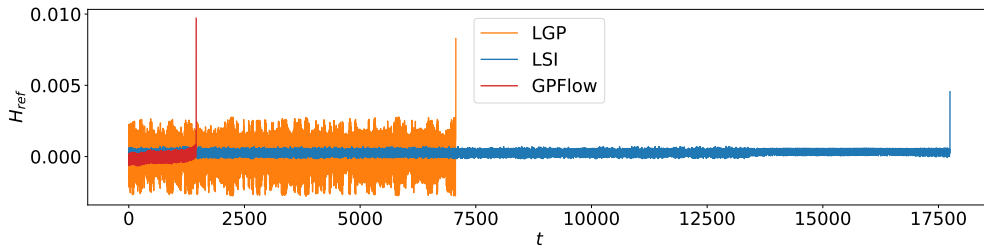

Figure 8: Energy behaviour of the trajectories from Figure 7. All trajectories diverge eventually but the LSI prediction diverges much later than GPFlow and LGP due to its better energy preservation properties.

$H(q_0, \dot{q}_0) = V(q_0) \approx 0.26041605$ is close to the critical value of $c_1 = 0.260416\overline{6}$ such that high energy accuracy is required in numerical simulations to prevent erroneous divergence. While GPFlow, LGP, and LSI all diverge eventually, the LSI prediction shows the correct behaviour (an orbit densely filling the area bounded by $V^{-1}(q_0)$) for much longer than LGP or GPFlow.

The energy error plot (Figure 8) confirms that LSI preserves energy most accurately, while the energy behaviour of GPFlow shows a drift as no Lagrangian or Hamiltonian structure has been enforced. The steady energy growth causes the trajectory to enter a section of the phase space where trajectories diverge. Although LGP does not suffer from a steady energy growth since variational structure is enforced, its energy error oscillations are wider than those of LSI because discretisation errors in the training process and by the application of variational midpoint integration do not get corrected. The wider oscillations cause the trajectory to diverge erroneously much earlier than the LSI prediction which conserves energy more accurately.

We again use backward error analysis and compute the modified system whose exact motion (up to truncation error) coincide at the snapshot times with the numerical predictions by LSI. As in the previous experiment, we apply backward error analysis to the learned inverse modified Lagrangian L_{invmod} (truncating after second order terms) and then calculate $H_{\mathrm{LSI}}^{[[2]]}$. A contour plot of $V_{\mathrm{LSI}}^{[[2]]}(q) = H_{\mathrm{LSI}}^{[[2]]}(q, (0, 0))$ approximating the contours of the exact potential V_{ref} are plotted in Figure 9. The plots confirm that LSI can successfully be used for system identification. Moreover, it verifies that the learned system is close to the exact system and that all contours away from the separatrix have the correct topology. This shows that motions initialised not too close to the separatrix qualitatively show the correct long-term behaviour. However, very close to the separatrix there are incorrectly connected contour lines leaving small channels for the numerical motions with close to critical energy to escape erroneously into another region of the phase space. This explains why the LSI prediction of the experiment shown in the previous figures diverges eventually.

4. Inverse modified Lagrangians

In this section, we review variational backward error analysis which was developed in [3]. Its methods are needed for the system identification part of Lagrangian Shadow Integration. Moreover, we develop an inverse version of
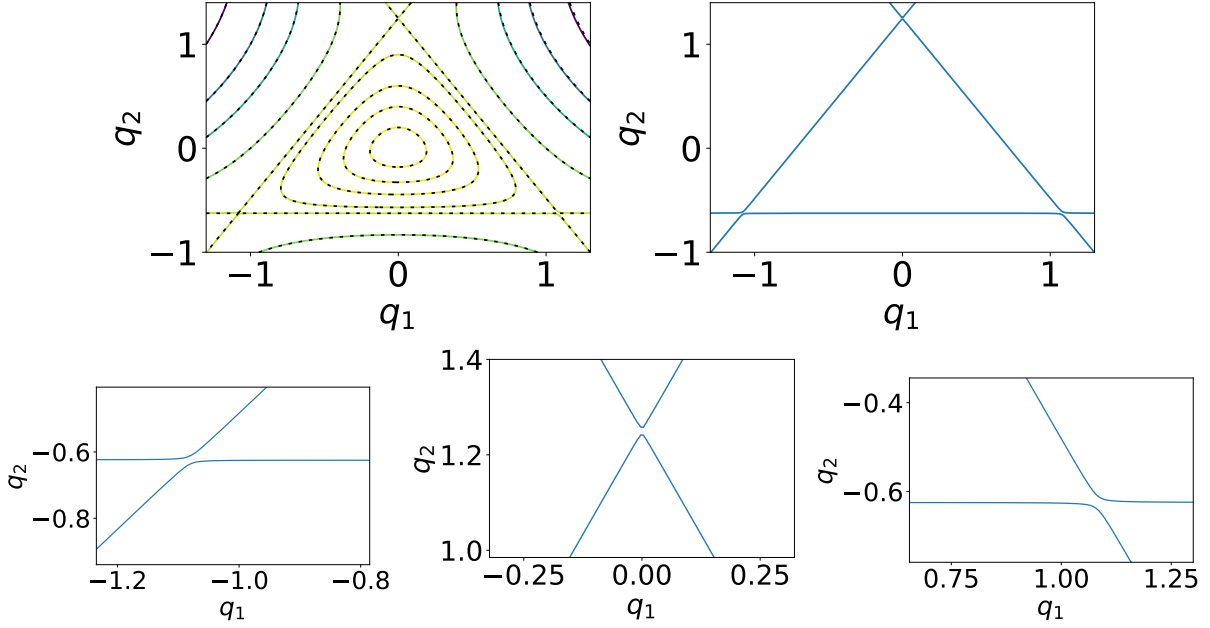
15

Figure 9: The contour lines of the numerical energy (solid lines) closely approximate the contour lines of the energy of the exact system (black dashed lines). This is visualised by viewing the section of the phase space where $\dot{q} = 0$ (top left). However, some level sets close to the separatrix of the exact system are incorrectly connected (top right). Close-up plots of the top right figure are displayed in the second row.

variational backward error analysis and prove the existence of inverse modified Lagrangians in the setting of formal power series. This provides a theoretical justification of why Lagrangian Shadow Integration works.

Let us fix notation and review and extend classical notions in Lagrangian mechanics and in the theory of variational integrators to formal power series. In the following, Q denotes an analytic manifold with tangent bundle TQ and $C^a(Q, \mathbb{R})$ denotes the linear space of analytic maps from Q to the real line $\mathbb{R}$. The space of k-jets over Q is denoted by $\text{Jet}^k(Q)$ for $k \in \mathbb{N} \cup \{\infty\}$. We will make use of a local identifications of Q with a subset of $\mathbb{R}^n$ and corresponding trivialisations of TQ and the jet-spaces. Local coordinates on Q are denoted by $q = (q_1, \ldots q_n)$ and coordinates on the jet space by $q^{[k]} = (q, \dot{q}, \ddot{q}, q^{(3)}, q^{(4)}, \ldots)$.

Definition 4.1 (Euler-Lagrange operator). *We define the linear operator $\mathcal{E}$ by*

$$\mathcal{E}(L) = \sum_{j=0}^{k} (-1)^j \frac{\mathrm{d}^j}{\mathrm{d}t^j} \frac{\partial L}{\partial q^{(j)}}$$

for $L \in C^a(\text{Jet}^k(Q), \mathbb{R})$. Here $\frac{\mathrm{d}^j}{\mathrm{d}t^j}$ denotes the total derivative. The equation $\mathcal{E}(L) = 0$ is called Euler-Lagrange equation *and $\mathcal{E}$ Euler-Lagrange operator. The definition of $\mathcal{E}$ can be linearly extended to formal power series $L = \sum_{j=0}^{\infty} s^j L_j \in C^a(\text{Jet}^k(Q), \mathbb{R})[[s]]$ in a formal variable s.*

Definition 4.2 (regular Lagrangian). *A Lagrangian $L \in C^a(TQ, \mathbb{R})$ is* regular *or* non-degenerate, *if $\left(\frac{\partial^2 L}{\partial \dot{q}^i \partial \dot{q}^j}\right)_{i,j}$ is invertible at $h = 0$. A Lagrangian $L = \sum_{j=0}^{\infty} s^j L_j \in \text{Jet}^\infty(Q)[[s]]$ with $L_0 \in C^a(TQ, \mathbb{R})$ is* regular *or* non-degenerate, *if L_0 constitutes a regular Lagrangian.*

Proposition 4.1. *Let $\sum_{j=0}^{\infty} s^j L_j \in C^a(\text{Jet}^\infty(Q), \mathbb{R})[[s]]$ be a regular Lagrangian, where $L_0 \in C^a(TQ, \mathbb{R})$ and $L_j \in C^a(\text{Jet}^{m_j}(Q), \mathbb{R})$ for $m_j < \infty$ for all j. There exists a formal power series $\sum_{j=0}^{\infty} s^j g_j \in C^a(TQ, \mathbb{R})$ such that for any truncation index k all solutions to the second order equation*

$$\ddot{q} = \sum_{j=0}^{k} s^j g_j(q, \dot{q}) \tag{22}$$

16

solve the Euler-Lagrange equation

$$\mathcal{E}\left(\sum_{j=0}^{k} s^j L_j\right) = 0 \tag{23}$$

up to an error of order $O(s^{k+1})$.

Proof. Let k be a truncation index and $m_{\max} = \max(m_0, m_1, \ldots, m_k)$. Since $L_0 \in C^a(TQ, \mathbb{R})$ is a regular Lagrangian, the Euler-Lagrange equations (23) are equivalent to a differential equation of the form

$$\ddot{q} = g_0(q, \dot{q}) + \sum_{j=1}^{k} s^j \bar{g}_j(q, \dot{q}, \ddot{q}, \ldots, q^{(m_{\max})}). \tag{24}$$

Time derivatives of (24) yield expressions for $\ddot{q}, \ldots, q^{(k_{\max})}$. Repeated substitutions of these expressions into the right hand side of (24) yield after finitely many steps

$$\ddot{q} = \sum_{j=0}^{k} s^j g_j(q, \dot{q}) + O(s^{k+1}). \tag{25}$$

Solutions to (22) solve (25) and thus the Euler–Lagrange equations (23) up to an error term of order $O(s^{k+1})$. $\qquad\square$

Definition 4.3 (variational integrator). *A variational integrator I is a bounded linear map*

$$I \colon C^a(TQ, \mathbb{R}) \to C^a(Q \times Q \times \mathbb{R}, \mathbb{R}).$$

Remark 4.1. *Images of variational integrators $L_\Delta = I(L) \in C^a(Q \times Q \times \mathbb{R}, \mathbb{R})$ are called* discrete Lagrangians. *Their last input argument is interpreted as a discretisation parameter or step-size. It is sometimes implicit in the notation $L_\Delta(q_0, q_1) = L_\Delta(q_0, q_1, h)$.*

Definition 4.4 (consistent variational discretisation/integrator). *A discrete Lagrangian $L_\Delta \in C^a(Q \times Q, \mathbb{R})$ is a consistent variational discretisation of a Lagrangian $L \in C^a(TQ, \mathbb{R})$ if it has variational order of at least 1, i.e. there exists an open set $U \subset TQ$ with compact closure and constants $c, \tau > 0$ such that for all solutions of the Euler–Lagrange equations*

$$0 = \mathcal{E}(L) = \frac{\partial L}{\partial q}(q, \dot{q}) - \frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial L}{\partial \dot{q}}(q, \dot{q})$$

with initial conditions $(q(0), \dot{q}(0)) \in U$ we have

$$\left\| I(L)(q(0), q(h)) - \int_0^h L(q(t), \dot{q}(t)) \mathrm{d}t \right\| \leq ch^2$$

for all $h \leq \tau$.

   A variational integrator I is consistent *if for all regular Lagrangians $L \in C^a(TQ, \mathbb{R})$ the discrete Lagrangian $I(L)$ is a consistent variational discretisation of L.*

Definition 4.5 ((consistent) variational integrator on power series). *Variational integrators can be defined on $C^a(TQ, \mathbb{R})[[s]]$ by a linear extension of Definition 4.3. Here $C^a(TQ, \mathbb{R})[[s]]$ denotes the ring of formal power series in the formal variable s over $C^a(TQ, \mathbb{R})$ (interpreted as a ring). A variational integrator I is* consistent *if for any $L = \sum_{j=0}^{\infty} s^j L_j \in C^a(TQ, \mathbb{R})[[s]]$ the discrete Lagrangian $I(L_j)$ is a consistent variational discretisation of L_j.*

Definition 4.6 (modified equation). *Consider a consistent variational discretisation $L_\Delta = I(L)$ of a regular Lagrangian L. An application of Proposition 4.1 to the series expansion of $L_\Delta(q(t), q(t+s))$ around $s = 0$ yields a formal power series $\sum_{j=0}^{\infty} s^j g_j$. The equation*

$$\ddot{q} = \sum_{j=0}^{\infty} s^j g_j(q, \dot{q}) \tag{26}$$

is called modified equation *for L_Δ.*

Remark 4.2. *Since L in Definition [4.6](#) is regular and the discretisation $L_\Delta = I(L)$ consistent, the zeroth order truncation of the modified equation ([26](#)) is equivalent to the Euler–Lagrange equation $\mathcal{E}(L) = 0$.*

Remark 4.3. *Solutions to optimal truncations of the modified equation ([26](#)) constitute solutions to the discrete Euler-Lagrange equations*

$$\nabla_2 L_\Delta(q(t-s), q(t)) + \nabla_1 L_\Delta(q(t), q(t+s)) = 0 \tag{27}$$

on exponentially long time intervals up to an exponentially small error in the step size s. In ([27](#)) $\nabla_1 L_\Delta$ and $\nabla_2 L_\Delta$ denote the derivatives with respect to the first or second input argument, respectively. See, for instance, [4, 26] for a rigorous analysis and optimal truncation techniques.

The following result was developed by Vermeeren in [3].

Theorem 4.2 (variational backward error analysis). *Consider a regular Lagrangian $L \in C^a(TQ, \mathbb{R})$ and a consistent variational discretisation $L_\Delta = I(L)$. There exists a formal power series L_{mod} in h, called* modified Lagrangian, *such that for any truncation $L_{\mathrm{mod}}^{[[k]]}$ after order $O(h^k)$ the Euler-Lagrange equations $\mathcal{E}(L_{\mathrm{mod}}^{[[k]]}) = 0$ solved for $\ddot{q}$ coincide with the modified equation to $I(L)$ up to an error of order $O(h^{k+1})$.*

Let us review the construction of L_{mod} from $L_\Delta = I(L)$ for a given maximal truncation index k.

Construction. Let $\mathcal{L}_\Delta^{[[k]]}(q^{[m]})$ be a series expansion around $h = 0$ of $L_\Delta(q(t-\frac{h}{2}), q(t+\frac{h}{2}), h)$ after order $O(h^k)$. Define the meshed Lagrangian

$$\mathcal{L}_{\mathrm{mesh}}^{[[k]]} = \sum_{j=0}^{\lfloor \frac{k}{2} \rfloor} \left(2^{(1-2j)} - 1\right) h^{2j} \frac{B_{2j}}{(2j)!} \frac{\mathrm{d}^{2j}}{\mathrm{d}t^{2j}} \mathcal{L}_\Delta^{[[k]]}(q^{[m]}), \tag{28}$$

where terms of order $O(h^{k+1})$ are truncated at the right hand side. Instances of $q^{(l)}$ with $l \geq 2$ are eliminated by iteratively substituting the modified equation for L_Δ and its derivatives and truncating terms of order $O(h^{k+1})$. The resulting expression is the modified Lagrangian $L_{\mathrm{mod}}^{[[k]]}$. $\qquad\square$

Remark 4.4. *The construction of $L_{\mathrm{mod}}^{[[k]]}$ is consistent with truncations, i.e. $L_{\mathrm{mod}}^{[[k]]} - L_{\mathrm{mod}}^{[[k+1]]} \in O(h^{k+1})$.*

We now introduce inverse modified Lagrangians L_{invmod} and prove their existence as formal power series in the discretisation parameter.

Theorem 4.3 (inverse variational backward error analysis). *Consider a regular Lagrangian $L \in C^a(TQ, \mathbb{R})$ and a consistent variational integrator I. There exists a formal power series L_{invmod} in the step size h such that for any truncation index k the following statements hold.*

- *The modified equation to $I(L_{\mathrm{invmod}}^{[[k]]})$ is equivalent to the exact Euler–Lagrange equations $\mathcal{E}(L) = 0$ up to terms of order $O(h^{k+1})$.*

- *The modified Lagrangian $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mod}}$ to the variational discretisation of $I(L_{\mathrm{invmod}}^{[[k]]})$ coincides with L up to terms of order $O(h^{k+1})$.*

Proof. The first two statements of the theorem follow directly from the last statement, which we now prove. Let k be a truncation index. We consider the ansatz

$$L_{\mathrm{invmod}}^{[[k]]} = L(q, \dot{q}) + sL_1(q, \dot{q}) + \ldots + s^k L_k(q, \dot{q}) \tag{29}$$

for an additional formal variable s, where $L_1, \ldots, L_k$ need to be determined. When applied to the ansatz $L_{\mathrm{invmod}}^{[[k]]}$, the construction of Theorem [4.2](#) yields $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mod}}^{[[k]]}$ which is a formal multivariate polynomial in h and s. We then set $s = h$, truncate terms of order $O(h^{k+1})$. We now show that one can derive analytic expressions for $L_1, \ldots, L_k$ by equating the coefficients of $h, h^2, \ldots, h^k$ successively with 0.

The meshed Lagrangian $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mesh}}^{[[k]]}$ is of the form

$$\mathfrak{L} := (L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mesh}}^{[[k]]} = I(L_{\mathrm{invmod}}^{[[k]]}) + h^2 \cdot O(s) = L(q, \dot{q}) + sL_1(q, \dot{q}) + \ldots + s^k L_k(q, \dot{q}) + (O(h) + h \cdot O(s))$$

since I is consistent. For $l \leq k$ let $\mathfrak{L}^{[[l]]}$ denote a truncation of the multivariate polynomial $\mathfrak{L}$ after terms of total degree l. Observe that the multivariate polynomial $\mathfrak{L}^{[[l]]}$ contains only one instance of L_l. More precisely, $L_l(q, \dot{q})$ is the coefficient of $s^l h^0$.

Using the regularity of L and consistency of I, the modified equation for $I(L_{\mathrm{invmod}}^{[[k]]})$ is of the form

$$\ddot{q} = \left(\frac{\partial^2 L}{\partial \dot{q} \partial \dot{q}}\right)^{-1} \left(\frac{\partial L}{\partial q} - \frac{\partial^2 L}{\partial q \partial \dot{q}}\dot{q}\right) + (O(h) + O(s)). \tag{30}$$

In this expression, L_l appears only in summands which have total degree at least l. All appearances of $\ddot{q}, q^{(3)}, \ldots$ in $\mathfrak{L}$ are in summands of total degree at least 1. When $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mod}}^{[[k]]}$ is obtained from $\mathfrak{L}$ by repeated substitution of the right hand side of (30) and its derivatives, then a truncation of $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mod}}^{[[k]]}$ to total degree l contains exactly one instance of $L_l(q, \dot{q})$ (the coefficient of $s^l h^0$). Now we set $s = h$ and truncate $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mod}}^{[[k]]}$ after terms of order $O(h^k)$. The discussion shows that equating the lth coefficient of $(L_{\mathrm{invmod}}^{[[k]]})_{\mathrm{mod}}^{[[k]]}$ with zero, we obtain an expression for $L_l(q, \dot{q})$ in terms of $L, L_1, \ldots, L_{l-1}$ and their derivatives. Setting $L_0 = L$, by an inductive argument, these expressions are analytic since L is analytic. $\qquad\square$

The exact discrete Lagrangian $L_\Delta^{\mathrm{exact}}$ to a regular Lagrangian L is defined as

$$L_\Delta^{\mathrm{exact}}(q_0, q_1) = \int_0^h L(q(t), \dot{q}(t)) \mathrm{d}t,$$

where q solves the Euler–Lagrange equations $\mathcal{E}(L) = 0$ with boundary conditions $q(0) = q_0$ and $q(h) = q_1$, where q_0, q_1 are sufficiently close such that the solution q exists and is unique.

A reformulation of Theorem 4.3 yields

Corollary 4.4. *To a consistent variational integrator I and a Lagrangian L there exists an inverse modified Lagrangian L_{invmod} such that the asymptotic expansion around 0 in the discretisation parameter h of $I(L_{\mathrm{invmod}})$ and the exact discrete Lagrangian to L coincide.*

Example 4.1. *As a variational integrator, consider the midpoint rule I defined by*

$$I(L)(q_0, q_1) = hL\left(\frac{q_0 + q_1}{2}, \frac{q_1 - q_0}{h}\right).$$

Using variational backward error analysis (Theorem 4.2) and inverse variational backward error analysis (Theorem 4.3) we obtain in the one-dimensional case

$$L_{\mathrm{mod}} = L + \frac{1}{24}h^2 \left(\frac{\left(L^{(1,0)} - L^{(1,1)}\dot{q}\right)^2}{L^{(0,2)}} - L^{(2,0)}\dot{q}^2\right) + O\left(h^4\right)$$

$$L_{\mathrm{invmod}} = L - \frac{1}{24}h^2 \left(\frac{\left(L^{(1,0)} - L^{(1,1)}\dot{q}\right)^2}{L^{(0,2)}} - L^{(2,0)}\dot{q}^2\right) + O\left(h^4\right).$$

Here $L^{(i,j)} = \frac{\partial^{i+j}}{(\partial q)^i (\partial \dot{q})^j}L(q, \dot{q})$. Higher order terms are more complicated and can be found along computational details in the corresponding GitHub repository.

Example 4.2. *Now consider the trapezoidal variational integrator I defined by*

$$I(L)(q_0, q_1) = \frac{h}{2}L\left(q_0, \frac{q_1 - q_0}{h}\right) + \frac{h}{2}L\left(q_1, \frac{q_1 - q_0}{h}\right).$$

The modified and inverse modified Lagrangians are given by

$$L_{\text{mod}} = L + \frac{1}{24}h^2\left(2L^{(2,0)}(\dot{q})^2 + \frac{\left(L^{(1,0)} - L^{(1,1)}\dot{q}\right)^2}{L^{(0,2)}}\right) + O\left(h^4\right)$$

$$L_{\text{invmod}} = L - \frac{1}{24}h^2\left(2L^{(2,0)}(\dot{q})^2 + \frac{\left(L^{(1,0)} - L^{(1,1)}\dot{q}\right)^2}{L^{(0,2)}}\right) + O\left(h^4\right)$$

Higher order terms are more complicated and can be found along computational details in the corresponding GitHub repository.

5. Conclusion

We have introduced the framework of Lagrangian Shadow Integration (LSI) and proved the existence of inverse modified Lagrangians in the sense of formal power series. The key novelty of LSI is that two types of discretisation errors are eliminated by design: (1) inverse modified Lagrangians can be learned directly from position data of trajectories. An approximation of velocities or acceleration data is not required. (2) Inverse modified Lagrangians compensate for the discretisation error of the integrator which is used to compute predictions of motions. Therefore, large step-sizes can be used while maintaining high accuracy. Further, LSI does not require prior knowledge about the form of the Lagrangian such that the method applies very generally. Moreover, the performance of LSI can be conveniently analysed using variational backward error analysis techniques. Additionally, system identification can be performed.

Future research directions are to embed symmetries into the learning process, for instance using GIM kernels or symmetric neural networks. An LSI prediction will then, by Noether's theorem, profit from exact conservation of conserved quantities related to variational symmetries, which should be beneficial especially for completely integrable systems. It appears plausible that an introduction of structure into learning approaches for dynamical systems increases its resilience to noise in the learning data, reduces data requirements, and that the compensation of discretisation errors, as performed by LSI, remains important for noisy data since discretisation errors tend to be biased in contrast to noise. A systematic investigation of these questions is part of future work. Moreover, the statistical framework of Gaussian Processes could be used to obtain uncertainty estimations of predictions obtained using LSI with GPs. Furthermore, Lagrangian neural networks [1] have successfully been applied to PDEs with variational structure. Practical applicability of LSI in this context can be investigated. A more theoretical research direction is to transfer optimal truncation results for modified Lagrangians and Hamiltonians to inverse modified Lagrangian.

References

[1] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, S. Ho, Lagrangian Neural Networks (2020). `arXiv:2003.04630`.

[2] T. Y. Takehiro Aoshima, Takashi Matsubara, Deep discrete-time Lagrangian mechanics, ICLR 2021 SimDL Workshop (2021). URL `https://simdl.github.io/files/49.pdf`

[3] M. Vermeeren, Modified equations for variational integrators, Numerische Mathematik 137 (4) (2017) 1001–1037. `doi:10.1007/s00211-017-0896-4`.

[4] E. Hairer, C. Lubich, G. Wanner, Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations, Springer Series in Computational Mathematics, Springer Berlin Heidelberg, 2013. `doi:10.1007/3-540-30666-8`.

[5] S. Ridderbusch, C. Offen, S. Ober-Blöbaum, P. Goulart, Learning ODE models with qualitative structure using Gaussian processes, in: 2021 60th IEEE Conference on Decision and Control (CDC), 2021, pp. 2896–2896. `doi:10.1109/CDC45484.2021.9683426`.

[6] T. S. Cohen, M. Geiger, M. Weiler, A General Theory of Equivariant CNNs on Homogeneous Spaces, Curran Associates Inc., Red Hook, NY, USA, 2019.

[7] T. Cohen, M. Welling, Group equivariant convolutional networks, in: M. F. Balcan, K. Q. Weinberger (Eds.), Proceedings of The 33rd International Conference on Machine Learning, Vol. 48 of Proceedings of Machine Learning Research, PMLR, New York, New York, USA, 2016, pp. 2990–2999. URL `https://proceedings.mlr.press/v48/cohenc16.html`

[8] R. Kondor, S. Trivedi, On the generalization of equivariance and convolution in neural networks to the action of compact groups, in: J. Dy, A. Krause (Eds.), Proceedings of the 35th International Conference on Machine Learning, Vol. 80 of Proceedings of Machine Learning Research, PMLR, 2018, pp. 2747–2755.
URL https://proceedings.mlr.press/v80/kondor18a.html

[9] N. Dehmamy, R. Walters, Y. Liu, D. Wang, R. Yu, Automatic symmetry discovery with lie algebra convolutional network, in: M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, J. W. Vaughan (Eds.), Advances in Neural Information Processing Systems, Vol. 34, Curran Associates, Inc., 2021, pp. 2503–2515.
URL https://proceedings.neurips.cc/paper/2021/file/148148d62be67e0916a833931bd32b26-Paper.pdf

[10] G. Evangelisti, S. Hirche, Physically consistent learning of conservative Lagrangian systems with Gaussian processes (submitted), in: 61st IEEE Conference on Decision and Control, 2022.

[11] C.-A. Cheng, H.-P. Huang, Learn the Lagrangian: A vector-valued RKHS approach to identifying Lagrangian systems, IEEE Transactions on Cybernetics 46 (12) (2016) 3247–3258. doi:10.1109/tcyb.2015.2501842.

[12] S. Greydanus, M. Dzamba, J. Yosinski, Hamiltonian Neural Networks, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Vol. 32, Curran Associates, Inc., 2019.
URL https://proceedings.neurips.cc/paper/2019/file/26cd8ecadce0d4efd6cc8a8725cbd1f8-Paper.pdf

[13] T. Bertalan, F. Dietrich, I. Mezić, I. G. Kevrekidis, On learning Hamiltonian systems from data, Chaos: An Interdisciplinary Journal of Nonlinear Science 29 (12) (2019) 121107. doi:10.1063/1.5128231.

[14] P. Jin, Z. Zhang, A. Zhu, Y. Tang, G. E. Karniadakis, Sympnets: Intrinsic structure-preserving symplectic networks for identifying hamiltonian systems, Neural Networks 132 (2020) 166–179. doi:10.1016/j.neunet.2020.08.017.
URL https://www.sciencedirect.com/science/article/pii/S0893608020303063

[15] K. Rath, C. G. Albert, B. Bischl, U. von Toussaint, Symplectic Gaussian process regression of maps in Hamiltonian systems, Chaos: An Interdisciplinary Journal of Nonlinear Science 31 (5) (2021) 053121. doi:10.1063/5.0048129.

[16] C. Offen, S. Ober-Blöbaum, Symplectic integration of learned Hamiltonian systems, Chaos: An Interdisciplinary Journal of Nonlinear Science 32 (1) (2022) 013122. doi:10.1063/5.0065913.

[17] P. Chartier, E. Hairer, G. Vilmart, Numerical integrators based on modified differential equations, Mathematics of Computation 76 (260) (2007) 1941–1954. doi:10.1090/s0025-5718-07-01967-9.

[18] M. David, F. Méhats, Symplectic learning for hamiltonian neural networks (2021). doi:10.48550/ARXIV.2106.11753.

[19] H. Qin, Machine learning and serving of discrete field theories, Scientific Reports 10 (2020) 19329. doi:10.1038/s41598-020-76301-0.

[20] R. I. McLachlan, C. Offen, Backward error analysis for conjugate symplectic methods (2022). arXiv:2201.03911.

[21] R. I. McLachlan, C. Offen, Backward error analysis for variational discretisations of PDEs, Journal of Geometric Mechanics (2022). doi:10.3934/jgm.2022014.

[22] J. E. Marsden, M. West, Discrete mechanics and variational integrators, Acta Numerica 10 (2001) 357–514. doi:10.1017/S096249290100006X.

[23] C. Offen, GitHub repository Christian-Offen/LagrangianShadowIntegration (2021).
URL https://github.com/Christian-Offen/LagrangianShadowIntegration

[24] C. E. Rasmussen, C. K. I. Williams, Gaussian Processes for Machine Learning, The MIT Press, 2005. doi:10.7551/mitpress/3206.001.0001.

[25] M. Henon, C. Heiles, The applicability of the third integral of motion: Some numerical experiments, The Astronomical Journal 69 (1964) 73. doi:10.1086/109234.

[26] B. Leimkuhler, S. Reich, Simulating Hamiltonian Dynamics, Cambridge University Press, 2005. doi:10.1017/cbo9780511614118.