
Variational multiple shooting for Bayesian ODEs with Gaussian processes

Pashupati Hegde¹

Çağatay Yıldız¹

Harri Lähdesmäki¹

Samuel Kaski¹

Markus Heinonen¹

¹ Department of Computer Science, Aalto University, Finland

Abstract

Recent machine learning advances have proposed black-box estimation of *unknown continuous-time system dynamics* directly from data. However, earlier works are based on approximative solutions or point estimates. We propose a novel Bayesian nonparametric model that uses Gaussian processes to infer posteriors of unknown ODE systems directly from data. We derive sparse variational inference with decoupled functional sampling to represent vector field posteriors. We also introduce a probabilistic shooting augmentation to enable efficient inference from arbitrarily long trajectories. The method demonstrates the benefit of computing vector field posteriors, with predictive uncertainty scores outperforming alternative methods on multiple ODE learning tasks.

1 INTRODUCTION

Ordinary differential equations (ODEs) are powerful models for continuous-time non-stochastic systems, which are ubiquitous from physical and life sciences to engineering (Hirsch et al., 2012). In this work, we consider non-linear ODE systems

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_0^t \mathbf{f}(\mathbf{x}(\tau)) d\tau \quad (1)$$

$$\dot{\mathbf{x}}(t) := \frac{d\mathbf{x}(t)}{dt} = \mathbf{f}(\mathbf{x}(t)), \quad (2)$$

where the state vector $\mathbf{x}(t) \in \mathbb{R}^D$ evolves over time $t \in \mathbb{R}_+$ from an initial state \mathbf{x}_0 following its time derivative $\dot{\mathbf{x}}(t)$, and τ is an auxiliary time variable. Our goal is to learn the differential function $\mathbf{f} : \mathbb{R}^D \mapsto \mathbb{R}^D$ from state observations, when the functional form of \mathbf{f} is unknown.

The conventional mechanistic approach involves manually defining the equations of dynamics and optimizing

their parameters (Butcher and Goodwin, 2008), or inferring their posteriors (Girolami, 2008) from data. However, the equations are unknown or ambiguous for many systems, such as human motion (Wang et al., 2008). Some early works explored fitting unknown ODEs with splines (Henderson and Michailidis, 2014), Gaussian processes (Äijö and Lähdesmäki, 2009) or kernel methods (Heinonen and d’Alché-Buc, 2014) by resorting to less accurate gradient matching approximations (Varah, 1982). Recently, Heinonen et al. (2018) proposed estimation of free-form non-linear dynamics using Gaussian processes without gradient matching. However, the approach is restricted to learning point estimates of the dynamics, limiting the uncertainty characterization and generalization. Chen et al. (2018) proposed modeling ODEs with neural networks and adjoints, which was later extended to the Bayesian setting by Dandekar et al. (2020). However, the gradient descent training in such approaches can be ill-suited for complex or long-horizon ODEs with typically highly non-linear integration maps (Diehl and Gros, 2017).

In this work, we introduce efficient Bayesian learning of unknown, non-linear ODEs. Our contributions are:

- We introduce a way of learning posteriors of vectorfields using Gaussian processes as flexible priors over differentials \mathbf{f} , and thereby build on the work by Heinonen et al. (2018). We adapt decoupled functional sampling to simulate ODEs from vector field posteriors.
- For the difficult problem of gradient optimizations of ODEs, we introduce a novel probabilistic shooting method. It is motivated by the canonical shooting methods from optimal control and makes inference stable and efficient on long trajectories.
- We empirically show the effectiveness of the proposed method even while learning from a limited number of observations. We demonstrate the ability to infer arbitrarily long trajectories efficiently with the shooting extension.

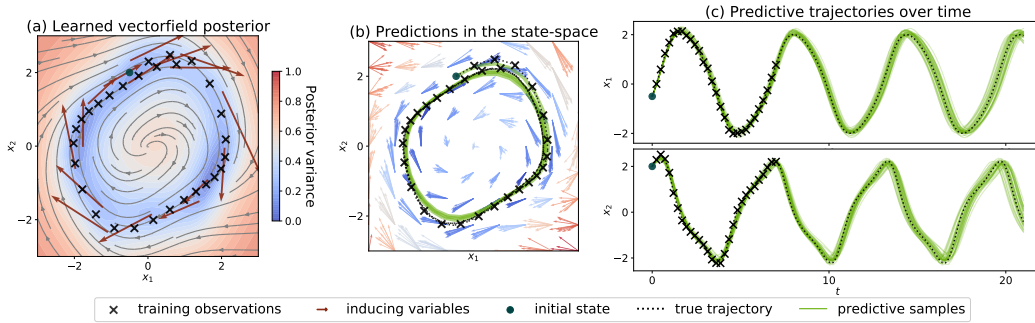


Figure 1: Illustration of GPODE: The model learns a GP posterior (a) of a vector field. Valid ODE trajectories are sampled from the posterior process as shown in (b) and (c).

2 RELATED WORKS

Mechanistic ODE models. In mechanistic modelling the equation f_θ is predefined with a set of coefficients θ to be fitted (Butcher and Goodwin, 2008). Several works have proposed embedding mechanistic models within Bayesian or Gaussian process models (Calderhead et al., 2008; Dondelinger et al., 2013; Wenk et al., 2020). Recently both Julia and Stan have introduced support for Bayesian analysis of parametric ODEs (Rackauckas and Nie, 2017; Stan, 2021). Since this line of work assumes a known dynamics model, we do not consider these methods in the experiments.

Free-form ODE models. Multiple works have proposed fitting unknown, non-linear and free-form ODE differentials with gradient matching using splines (Ramsay et al., 2007), Gaussian processes (Äijö and Lähdesmäki, 2009) or kernel methods (Heinonen and d’Alché-Buc, 2014). Recently, Heinonen et al. (2018) proposed accurate *maximum a posteriori* (MAP) optimisation of vector fields with sensitivity equation gradients (Kokotovic and Heller, 1967). Neural ODEs (Chen et al., 2018) introduced adjoint gradients (Pontryagin et al., 1962) along with flexible black-box neural network vector fields. Several extensions to learning latent ODEs have been proposed (Yildiz et al., 2019; Rubanova et al., 2019).

Discrete-time state-space models. There is a large literature on Markovian state-space models that operate over discrete time increments (Wang et al., 2005; Turner et al., 2010; Frigola et al., 2014). Typically nonlinear state transition functions are modeled with Gaussian processes and applied to latent state estimation or system identification problems with dynamical systems (Eleftheriadis et al., 2017; Doerr et al., 2018; Ialongo et al., 2019). In this paper, we focus strictly on continuous-time models and leave the study of discrete vs. continuous formulations for future work.

Stochastic differential equations. As an alternative formulation of inferring unknown dynamics from observational data, one can assume stochastic transitions and learn models

of stochastic differential equations (SDEs). Existing works have utilized Gaussian processes (Archambeau et al., 2007; Duncker et al., 2019; Jørgensen et al., 2020) and neural networks (Tzen and Raginsky, 2019; Li et al., 2020) to model non-linear SDEs. However, since they assume a different model (i.e. deterministic transitions vs stochastic transitions), we will restrict the experimental comparisons to other ODE-based approaches.

3 METHODS

We consider the problem of learning ODEs (2) with GPs and propose a Bayesian model to infer posteriors over the differential $f(\cdot)$.

3.1 BAYESIAN MODELING OF ODES USING GPs

We assume a sequence of N observations $\mathbf{Y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N)^T \in \mathbb{R}^{N \times D}$ along a trajectory, with $\mathbf{y}_i \in \mathbb{R}^D$ representing the noisy observation of the unknown state $\mathbf{x}(t_i) \in \mathbb{R}^D$ at time t_i . Similar to Heinonen et al. (2018), we assume a zero mean vector-valued Gaussian process prior over \mathbf{f} ,

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

which defines a distribution of differentials $\mathbf{f}(\mathbf{x})$ with covariance $\text{cov}[\mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{x}')] = K(\mathbf{x}, \mathbf{x}')$, where $K(\mathbf{x}, \mathbf{x}') \in \mathbb{R}^{D \times D}$ is a stationary matrix-valued kernel. We follow the commonly used sparse inference framework for GPs using inducing variables (Titsias, 2009), and augment the full model with inducing values $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_M)^T \in \mathbb{R}^{M \times D}$ and inducing locations $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_M)^T \in \mathbb{R}^{M \times D}$ such that $\mathbf{u}_m = \mathbf{f}(\mathbf{z}_m)$. The inducing variables are trainable ‘landmark’ state-differential pairs, from which the rest of the differential field is interpolated (See Figure 1, where arrow locations are the \mathbf{z}_m and arrow end-points are the \mathbf{u}_m). The inducing augmentation leads to the following prior and

conditionals (Hensman et al., 2013):

$$p(\mathbf{U}) = \mathcal{N}(\mathbf{U}|\mathbf{0}, \mathbf{K}_{\mathbf{ZZ}}), \quad (4)$$

$$p(\mathbf{f}|\mathbf{U}; \mathbf{Z}) = \mathcal{N}(\mathbf{f}|\mathbf{A}\text{vec}(\mathbf{U}), \mathbf{K}_{\mathbf{XX}} - \mathbf{A}\mathbf{K}_{\mathbf{ZZ}}\mathbf{A}^T), \quad (5)$$

where $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N'})^T \in \mathbb{R}^{N' \times D}$ collects all the intermediate state evaluations $\mathbf{x}(t_i)$ encountered along a numerical approximation of the true continuous ODE integral (1), $\mathbf{f} = (\mathbf{f}(\mathbf{x}_1)^T, \dots, \mathbf{f}(\mathbf{x}_{N'})^T)^T \in \mathbb{R}^{N' \times D \times 1}$, $\mathbf{K}_{\mathbf{XX}}$ is a block-partitioned matrix of size $N'D \times N'D$ with $D \times D$ blocks, so that block $(\mathbf{K}_{\mathbf{XX}})_{i,j} = K(\mathbf{x}_i, \mathbf{x}_j)$, and $\mathbf{A} = \mathbf{K}_{\mathbf{XZ}}\mathbf{K}_{\mathbf{ZZ}}^{-1}$. For notational simplicity, we assume that the measurement time points are among the time points of the intermediate state evaluations of a numerical ODE solver.

The joint probability distribution follows

$$p(\mathbf{Y}, \mathbf{f}, \mathbf{U}, \mathbf{x}_0) = p(\mathbf{Y}|\mathbf{f}, \mathbf{x}_0)p(\mathbf{f}, \mathbf{U})p(\mathbf{x}_0) \quad (6)$$

$$= \prod_{i=1}^N p(\mathbf{y}_i|\mathbf{f}, \mathbf{x}_0)p(\mathbf{f}|\mathbf{U})p(\mathbf{U})p(\mathbf{x}_0), \quad (7)$$

where the conditional distribution $p(\mathbf{y}_i|\mathbf{f}, \mathbf{x}_0) = p(\mathbf{y}_i|\mathbf{x}_i)$ computes the likelihood over ODE state solutions $\mathbf{x}_i = \mathbf{x}_0 + \int_0^{t_i} \mathbf{f}(\mathbf{x}(\tau))d\tau$.

3.2 VARIATIONAL INFERENCE FOR GP-ODES

In contrast to earlier approach that estimates MAP solutions (Heinonen et al., 2018), our goal is to infer the posterior distribution $p(\mathbf{f}, \mathbf{x}_0|\mathbf{Y})$ of the vector field \mathbf{f} and initial value \mathbf{x}_0 from observations \mathbf{Y} . The posterior is intractable due to the non-linear integration map $\mathbf{x}_0 \xrightarrow{\mathbf{f}} \mathbf{x}(t)$.

We use the stochastic variational inference (SVI) formulation for sparse GPs (Hensman et al., 2013) in this work. We introduce a factorized Gaussian posterior approximation for the inducing variables across state dimensions $q(\mathbf{U}) = \prod_{d=1}^D \mathcal{N}(\mathbf{u}_d|\mathbf{m}_d, \mathbf{Q}_d)$, $\mathbf{u}_d \in \mathbb{R}^M$ where $\mathbf{m}_d \in \mathbb{R}^M$, $\mathbf{Q}_d \in \mathbb{R}^{M \times M}$ are the mean and covariance parameters of the variational Gaussian posterior approximation for the inducing variables. We treat the inducing locations \mathbf{Z} as optimized hyperparameters. The posterior distribution for the variational approximation can be written as

$$q(\mathbf{f}) = \int p(\mathbf{f}|\mathbf{U})q(\mathbf{U})d\mathbf{U} \quad (8)$$

$$= \int \mathcal{N}(\mathbf{f}|\mathbf{A}\text{vec}(\mathbf{U}), \mathbf{K}_{\mathbf{XX}} - \mathbf{A}\mathbf{K}_{\mathbf{ZZ}}\mathbf{A}^T)q(\mathbf{U})d\mathbf{U}. \quad (9)$$

The posterior inference goal then translates to estimating the posterior $p(\mathbf{f}, \mathbf{U}, \mathbf{x}_0|\mathbf{Y})$ of the inducing points \mathbf{U} and initial state \mathbf{x}_0 . Under variational inference this learning objective

$$\arg \min_q \text{KL} [q(\mathbf{f}, \mathbf{U}, \mathbf{x}_0) || p(\mathbf{f}, \mathbf{U}, \mathbf{x}_0|\mathbf{Y})] \quad (10)$$

translates into maximizing the evidence lowerbound (ELBO),

$$\log p(\mathbf{Y}) \geq \sum_{i=1}^N \underbrace{\mathbb{E}_{q(\mathbf{f}, \mathbf{x}_0)} \log p(\mathbf{y}_i|\mathbf{f}, \mathbf{x}_0)}_{\text{variational likelihood}} - \underbrace{\text{KL}[q(\mathbf{U})||p(\mathbf{U})]}_{\text{inducing KL}} - \underbrace{\text{KL}[q(\mathbf{x}_0)||p(\mathbf{x}_0)]}_{\text{initial state KL}}, \quad (11)$$

where we also assume variational approximation $q(\mathbf{x}_0) = \mathcal{N}(\mathbf{a}_0, \Sigma_0)$ for the initial state \mathbf{x}_0 . See supplementary section 1.1 for detailed derivations of the above equations.

3.3 SAMPLING ODES FROM GAUSSIAN PROCESSES

The Picard-Lindelöf theorem (Lindelöf, 1894) ensures valid ODE systems define unique solutions to the initial value problem (IVP) (1). In order to sample valid state trajectories for the IVP, we need to efficiently sample GP functions $\mathbf{f}(\cdot) \sim q(\mathbf{f})$ (9). This way, we can evaluate the sample function $\mathbf{f}(\mathbf{x}(t))$ at arbitrary states $\mathbf{x}(t)$ encountered during ODE forward integration, while accounting for both the inducing and interpolation distributions of Equation (9). Unfortunately, function-space sampling of such GPs has prohibitive cubic complexity (Rasmussen and Williams, 2006; Ustyuzhaninov et al., 2020), while the more efficient weight-space sampling with Fourier bases cannot accurately express the posterior (9) (Wilson et al., 2020).

We use the decoupled sampling that decomposes the posterior into two parts (Wilson et al., 2020),

$$\underbrace{\mathbf{f}(\mathbf{x})|\mathbf{U}}_{\text{posterior}} = \underbrace{\mathbf{f}(\mathbf{x})}_{\text{prior}} + \underbrace{K(\mathbf{x}, \mathbf{Z})K(\mathbf{Z}, \mathbf{Z})^{-1}(\mathbf{U} - \mathbf{f}_{\mathbf{Z}})}_{\text{update}}. \quad (12)$$

$$\approx \sum_{i=1}^F \mathbf{w}_i \phi_i(\mathbf{x}) + \sum_{j=1}^M \nu_j K(\mathbf{x}, \mathbf{z}_j), \quad (13)$$

where we use F Fourier bases $\phi_i(\cdot)$ with $\mathbf{w}_i \sim \mathcal{N}(\mathbf{0}, I)$ (Rahimi and Recht, 2007) to represent the stationary prior, and function basis $K(\cdot, \mathbf{z}_j)$ for the posterior update with $\nu = K(\mathbf{Z}, \mathbf{Z})^{-1}(\mathbf{U} - \Phi\mathbf{W})$, $\Phi = \phi(\mathbf{Z}) \in \mathbb{R}^{M \times F}$, $\mathbf{W} \in \mathbb{R}^{F \times D}$. By combining these two steps, we can accurately evaluate functions from the posterior (9) in linear time at arbitrary locations. We refer the reader to the supplementary section 1.2 for more details. We note that concurrent works by Mikheeva et al. (2021) and Ensinger et al. (2021) also utilize the decoupled-sampling to infer ODE posteriors with GPs.

3.4 AUGMENTING THE ODE MODEL WITH SHOOTING SYSTEM

A key bottleneck in ODE modeling is the poor gradient descent performance over long integration times $\mathbf{x}_{0:T}$, which

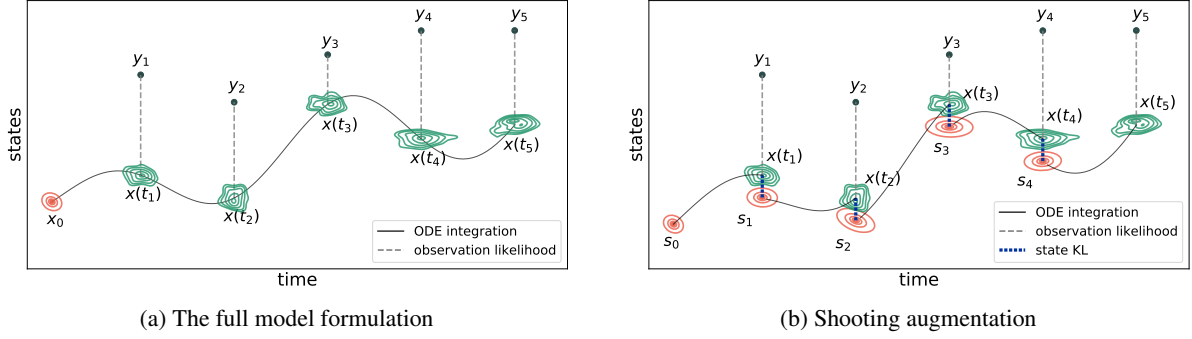


Figure 2: Illustrations of GPODE formulations: the full model formulation (a) follows the long trajectory integration, whereas the shooting version (b) splits the long trajectory into multiple short subintervals.

can exhibit vanishing or exploding gradients (Haber and Ruthotto, 2017; Choromanski et al., 2020). Earlier approaches tackled this issue mainly with more accurate numerical solvers (Zhuang et al., 2020, 2021). The nonlinearity of the integration map $\mathbf{x}_0 \xrightarrow{\mathbf{f}} \mathbf{x}_t$ motivates us to instead segment the full integration $\mathbf{x}_{0:T}$ into short segments, which are easier to optimize and can be trivially parallelized. This is called the *multiple shooting* method in optimal control literature (Osborne, 1969; Bock and Plitt, 1984), in the context of parameter estimation of ODEs (vanDomselaar and Hemker, 1975; Bock, 1983). Recently, Massaroli et al. (2021); Turan and Jäschke (2021) also introduced a multiple-shooting framework within the context of deterministic neural ODEs. We introduce probabilistic shooting for the Gaussian process posterior inference of ODEs.

We begin by introducing shooting state variables $\mathbf{S} = (\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_{N-1})$, $\mathbf{s}_i \in \mathbb{R}^D$, and segment the continuous state function $\mathbf{x}(t; \mathbf{x}_0)$ (1) into N segments $\{(\mathbf{s}_{i-1}, \mathbf{x}(t_i; \mathbf{s}_{i-1}))\}_{i=1}^N$ that branch from the shooting variables \mathbf{s}_{i-1} (See Figure 2);

$$\mathbf{x}(t_i; \mathbf{s}_{i-1}) = \mathbf{s}_{i-1} + \int_{t_{i-1}}^{t_i} \mathbf{f}(\mathbf{x}(\tau)) d\tau. \quad (14)$$

In addition, every shooting variable is approximately matched with the ODE state evolution from the previous shooting state,

$$\mathbf{s}_i = \mathbf{x}(t_i; \mathbf{s}_{i-1}) + \boldsymbol{\xi}, \quad (15)$$

where $\boldsymbol{\xi} \in \mathbb{R}^D$ represents the tolerance parameter controlling the shooting approximation. The augmented system is equivalent to the original ODE system in case the constraints $\mathbf{s}_i = \mathbf{x}(t_i; \mathbf{s}_{i-1})$ are satisfied exactly at the limit $\boldsymbol{\xi} \rightarrow \mathbf{0}$. We place a Gaussian prior over the tolerance parameter $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \sigma_{\boldsymbol{\xi}}^2 \mathbf{I})$, which translates into the following prior over shooting variables

$$p(\mathbf{s}_i | \mathbf{s}_{i-1}) = \mathcal{N}(\mathbf{s}_i | \mathbf{x}(t_i; \mathbf{s}_{i-1}), \sigma_{\boldsymbol{\xi}}^2 \mathbf{I}). \quad (16)$$

Further, the joint probability of the augmented model after

placing a GP prior over the vectorfield \mathbf{f} can be written as

$$p(\mathbf{Y}, \mathbf{S}, \mathbf{f}) = \prod_{i=1}^N p(\mathbf{y}_i | \mathbf{s}_{i-1}, \mathbf{f}) \prod_{i=1}^{N-1} p(\mathbf{s}_i | \mathbf{s}_{i-1}, \mathbf{f}) p(\mathbf{s}_0) p(\mathbf{f}). \quad (17)$$

3.5 VARIATIONAL INFERENCE FOR THE AUGMENTED MODEL

To infer the augmented posterior $p(\mathbf{f}, \mathbf{U}, \mathbf{S} | \mathbf{Y})$ we introduce variational approximation for the shooting variables $q(\mathbf{S}) = q(\mathbf{s}_0) \cdots q(\mathbf{s}_{N-1})$, where each distribution $q(\mathbf{s}_i) = \mathcal{N}(\mathbf{s}_i | \mathbf{a}_i, \Sigma_i)$ is a Gaussian. This results in the joint variational approximation

$$q(\mathbf{S}, \mathbf{f}, \mathbf{U}) = \prod_{i=0}^{N-1} q(\mathbf{s}_i) p(\mathbf{f} | \mathbf{U}) q(\mathbf{U}), \quad (18)$$

and the following evidence lower bound for the shooting model,

$$\begin{aligned} \mathcal{L}_{\text{shooting}} &= \sum_{i=1}^N \mathbb{E}_{q(\mathbf{s}_{i-1}, \mathbf{f})} \left[\log p(\mathbf{y}_i | \mathbf{s}_{i-1}, \mathbf{f}) \right] \\ &+ \sum_{i=1}^{N-1} \mathbb{E}_{q(\mathbf{s}_i, \mathbf{s}_{i-1}, \mathbf{f})} \left[\log p(\mathbf{s}_i | \mathbf{s}_{i-1}, \mathbf{f}) \right] - \mathbb{E}_{q(\mathbf{s}_i)} \left[\log q(\mathbf{s}_i) \right] \\ &- \text{KL}[q(\mathbf{s}_0) || p(\mathbf{s}_0)] - \text{KL}[q(\mathbf{U}) || p(\mathbf{U})]. \end{aligned} \quad (19)$$

The ELBO consists of an expected log-likelihood term, which matches the state evolution (14) from every shooting variable to the corresponding observation. In addition, the posterior approximation for every shooting variable is also matched with the ODE evolution of the approximated posterior of the previous shooting state, leading to corresponding cross-entropy and entropy terms.

The ELBO for the augmented shooting model requires solving only the short segments (14) with simpler integration maps, thus great at mitigating problems with vanishing/exploding gradients. Since the involved numerical ODE integration can be done in parallel, the shooting model is also

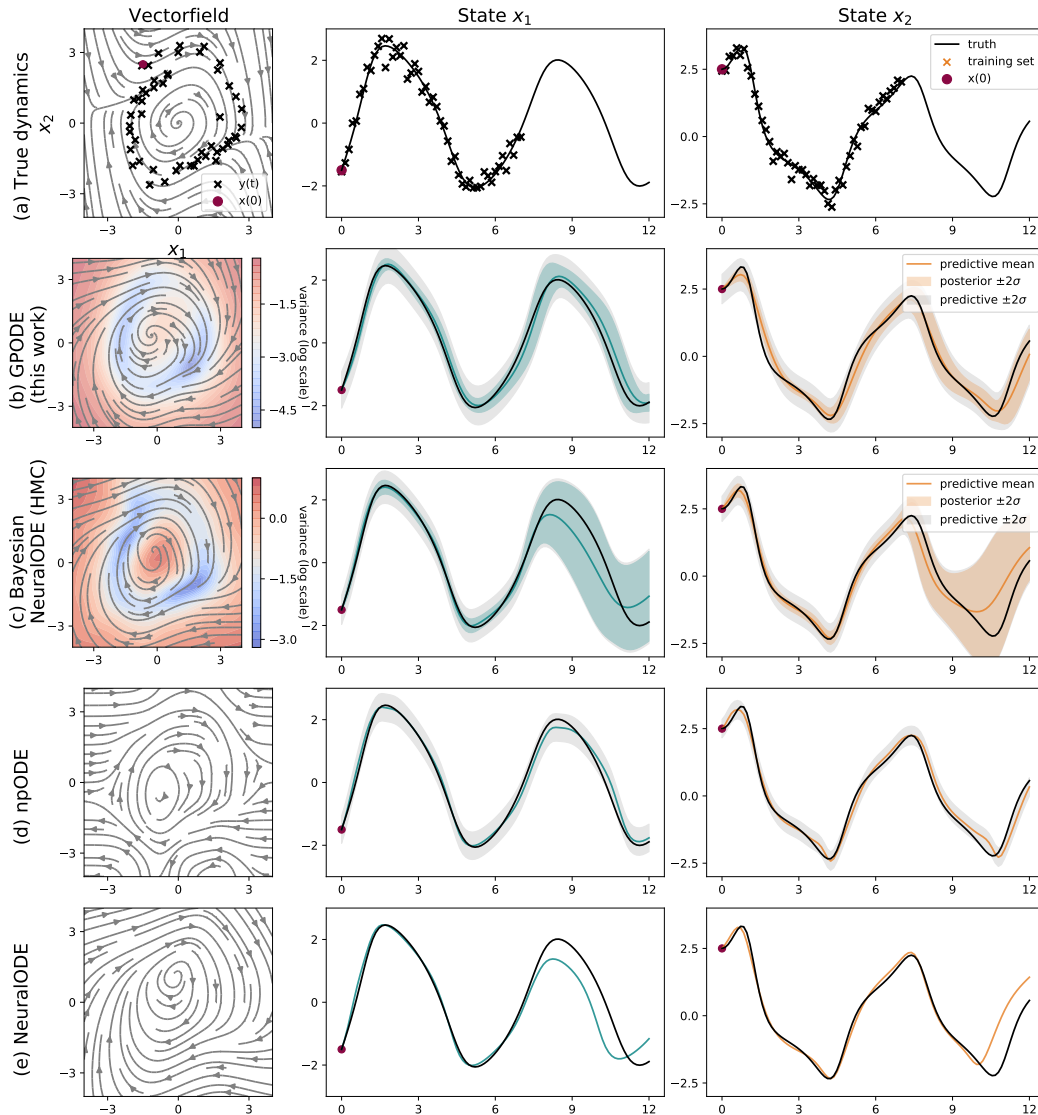


Figure 3: Learning the 2D Van der Pol dynamics (a) with alternative methods (b-d). Column 1 shows the vector fields while columns 2 and 3 show the state trajectories $x_1(t)$ and $x_2(t)$. GPODE learns the posterior accurately.

computationally faster than the full model in practice. See supplementary section 1.3 for a plate diagram and detailed derivation of the approach.

4 EXPERIMENTS

We validate the proposed method on Van der Pol (VDP) and FitzHughNagumo (FHN) systems and on the task of learning human motion dynamics (MoCap). The predictive performance of the proposed GPODE is compared against npODE (Heinonen et al., 2018), NeuralODE (Chen et al., 2018) and Bayesian version of NeuralODE (Dandekar et al., 2020). We use 16 inducing points in VDP and FHN experiments and 100 inducing points for the MoCap experiments. Except for the NeuralODE model, we assume Gaussian

observation likelihood, and infer the unknown noise scale parameter from the training data. All the experiments use squared exponential kernel with automatic relevance determination (ARD) along with 256 Fourier basis functions for decoupled GP sampling. Along with the variational parameters, kernel lengthscales, signal variance, noise scale, and inducing locations are jointly optimized against the model ELBO while training. In addition, for the shooting model, we fix the constraint tolerance parameter to a small value $\sigma_\xi^2 = 1e^{-6}$ consistently across all the experiments. In all the shooting experiments, we considered the number of shooting segments to be the same as the number of observation segments in the dataset. A codebase for implementing the proposed methods is provided <https://github.com/hegdepashupati/gaussian-process-odes>.

We use the `dopri5` solver with tolerance parameters

Table 1: VDP system learning performance on extrapolation task with observations on regular (task 1) and irregular time intervals (task 2). We report mean \pm standard error over 5 runs from different random initialization, the best values bolded. (\uparrow): higher is better, (\downarrow) lower is better

	Task 1: Regular time-grid		Task 2: Irregular time-grid	
	MNLL (\downarrow)	MSE (\downarrow)	MNLL (\downarrow)	MSE (\downarrow)
Bayesian NeuralODE (HMC)	0.82 \pm 0.01	1.45 \pm 0.04	0.88 \pm 0.01	1.68 \pm 0.04
NeuralODE	-	0.29 \pm 0.11	-	0.55 \pm 0.07
npODE	1.47 \pm 0.59	0.16 \pm 0.05	8.89 \pm 3.06	2.08 \pm 0.78
GPODE	0.60 \pm 0.03	0.13 \pm 0.01	0.41 \pm 0.18	0.21 \pm 0.07

$rtol=1e^{-5}$ and $atol=1e^{-5}$, and use the adjoint method for computing loss gradients with `torchdiffeq`¹ package (Chen et al., 2018). All the experiments are repeated 5 times with random initialization, and means and standard errors are reported over multiple runs. The predictive performance of different models are measured with mean squared error (MSE) and mean negative log likelihood (MNLL) metrics.

4.1 LEARNING VAN DER POL DYNAMICS

We first illustrate the effectiveness of the proposed method by inferring the vector field posterior on a two-dimensional VDP (see Figure 3),

$$\begin{aligned} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= -x_1 + 0.5x_2(1 - x_1^2). \end{aligned} \quad (20)$$

We simulate a trajectory of 50 states following the true system dynamics from the initial state $(x_1(0), x_2(0)) = (-1.5, 2.5)$, and add Gaussian noise with $\sigma^2 = 0.05$ to generate the training data. We explore two scenarios with training time interval $t \in [0, 7]$ and forecasting interval $t \in [7, 14]$: (1) over a regularly sampled time grid, (2) over an irregular grid using uniform random sampling of time points. Task (2) demonstrates one of the key advantages of continuous-time models with the ability to handle irregular data.

Figure 3(b) shows that both GPODE and Bayesian NeuralODE learn a vector field posterior whose posterior mean closely matches the ground truth, with low variance (blue regions) near the observed data. The posterior variance increases away from the observed data (orange regions), indicating a good uncertainty characterization, while the npODE with MAP estimation seems to overfit. NeuralODE learns an appropriate vector field, but requires careful tuning of regularization and hyperparameters for a good fit with a limited number of observations. A quantitative evaluation of the model fits in Table 1 indicates the better performance of GPODE as compared to the other methods under comparison.

¹<https://github.com/rtqichen/torchdiffeq>

Table 2: Imputation results on the FHN system.

	MNLL (\downarrow)	MSE (\downarrow)
Bayesian NeuralODE (HMC)	0.77 \pm 0.12	0.24 \pm 0.03
NeuralODE	-	0.18 \pm 0.00
npODE	6.49 \pm 1.49	0.08 \pm 0.01
GPODE	0.09 \pm 0.05	0.07 \pm 0.02

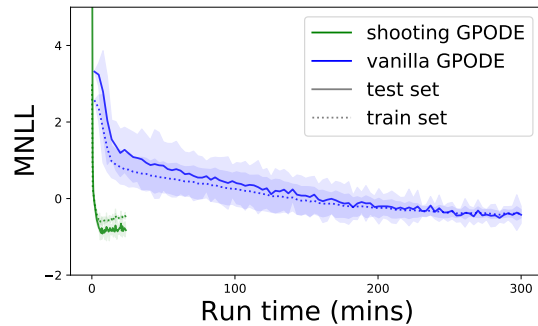


Figure 4: Optimization efficiency with GPODE models.

4.2 LEARNING WITH MISSING OBSERVATIONS

We illustrate the usefulness of learning Bayesian ODE posteriors under missing data with the FHN oscillator

$$\begin{aligned} \dot{x}_1 &= 3(x_1 - x_1^3/3 + x_2), \\ \dot{x}_2 &= (0.2 - 3x_1 - 0.2x_2)/3. \end{aligned} \quad (21)$$

We generate a training sequence by simulating 25 regularly-sampled time points from $t \in [0, 5.0]$ with added Gaussian noise with $\sigma^2 = 0.025$. We remove all observations at the quadrant $x_1 > 0, x_2 < 0$ and evaluate model accuracy in this region. The interpolation performance for different models is shown in Table 2. The point estimates of npODE and NeuralODE have biases, while the Bayesian variants of GPODE and NeuralODE provide good uncertainty estimates corresponding to their better predictive performance.

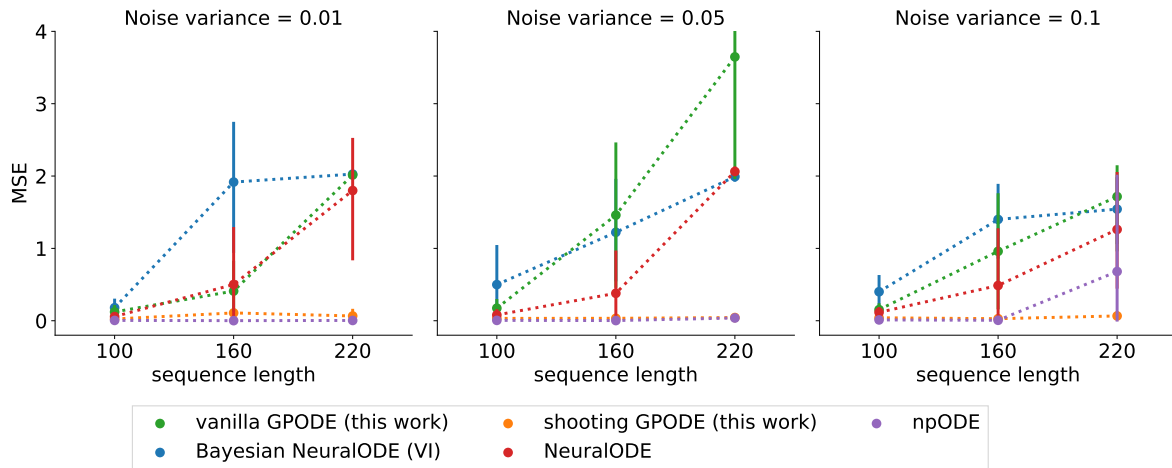


Figure 5: Varying sequence length and observation noise: shooting formulation makes GPODE feasible for long sequences, outperforming the non-shooting version and competing methods. We report the results for different levels of observation noise and training sequence length on the VDP system.

4.3 LEARNING LONG TRAJECTORIES WITH THE SHOOTING FORMULATION

We demonstrate the necessity of the shooting formulation for working with long training trajectories. We use the VDP system with four observations per unit of time for $T = (25, 40, 55)$ corresponding to $N = (100, 160, 220)$ observed states. We also vary the observation variance as $\sigma^2 = (0.01, 0.05, 0.1)$ and test the model for forecasting additional 50 time points.

Figure 5 demonstrates that vanilla-GPODE and NeuralODE, and Bayesian NeuralODE fail to fit the data with long sequences on all noise levels. In contrast, inference for the shooting model is successful in all settings. The npODE is remarkably robust to long trajectories. We believe the robustness of npODE mainly stems from the excellent parameter initialization strategy (see supplementary section 2.2) coupled with the fully deterministic optimization setup (no reparametrization gradients).

Figure 4 shows a runtime trace comparison between vanilla GPODE and the shooting variant in wall-clock time for a fixed budget of 15000 optimization steps on the VDP system with $N = 100$, $T = 25$ and $\sigma^2 = 0.01$. The shooting model converges approximately 10 times faster. The speedup stems from the parallelization of the shooting ODE solver, since the shooting method splits the full IVP problem into numerous short and less non-linear IVPs. In addition, the shooting method relaxes the inference problem with its auxiliary augmentation. This experiment was conducted on a system with AMD Ryzen 5 3600 processor and Nvidia GeForce GTX 1660S GPUs.

4.4 LEARNING HUMAN MOTION DYNAMICS

We learn the dynamics of human motion from noisy experimental data from CMU MoCap database for three subjects, 09, 35 and 39. The dataset consists of 50 sensor readings from different parts of the body while walking or running. We follow the preprocessing of Wang et al. (2008) and center the data. The dataset was further split into train, test, and validation sequences. We observed that the NeuralODE, the Bayesian NeuralODE version with VI, and npODE models suffer from over-fitting, and we remedy this by applying early stopping by monitoring the validation loss during optimization.

We project the original 50-dimensional data into a 5-dimensional latent space using PCA and learn the dynamics in the latent space. To compute the data likelihood, we project the latent dynamics back to the original data space by inverting the PCA. We divide the experiment into sub-tasks MoCap-short and MoCap-long, based on the length of the sequence considered for model training (see the supplementary section for more details on the dataset and experimental setup). The model predictive performance is measured on unseen test sequences in both tasks.

Table 3 indicates that GPODE outperforms the competing npODE and NeuralODE model variants. Figure 6 visualizes the predicted dynamics for a test sequence. The GPODE variants have reasonable posterior uncertainties, while NeuralODE variants and npODE tend to be overconfident and make more mistakes (see Figure 6 (b), sensors 05, 41 and 47). We note that some variations in the data space cannot be accurately estimated due to the low-dimensional PCA projection.

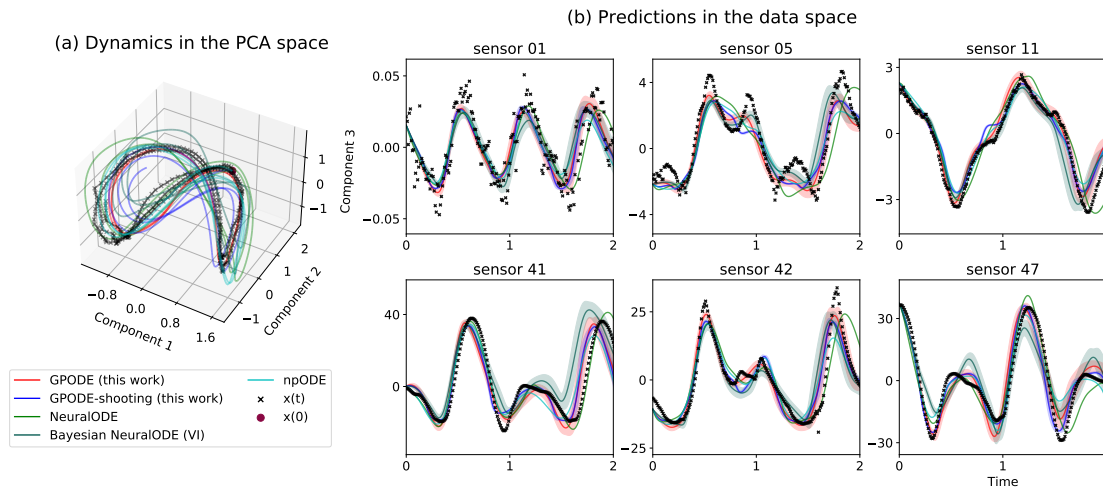


Figure 6: Learning the walking dynamics of subject 39: The true dynamics and predicted dynamics (mean) for the first three components in PCA space are shown in (a). Corresponding trajectories in the observation space for 6 different sensors are shown in (b) (We do not plot the observation noise variance)

Table 3: Test MNLL and MSE metrics for dynamics prediction task on CMU MoCap dataset.

Metric	Method	Subject 09		Subject 35		Subject 39	
		short	long	short	long	short	long
MNLL(\downarrow)	Bayesian NeuralODE (VI)	2.03 ± 0.10	1.50 ± 0.05	1.42 ± 0.05	1.37 ± 0.06	1.61 ± 0.07	1.45 ± 0.03
	npODE	2.09 ± 0.01	1.78 ± 0.08	1.67 ± 0.02	1.66 ± 0.04	2.06 ± 0.05	1.78 ± 0.04
	GPODE-vanilla	1.30 ± 0.02	1.26 ± 0.02	1.27 ± 0.04	1.39 ± 0.04	1.29 ± 0.01	1.13 ± 0.01
	GPODE-shooting	1.19 ± 0.02	1.14 ± 0.02	1.25 ± 0.06	1.08 ± 0.02	1.25 ± 0.01	1.36 ± 0.02
MSE(\downarrow)	Bayesian NeuralODE (VI)	25.50 ± 1.70	21.32 ± 2.58	23.09 ± 3.95	20.86 ± 2.95	53.34 ± 5.31	39.66 ± 6.82
	NeuralODE	27.53 ± 2.87	33.83 ± 2.46	36.50 ± 3.86	23.54 ± 0.56	115.38 ± 10.96	53.51 ± 2.98
	npODE	17.91 ± 1.62	19.76 ± 4.29	26.24 ± 2.88	22.83 ± 3.91	92.80 ± 15.74	55.94 ± 4.63
	GPODE-vanilla	15.78 ± 0.67	12.62 ± 1.14	16.14 ± 0.99	15.53 ± 0.76	20.71 ± 1.25	23.64 ± 1.94
	GPODE-shooting	9.11 ± 0.37	8.38 ± 1.23	10.11 ± 0.79	11.66 ± 0.73	26.72 ± 0.63	21.17 ± 2.88

5 CONCLUSION AND DISCUSSION

We proposed a novel model for Bayesian inference of ODEs using Gaussian processes. With this approach, one can model unknown ODE systems directly from the observational data and learn posteriors of the continuous-time vector fields. In contrast, earlier works produce point estimate solutions. We believe this to be a significant addition to the data-descriptive ODE modeling methods, especially for applications where uncertainty quantification is critical. Many conventional machine learning algorithms have been interpreted and modeled as continuous-time dynamical systems, with applications to generative modeling (Grathwohl et al., 2019) and probabilistic alignment (Ustyuzhaninov et al., 2020), among others. However, scaling GPs to high-dimensional datasets (such as images) can be a bottleneck. The applicability of the proposed model as a plug-in extension for these applications can be studied as part of future work.

We also highlighted a problem of learning black-box ODE models on long trajectories and proposed a probabilistic shooting framework enabling efficient inference on such

tasks. This framework can be applied to other existing approaches, such as NeuralODEs. However, the proposed shooting augmentation introduces model approximation and involves approximating inference over auxiliary shooting variables. Hence the benefits of the shooting augmentation can be task specific, especially on short sequences. Comprehensive empirical studies across different types of tasks can be considered in future work.

References

- Tarmo Äijö and Harri Lähdesmäki. Learning gene regulatory networks from gene expression measurements using non-parametric molecular kinetics. *Bioinformatics*, 25 (22):2937–2944, 2009.
- Cedric Archambeau, Dan Cornford, Manfred Opper, and John Shawe-Taylor. Gaussian process approximations of stochastic differential equations. In *Gaussian Processes in Practice*, pages 1–16. PMLR, 2007.
- Hans Bock and Karl Plitt. A multiple shooting algorithm

- for direct solution of optimal control problems. In *IFAC World Congress*, pages 242–247, 1984.
- Hans Georg Bock. Recent advances in parameter identification techniques for ode. *Numerical treatment of inverse problems in differential and integral equations*, pages 95–121, 1983.
- John Butcher and Nicolette Goodwin. *Numerical methods for ordinary differential equations*. Wiley Online Library, 2nd edition, 2008.
- Ben Calderhead, Mark Girolami, and Neil Lawrence. Accelerating Bayesian inference over nonlinear differential equations with Gaussian processes. In *Advances in Neural Information Processing Systems*, 2008.
- Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pages 6571–6583, 2018.
- Krzysztof Choromanski, Jared Davis, Valerii Likhoshershtov, Xingyou Song, Jean-Jacques Slotine, Jacob Varley, Honglak Lee, Adrian Weller, and Vikas Sindhwani. An Ode to an ODE. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Raj Dandekar, Karen Chung, Vaibhav Dixit, Mohamed Tarek, Aslan Garcia-Valadez, Krishna Vishal Vemula, and Chris Rackauckas. Bayesian neural ordinary differential equations. *arXiv preprint arXiv:2012.07244*, 2020.
- Moritz Diehl and Sebastian Gros. *Numerical Optimal Control*. University of Freiburg, 2017.
- Andreas Doerr, Christian Daniel, Martin Schiegg, Nguyen-Tuong Duy, Stefan Schaal, Marc Toussaint, and Trimpe Sebastian. Probabilistic recurrent state-space models. In *International Conference on Machine Learning*, pages 1280–1289, 2018.
- Frank Dondelinger, Dirk Husmeier, Simon Rogers, and Maurizio Filippone. ODE parameter inference using adaptive gradient matching with Gaussian processes. In *Artificial Intelligence and Statistics*, pages 216–228, 2013.
- Lea Duncker, Gergo Böhner, Julien Boussard, and Maneesh Sahani. Learning interpretable continuous-time models of latent stochastic dynamical systems. In *International Conference on Machine Learning*, pages 1726–1734. PMLR, 2019.
- Stefanos Eleftheriadis, Tom Nicholson, Marc Peter Deisenroth, and James Hensman. Identification of Gaussian process state space models. In *Advances in Neural Information Processing Systems*, pages 5309–5319, 2017.
- Katharina Ensinger, Friedrich Solowjow, Michael Tiemann, and Sebastian Trimpe. Symplectic gaussian process dynamics. *arXiv preprint arXiv:2102.01606*, 2021.
- Roger Frigola, Yutian Chen, and Carl Rasmussen. Variational Gaussian process state-space models. In *Advances in Neural Information Processing Systems*, 2014.
- Mark Girolami. Bayesian inference for differential equations. *Theoretical Computer Science*, 408:4–16, 2008.
- Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD: Free-form continuous dynamics for scalable reversible generative models. In *ICLR*, 2019.
- Eldad Haber and Lars Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 34:014004, 2017.
- Markus Heinonen and Florence d’Alché-Buc. Learning nonparametric differential equations with operator-valued kernels and gradient matching. Technical report, Université d’Evry, 2014.
- Markus Heinonen, Cagatay Yildiz, Henrik Mannerström, Jukka Intosalmi, and Harri Lähdesmäki. Learning unknown ODE models with Gaussian processes. In *International Conference on Machine Learning*, pages 1959–1968, 2018.
- James Henderson and George Michailidis. Network reconstruction using nonparametric additive ODE models. *PLOS ONE*, 9(4):e94003, 2014.
- James Hensman, Nicolò Fusi, and Neil Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, pages 282–290, 2013.
- Morris Hirsch, Stephen Smale, and Robert Devaney. *Differential equations, dynamical systems, and an introduction to chaos*. Academic press, 2012.
- Alessandro Davide Ialongo, Mark Van Der Wilk, James Hensman, and Carl Edward Rasmussen. Overcoming mean-field approximations in recurrent gaussian process models. In *International Conference on Machine Learning*, pages 2931–2940, 2019.
- Martin Jørgensen, Marc Deisenroth, and Hugh Salimbeni. Stochastic differential equations with variational wishart diffusions. In *International Conference on Machine Learning*, pages 4974–4983. PMLR, 2020.
- Petar Kokotovic and James Heller. Direct and adjoint sensitivity equations for parameter optimization. *IEEE Transactions on Automatic Control*, 12(5):609–610, 1967.
- Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pages 3870–3882. PMLR, 2020.

- Ernst Lindelöf. Sur l'application de la méthode des approximations successives aux équations différentielles ordinaires du premier ordre. *Comptes rendus hebdomadaires des séances de l'Académie des sciences*, 116:454–457, 1894.
- Stefano Massaroli, Michael Poli, Sho Sonoda, Taiji Suzuki, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Differentiable multiple shooting layers. *Advances in Neural Information Processing Systems*, 34, 2021.
- Olga Mikheeva, Ieva Kazlauskaitė, Adam Hartshorne, Hedvig Kjellström, Carl Henrik Ek, and Neill DF Campbell. Aligned multi-task gaussian process. *arXiv preprint arXiv:2110.15761*, 2021.
- Michael Osborne. On shooting methods for boundary value problems. *Journal of Mathematical Analysis and Applications*, 27:417–433, 1969.
- Lev Pontryagin, Evgenii Mishchenko, Vladimir Boltyanskii, and Revas Gamkrelidze. *The mathematical theory of optimal processes*. Interscience Publishers, 1962. Translation KN Tririgoff.
- Christopher Rackauckas and Qing Nie. Differentialequations.jl – a performant and feature-rich ecosystem for solving differential equations in Julia. *Journal of Open Research Software*, 5(1), 2017.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2007.
- Jim Ramsay, Giles Hooker, David Campbell, and Jiguo Cao. Parameter estimation for differential equations: A generalized smoothing approach. *Journal of the Royal Statistical Society: Series B*, 69(5):741–796, 2007.
- Carl Rasmussen and Christopher Williams. *Gaussian processes for machine learning*. The MIT Press, 2006.
- Yulia Rubanova, Ricky T. Q. Chen, and David Duvenaud. Latent ODEs for irregularly-sampled time series. In *Advances in Neural Information Processing Systems*, 2019.
- Development Team Stan. Stan modeling language users guide and reference manual, mc-stan.org. 2021.
- Michalis Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574, 2009.
- Evren Mert Turan and Johannes Jäschke. Multiple shooting with neural differential equations. *arXiv preprint arXiv:2109.06786*, 2021.
- Ryan Turner, Marc Deisenroth, and Carl Rasmussen. State-space inference and learning with Gaussian processes. In *Artificial Intelligence and Statistics*, pages 868–875, 2010.
- Belinda Tzen and Maxim Raginsky. Neural stochastic differential equations: Deep latent gaussian models in the diffusion limit. *arXiv preprint arXiv:1905.09883*, 2019.
- Ivan Ustyuzhaninov, Ieva Kazlauskaitė, Carl Henrik Ek, and Neill Campbell. Monotonic Gaussian process flows. In *Artificial Intelligence and Statistics*, volume 108, pages 3057–3067, 2020.
- B vanDomselaar and Piet W Hemker. Nonlinear parameter estimation in initial value problems. *Stichting Mathematisch Centrum. Numerieke Wiskunde*, (NW 18/75), 1975.
- James Varah. A spline least squares method for numerical parameter estimation in differential equations. *SIAM Journal on Scientific and Statistical Computing*, 3(1): 28–46, 1982.
- Jack Wang, Aaron Hertzmann, and David Fleet. Gaussian process dynamical models. In *Advances in Neural Information Processing Systems*, 2005.
- Jack Wang, David Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):283–298, 2008.
- Philippe Wenk, Gabriele Abbati, Michael Osborne, Bernhard Schölkopf, Andreas Krause, and Stefan Bauer. ODIN: ODE-informed regression for parameter and state inference in time-continuous dynamical systems. In *AAAI Conference on Artificial Intelligence*, volume 34, pages 6364–6371, 2020.
- James Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *International Conference on Machine Learning*, pages 10292–10302, 2020.
- Cagatay Yildiz, Markus Heinonen, and Harri Lähdesmäki. ODE2VAE: Deep generative second order ODEs with Bayesian neural networks. In *Advances in Neural Information Processing Systems*, pages 13412–13421, 2019.
- Juntang Zhuang, Nicha Dvornek, Xiaoxiao Li, Sekhar Tatikonda, Xenophon Papademetris, and James Duncan. Adaptive checkpoint adjoint method for gradient estimation in neural ODE. In *International Conference on Machine Learning*, pages 11639–11649, 2020.
- Juntang Zhuang, Nicha Dvornek, Sekhar Tatikonda, and James Duncan. MALI: A memory efficient and reverse accurate integrator for neural ODEs. In *ICLR*, 2021.