Symplectic Spectrum Gaussian Processes: Learning Hamiltonians from Noisy and Sparse Data

Yusuke Tanaka **Tomoharu Iwata** **Naonori Ueda**
NTT Communication Science Laboratories
{yusuke.tanaka.rh,tomoharu.iwata.gy,naonori.ueda.fr}@hco.ntt.co.jp

Abstract

Hamiltonian mechanics is a well-established theory for modeling the time evolution of systems with conserved quantities (called *Hamiltonian*), such as the total energy of the system. Recent works have parameterized the Hamiltonian by machine learning models (e.g., neural networks), allowing Hamiltonian dynamics to be obtained from state trajectories without explicit mathematical modeling. However, the performance of existing models is limited as we can observe only noisy and sparse trajectories in practice. This paper proposes a probabilistic model that can learn the dynamics of conservative or dissipative systems from noisy and sparse data. We introduce a Gaussian process that incorporates the symplectic geometric structure of Hamiltonian systems, which is used as a prior distribution for estimating Hamiltonian systems with additive dissipation. We then present its spectral representation, *Symplectic Spectrum Gaussian Processes (SSGPs)*, for which we newly derive random Fourier features with symplectic structures. This allows us to construct an efficient variational inference algorithm for training the models while simulating the dynamics via ordinary differential equation solvers. Experiments on several physical systems show that SSGP offers excellent performance in predicting dynamics that follow the energy conservation or dissipation law from noisy and sparse data.

1 Introduction

There is great interest in the data-driven approach for learning the dynamics of physical systems. Modeling, simulating, and forecasting dynamics from data are fundamental in engineering and physical sciences; the data-driven approach also has the potential to discover new laws in physics [17, 38]. Many real-world systems can be described by ordinary differential equations (ODEs). Classical data-driven approaches make strong assumptions about the form of the equations; thus, they are not applicable if the system's equation is unknown [8, 32, 33, 37, 42]. The pioneering work of neural ordinary differential equation (NODE) [4] has been presented as a general *black-box* model for learning *vector fields*, represented by functions, to output time derivatives of the system's state; the state of the system is continuously transformed along the vector fields (see the right part of Figure 1). The Hamiltonian neural network (HNN) [14] and variants (e.g., [6, 12, 49]) allow one to learn vector fields such that the total energy of the system (called *Hamiltonian*) is conserved. This formulation is advantageous for learning dynamics that follow the basic laws of physics (i.e., the energy conservation law). Several extended models cover Hamiltonian systems with additive dissipation [26, 48]. However, neural network based models implicitly assume that a large amount of training data with high temporal resolution is available.

This paper addresses the problem of learning Hamiltonian dynamics from noisy and sparse trajectories and predicting the dynamics from an arbitrary initial condition. We illustrate our problem in Figure 1.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

The state trajectories are assumed to be *sparse* in the following senses: 1) Limited observation trials and 2) low temporal resolution (see Figure 1). One promising approach to alleviate overfitting in sparse settings is a Gaussian process (GP) [35], which allows for model learning while considering that the data contains uncertainties. GP models have been proposed for inferring unknown dynamics from trajectories [16, 28, 29], in which GPs are used for modeling the vector fields. They utilize a GP approximation with random Fourier features (RFFs) [34] to learn the dynamics via ODE solvers. The approximation is needed for the following reasons. Since we can obtain only trajectory data, not derivative observations (i.e., direct observations of vector fields), we cannot use the standard GP posterior conditioned on data points; thus, we require posterior approximation. The RFF-based approximation avoids the prohibitive computational complexity of generating sample paths (i.e., realizations of the vector field) from the GP posterior (details are discussed in [45, 46]). This advantage in computation is essential for ODE-based learning. Since they do not, however, consider Hamiltonian mechanics, it is difficult to accurately capture dynamics that follow physical laws such as energy conservation and dissipation.

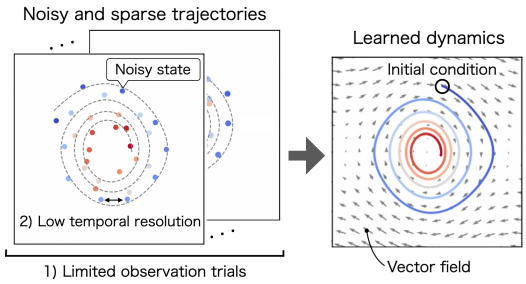

Figure 1: Problem setting. Color indicates time-evolution, starting at blue and ending at red.

GP regression models for derivative observations have been proposed [39], in which the covariance function is given by matrix-valued kernels for estimating vector fields. Recently, symplectic Gaussian process regression (SympGPR) [36] has been proposed; the covariance function for learning conservative vector fields is derived by incorporating the theory of Hamiltonian mechanics into GP modeling. However, SympGPR assumes that one can obtain derivative observations, not trajectories, for training. Although it can be naively applicable to our problem by using finite differences, it is difficult to learn the dynamics accurately from noisy state observations with low temporal resolution. Also, it is not applicable to Hamiltonian systems with dissipation.

This paper proposes a GP framework that can infer Hamiltonian systems with additive dissipative terms from noisy and sparse trajectories. We first extend the GP prior for Hamiltonian dynamics proposed in [36] to handle additive dissipation. The conserved quantity of the systems (i.e., Hamiltonian) is assumed to be a single-output GP. By employing Hamilton's equations, the vector fields are derived as a multi-output GP whose covariance function incorporates the *symplectic structure* for the energy conservation law. Moreover, one can handle Hamiltonian dynamics with friction by introducing the dissipation matrix. The most significant contribution of this work is the derivation of the RFFs that encode the symplectic structure to obtain the GP approximation for dissipative Hamiltonian systems, which we call *Symplectic Spectrum Gaussian Processes (SSGPs)*. This can efficiently generate samples of the Hamiltonian vector fields from the GP posterior and construct the scalable learning algorithm based on ODE solvers. Inference in SSGP is based on variational Bayes, which allows for learning the Hamiltonian vector fields while considering the uncertainties posed by noisy and sparse data. Another benefit of SSGP is that it can be used for decomposing the dynamics into conservative and dissipative terms allowing the dynamics for unseen friction coefficients to be predicted. Such a task cannot be handled with models (e.g., NODE) that do not use prior knowledge available in physics.

The main contributions of this work are as follows:

- We introduce a GP prior for modeling Hamiltonian systems with additive dissipation.
- We propose its spectral representation (called SSGP) by deriving RFFs that incorporate the symplectic structure of Hamiltonian systems.
- We develop a variational inference procedure for SSGP that offers numerical integration by ODE solvers as a subroutine.
- Experiments on several physical systems show that SSGP can accurately predict the dynamics that follow the conservation or dissipation laws from noisy and sparse trajectories [1].

[1] Code is available at `https://github.com/yusuk-e/SSGP`

Table 1: Comparison of the proposed model (SSGP) with the existing models. We compare the models for four items. The first and second items represent that the models can use the physics priors, i.e., (a) energy conservation law and (b) energy dissipation law, respectively, as an inductive bias for training. The third item (c) states that the models can be learned using ODE solvers from trajectory data; the models without a checkmark require derivative observations for training. The last item (d) shows that the models can handle uncertainties present in data.

	HNN [15]	D-HNN [9, 40]	NODE [4]	SymODEN [49]	D-SymODEN [48]	SympGPR [36]	SSGP
(a) Energy conservation law	✓	✓		✓	✓	✓	✓
(b) Energy dissipation law		✓			✓		✓
(c) Learning with ODE solver			✓	✓	✓		✓
(d) Uncertainty						✓	✓

2 Related Work

This work is positioned in the series of studies for learning *black-box* models that can estimate Hamiltonian dynamics from data. Here, black-box means that explicit mathematical modeling of differential equations is not required. We describe below prior works in this research line. Table 1 shows the comparison of the proposed model with the existing representative models.

Neural network models. Time-evolution of system states is generally described by differential equations. The neural ordinary differential equation (NODE) [4] and its variants [3, 20] have been proposed for learning the continuous-time evolution of the states from data. In these studies, the vector fields that determine the state evolution are parameterized by deep neural networks (DNNs) and estimated by back-propagating errors of observed trajectories via the ODE solver. The Hamiltonian neural network (HNN) [15] introduced prior knowledge of Hamiltonian mechanics as an inductive bias for training DNNs; the core concept is to parameterize the Hamiltonian (i.e., energy function) using DNNs. This formulation allows one to obtain dynamics that follow the energy conservation laws. Although the HNN assumes derivative observations for training, a learning procedure with ODE solvers from state trajectories was proposed in [6, 49]. Several models have been extended to handle Hamiltonian systems with additive dissipation [40, 48]. Most recent studies have further expanded the scope of application, including Hamiltonian systems with controllable inputs [9], stiff Hamiltonian dynamics [24], energy-conserving partial differential equation systems [26], odd-dimensional chaotic systems [7], and Poisson systems [5, 18]. Another approach, SympNets [19, 47], models the symplectic map using neural networks, which are shown to be universal approximators. However, existing DNN-based models implicitly assume a situation wherein a large amount of training data with high temporal resolution is available; they might fail to capture dynamics in the noisy and sparse settings this work focuses on.

Gaussian process models. Gaussian process (GP) modeling has the advantage that it makes it possible to learn models while considering uncertainties present in data [35]. Solak et al. [39] proposed GP regression models for derivative observations, in which the covariance function is given by matrix-valued kernels for inferring vector fields. Curl-free and the divergence-free kernels have been introduced for learning conservative vector fields [1, 25]; however, they lack the flexibility to be extended to cover various physical dynamics as described in the previous paragraph. Rath et al. extended the model in [39] to learn Hamiltonian vector fields, which is called symplectic Gaussian process regression (SympGPR) [36]. In SympGPR, the conserved quantity (i.e., Hamiltonian) is defined by a single-output GP; the covariance function for Hamiltonian vector fields can then be naturally derived on the basis of Hamilton's equations. However, these GP regression models assume derivative observations, not state trajectories. Although the use of finite differences allows for naively applying them to the problem of learning from trajectories, they might fail to capture dynamics, especially when the temporal resolution is lower, because each finite difference is computed from state pairs at neighboring discrete time steps. Most recently, a few models combined GPs with symplectic integrators [11, 30]; however, all the GP models [11, 30, 36] are not applicable to Hamiltonian systems with additive dissipation.

Learning with ODE solvers. Numerical integration methods play an essential role in learning dynamics from trajectory data. Chen et al. [4] provide a convenient tool for solving ODEs and

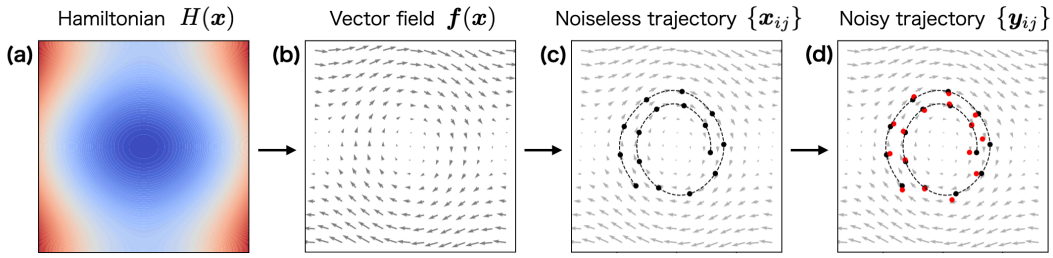

Figure 2: Schematic diagram of SSGP: Generative processes of noisy trajectories. (a) We model the unknown Hamiltonian $H(\boldsymbol{x})$ by a single-output GP. Here, the color represents the magnitude of energy. (b) The vector field $\boldsymbol{f}(\boldsymbol{x})$ is calculated by applying the differential operator $\mathcal{L}$ to $H(\boldsymbol{x})$. (c) We sample the initial condition from the standard Gaussian distribution and solve the ODE defined by $\boldsymbol{f}(\boldsymbol{x})$ to obtain the noiseless trajectory $\{\boldsymbol{x}_{ij}\}$ depicted by black dots. (d) The noisy trajectory $\{\boldsymbol{y}_{ij}\}$ depicted by red dots is observed by adding Gaussian noise.

learning parameters via back-propagation or memory-efficient adjoint methods, which are widely used for learning Hamiltonian dynamics (e.g., [9, 12, 49]). One can also use advanced learning methods based on symplectic integrators [10, 27, 50] or discrete gradients [26], which can evaluate the loss function and its gradient while preserving energy accurately. These studies are related to ours but different in the research focus. They focus on how to estimate the parameters given the models (e.g., DNNs); whereas our proposal is a model whose learning method can be chosen as needed. It might be possible to employ the above techniques for learning our model. Different from the above studies, this work can also consider uncertainties to learn the model from noisy and sparse data.

3 Hamiltonian Mechanics

In this section, we briefly review Hamiltonian mechanics [13]. Let us consider a system with N degrees of freedom. In the Hamiltonian formalism, the continuous-time evolution of the system is described in phase space, that is, the product space of generalized coordinates $\boldsymbol{x}^{\mathrm{q}} = (x_1^{\mathrm{q}}, \ldots, x_N^{\mathrm{q}})$ and generalized momenta $\boldsymbol{x}^{\mathrm{p}} = (x_1^{\mathrm{p}}, \ldots, x_N^{\mathrm{p}})$. Let $\boldsymbol{x} = (\boldsymbol{x}^{\mathrm{q}}, \boldsymbol{x}^{\mathrm{p}}) \in \mathbb{R}^D$ be a state of the system, where $D = 2N$. The system's evolution is determined by the Hamiltonian $H(\boldsymbol{x}) : \mathbb{R}^D \to \mathbb{R}$, which denotes the system's total energy. Traditionally, the Hamiltonian is *manually designed* to suit the system. The dynamics of a Hamiltonian system with additive dissipative terms is given by

$$\frac{d\boldsymbol{x}}{dt} = (\mathbf{S} - \mathbf{R})\nabla H(\boldsymbol{x}) =: \boldsymbol{f}(\boldsymbol{x}), \quad \text{where} \quad \mathbf{S} = \begin{pmatrix} \mathbf{O} & \mathbf{I} \\ -\mathbf{I} & \mathbf{O} \end{pmatrix}. \tag{1}$$

Here, $\nabla H(\boldsymbol{x}) : \mathbb{R}^D \to \mathbb{R}^D$ is the gradient of the Hamiltonian with respect to state $\boldsymbol{x}$, $\mathbf{S} \in \mathbb{R}^{D \times D}$ is the skew-symmetric matrix, $\mathbf{R} \in \mathbb{R}^{D \times D}$ is the positive semi-definite dissipation matrix, $\mathbf{I}$ is the identity matrix, and $\mathbf{O}$ is the zero matrix. In (1), we define the time derivatives of the state by the function $\boldsymbol{f}(\boldsymbol{x}) : \mathbb{R}^D \to \mathbb{R}^D$, which is a special kind of vector field that has a symplectic geometric structure (called Hamiltonian vector field or symplectic gradient). The dynamics on this vector field conserve the total energy when $\mathbf{R} = \mathbf{O}$. One example of the dissipation matrix is $\mathbf{R} = \mathrm{diag}(0, \ldots, 0, r_1, \ldots, r_N)$, representing a dissipative system with friction coefficient $r_n \geq 0$. Although we assume this kind of dissipation matrix in the following, this work is extendable to the general dissipation matrix, such as the state-dependent damping term [9]. Given vector field $\boldsymbol{f}(x)$ and initial condition $\boldsymbol{x}_1$ at time t_1, one can predict state $\boldsymbol{x}_t$ at time t by integrating $\boldsymbol{f}(\boldsymbol{x})$ from t_1 to t, as follows: $\boldsymbol{x}_t = \boldsymbol{x}_1 + \int_{t_1}^{t} \boldsymbol{f}(\boldsymbol{x})\, dt$.

4 Model

We propose SSGP (Symplectic Spectrum Gaussian Process), a probabilistic model for learning Hamiltonian systems with additive dissipation from noisy and sparse trajectories. Figure 2 shows a schematic diagram of the proposed generative processes of noisy trajectories. In the following, we first introduce a Gaussian process (GP) prior for modeling conservative and dissipative vector fields

by incorporating the theory of Hamiltonian mechanics. We then derive its spectral representation, for which we propose random Fourier features that encode symplectic structures. This spectral representation is used for constructing a scalable inference algorithm described in Section 5. Finally, we describe the generative processes of noisy state observations.

GP priors for Hamiltonian systems with additive dissipation. In the proposed model, the *unknown* Hamiltonian $H(\boldsymbol{x})$ is assumed to be a single-output GP with zero mean. Let $\mathcal{L} := (\mathbf{S} - \mathbf{R})\nabla$ denote a differential operator. According to (1), the vector field can be represented using $\mathcal{L}$, as follows:

$$\boldsymbol{f}(\boldsymbol{x}) = \mathcal{L}\, H(\boldsymbol{x}), \quad \text{where} \quad H(\boldsymbol{x}) \sim \mathcal{GP}(0, \gamma(\boldsymbol{x}, \boldsymbol{x}')). \tag{2}$$

Here, $\gamma(\boldsymbol{x}, \boldsymbol{x}') : \mathbb{R}^D \times \mathbb{R}^D \to \mathbb{R}$ is a covariance function. Since differentiation is a linear operator, the derivative of a GP is again a GP [39]; thus, $\boldsymbol{f}(\boldsymbol{x})$ is given by a multi-output GP,

$$\boldsymbol{f}(\boldsymbol{x}) \sim \mathcal{GP}(\mathbf{0}, \mathbf{K}(\boldsymbol{x}, \boldsymbol{x}')), \tag{3}$$

where $\mathbf{0}$ is a column vector of 0's, and $\mathbf{K}(\boldsymbol{x}, \boldsymbol{x}') : \mathbb{R}^D \times \mathbb{R}^D \to \mathbb{R}^{D \times D}$ is the matrix-valued covariance function represented by

$$\mathbf{K}(\boldsymbol{x}, \boldsymbol{x}') = \mathcal{L}\mathcal{L}^\top \gamma(\boldsymbol{x}, \boldsymbol{x}') = (\mathbf{S} - \mathbf{R})\nabla^2 (\mathbf{S} - \mathbf{R})^\top \gamma(\boldsymbol{x}, \boldsymbol{x}'). \tag{4}$$

Here, ∇^2 is the Hessian operator. As described above, we can obtain the GP prior (3) with the covariance function (4) that encodes the geometric structure of Hamiltonian systems with additive dissipation. This formulation can be regarded as a generalization of the GP for non-dissipative Hamiltonian systems [36]. Though one can opt for any positive definite kernel as $\gamma(\boldsymbol{x}, \boldsymbol{x}')$, we assume that $\gamma(\boldsymbol{x}, \boldsymbol{x}')$ is shift-invariant in the approximation described in the next paragraph. One of the most widely used shift-invariant kernels is the ARD (Automatic Relevance Determination) Gaussian kernel,

$$\gamma(\boldsymbol{x}, \boldsymbol{x}') = \sigma_0^2 \exp\left(-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}')^\top \boldsymbol{\Lambda}^{-1} (\boldsymbol{x} - \boldsymbol{x}')\right), \quad \text{where} \quad \boldsymbol{\Lambda} = \operatorname{diag}\left(\lambda_1^2, \ldots, \lambda_D^2\right). \tag{5}$$

Here, $\sigma_0^2 \in \mathbb{R}_{>0}$ and $\lambda_d^2 \in \mathbb{R}_{>0}$ are the signal variance and the length scale, respectively.

Spectral representations. We present the approximation of the GP prior (3), for which we newly derive random Fourier features (RFF) that encode symplectic structures of Hamiltonian systems. The RFF-based approximation is advantageous for 1) estimating the GP posterior for vector fields without derivative observations and 2) efficiently sampling the vector field from the GP posterior, which has been suggested in recent studies on learning ODEs [16, 28, 29, 45]. These advantages make it feasible to train our GP model for Hamiltonian systems with dissipation by utilizing ODE solvers, even when only trajectory data, not derivative observations, are available (see also the **On computational complexity** paragraph in Section 5 and Appendix E). We begin by modeling the Hamiltonian $H(\boldsymbol{x})$ as a GP approximation,

$$H(\boldsymbol{x}) = \sum_{m=1}^{M} \boldsymbol{w}_m \boldsymbol{\phi}_m(\boldsymbol{x}), \quad \text{where} \quad \boldsymbol{w}_m \sim \mathcal{N}\left(\mathbf{0}, \frac{\sigma_0^2}{M}\mathbf{I}\right). \tag{6}$$

Here, we adopt the M pairs of basis functions[2], $\boldsymbol{\phi}_m(\boldsymbol{x}) = \left[\cos(2\pi \boldsymbol{s}_m^\top \boldsymbol{x}), \sin(2\pi \boldsymbol{s}_m^\top \boldsymbol{x})\right]^\top$, parameterized by a D-dimensional column vector $\boldsymbol{s}_m$ of spectral points. The point $\boldsymbol{s}_m$ is sampled from the kernel's spectral density, which is given by

$$p(\boldsymbol{s}) = \mathcal{N}\left(\mathbf{0}, (4\pi^2 \boldsymbol{\Lambda})^{-1}\right), \tag{7}$$

in the case in which the ARD Gaussian kernel (5) is used [23]. In (6), $\boldsymbol{w}_m \in \mathbb{R}^2$ is a row vector of weights for the mth pair of basis. We apply the differential operator $\mathcal{L}$ to the approximation (6); which yields the spectral representation (i.e., SSGP) of (3), represented by

$$\boldsymbol{f}(\boldsymbol{x}) = \mathcal{L}\, H(\boldsymbol{x}) =: \boldsymbol{\Psi}(\boldsymbol{x}) \boldsymbol{w}^\top, \tag{8}$$

where $\boldsymbol{w} = (\boldsymbol{w}_1, \ldots, \boldsymbol{w}_M)$, and the resulting feature maps are defined by $\boldsymbol{\Psi}(\boldsymbol{x}) = (\boldsymbol{\Psi}_1(\boldsymbol{x}), \ldots, \boldsymbol{\Psi}_M(\boldsymbol{x}))$; the mth feature map $\boldsymbol{\Psi}_m(\boldsymbol{x}) : \mathbb{R}^D \to \mathbb{R}^{D \times 2}$ is given by

$$\boldsymbol{\Psi}_m(\boldsymbol{x}) = 2\pi(\mathbf{S} - \mathbf{R})\boldsymbol{s}_m \left[-\sin(2\pi \boldsymbol{s}_m^\top \boldsymbol{x}), \cos(2\pi \boldsymbol{s}_m^\top \boldsymbol{x})\right], \tag{9}$$

[2] An alternative version of basis functions [41] is known, and our proposal can easily adopt the alternative.

which we call *symplectic random Fourier features (S-RFF)*. The derivation of (8) is described in Appendix A. S-RFF (9) is the first to incorporate symplectic structures into the random features for modeling vector fields by leveraging the knowledge of Hamiltonian mechanics. This result is not trivial in that we bridge two clearly distinct research topics, random features for kernel machines (i.e., GPs) and Hamiltonian mechanics. Let us explore the connection between the exact GP (3) and its spectral approximation (8). Denoting Dirac's delta function by $p(\boldsymbol{f}|\boldsymbol{w})$, the distribution of $\boldsymbol{f}$ is given by integrating out $\boldsymbol{w}$ from (8), as follows[3]:

$$p(\boldsymbol{f}) = \int p(\boldsymbol{f}|\boldsymbol{w})p(\boldsymbol{w})d\boldsymbol{w} = \mathcal{N}\left(\mathbf{0}, \tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}')\right), \quad \text{where} \quad \tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}') = \frac{\sigma_0^2}{M}\boldsymbol{\Psi}(\boldsymbol{x})\boldsymbol{\Psi}(\boldsymbol{x}')^\top. \quad (10)$$

Here, $\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}') : \mathbb{R}^D \times \mathbb{R}^D \to \mathbb{R}^{D\times D}$ is the covariance function for GP approximation. We observed that approximation quality improves with the number of spectral points (i.e., the number of basis functions) by comparing the gram matrix of $\mathbf{K}(\boldsymbol{x}, \boldsymbol{x}')$ (4) with that of $\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}')$ in (10), where the spectral points $\boldsymbol{s}_m$ are sampled from (7). We show the gram matrices in Appendix B.

Generative processes of noisy observations. Suppose that we have a collection of I trajectories $\{(t_{ij}, \boldsymbol{y}_{ij})|i = 1, \ldots, I; j = 1, \ldots, J_i\}$, where J_i is the number of samples in the ith trajectory. Each sample is specified by a pair $(t_{ij}, \boldsymbol{y}_{ij})$, which represents the observation of noisy state $\boldsymbol{y}_{ij}$ at time t_{ij}. We treat the noiseless state $\boldsymbol{x}_{ij}$, the counterpart of $\boldsymbol{y}_{ij}$, as a latent variable. We assume that the observation model of $\boldsymbol{y}_{ij}$ is a Gaussian distribution with a variance of σ^2. Letting $\mathbf{Y} = \{\boldsymbol{y}_{ij}\}$, the marginal likelihood (i.e., evidence) is given by

$$p(\mathbf{Y}) = \int p(\boldsymbol{f}) \prod_{i=1}^{I} \left[\int p(\boldsymbol{y}_{i1}|\boldsymbol{x}_{i1})p(\boldsymbol{x}_{i1}) \prod_{j=2}^{J_i} p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij})p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1})d\boldsymbol{x}_{i1} \right] d\boldsymbol{f}, \quad (11)$$

where $p(\boldsymbol{f})$ is the GP prior (10) of the vector field, and $p(\boldsymbol{x}_{i1})$ is the prior distribution[4] of the initial condition $\boldsymbol{x}_{i1}$. Given $\boldsymbol{f}$ and $\boldsymbol{x}_{i1}$, the state $\boldsymbol{x}_{ij}$ is deterministically given by solving the ODE; thus, we can write the conditional distribution $p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1})$ in (11) using Dirac's delta function, as follows:

$$p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) = \delta\left(\boldsymbol{x}_{ij} - \left[\boldsymbol{x}_{i1} + \int_{t_{i1}}^{t_{ij}} \boldsymbol{f}(\boldsymbol{x})\, dt\right]\right). \quad (12)$$

Note that, although we omit the observation time points $\{t_{ij}\}$ in (11), it is actually conditioned on $\{t_{ij}\}$. For simplicity, we adopt this notation hereinafter.

5 Inference

Although we would like to estimate the model parameters $\boldsymbol{w}$, $\boldsymbol{\Lambda}$, σ_0^2, σ^2 and $\mathbf{R}$ by maximizing the logarithm of (11), we cannot calculate the exact marginal likelihood (11) analytically as it includes the process of solving the ODEs. We then present a variational inference procedure for learning Hamiltonian systems with additive dissipation from noisy and sparse trajectories. The spectral representation of SSGP is used for efficient sampling of the conservative or dissipative vector fields from the GP posterior, allowing the learning to proceed with ODE solvers while handling uncertainties in the vector fields.

Parameter learning. We consider the following evidence lower bound (ELBO),

$$\log p(\mathbf{Y}) \geq \sum_{i=1}^{I} \left[\sum_{j=1}^{J_i} \mathbb{E}_{q(\boldsymbol{x}_{ij})} \left[\log p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij})\right] - \mathrm{KL}\left[q(\boldsymbol{x}_{i1})||p(\boldsymbol{x}_{i1})\right] \right] - \mathrm{KL}\left[q(\boldsymbol{w})||p(\boldsymbol{w})\right], \quad (13)$$

where $q(\boldsymbol{x}_{ij})$ and $q(\boldsymbol{w})$ are the variational distributions of the noiseless state $\boldsymbol{x}_{ij}$ and the weights $\boldsymbol{w}$, respectively. We assume that the variational distribution of $\boldsymbol{w}$ is given by a Gaussian distribution, $q(\boldsymbol{w}) = \mathcal{N}(\boldsymbol{b}, \mathbf{C})$, where $\boldsymbol{b} \in \mathbb{R}^{2M}$ and $\mathbf{C} \in \mathbb{R}^{2M\times 2M}$ are the mean and the covariance matrix, respectively. For $j = 1$, we assume that the variational distribution of initial condition $\boldsymbol{x}_{i1}$ is given

[3]The derivation procedure is similar to that described in [2, Chapter 4.5.2].

[4]For simplicity, the prior is set to the standard Gaussian distribution in the implementation.

by a Gaussian distribution with the mean being the observed state $\boldsymbol{y}_{i1}$, $q(\boldsymbol{x}_{i1}) = \mathcal{N}(\boldsymbol{y}_{i1}, \mathbf{A})$, where $\mathbf{A} \in \mathbb{R}^{D \times D}$ is the covariance matrix. For $j \geq 2$, the variational distribution of $\boldsymbol{x}_{ij}$ is given by

$$q(\boldsymbol{x}_{ij}) = \iint p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) \left[\int p(\boldsymbol{f}|\boldsymbol{w}) q(\boldsymbol{w}) d\boldsymbol{w} \right] q(\boldsymbol{x}_{i1}) d\boldsymbol{x}_{i1} d\boldsymbol{f}, \tag{14}$$

where the factor in square brackets is the variational distribution of $\boldsymbol{f}$. The derivation of (13) is described in Appendix C. We approximate the expectation in (13) using Monte Carlo integration,

$$\mathbb{E}_{q(\boldsymbol{x}_{ij})} \left[\log p\left(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij}\right) \right] \approx \frac{1}{K} \sum_{k=1}^{K} \log p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij}^{(k)}), \tag{15}$$

where the Monte Carlo samples are given by

$$\boldsymbol{x}_{i1}^{(k)} = \boldsymbol{y}_{i1} + \sqrt{\mathbf{A}}\, \boldsymbol{\epsilon}_i^{(k)}, \quad \text{where} \quad \boldsymbol{\epsilon}_i^{(k)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{16}$$

$$\boldsymbol{x}_{i2}^{(k)}, \ldots, \boldsymbol{x}_{iJ_i}^{(k)} = \text{ODESolve}\left(\boldsymbol{x}_{i1}^{(k)}, \boldsymbol{f}^{(k)}(\boldsymbol{x}), t_{i2}, \ldots, t_{iJ_i}\right), \tag{17}$$

where we use the reparameterization trick [21] in (16), and we perform the numerical integration by the ODE solvers (e.g., the Runge-Kutta method) in (17). Here, K is the number of Monte Carlo samples to obtain (15), and $\sqrt{\cdot}$ denotes a matrix square root. $\boldsymbol{f}^{(k)}(\boldsymbol{x})$ in (17) is the sample of the vector field generated from the variational posterior of $\boldsymbol{f}$,

$$\boldsymbol{f}^{(k)}(\boldsymbol{x}) = \boldsymbol{\Psi}(\boldsymbol{x}) \left[\frac{1}{L} \sum_{l=1}^{L} \boldsymbol{w}^{(k,l)} \right]^{\top}, \quad \text{where} \quad \boldsymbol{w}^{(k,l)} = \boldsymbol{b} + \sqrt{\mathbf{C}}\, \boldsymbol{\epsilon}^{(k,l)}. \tag{18}$$

Here, L is the number of Monte Carlo samples to obtain (18), and $\boldsymbol{\epsilon}^{(k,l)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. In the spectral representation, the randomness of function $\boldsymbol{f}$ is totally controlled by the distribution of weights $\boldsymbol{w}$. Accordingly, we can optimize both model parameters and variational parameters while considering uncertainties in the vector field by sampling $\boldsymbol{w}$ at each training iteration. The inference procedure is shown in Appendix D. We can easily extend our proposed approach to learn Hamiltonians from high-dimensional data (such as images) by combining an autoencoder with an SSGP, as in [14, 43]. The states $\boldsymbol{x}$ and $\boldsymbol{x}'$ in the covariance function $\mathbf{K}(\boldsymbol{x}, \boldsymbol{x}')$ (4) are modeled as the latent vectors of the autoencoder. This formulation can be regarded as an extension of the SSGP based on deep kernel learning [44].

Prediction. The variational posterior, $q(\boldsymbol{f}) = \int p(\boldsymbol{f}|\boldsymbol{w}) q(\boldsymbol{w}) d\boldsymbol{w}$, has a closed-form solution, a Gaussian distribution whose mean function and covariance function are given by $\tilde{\boldsymbol{m}}^*(\boldsymbol{x}) = \boldsymbol{\Psi}(\boldsymbol{x})\boldsymbol{b}^{\top}$ and $\tilde{\mathbf{K}}^*(\boldsymbol{x}, \boldsymbol{x}') = \boldsymbol{\Psi}(\boldsymbol{x})\mathbf{C}\boldsymbol{\Psi}^{\top}(\boldsymbol{x}')$, respectively. One can predict the dynamics from an arbitrary initial condition by numerically integrating the mean function $\tilde{\boldsymbol{m}}^*(\boldsymbol{x})$. Also, one can evaluate the predictive uncertainty for the vector field by using the covariance function $\tilde{\mathbf{K}}^*(\boldsymbol{x}, \boldsymbol{x}')$.

On computational complexity. The computational bottleneck is the sampling of $\boldsymbol{w}$; the cost of computing $\sqrt{\mathbf{C}}$ in (18) is $\mathcal{O}(M^3)$. The cost increases with the number of basis functions; however, the most important thing is that the bottleneck is outside of the ODE solver. The sampling of $\boldsymbol{w}$ makes it possible to efficiently evaluate function values at arbitrary point $\boldsymbol{x}$, as in (18), and to use it in the ODE solver in (17). The advantage in computational costs is discussed in Appendix E. In addition, previous studies have shown experimentally that the GP models approximated by RFF have, in many cases, high predictive performance even when setting M lower than 10^3 (e.g., [23]). Actually, our experiments show that SSGP yields an accurate prediction using a small set of RFFs. The average training time when setting $M = 250$ was 2943.0 seconds for the dataset of the pendulum with friction; the experiments were conducted on the AMD EPYC 7313 CPU (3.0GHz).

6 Experiments

Data. We evaluated the proposed model, SSGP, using two physical systems: *pendulum*, and *Duffing oscillator*. The Hamiltonian of the pendulum system is

$$H(\boldsymbol{x}) = 2mgl(1 - \cos x^{\mathrm{q}}) + \frac{l^2 (x^{\mathrm{p}})^2}{2m}, \tag{19}$$

where we denote the gravitational constant by g, the mass constant by m, and the length of the pendulum by l. In the experiments, we set $g = 3$ and $m = l = 1$. The Hamiltonian of the Duffing oscillator system is

$$H(\boldsymbol{x}) = \frac{1}{2}(x^{\mathrm{p}})^2 + \frac{\alpha}{2}(x^{\mathrm{q}})^2 + \frac{\beta}{4}(x^{\mathrm{q}})^4, \tag{20}$$

where we set the parameters, $\alpha = \beta = 1$. For each system, the dissipation matrix was set to $\mathbf{R} = \mathbf{O}$ (energy conservation) or $\mathbf{R} = \mathrm{diag}(0, 0.05)$ (energy dissipation). We generated trajectory data by employing a numerical integrator, i.e., the Dormand–Prince method with adaptive time-stepping, implemented in torchdiffeq[5] [3, 4]. We sampled initial conditions with total energies uniformly distributed across a predefined range. The energy ranges for pendulum and Duffing oscillator were $[1.3, 2.3]$ and $[0.5, 1.5]$, respectively. To evaluate the robustness of our model against the degree of sparsity, we prepared $\{10, 15, 20, 30, 50\}$ trajectories sampled at a frequency of $\{3, 5, 10\}$ Hz for 10 seconds, and added Gaussian noise with variance $\sigma^2 = 0.1$ to each sample. We randomly split the trajectory data and used 70% for training and 30% for validation. Here, if the number of trajectories for validation was less than five, we used five trajectories for validation and the rest for training. We generated a test set of 25 trajectories independently from training and validation sets; each test trajectory was sampled at a frequency of 100 Hz for 15 seconds. The experiments were conducted five times by resampling the training and validation sets.

Task 1: Normal prediction. We evaluated performance by comparing the predicted state trajectories $\{\hat{\boldsymbol{x}}_{ij}\}$ with the ground truth $\{\boldsymbol{x}_{ij}^{\mathrm{true}}\}$ (i.e., the test set described above). The evaluation metric is the mean squared error (MSE), $\frac{1}{I}\sum_{i=1}^{I}(\frac{1}{J_i}\sum_{j=1}^{J_i} \parallel \hat{\boldsymbol{x}}_{ij} - \boldsymbol{x}_{ij}^{\mathrm{true}} \parallel^2)$, where $\parallel \cdot \parallel^2$ is the Euclidean norm. In the evaluation, we used the datasets from four settings: pendulum and Duffing oscillator with or without energy dissipation.

Task 2: Predicting the dynamics for unseen friction coefficients. One benefit of SSGP is that it can decompose the dynamics into its conservative and dissipative terms, thus predicting the dynamics for arbitrary friction coefficients. To show this, we evaluated SSGP by the following procedure: 1) we trained the model using the dataset from pendulum or Duffing oscillator with the friction coefficients, $\mathbf{R} = \mathrm{diag}(0, 0.05)$; 2) we predicted its conservative system by using the learned Hamiltonian, where we set the friction coefficients, $\mathbf{R} = \mathbf{O}$. The evaluation metric is the same as that in **Task 1**.

Setup of the SSGP. We trained the model using the Adam optimizer [22] with learning rate of 10^{-3} for 10^4 epochs, implemented in PyTorch [31]. We performed numerical integration by the adaptive Dormand–Prince method [3, 4] with the relative and absolute tolerances of 10^{-8}. We set the numbers of Monte Carlo samples to $K = 1$ and $L = 100$. The number M of spectral points was chosen from $\{100, 250, 500\}$ based on the validation error. We used a block-diagonal approximation of $\mathbf{C}$, namely, $q(\boldsymbol{w})$ is factorized as two factors separating the weights for the sines and cosines, where the computational complexity is the same as that of the original. We fixed the friction coefficients $\mathbf{R} = \mathbf{O}$ when training the conservative systems in **Task 1**.

Baselines. We compared the SSGP with the existing models (see Table 1): Hamiltonian neural network (HNN) [15], dissipative HNN (D-HNN) [40], neural ordinary differential equation (NODE) [4], symplectic ODE-Net (SymODEN) [49], dissipative SymODEN (D-SymODEN) [48] and symplectic Gaussian process regression (SympGPR) [36]. Since HNN, D-HNN and SympGPR require derivative observations for training, we used the finite difference instead. Another baseline is the GP model using standard random Fourier features (RFF), which corresponds to the case that the vector field $\boldsymbol{f}(\boldsymbol{x})$ in SSGP is modeled as a multi-output GP approximated by the standard RFF (not considering Hamiltonian mechanics). In the experiments, we call it *RFF*. The details of each baseline are described in Appendix F.

Results. We present the experimental results that are picked up from the case where the sampling frequency was 5 Hz. Appendix G shows all the results, including those for the frequencies of 3 and 10 Hz. Figure 3 shows MSE and standard errors for SSGP and the baselines in **Task 1** and **Task 2**. As expected, in many cases, MSE decreased as the number of trajectories increased. The performance of HNN, D-HNN and SympGPR was limited because their training was based on finite differences. NODE and RFF had large errors except for the case of Figure 3 **(b)** because they cannot use the prior knowledge of Hamiltonian mechanics as the inductive bias. SymODEN and D-SymODEN were more competitive with SSGP; nevertheless, SSGP matched or bettered their predictive performance.

[5] https://github.com/rtqichen/torchdiffeq (MIT License)

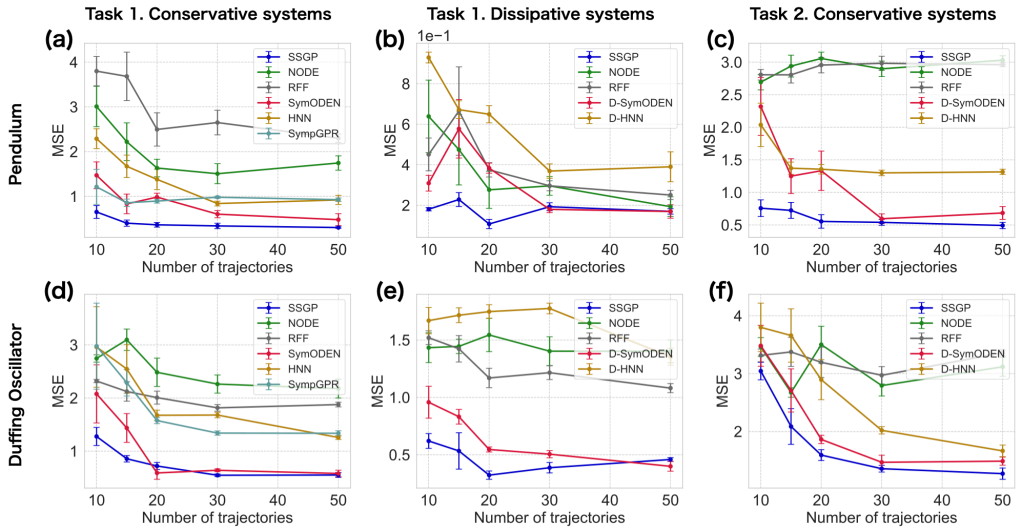

Figure 3: MSE and standard errors for the predicted trajectories when the sampling frequency was 5 Hz. The first and second columns are the results for conservative and dissipative systems, respectively, in **Task 1**. The third column holds the results predicted for conservative systems in **Task 2**.

Table 2: MSE of the friction coefficient. All values are multiplied by 10^3.

	pendulum	Duffing
SSGP	**0.519**	**0.506**
D-SymODEN	1.010	0.630
D-HNN	2.486	2.861

Moreover, SSGP improved the performance rather than all baselines, especially when the number of trajectories was small. These results show that SSGP is advantageous in sparse settings. In Figure 3 **(b)**, NODE and RFF yielded relatively low errors; however, since they cannot distinguish between the conservative and dissipative terms, their performance was worse in **Task 2**, as shown in Figure 3 **(c)**. In contrast, SSGP can estimate the conservative and dissipative terms separately by incorporating knowledge of Hamiltonian mechanics, which yields better performance in both **Task 1** and **Task 2** than all baselines. Table 2 shows MSE of the friction coefficients estimated by SSGP, D-SymODEN and D-HNN when inferring the dissipative systems. Here, MSE was averaged for all experimental settings. This result shows that SSGP can accurately estimate the friction coefficients.

In the following, we deeply explore the results when the number of observed trajectories was 20. Figure 4 shows the visualization of the trajectories predicted by SSGP. One observes that SSGP appropriately captured the dynamics of the dissipative Hamiltonian systems (the third column of Figure 4) and better discerned the conservative and dissipative terms (the fourth and fifth columns of Figure 4). Figure (5) shows the cumulative errors for the trajectories and energies for the pendulum data. Let I^{test} denote the number of time steps in the test set, and j_t denote a time step at time t. We calculated the cumulative error of trajectories at time t: $\sum_{j=1}^{j_t}(\frac{1}{I^{\mathrm{test}}}\sum_{i=1}^{I^{\mathrm{test}}} \parallel \hat{\boldsymbol{x}}_{ij} - \boldsymbol{x}_{ij}^{\mathrm{true}} \parallel^2)$. The cumulative error of energies was obtained similarly by calculating the squared error of the energy, $H^{\mathrm{true}}(\hat{\boldsymbol{x}}_{ij})$, evaluated using the predicted trajectories and the true energy, $H^{\mathrm{true}}(\boldsymbol{x}_{ij}^{\mathrm{true}})$. Here, $H^{\mathrm{true}}(\cdot)$ is the true Hamiltonian of each system. As shown in Figure 5, SSGP achieved lower errors than the baselines for both trajectories and energies at all time points. The performance improvement of SSGP was more significant in the period $[10, 15]$, which is out of the simulation period (10 seconds) when generating the training data. These results show that SSGP can accurately predict the dynamics with energy conservation and dissipation in both short-term and long-term simulations.

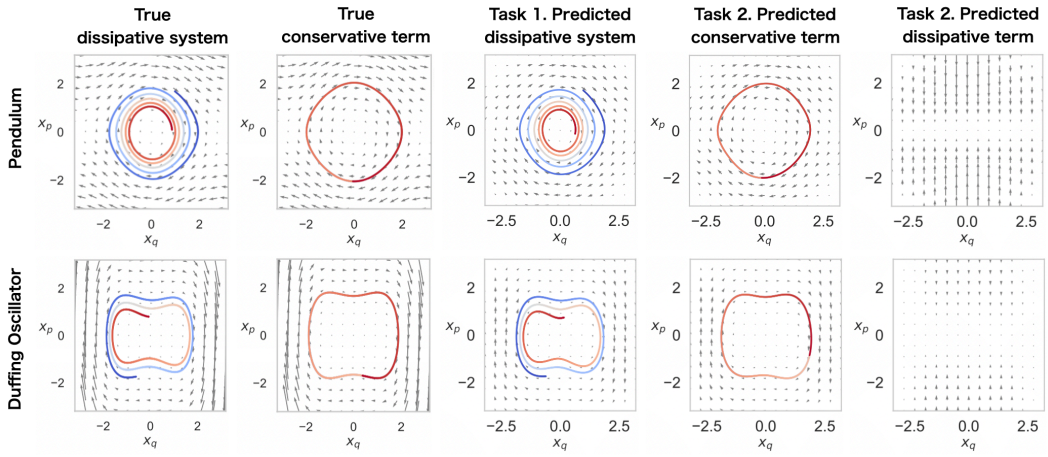

Figure 4: Prediction results of SSGP. The color indicates time-evolution, starting at blue and ending at red. The first and second columns are the true trajectories for the dissipative systems and their conservative terms, respectively. The third column is the prediction for the dissipative systems in **Task 1**. The fourth and fifth columns are the predicted conservative and dissipative terms in **Task 2**, respectively. Here, the dissipative terms are multiplied by 30 for enhanced clarity. Comparisons with other models are shown in Appendix G.

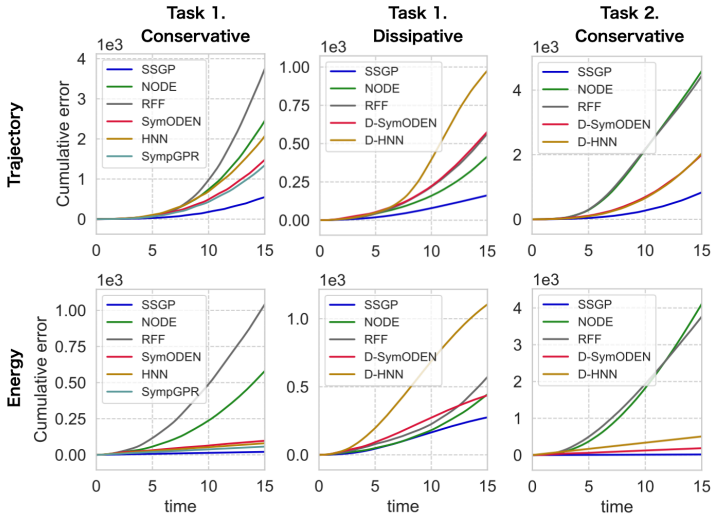

Figure 5: Cumulative error for the predicted trajectories and energies when using pendulum data. The horizontal axis is simulation time in the test phase.

7 Conclusion

We have proposed the Symplectic Spectrum Gaussian Process (SSGP), which allows one to predict systems whose dynamics follow energy conservation and dissipation laws from noisy and sparse data. Our result, the symplectic random Fourier feature, is a general tool and has the potential to use the design of kernel machines with prior knowledge in physics. As described in Section 5, the SSGP can easily extend to learn Hamiltonians from high-dimensional data (e.g., images). Our future work is to evaluate the effectiveness of its extension.

Acknowledgments and Disclosure of Funding

This work was supported by JST, ACT-X Grant Number JPMJAX210D, Japan.

References

[1] Luca Baldassarre, Lorenzo Rosasco, Annalisa Barla, and Alessandro Verri. Multi-output learning via spectral filtering. *Mach. Learn.*, 87(3):259–301, 2012.

[2] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[3] Ricky T. Q. Chen, Brandon Amos, and Maximilian Nickel. Learning neural event functions for ordinary differential equations. In *International Conference on Learning Representations*, 2021.

[4] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[5] Yuhan Chen, Takashi Matsubara, and Takaharu Yaguchi. Neural symplectic form: Learning Hamiltonian equations on general coordinate systems. In *Advances in Neural Information Processing Systems*, 2021.

[6] Zhengdao Chen, Jianyu Zhang, Martin Arjovsky, and Léon Bottou. Symplectic recurrent neural networks. In *International Conference on Learning Representations*, 2020.

[7] Kevin Course, Trefor Evans, and Prasanth Nair. Weak form generalized Hamiltonian learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 18716–18726, 2020.

[8] P. Courtier, Jean-Noël Thépaut, and A. Hollingsworth. A strategy for operational implementation of 4d-var. In *Workshop on Variational Assimilation, with special emphasis on Three-dimensional Aspects, 9-12 November 1992*, pages 437–464, Shinfield Park, Reading, 1992. ECMWF, ECMWF.

[9] Shaan A. Desai, Marios Mattheakis, David Sondak, Pavlos Protopapas, and Stephen J. Roberts. Port-Hamiltonian neural networks for learning explicit time-dependent dynamical systems. *Phys. Rev. E*, 104:034312, Sep 2021.

[10] Daniel DiPietro, Shiying Xiong, and Bo Zhu. Sparse symplectically integrated neural networks. In *Advances in Neural Information Processing Systems*, volume 33, pages 6074–6085, 2020.

[11] Katharina Ensinger, Friedrich Solowjow, Sebastian Ziesche, Michael Tiemann, and Sebastian Trimpe. Structure-preserving Gaussian process dynamics. In *arXiv*, 2021.

[12] Marc Finzi, Ke Alexander Wang, and Andrew G Wilson. Simplifying Hamiltonian and Lagrangian neural networks via explicit constraints. In *Advances in Neural Information Processing Systems*, volume 33, pages 13880–13889, 2020.

[13] Herbert Goldstein. *Classical Mechanics*. Addison-Wesley, 1980.

[14] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[15] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[16] Pashupati Hegde, Çağatay Yıldız, Harri Lähdesmäki, Samuel Kaski, and Markus Heinonen. Variational multiple shooting for Bayesian ODEs with Gaussian processes. In *arXiv*, 2021.

[17] Raban Iten, Tony Metger, Henrik Wilming, Lídia del Rio, and Renato Renner. Discovering physical concepts with neural networks. *Phys. Rev. Lett.*, 124:010508, Jan 2020.

[18] Pengzhan Jin, Zhen Zhang, Ioannis G. Kevrekidis, and George Em Karniadakis. Learning Poisson systems and trajectories of autonomous systems via Poisson neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–13, 2022.

[19] Pengzhan Jin, Zhen Zhang, Aiqing Zhu, Yifa Tang, and George Em Karniadakis. SympNets: Intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems. *Neural Networks*, 132:166–179, 2020.

[20] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. In *Advances in Neural Information Processing Systems*, volume 33, pages 6696–6707, 2020.

[21] D. Kingma and M. Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.

[22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

[23] Miguel Lázaro-Gredilla, Joaquin Quiñnero-Candela, Carl Edward Rasmussen, and Aníbal R. Figueiras-Vidal. Sparse spectrum Gaussian process regression. *Journal of Machine Learning Research*, 11(63):1865–1881, 2010.

[24] SENWEI Liang, Zhongzhan Huang, and Hong Zhang. Stiffness-aware neural network for learning Hamiltonian systems. In *International Conference on Learning Representations*, 2022.

[25] Y. Macedo and R. Castro. Learning div-free and curl-free vector fields by matrix-valued kernels. In *Technical report, Preprint A 679/2010 IMPA*, 2008.

[26] Takashi Matsubara, Ai Ishikawa, and Takaharu Yaguchi. Deep energy-based modeling of discrete-time physics. In *Advances in Neural Information Processing Systems*, volume 33, pages 13100–13111, 2020.

[27] Takashi Matsubara, Yuto Miyatake, and Takaharu Yaguchi. Symplectic adjoint method for exact gradient of neural ODE with minimal memory. In *Advances in Neural Information Processing Systems*, volume 34, pages 20772–20784, 2021.

[28] Thomas McDonald and Mauricio Álvarez. Compositional modeling of nonlinear dynamical systems with ODE-based random features. In *Advances in Neural Information Processing Systems*, volume 34, pages 13809–13819, 2021.

[29] Hossein Mohammadi, Peter Challenor, and Marc Goodfellow. Emulating complex dynamical simulators with random Fourier features. In *arXiv*, 2021.

[30] C. Offen and S. Ober-Blöbaum. Symplectic integration of learned Hamiltonian systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 32(1):013122, 2022.

[31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035. 2019.

[32] Pascal Pernot and Fabien Cailliez. A critical review of statistical calibration/prediction models handling data inconsistency and model inadequacy. *AIChE Journal*, 63(10):4642–4665, 2017.

[33] Dimitris C. Psichogios and Lyle H. Ungar. A hybrid neural network-first principles approach to process modeling. *AIChE Journal*, 38(10):1499–1511, 1992.

[34] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, volume 20, 2007.

[35] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[36] Katharina Rath, Christopher G. Albert, Bernd Bischl, and Udo von Toussaint. Symplectic Gaussian process regression of maps in Hamiltonian systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(5):053121, 2021.

[37] R. Rico-Martinez, J.S. Anderson, and I.G. Kevrekidis. Continuous-time nonlinear signal processing: A neural network based approach for gray box identification. In *Proceedings of IEEE Workshop on Neural Networks for Signal Processing*, pages 596–605, 1994.

[38] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 2009.

[39] E. Solak, R. Murray-smith, W. Leithead, D. Leith, and Carl Rasmussen. Derivative observations in Gaussian process models of dynamic systems. In *Advances in Neural Information Processing Systems*, volume 15, 2003.

[40] Andrew Sosanya and Sam Greydanus. Dissipative Hamiltonian neural networks: Learning dissipative and conservative dynamics separately. *CoRR*, abs/2201.10085, 2022.

[41] Danica J. Sutherland and Jeff Schneider. On the error of random Fourier features. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 862–871, 2015.

[42] Michael L. Thompson and Mark A. Kramer. Modeling chemical processes using prior knowledge and neural networks. *AIChE Journal*, 40(8):1328–1340, 1994.

[43] Peter Toth, Danilo J. Rezende, Andrew Jaegle, Sébastien Racanière, Aleksandar Botev, and Irina Higgins. Hamiltonian generative networks. In *International Conference on Learning Representations*, 2020.

[44] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. In *International Conference on Artificial Intelligence and Statistics*, volume 51, pages 370–378. PMLR, 2016.

[45] James Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *International Conference on Machine Learning*, volume 119, pages 10292–10302. PMLR, 2020.

[46] James T. Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Pathwise conditioning of Gaussian processes. *Journal of Machine Learning Research*, 22(105):1–47, 2021.

[47] Shiying Xiong, Yunjin Tong, Xingzhe He, Shuqi Yang, Cheng Yang, and Bo Zhu. Nonseparable symplectic neural networks. In *International Conference on Learning Representations*, 2021.

[48] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Dissipative SymODEN: Encoding Hamiltonian dynamics with dissipation and control into deep learning. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.

[49] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Symplectic ODE-Net: Learning Hamiltonian dynamics with control. In *International Conference on Learning Representations*, 2020.

[50] Juntang Zhuang, Nicha C Dvornek, sekhar tatikonda, and James s Duncan. MALI: A memory efficient and reverse accurate integrator for neural ODEs. In *International Conference on Learning Representations*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] The main contributions of this work have been summarized in the last paragraph of Section 1.
   (b) Did you describe the limitations of your work? [Yes] The limitations have been described in Appendix H.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] We have discussed the societal impacts in Appendix H.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We have read the guidelines. No ethical impact is supposed.
2. If you are including theoretical results...

(a) Did you state the full set of assumptions of all theoretical results? [Yes] All the assumptions have been described in Sections 4 and 5.

(b) Did you include complete proofs of all theoretical results? [Yes] Derivation procedures of our models have been included in Appendices A and C.

3. If you ran experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code is available at `https://github.com/yusuk-e/SSGP`.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We have described how to generate datasets in the **Data** paragraph of Section 6. The setups of the proposed model and the baselines were specified in the **Setup of the SSGP** paragraph of Section 6 and Appendix F, respectively.

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We have provided the standard errors of the performance in Figure 3.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We used AMD EPYC 7313 CPU (3.0GHz) in the experiments, as mentioned in the **On computational complexity** paragraph of Section 5.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] We used the python library, torchdiffeq [3, 4], which we have cited in Section 6.

(b) Did you mention the license of the assets? [Yes] We have mentioned the license information in the footnote on page 7.

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code is available at `https://github.com/yusuk-e/SSGP`.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Derivation of the spectral representation

The spectral representation of the GP prior (3) for Hamiltonian systems with additive dissipation can be derived by applying the differential operator $\mathcal{L} = (\mathbf{R} - \mathbf{S})\nabla$ to the $2M$-dimensional GP

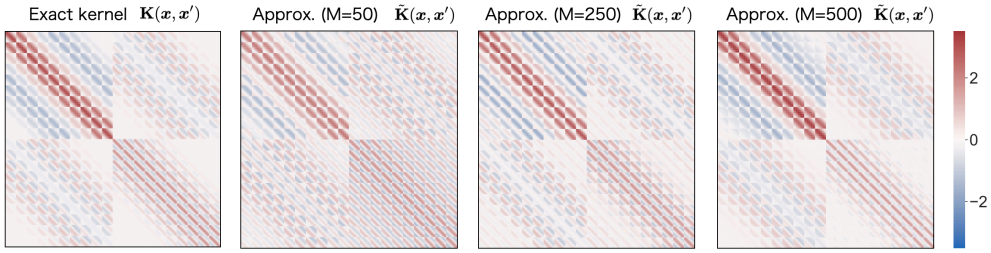

Figure 6: Gram matrices for the exact kernel and for its approximation with $M = \{50, 250, 500\}$ spectral points. The covariances are computed using 10×10 regular grid points $\{(x^{\mathrm{q}}, x^{\mathrm{p}})\}$ on a two-dimensional phase space defined by $[-1, 1]^2$. The upper left and lower right corners of each matrix are the covariances of $(x^{\mathrm{q}}, x^{\mathrm{q}'})$ and $(x^{\mathrm{p}}, x^{\mathrm{p}'})$, respectively; the remaining corners are the cross-covariances of $(x^{\mathrm{q}}, x^{\mathrm{p}})$. The parameters are set as follows: $\sigma_0^2 = 1, \boldsymbol{\Lambda} = \mathrm{diag}(0.5, 0.3), \mathbf{R} = \mathrm{diag}(0, 0.1)$.

approximation (6) of the Hamiltonian $H(\boldsymbol{x})$, as follows:

$$\boldsymbol{f}(\boldsymbol{x}) = \mathcal{L}\, H(\boldsymbol{x}) \tag{21}$$

$$= (\mathbf{S} - \mathbf{R})\nabla \left[\sum_{m=1}^{M} \boldsymbol{w}_m \boldsymbol{\phi}_m(\boldsymbol{x})\right] \tag{22}$$

$$= (\mathbf{S} - \mathbf{R}) \sum_{m=1}^{M} \boldsymbol{w}_m \nabla \left[\cos\left(2\pi \boldsymbol{s}_m^\top \boldsymbol{x}\right), \sin\left(2\pi \boldsymbol{s}_m^\top \boldsymbol{x}\right)\right] \tag{23}$$

$$= \sum_{m=1}^{M} 2\pi(\mathbf{S} - \mathbf{R})\boldsymbol{s}_m \left[-\sin(2\pi \boldsymbol{s}_m^\top \boldsymbol{x}), \cos(2\pi \boldsymbol{s}_m^\top \boldsymbol{x})\right] \boldsymbol{w}_m^\top \tag{24}$$

$$= \sum_{m=1}^{M} \boldsymbol{\Psi}_m(\boldsymbol{x}) \boldsymbol{w}_m^\top \tag{25}$$

$$= \boldsymbol{\Psi}(\boldsymbol{x}) \boldsymbol{w}^\top. \tag{26}$$

As shown in (10), the spectral representation of $\boldsymbol{f}(\boldsymbol{x})$ corresponds to the GP approximation equipped with covariance function $\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}')$ arising from *symplectic* random Fourier features (S-RFFs) $\boldsymbol{\Psi}(\boldsymbol{x})$. The covariance function $\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}')$ is given by

$$\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}') = \frac{\sigma_0^2}{M} \boldsymbol{\Psi}(\boldsymbol{x}) \boldsymbol{\Psi}(\boldsymbol{x}')^\top \tag{27}$$

$$= \frac{\sigma_0^2}{M} \sum_{m=1}^{M} \boldsymbol{\Psi}_m(\boldsymbol{x}) \boldsymbol{\Psi}_m(\boldsymbol{x}')^\top \tag{28}$$

$$= \frac{(2\pi)^2 \sigma_0^2}{M} \sum_{m=1}^{M} (\mathbf{S} - \mathbf{R})\boldsymbol{s}_m \left[(\mathbf{S} - \mathbf{R})\boldsymbol{s}_m\right]^\top \cos\left(2\pi \boldsymbol{s}_m^\top (\boldsymbol{x} - \boldsymbol{x}')\right). \tag{29}$$

One observes that $\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}')$ is an extention of the covariance function arising from *standard* RFFs [23, p. 1870] to encode the geometric structure of Hamiltonian systems with additive dissipation.

B Comparisons of gram matrices

To investigate the connection between the exact GP (3) for Hamiltonian systems with additive dissipation and its spectral approximation (8), we show the comparisons of the gram matrix of $\mathbf{K}(\boldsymbol{x}, \boldsymbol{x}')$ (4) and that of $\tilde{\mathbf{K}}(\boldsymbol{x}, \boldsymbol{x}')$ (10) in Figure 6. We observe that approximation quality improves as the number M of spectral points increases.

C Derivation of the ELBO

In this appendix, we derive the evidence lower bound (ELBO) (13). Let $\mathbf{X} = \{\boldsymbol{x}_{ij}\}$ denote the set of the noiseless states. We assume that the variational distribution is given by

$$q\left(\mathbf{X}, \boldsymbol{f}, \boldsymbol{w}\right) = p(\boldsymbol{f}|\boldsymbol{w})q(\boldsymbol{w}) \prod_{i=1}^{I} \left[q(\boldsymbol{x}_{i1}) \prod_{j=2}^{J_i} p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) \right]. \tag{30}$$

As described in Section 5, $p(\boldsymbol{f}|\boldsymbol{w})$ is Dirac's delta function; $q(\boldsymbol{x}_{i1})$ and $q(\boldsymbol{w})$ are the variational distributions of the initial condition $\boldsymbol{x}_{i1}$ and the weights $\boldsymbol{w}$, respectively. $p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1})$ is Dirac's delta function represented by (12). The ELBO is derived from Jensen's inequality as follows:

$$\begin{aligned}
&\log p\left(\mathbf{Y}\right) \\
&\geq \iiint q\left(\mathbf{X}, \boldsymbol{f}, \boldsymbol{w}\right) \log \frac{p\left(\mathbf{Y}, \mathbf{X}, \boldsymbol{f}, \boldsymbol{w}\right)}{q\left(\mathbf{X}, \boldsymbol{f}, \boldsymbol{w}\right)} d\boldsymbol{w} d\boldsymbol{f} d\mathbf{X} && (31) \\
&= \iiint p(\boldsymbol{f}|\boldsymbol{w})q(\boldsymbol{w}) \prod_{i=1}^{I} \left[q(\boldsymbol{x}_{i1}) \prod_{j=2}^{J_i} p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) \right] \\
&\times \log \frac{p(\boldsymbol{f}|\boldsymbol{w})p(\boldsymbol{w}) \prod_{i=1}^{I} \left[p(\boldsymbol{y}_{i1}|\boldsymbol{x}_{i1})p(\boldsymbol{x}_{i1}) \prod_{j=2}^{J_i} \left\{ p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij})p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) \right\} \right]}{p(\boldsymbol{f}|\boldsymbol{w})q(\boldsymbol{w}) \prod_{i=1}^{I} \left[q(\boldsymbol{x}_{i1}) \prod_{j=2}^{J_i} p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) \right]} d\boldsymbol{w} d\boldsymbol{f} d\mathbf{X} && (32) \\
&= \sum_{i=1}^{I} \left[\int \left\{ q(\boldsymbol{x}_{i1}) \log p(\boldsymbol{y}_{i1}|\boldsymbol{x}_{i1}) + \sum_{j=2}^{J_i} \int p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1}) q(\boldsymbol{f}) q(\boldsymbol{x}_{i1}) \log p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij}) d\boldsymbol{f} \right\} d\mathbf{X} \right] \\
&- \sum_{i=1}^{I} \mathrm{KL}\left[q(\boldsymbol{x}_{i1}) || p(\boldsymbol{x}_{i1}) \right] - \mathrm{KL}\left[q(\boldsymbol{w}) || p(\boldsymbol{w}) \right] && (33) \\
&= \sum_{i=1}^{I} \left[\sum_{j=1}^{J_i} \mathbb{E}_{q(\boldsymbol{x}_{ij})} \left[\log p\left(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij}\right) \right] - \mathrm{KL}\left[q(\boldsymbol{x}_{i1}) || p(\boldsymbol{x}_{i1}) \right] \right] - \mathrm{KL}\left[q(\boldsymbol{w}) || p(\boldsymbol{w}) \right], && (34)
\end{aligned}$$

where $p(\boldsymbol{f}|\boldsymbol{w})$ and $p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1})$ are cancelled in (32), and we define $q(\boldsymbol{f}) = \int p(\boldsymbol{f}|\boldsymbol{w})q(\boldsymbol{w})d\boldsymbol{w}$ in (33). For $j \geq 2$, the variational distribution $q(\boldsymbol{x}_{ij})$ in (34) is given by

$$q(\boldsymbol{x}_{ij}) = \iint q(\boldsymbol{f})p(\boldsymbol{x}_{ij}|\boldsymbol{f}, \boldsymbol{x}_{i1})q(\boldsymbol{x}_{i1})d\boldsymbol{x}_{i1}d\boldsymbol{f}. \tag{35}$$

It should be noted that, given $\boldsymbol{f}$ and $\boldsymbol{x}_{i1}$, $\boldsymbol{x}_{ij} (j \geq 2)$ is deterministically generated by the ODE.

D Inference procedure of SSGP

The inference procedure of SSGP is shown in Algorithm 1. Here, $\mathrm{unif}\{1, J_i - S\}$ in Algorithm 1 is a discrete uniform distribution with support $\{1, \ldots, J_i - S\}$. Although we assume that, for simplicity, the sampling frequency G_s is constant, our algorithm can be easily extended to irregularly sampled trajectories. In addition, this algorithm requires a hyperparameter, *integration time window* τ, which is often used for effectively learning ODEs when the trajectory length J_i is large. In the implementation, we randomly re-select a state taken to be an initial condition from the observed samples in each iteration and calculate the ELBO by generating trajectories up to τ seconds later, as in previous works (e.g., [7, 49]). In the experiments, we set the integration time window $\tau = 1$.

E On computational complexity

As described in Section 5, our inference algorithm is computationally efficient because the proposed model is based on symplectic random Fourier features (S-RFFs) and the bottleneck (i.e., sampling of weights) is outside of the ODE solver. To clarify it, let us consider the case where we do not use the

Algorithm 1 Inference procedure for SSGP.

Input: Trajectory data $\{(t_{ij}, \boldsymbol{y}_{ij}) | i = 1 \ldots, I; j = 1 \ldots, J_i\}$, sampling frequency G_s, number of spectral points M, numbers of Monte Carlo samples K, L, integration time window τ, number of epochs E

Output: Parameters $\boldsymbol{\Lambda}$, σ_0^2, σ^2, $\mathbf{R}$, $\mathbf{A}$, $\boldsymbol{b}$, $\mathbf{C}$

Initialize the parameters
$e \leftarrow 1$
repeat
  /* Sample mini-batches $\mathcal{D}$ */
  $S \leftarrow \tau / G_\mathrm{s}$
  $\mathcal{D} \leftarrow \emptyset$
  for $i = 1, \ldots I$ **do**
    $s_{i1} \sim \mathrm{unif}\{1, J_i - S\}$
    $\mathcal{D} \leftarrow \mathcal{D} \cup \{(t_{i,s_{i1}}, \boldsymbol{y}_{i,s_{i1}}), \ldots, (t_{i,s_{i1}+S}, \boldsymbol{y}_{i,s_{i1}+S})\}$
  end for

  /* Generate trajectories */
  for $k = 1, \ldots, K$ **do**
    for $l = 1, \ldots, L$ **do**
      $\boldsymbol{w}^{(k,l)} \leftarrow \boldsymbol{b} + \sqrt{\mathbf{C}}\, \boldsymbol{\epsilon}^{(k,l)} \quad \text{where} \quad \boldsymbol{\epsilon}^{(k,l)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
    end for
    Construct the vector field $\boldsymbol{f}^{(k)}(\boldsymbol{x})$ by (18).
    for $i = 1, \ldots, I$ **do**
      $\boldsymbol{x}_{i,s_{i1}}^{(k)} \leftarrow \boldsymbol{y}_{i,s_{i1}} + \sqrt{\mathbf{A}}\, \boldsymbol{\epsilon}_i^{(k)}, \quad \text{where} \quad \boldsymbol{\epsilon}_i^{(k)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
      $\boldsymbol{x}_{i,s_{i1}+1}^{(k)}, \ldots, \boldsymbol{x}_{i,s_{i1}+S}^{(k)} \leftarrow \mathrm{ODESolve}\left(\boldsymbol{x}_{i,s_{i1}}^{(k)}, \boldsymbol{f}^{(k)}(\boldsymbol{x}), t_{i,s_{i1}+1}, \ldots, t_{i,s_{i1}+S}\right)$
    end for
  end for

  Update the parameters by maximizing the ELBO (13) evaluated using $\mathcal{D}$.
  $e \leftarrow e + 1$
until $e > E$

RFF-based approximation. Since the derivative observations are not available, we should approximate the GP posterior for vector fields; one possible way is to introduce latent variables (similar to inducing variables) for the derivative observations, as in [A1]. Unfortunately, this technique does not reduce the cost of sampling paths from the GP posterior which scales cubically with the number of test points [45, 46]. Note that sampling paths is different from sampling function values independently at each test point. Denote the number of test points by N_t, which is proportional to the inverse of step size Δ in the ODE solver, as follows: $N_\mathrm{t} \propto 1/\Delta$, where Δ should be set to a sufficiently small value (often lower than 10^{-3}). Moreover, the GP posterior should be updated at each numerical integration step by conditioning all data points (i.e., test points) generated by the ODE solver. The cost of this approach results in $\mathcal{O}(N_\mathrm{t}^4)$. On the other hand, the cost of our inference algorithm is $\mathcal{O}(M^3)$, where M is the number of spectral points (often lower than 10^3). Accordingly, our proposed model based on RFF is also beneficial in terms of computational complexity.

F Baselines

In this appendix, we describe baseline models for the experiments in Section 6. In all experiments, we trained the model using the Adam optimizer with a learning rate of 10^{-3} and a weight decay of 10^{-4} for 10^4 epochs. We also used the ODE solver, the Dormand–Prince method with adaptive time-stepping, with the relative and absolute tolerances of 10^{-8}.

Neural network models. HNN can utilize the prior knowledge of energy conservation as an inductive bias for training neural networks; D-HNN is its extension to handle not only energy conservation but also energy dissipation. These models require derivative observations, not trajectories, for training; thus, we used finite differences instead in the experiments. NODE can infer unknown dynamics

from trajectory data by utilizing ODE solvers, which does not consider prior knowledge in physics. SymODEN is the extension of HNN to allow the model learning from trajectory data by using ODE solvers, as in NODE. D-SymODEN can also apply to the dissipative systems. For all the models, we used a fully connected neural network with 3 layers, 200 hidden units, and tanh activations.

Gaussian process models. SympGPR can estimate conservative vector fields from derivative observations by considering Hamiltonian mechanics; we used finite differences for training. RFF is a baseline that uses standard random Fourier features for modeling vector fields (not considering Hamiltonian mechanics); its training is based on ODE solvers. This baseline is equivalent to [16] except for using the decoupled sampling. SSGP is our proposal, which can handle conservative and dissipative systems by considering Hamiltonian mechanics and train the model with uncertainties from trajectories using ODE solvers. For all the GP models, we used the ARD Gaussian kernel.

G Experimental results

This appendix contains additional results for our experiments in Section 6. Figures 7(a) and 7(b) show MSEs and standard errors for SSGP and the baselines when the sampling frequency was 3 or 10 Hz. We obtained similar results to those for 5 Hz shown in Figure 3. SSGP had a more significant performance improvement when the number of trajectories was smaller and the sampling frequency was lower (see Figure 7(a)). On the other hand, the performance of SSGP was closer to those of the baselines when the trajectory data with relatively high time resolution was available (see Figure 7(b)). These results are convincing to clarify that SSGP is more advantageous in the sparse setting.

We provide the visualization comparisons of the predicted trajectories when the sampling frequency was 5 Hz. Figures 8 and 9 show the visualization results for predicting the conservative and dissipative systems, respectively, in **Task 1**. As shown in Figure 8, NODE and RFF do not capture the energy conservation law. As one can see from the results of Duffing oscillator (see Figures 8(b) and 9(b)), HNN, D-HNN, SymODEN, and D-SymODEN failed to estimate the vector fields precisely because they suffered from overfitting to noisy observations; thus, they output false trajectories. Compared with the baselines, SSGP estimated the vector fields without overfitting and accurately predicted the trajectories with energy conservation and dissipation. Figure 10 shows the visualizations of the predicted conservative and dissipative terms in **Task 2**. One observes that SSGP accurately decomposed the dynamics into conservative and dissipative terms rather than the baseline models.

H Discussion

Limitations. Finite-dimensional approximation of Gaussian processes (GPs) via a set of random Fourier features (RFFs) is helpful for efficient sampling from GP posterior. In the proposed approach, we introduce the RFFs that encode symplectic geometric structures of Hamiltonian systems to construct the efficient variational inference algorithm for learning dynamics with energy conservation or dissipation. However, the RFF-based approximation is well-known to yield undesirable extrapolatory predictions as the number of samples increases, which is called *variance starvation* [A2, A3]. This might prevent the proposed model from learning Hamiltonian dynamics effectively. It would be promising to employ the *decoupled sampling* technique [45], as in [11, 16], which might be useful to infer unknown dynamics while avoiding the variance starvation.

Although our experimental results are encouraging, they are the first step to investigating the empirical properties of the proposed model. We will apply it to various types of systems. In addition, the proposed model is easily extendable to learning Hamiltonians from high-dimensional data (such as images), as described in Section 5. We will empirically evaluate the effectiveness of its extension.

Societal impact. This paper presents a framework for data-driven modeling that incorporates prior knowledge of Hamiltonian mechanics to improve predictive performance in noisy and sparse settings. Generally, a data-driven approach can have a significant impact on a wide range of applications, including climate and robotics. However, since we have not evaluated our proposed framework for such real-world applications, we need to carefully assess the prediction and uncertainty estimation performance when applying it to them.

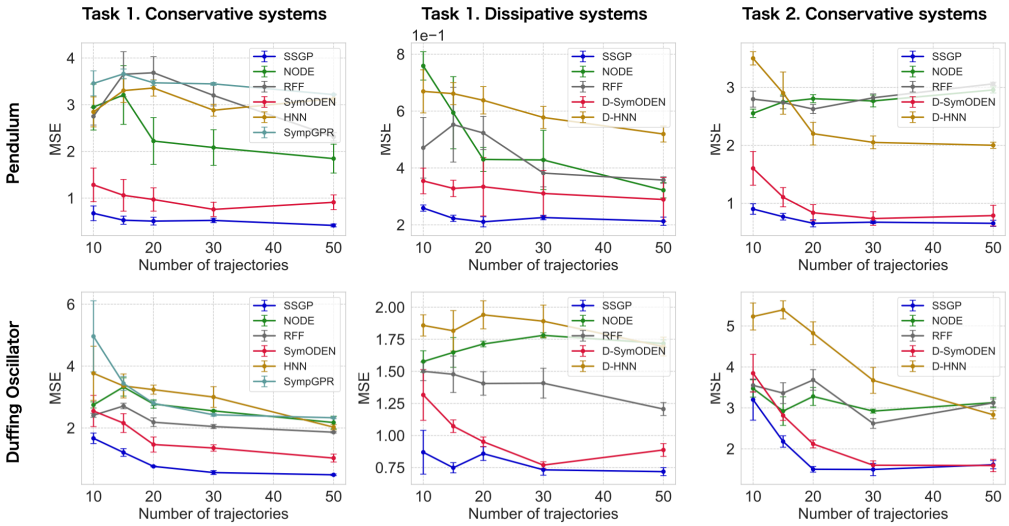

(a) Results when setting the sampling frequency at 3 Hz.

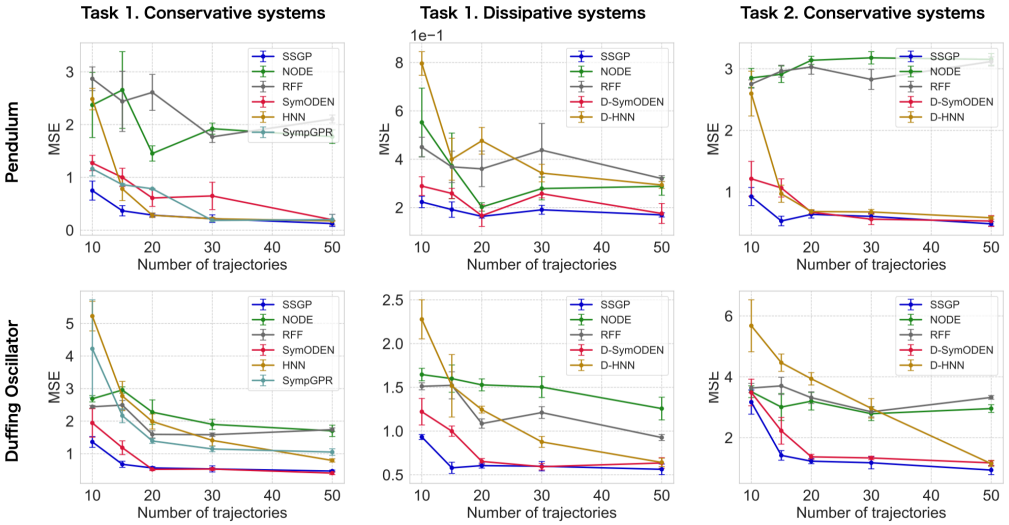

(b) Results when setting the sampling frequency at 10 Hz.

Figure 7: MSEs and standard errors for the predicted trajectories of pendulum and Duffing oscillator. We set the sampling frequency to (a) 3 Hz and (b) 10 Hz. The first and second columns are the results for the predictions of conservative and dissipative systems, respectively, in **Task 1**. The third column holds the results for predicting conservative systems in **Task 2**.

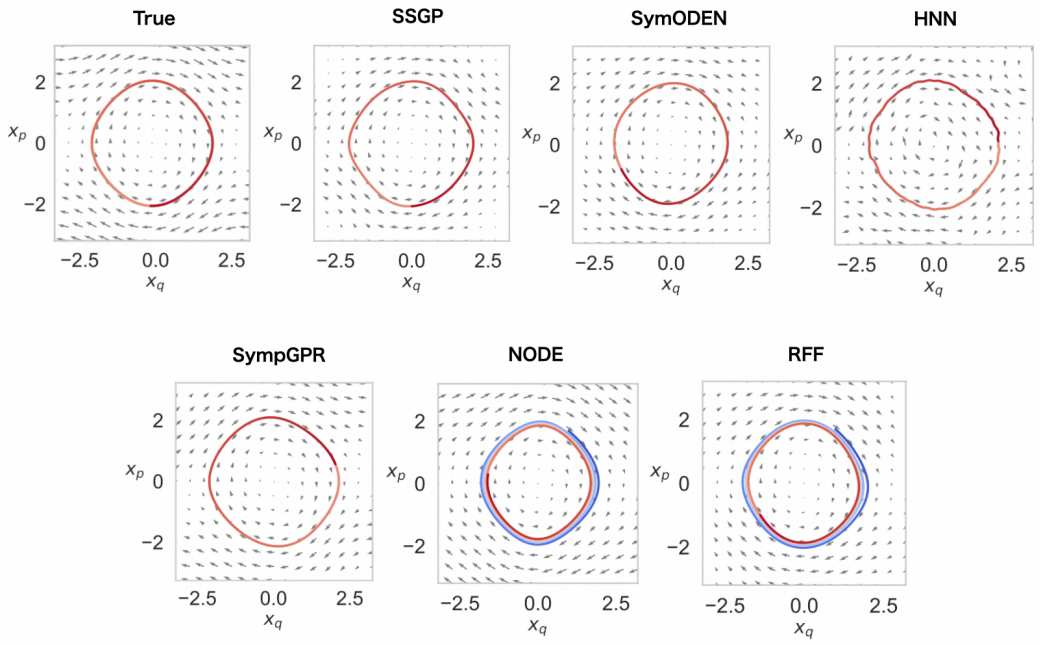

(a) Prediction results for pendulum (without dissipation)

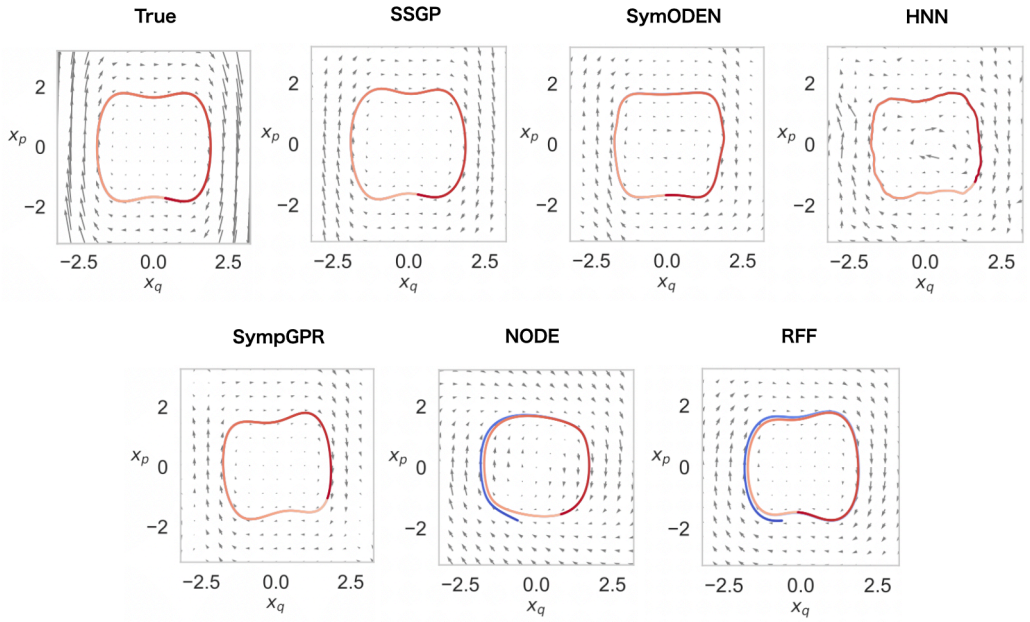

(b) Prediction results for Duffing oscillator (without dissipation)

Figure 8: Prediction results for pendulum and Duffing oscillator without dissipation in **Task 1**. The color indicates time-evolution, starting at blue and ending at red.

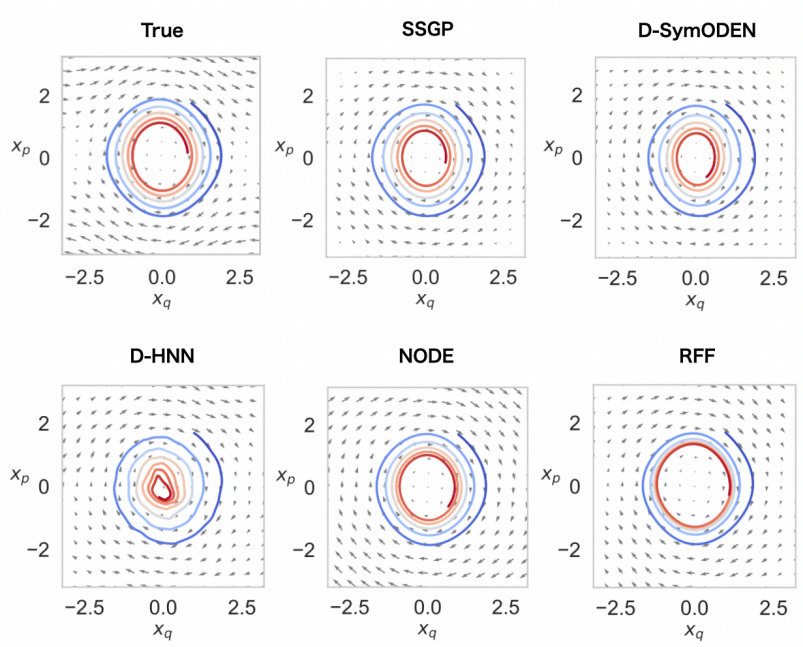

(a) Prediction results for pendulum (with dissipation)

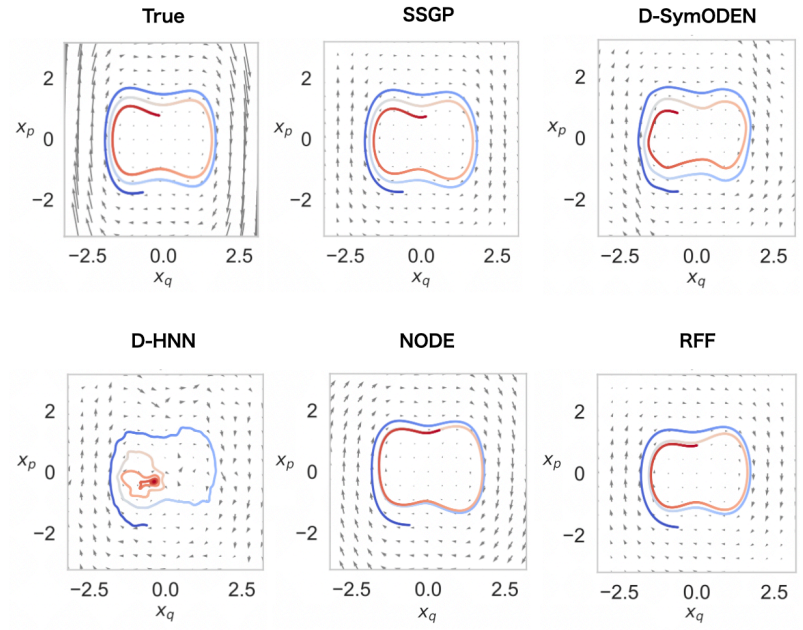

(b) Prediction results for Duffing oscillator (with dissipation)

Figure 9: Prediction results for pendulum and Duffing oscillator with dissipation in **Task 1**. The color indicates time-evolution, starting at blue and ending at red.

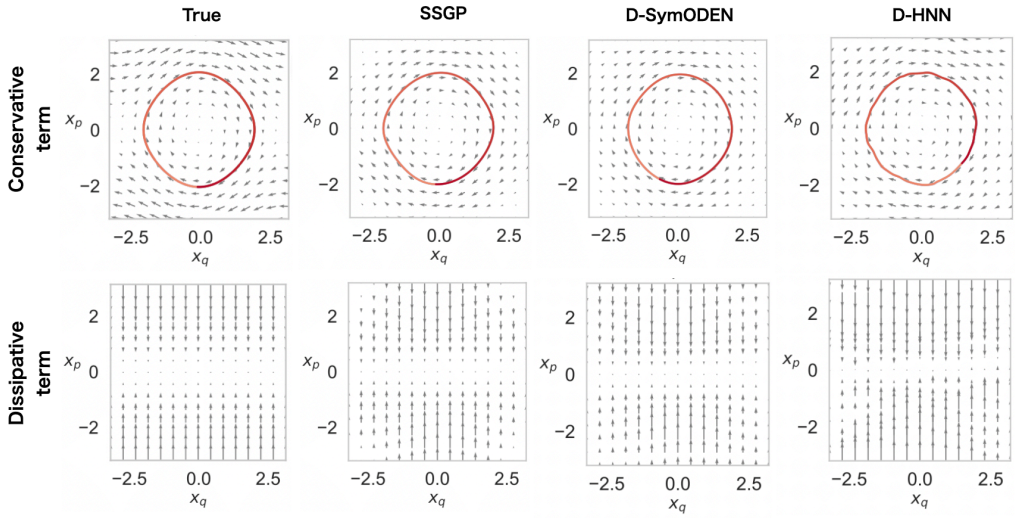

(a) Predicted conservative and dissipative terms for pendulum

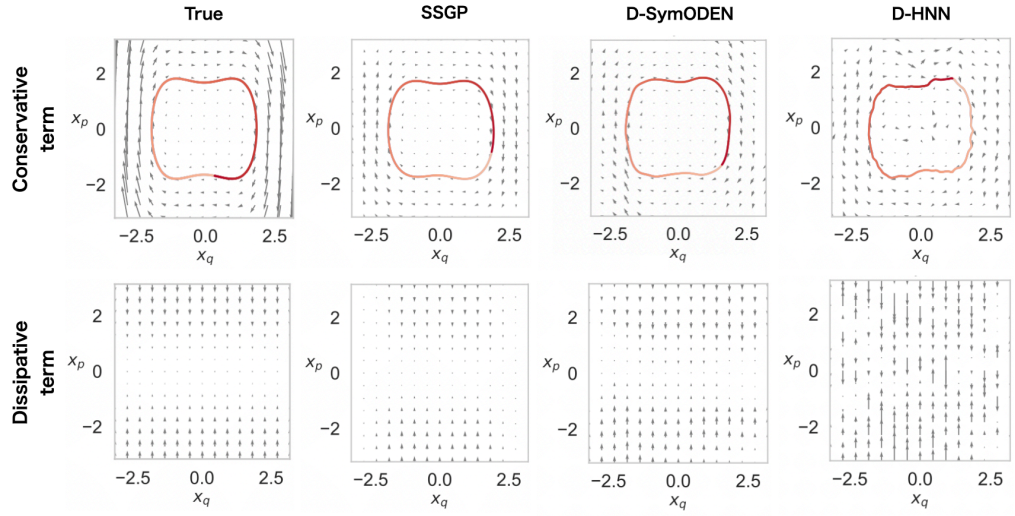

(b) Predicted conservative and dissipative terms for Duffing oscillator

Figure 10: Prediction results of the conservative and dissipative terms for pendulum and Duffing oscillator in **Task 2**. The dissipative terms are multiplied by 30 for enhanced clarity.

References

[A1] P. Hegde, M. Heinonen, H. Lähdesmäki, and H. Kaski. Deep learning with differential Gaussian process flows. In *Artificial Intelligence and Statistics*, volume 31, pages 1812–1831, 2019.

[A2] Mojmir Mutny and Andreas Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature Fourier features. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[A3] Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale Bayesian optimization in high-dimensional spaces. In *International Conference on Artificial Intelligence and Statistics*, volume 84, pages 745–754. PMLR, 2018.