

System-Aware Neural ODE Processes for Few-Shot Bayesian Optimization

Anonymous authors

Paper under double-blind review

Abstract

We consider the problem of optimizing initial conditions and termination time in dynamical systems governed by unknown ordinary differential equations (ODEs), where evaluating different initial conditions is costly and the state’s value can not be measured in real-time but only with a delay while the measuring device processes the sample. To identify the optimal conditions in limited trials, we introduce a few-shot Bayesian Optimization (BO) framework based on the system’s prior information. At the core of our approach is the System-Aware Neural ODE Processes (SANODEP), an extension of Neural ODE Processes (NODEP) designed to meta-learn ODE systems from multiple trajectories using a novel context embedding block. We further develop a two-stage BO framework to effectively incorporate search space constraints, enabling efficient optimization of both initial conditions and observation timings. We conduct extensive experiments showcasing SANODEP’s potential for few-shot BO within dynamical systems. We also explore SANODEP’s adaptability to varying levels of prior information, highlighting the trade-off between prior flexibility and model fitting accuracy.

1 Introduction

This paper studies a widely encountered, yet less investigated, problem: optimizing the initial conditions and termination time in unknown dynamical systems where evaluations are computationally expensive. We assume that the primary evaluation cost comes from switching initial conditions and wish to use as few trajectories as possible.

This issue is prevalent in multiple fields, including biotechnology, epidemiology, ecology, and chemistry. For instance, consider the optimization of a (bio)chemical reactor. Here, the objective is to determine the optimal “recipe,” i.e., the set of initial reactant concentrations and reaction times that maximize yield, enhance selectivity, and/or minimize waste (Taylor et al., 2023; Schoepfer et al., 2024; Schilter et al., 2024; Thebelt et al., 2022). Specifically, the high costs associated with changing reactants for multiple experimental runs highlight the need for developing efficient optimization algorithms (Paulson & Tsay, 2024).

Bayesian Optimization (BO) (Frazier, 2018; Garnett, 2023) is a well-established method for optimizing expensive-to-evaluate black-box objective functions. It relies on probabilistic surrogate models built from limited function evaluations to guide the optimization efficiently. However, standard Gaussian Processes (GPs) (Williams & Rasmussen, 2006) (the *de facto* probabilistic surrogate model in BO) with conventional kernels do not capture the dynamics of these systems effectively. Recent efforts to introduce Bayesian ODE models aim to incorporate suitable assumptions; however, their inference often involves time-consuming computations (Heinonen et al., 2018; Dandekar et al., 2020), making them impractical for a time-sensitive optimization scenario. Some Bayesian ODE models attempt to alleviate computational burdens by leveraging crude approximation inference (Ott et al., 2023) or approximating numerical integrators (Hegde et al., 2022; Ensinger et al., 2024). While these approaches reduce computation time, the necessary approximations may degrade model performance. Consequently, the widely encountered yet under-investigated problem of performing BO in dynamical systems remains open, primarily due to the lack of an appropriate model.




Figure 1: Illustration of meta-learning-based few-shot Bayesian Optimization (BO) with time-delay constraints (detailed in Section. 4) on the Lotka-Volterra (LV) system on different models: Gaussian Process (GP), Neural Process (NP) and System Aware Neural ODE Process (SANODEP). We start with one trajectory (marked as \circ). Each trajectory ensures that adjacent observations respect a minimum time delay Δt constraint. The figure clearly shows that after evaluating just one trajectory (the first, left-most column of three rows), the meta-learned SANODEP model already resembles the original system more closely than the non-meta-learned GP, demonstrating its efficiency in guiding the search toward promising areas.

This work attempts to tackle this optimization problem by conceptualizing it as a *few-shot* optimization, through leveraging *prior information* about dynamical systems to formulate the *meta tasks* to train a learning model, which is then able to adapt to new problems with very few data (a.k.a., shots) and served as our BO surrogate model. To achieve this, we focus on the Neural ODE Process (NODEP) (Norcliffe et al., 2021) as our learning model due to several useful model properties. First, NODEP combines Neural ODEs (Chen et al., 2018)—emphasizing the dynamical system perspective that may provide a more informative inductive bias—with Neural Processes (NP) (Garnelo et al., 2018b), enhancing the meta-learning aspect for few-shot optimization. Second, meta-learning endows NODEP with fast adaptation to new observations, mitigating the potential concern of model training time with incoming measurement results. Third, NODEP enables continuous gradient-based optimization over time, compared with discrete-time meta-learn models (e.g., Foong et al. (2020)), integrating seamlessly within the BO framework. Finally, compared to other continuous-time models involving computationally intensive operations (e.g., time domain attention (Chen et al., 2024)), NODEP is computationally efficient, making it suitable for BO.

However, building a meta-learning-based few-shot BO framework for unknown dynamical systems is not a straightforward downstream application. An obstacle is that NODEP’s functionality is limited by its original learning objective: predicting a single *trajectory* state value given *context* data for that trajectory. For initial condition optimization of unknown ODE systems, it is essential to learn to be aware of the underlying governing systems using context data consisting of one or several different trajectory observations, as this benefits for forecasting trajectories starting from arbitrary initial conditions. Moreover, no BO frameworks are proposed for such problem settings.

In line with our optimization requirements, we introduce the System-Aware Neural ODE Process (SANODEP), a generalization of NODEP in terms of meta-learning of ODE systems to plug-in our newly developed BO framework for few-shot BO in dynamical systems. Fig. 1 shows, through example, SANODEP outcomes in comparison to other approaches. Our contributions are as follows:

1. We propose a novel context embedding mechanism enabling SANODEP to meta-learn from batch of trajectories with minimal model structure adjustments based on NODEP.
2. We developed a model-agnostic, time-delay constraint process BO framework for optimizing initial conditions and termination time in dynamical systems.
3. We compare SANODEP and other meta-learned models under our BO framework, validating the benefits of the ODE-aware model structure and demonstrating the effectiveness of the few-shot BO framework’s model agnosticity property.

4. We investigate how the different levels of prior information can be utilized for SANODEP. Strong prior information can enable a physically-informed model structure with extended loss considering parameter inference, enabling a novel *few-shot parameter estimation* functionality, while weak prior information may still be useful through a properly designed *task distribution*, albeit at a detriment to model fitting capability.

The rest of the paper is organized as follows: In Section 2, we describe the preliminaries. Section 3 introduces the System-Aware Neural ODE Processes (SANODEP). In Section 4, we develop the optimization framework based on SANODEP specific to our optimization problem. Section 5 discusses related work. Section 6 presents the numerical experiments we conducted. In Section 7, we investigate the impact of different levels of prior information.

2 Problem Statement and Background

2.1 State Optimization

Consider a dynamical system whose evolution is described by the following ordinary differential equation (ODE) system in the time domain, denoted as t :

$$\begin{cases} \frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t) \\ \mathbf{x}(t = t_0) := \mathbf{x}_0 \end{cases} \quad (1)$$

where $\mathbf{x}(t) \in \mathcal{X}_{\mathbf{x}} \subset \mathbb{R}^{d_{\mathbf{x}}}$ represents the state value of the system at time t , $\mathbf{f}(\cdot) : \mathbb{R}^{d_{\mathbf{x}}} \times \mathbb{R} \rightarrow \mathbb{R}^{d_{\mathbf{x}}}$ is an unknown function (*vector field*) governing the system dynamics, and $\mathbf{x}_0 \in \mathcal{X}_0 \subseteq \mathcal{X}_{\mathbf{x}} \subset \mathbb{R}^{d_{\mathbf{x}}}$ is the initial system state. Define $\mathbf{f}_{\text{evolve}}(\mathbf{x}_0, t, \mathbf{f}) : \mathcal{X}_0 \times \tau \times (\mathbb{R}^{d_{\mathbf{x}}} \times \mathbb{R} \rightarrow \mathbb{R}^{d_{\mathbf{x}}}) \rightarrow \mathbb{R}^{d_{\mathbf{x}}}$, where $\tau \in [t_0, t_{\max}]$, to obtain the system state at time t from an initial condition \mathbf{x}_0 as: $\mathbf{f}_{\text{evolve}}(\mathbf{x}_0, t, \mathbf{f}) = \mathbf{x}_0 + \int_{t_0}^t \mathbf{f}(\mathbf{x}(h), h) dh$. Let $g(\cdot) : \mathbb{R}^{d_{\mathbf{x}}} \rightarrow \mathbb{R}$ be a practitioner-specified (known) function that aggregates the state values into a scalar, we consider the following multi-objective problem:

$$\text{maximize}_{\{t, \mathbf{x}_0\} \in \tau \times \mathcal{X}_0} \quad g(\mathbf{f}_{\text{evolve}}(\mathbf{x}_0, t, \mathbf{f})), -t \quad (2)$$

The goal is to identify the initial conditions \mathbf{x}_0 along with the corresponding evolution termination time t that provide the optimal trade-off between the objective g applied to the state values $\mathbf{x}(t)$ and the amount of time to reach this state. Simply put, we wish to maximize objective g early.

2.2 Neural ODE Processes (NODEP)

NODEP (Norcliffe et al., 2021) is a latent variable-based Bayesian meta-learning model with generative processes that can be summarized as follows: given a set of N observations of state values at different times from a single trajectory, represented by context set $\mathbb{C} := \{(t_i^{\mathbb{C}}, \mathbf{x}_i^{\mathbb{C}})\}_{i=1}^N$, NODEP assumes the conditional prediction has been generated from a latent controlled ODE of state dimensionality d_{ℓ} . The stochasticity of the model is induced by stochastic latent initial condition L_0 and the stochastic dynamics representation, which in practice is implemented via a time-invariant control L_D . These two terms are obtained through first encoding the context elements $\phi_r([t_i^{\mathbb{C}}, \mathbf{x}_i^{\mathbb{C}}])$ and then applying average pooling (Zaheer et al., 2017) to produce a single representation vector $\mathbf{r} = \frac{1}{N} \sum_{i=1}^N \phi_r([t_i^{\mathbb{C}}, \mathbf{x}_i^{\mathbb{C}}])$, which is then mapped to corresponding distributions L_0 and L_D . Once we sample the latent initial condition and the trajectory dynamics representation \mathbf{u}^1 , the evolution of the latent ODE is:

$$\mathbf{l}(t) = \mathbf{f}_{\text{evolve}}(\mathbf{l}_0, t, \mathbf{f}_{nn}(\mathbf{u}, \theta_{ode})) := \mathbf{l}(t_0) + \int_{t_0}^t \mathbf{f}_{nn}(\mathbf{l}(h), \mathbf{u}, h, \theta_{ode}) dh \quad (3)$$

Here \mathbf{f}_{nn} represents the vector field as a neural network² parameterized by θ_{ode} . We compute $\mathbf{f}_{\text{evolve}}$ using numerical ODE solvers. To predict the state values at time t , after solving the latent ODE, the latent state

¹With slight abuse of notation, we follow the common notation in control to use \mathbf{u} (instead of \mathbf{d}) to represent the realization of time-invariant control term L_D .

²We note that the latent ODE \mathbf{f}_{nn} in Eq. 3 does not have to align with the form of the unknown \mathbf{f} in Eq. 2. However, in section. 7, we also show that if explicitly know the form of \mathbf{f} , this is still possible to be leveraged as \mathbf{f}_{nn} .

$\mathbf{l}(t) \in \mathbb{R}^{d_l}$ can be decoded back to real state predictions through a decoder $p(\mathbf{x}(t)|\mathbf{l}(t), t)$. Thus, the overall joint probability of the model can be represented as:

$$p(\mathbf{u}, \mathbf{l}_0, \mathbf{x}|\mathbb{C}, t) = p(\mathbf{x}|\mathbf{f}_{evolve}(\mathbf{l}_0, t, \mathbf{f}_{nn}(\mathbf{u}, \theta)), t) p(\mathbf{u}|\mathbb{C}) p(\mathbf{l}_0|\mathbb{C}) \quad (4)$$

The latent state’s dimensionality d_l is typically chosen to be larger than d_x as additional state variables benefit the model flexibility (Dupont et al., 2019) and enable learning of higher-order dynamics implicitly (Norcliffe et al., 2020).

Given a context set \mathbb{C} and a target set $\mathbb{T} := \{(t_k^{\mathbb{T}}, \mathbf{x}_k^{\mathbb{T}})\}_{k=1}^J$ ³, to calculate the log-likelihood on \mathbb{T} for inference, the intractable posterior of $p(\mathbf{l}_0|\mathbb{C}) p(\mathbf{u}|\mathbb{C})$ has been approximated through the mean field approximation $q(\mathbf{l}_0|\mathbb{T}) q(\mathbf{u}|\mathbb{T})$, which in practice has been implemented through encoder $q_L(\mathbf{l}_0|\mathbb{C})$, $q_D(\mathbf{u}|\mathbb{C})$ in an amortized fashion, eventually leading to the following evidence lower bound (ELBO):

$$\begin{aligned} & \log p(\{(x_k^{\mathbb{T}})\}_{k=1}^J | \{(t_i^{\mathbb{C}}, \mathbf{x}_i^{\mathbb{C}})\}_{i=1}^N, \mathbf{T}^{\mathbb{T}}) \\ & \approx \mathbb{E}_{q(\mathbf{l}_0, \mathbf{u}|\mathbb{T})} \left[\sum_{k=1}^J \log p(\mathbf{x}_k|\mathbf{u}, \mathbf{l}_0, t_k) \right] - \text{KL}[q_L(\mathbf{l}_0|\mathbb{T}) || q_L(\mathbf{l}_0|\mathbb{C})] - \text{KL}[q_D(\mathbf{u}|\mathbb{T}) || q_D(\mathbf{u}|\mathbb{C})] \end{aligned} \quad (5)$$

where $\mathbf{T}^{\mathbb{T}} := \{(t_k^{\mathbb{T}})\}_{k=1}^J$ represents the (irregularly sampled) target times⁴. NODEP has competitive performance in predicting single trajectory, and we therefore consider extending its functionality to meta-learn ODE system distributions for our subsequent optimization purpose.

3 System Aware Neural ODE Process (SANODEP)

Few-Shot Optimization with Prior Information

Given that evaluating $\mathbf{f}_{evolve}(\mathbf{x}_0, t, \mathbf{f})$ is limited by cost to a small number of different initial conditions. To enable fast adaptation with few evaluations, we take a meta-learning approach by assuming that \mathbf{f} is a realization of a stochastic function \mathcal{F} , and we can access data from its different realizations to formulate the *meta-task distributions*,

More precisely, consider a random function $\mathcal{F} : \mathcal{X}_0 \times \tau \rightarrow \mathbb{R}^{d_x}$ representing the distribution of dynamical systems. From a specific realization of \mathcal{F} , suppose we have observed M distinct trajectories, the context set then encompasses observations from all M trajectories, denoted as $\mathcal{T}^{\mathbb{C}} = \{\mathcal{T}_1^{\mathbb{C}}, \dots, \mathcal{T}_M^{\mathbb{C}}\}$. Each trajectory, labeled as $\mathcal{T}_l^{\mathbb{C}}$, includes its own set of N_l context observations: $\mathcal{T}_l^{\mathbb{C}} = \{(t_{li}^{\mathbb{C}}, \mathbf{x}_{li}^{\mathbb{C}})\}_{i=1}^{N_l}$. Furthermore, consider a *new* trajectory \mathcal{T}_{new} . After observing additional context observations on this trajectory, $\mathcal{T}_{new}^{\mathbb{C}}$ (e.g., the initial condition of this new trajectory), we can define the extended context set as $\mathbb{C} = \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}$.⁵ Additionally, the target set is defined as $\mathbb{T} = \mathcal{T}^{\mathbb{T}} \cup \mathcal{T}_{new}^{\mathbb{T}}$, a visual elaboration is provided in Fig. 6 in Appendix A.

It is straightforward to see why NODEP is suboptimal in such scenario as it makes predictions based on single time series ($\mathbb{C} = \mathcal{T}_{new}^{\mathbb{C}}$, $\mathbb{T} = \mathcal{T}_{new}^{\mathbb{T}}$), unable to leverage M trajectories’ information. Below, we demonstrate how SANODEP is efficiently enabled through the set-based representation (Zaheer et al., 2017).

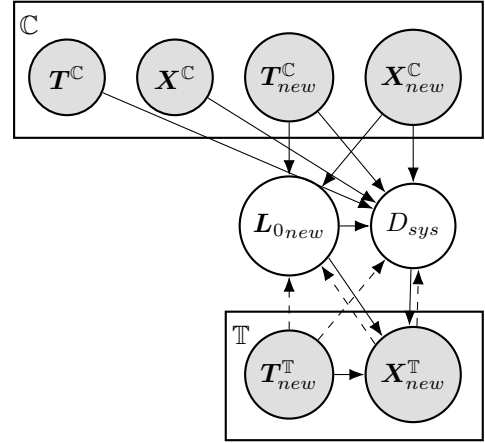


Figure 2: Graphical Model of SANODEP, the model predicts any time point in the new trajectory by knowing both observations from new trajectories and from past trajectories. Depending on whether the $X_{new}^{\mathbb{C}}$ and $T_{new}^{\mathbb{C}}$ consists of more than the initial condition, the model focuses on forecasting or interpolating tasks. The solid and dashed lines represent the generative and inference processes, respectively.

³We follow the Garnelo et al. (2018b) and Norcliffe et al. (2021) convention by assuming \mathbb{T} is a superset of \mathbb{C} .

⁴We denote context trajectory data either as a set of pairs $\{(t_i^{\mathbb{C}}, \mathbf{x}_i^{\mathbb{C}})\}_{i=1}^N$ or as vectors $\{\mathbf{T}^{\mathbb{C}}, \mathbf{X}^{\mathbb{C}}\}$, where $\mathbf{T}^{\mathbb{C}} = [t_1^{\mathbb{C}}, \dots, t_N^{\mathbb{C}}]^T$, $\mathbf{X}^{\mathbb{C}} = [\mathbf{x}_1^{\mathbb{C}}, \dots, \mathbf{x}_N^{\mathbb{C}}]^T$. We use similar notation $\{\mathbf{T}^{\mathbb{T}}, \mathbf{X}^{\mathbb{T}}\}$ for the target set.

⁵Note that again $\mathcal{T}_{new}^{\mathbb{C}} \subset \mathcal{T}_{new}^{\mathbb{T}}$.

3.1 Set-based Dynamical System Representation

Similar to the latent variable \mathbf{u} (Eq. 3) capturing trajectory dynamics in NODEP, we adapt a latent variable $D_{sys} \sim q(\mathbf{u}_{sys}|\mathbb{C})$ in SANODEP, effectively replacing \mathbf{u} in the model structure but with an enhanced conditioning on context observations from $M + 1$ trajectories, to capture the dynamical systems properties. Follow this mechanism, we will have a feature extraction for **multi-start multivariate irregular time series**, in addition with an efficiency requirement for our optimization purposes.

Such questions are less considered in contemporary time series models, as existing approaches (Shukla & Marlin, 2021; Schirmer et al., 2022) are mainly developed for single initial condition start trajectories. Specifically for multi-start scenario, Jiang et al. (2023) proposes first extracting a trajectory-wise aggregated feature vector, and then averaging feature vectors as a final representation. However, both extraction are implemented through a convolution operation, which poses challenges with irregular time series without additional modifications.

A straightforward thinking would be still using average pooling across all context elements as a set-based representation. While the Picard-Lindelöf theorem (Lindelöf, 1894) guarantees the uniqueness of state values in an initial value problem when $\mathbf{f}(\mathbf{x}, t)$ is Lipschitz continuous, in case when only part of the underline system states can be measured as $\mathbf{x}(t)$, different trajectories might still share identical $\mathbf{x}(t)$ values, leading to identifiability issues (duplicate context elements come from different initial conditions). Consequently, we propose an augmented sets-based approach: we augment each observation with its corresponding initial condition: $(t_i, \mathbf{x}_0, \mathbf{x}_i)$, to enhance the model’s ability to differentiate between trajectories that might otherwise appear identical. Then we perform the average pooling on the flattened context set to obtain the context representation $\mathbf{r}_{sys} = \frac{1}{\sum_{l=1}^{M+1} N_l} \sum_{l=1}^{M+1} \sum_{i=1}^{N_l} \phi_{sys}([t_{l_i}^{\mathbb{C}}, \mathbf{x}_{l_0}^{\mathbb{C}}, \mathbf{x}_{l_i}^{\mathbb{C}}])$, which is then mapped to a distribution representing possible system realizations from \mathcal{F} that has generated \mathbb{C} . Aside from the theoretical intuition, we also empirically compare the average pooling without initial condition augmentation in the section. 6.1, and find that the augmentation leads to more robust performance even \mathcal{F} only represents first-order systems with fully observable states.

For the rest of the model, SANODEP follows the NODEP structure, hence requiring minimal adjustment, only with additional care on activation function choices to enable differentiability (Appendix. A.3). Appendix A.2 provides a detailed description of the model structure.

3.2 Bi-scenario Loss Function

In line with the principles of episode learning (Vinyals et al., 2016), SANODEP’s training is structured through multiple episodes. Motivated by our optimization problem that will be detailed in Section. 4, the design of each episode’s problem ensures good model performance under two primary scenarios:

1. **Forecasting:** Using M context trajectories $\mathcal{T}^{\mathbb{C}}$, the model predicts future state values $\mathbf{X}_{new}^{\mathbb{T}}$ for a new trajectory initiated from $\mathcal{T}_{new}^{\mathbb{C}} = \{(t_{new_0}^{\mathbb{C}}, \mathbf{x}_{new_0}^{\mathbb{C}})\}$ at designated target times $\mathbf{T}_{new}^{\mathbb{T}}$.
2. **Interpolating:** From the same M trajectories, the model interpolates and extrapolates state values for a new trajectory that already includes $K > 1$ observations: $\mathcal{T}_{new}^{\mathbb{C}} = \{(t_{new_i}^{\mathbb{C}}, \mathbf{x}_{new_i}^{\mathbb{C}})\}_{i=0}^K$, predicting the states at times $\mathbf{T}_{new}^{\mathbb{T}}$.

Both scenarios can be considered under a bi-scenario loss function, designed to enhance model’s accuracy for predicting new trajectory states $\mathbf{X}_{new}^{\mathbb{T}}$ at $\mathbf{T}_{new}^{\mathbb{T}}$:

$$\mathcal{L}_{\theta} = \mathbb{E}_{\mathbf{f} \sim \mathcal{F}, M, \mathcal{T}^{\mathbb{C}}, \mathcal{T}_{new}^{\mathbb{C}}, \mathbb{1}_{forecast}, \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast})} \log p_{\theta}(\mathbf{X}_{new}^{\mathbb{T}} | \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}), \mathbf{T}_{new}^{\mathbb{T}}) \quad (6)$$

where $p_{\theta}(\cdot)$ represents the SANODEP prediction when parameterized by θ .

During training, once a dynamical system realization \mathbf{f} from \mathcal{F} has been drawn, we randomly sample the number of context trajectories M that we have observed, where the number of context elements, the initial condition of the trajectory and the observation time are all sampled. For the new trajectory to be predicted, besides sampling the target set $\{\mathbf{T}_{new}^{\mathbb{T}}, \mathbf{X}_{new}^{\mathbb{T}}\}$, a Bernoulli indicator $\mathbb{1}_{forecast} \sim \text{Bernoulli}(\lambda)$ parameterize

by λ is sampled to determine whether the episode will address a forecasting or interpolating scenario. This indicator directly influences the part of the context set sampled from the new trajectory⁶:

$$\mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}) = \begin{cases} \{(t_{new0}^{\mathbb{C}}, \mathbf{x}_{new0}^{\mathbb{C}})\} & \text{if } \mathbb{1}_{forecast} = 1, \\ \{(t_{newi}^{\mathbb{C}}, \mathbf{x}_{newi}^{\mathbb{C}})\}_{i=0}^K & \text{if } \mathbb{1}_{forecast} = 0. \end{cases} \quad (7)$$

For each system realization, we train in a mini-batch way by making prediction on a batch of different new trajectories.

The intractable log-likelihood in Eq. 6 can be approximated via the following evidence lower bound (ELBO):

$$\begin{aligned} & \log p_{\theta}(\mathbf{X}_{new}^{\mathbb{T}} | \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}), \mathbf{T}_{new}^{\mathbb{T}}) \\ & \approx \mathbb{E}_{q(\mathbf{u}_{sys} | \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}) \cup \mathcal{T}_{new}^{\mathbb{T}}) q(\mathbf{L}_{0new}^{\mathbb{T}} | (t_0^{\mathbb{C}}, \mathbf{x}_{0new}^{\mathbb{C}}))} \log p(\mathbf{X}_{new}^{\mathbb{T}} | \mathbf{T}_{new}^{\mathbb{T}}, \mathbf{u}_{sys}, \mathbf{L}_{0new}^{\mathbb{T}}) \\ & \quad - \text{KL}[q(\mathbf{u}_{sys} | \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}) \cup \mathcal{T}_{new}^{\mathbb{T}}) \parallel q(\mathbf{u}_{sys} | \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}))] \\ & \quad - \text{KL}[q(\mathbf{L}_{0new}^{\mathbb{T}} | (t_0^{\mathbb{C}}, \mathbf{x}_{0new}^{\mathbb{C}})) \parallel p(\mathbf{L}_{0new}^{\mathbb{T}})] \end{aligned} \quad (8)$$

where $q(\mathbf{u}_{sys} | \mathcal{T}^{\mathbb{C}} \cup \mathcal{T}_{new}^{\mathbb{C}}(\mathbb{1}_{forecast}) \cup \mathcal{T}_{new}^{\mathbb{T}})$ and $q(\mathbf{L}_{0new}^{\mathbb{T}} | (t_0^{\mathbb{C}}, \mathbf{x}_{0new}^{\mathbb{C}}))$ has been obtained through the encoder in an amortized fashion, the prior $p(\mathbf{L}_{0new}^{\mathbb{T}})$ is isotropic Gaussian. Appendix. B derives the ELBO and provides implementation details of the model inference procedure.

4 Time Delay Constraint Process Bayesian Optimization

Through maximizing the ELBO for θ , we can obtain SANODEP’s predictive distribution $p_{\theta}(\mathbf{X}^{\mathbb{T}} | \mathbb{C}, \mathbf{T}^{\mathbb{T}}, \mathbf{x}_0)$ for batch of time $\mathbf{T}^{\mathbb{T}} = \{t_1, \dots, t_N\}$ at specified initial condition \mathbf{x}_0 , which is sufficient for few-shot learning tasks. Specifically for our practically motivated termination time optimization problem that incorporates one additional time delay constraint, we further propose an optimization framework in this section, leveraging SANODEP for few-shot BO in ODE and benchmark in Section. 6.2.

Minimum Observation Delay Constraint When optimizing Eq. 2, we assume that, while one can observe state values at any chosen time t on a specific initial condition \mathbf{x}_0 , as illustrated in Fig. 1, the next observation can only commence after a fixed known time duration Δt . In practice, these delays stem from the need to sequentially conduct separate, smaller experiments for *state value* measurements, each requiring a known period to complete before the next can begin.

The optimization framework consists of the following two steps:

Initial Condition Identification identifies the optimal initial conditions necessary for starting experiments. The optimality of the initial conditions is defined as achieving the maximum expected reward after the completion of observations starting at this location. Inspired by the batch strategy to achieve a similar *non-myopic* objective (González et al., 2016; Jiang et al., 2020), we propose an **adaptive batch size** based optimization strategy for the initial condition identification:

$$\begin{aligned} & \text{maximize}_{\mathbf{x}_0 \in \mathcal{X}_0, \{t_1, t_2, \dots, t_N \in \tau\}, N \in \mathcal{N}_{opt}} \alpha(\mathbf{x}_0, t_1, \dots, t_N, p(\mathbf{X}^{\mathbb{T}} | \mathbb{C}, \mathbf{T}^{\mathbb{T}}, \mathbf{x}_0)) \\ & \text{s.t. } \forall i \in \{1, \dots, N\} : t_i - t_{i-1} \geq \Delta t \end{aligned} \quad (9)$$

where α is the batch acquisition function to be maximized, the search space of the total number of observation queries $\mathcal{N}_{opt} := \{1, 2, \dots, N_{max}\}$, $\mathbf{T}^{\mathbb{T}}$ now consists of $\{t_1, \dots, t_N\}$ and the maximum trajectory-wise observations number is $N_{max} := \lfloor (t_{max} - t_0) / \Delta t \rfloor$, we omit the dependence of $p_{\theta}(\mathbf{x} | \mathbb{C}, t, \mathbf{x}_0)$ in $\alpha(\cdot)$ for brevity thereafter.

Choose the next query time Suppose the initial condition \mathbf{x}_0 has been chosen, and our last observation query made is at time t_n , when still have query opportunity (i.e., $t_{max} > t_n + \Delta t$), redefine the maximum remaining trajectory observation N_{max} as $\lfloor \frac{t_{max} - t_n}{\Delta t} \rfloor$, the batch size search space \mathcal{N}_{opt} and time search space $\tau = [t_n + \Delta t, t_{max}]$, we choose the next query time recurrently via the following optimization problem:

⁶For notation simplicity, we always use subscript 0 to represent the initial condition of a trajectory.

$$\begin{aligned}
& \text{maximize}_{\{t_1, t_2, \dots, t_N \in \tau\}, N \in \mathcal{N}_{opt}} \alpha(\mathbf{x}_0, t_1, \dots, t_N) \\
& \text{s.t. } \forall i \in \{1, \dots, N\} : t_i - t_{i-1} \geq \Delta t
\end{aligned} \tag{10}$$

Search Space Reduction The integer variable N 's search space \mathcal{N}_{opt} , though one dimensional, can be cumbersome to optimize in practice. However, for the batch expected hypervolume improvement acquisition function (qEHVI) (Daulton et al., 2020) that we will use as α , the search space can be reduced:

Proposition 1. *For acquisition functions defined as $\alpha() = \mathbb{E}_{p(\mathbf{X}^\mathbb{T}|\mathbb{C}, \mathbf{T}^\mathbb{T}, \mathbf{x}_0)} [\text{HVI}(\mathbf{X}^\mathbb{T}, \mathbf{T}^\mathbb{T}, \mathcal{F}^*|\mathbb{C})]$:*

$$\begin{aligned}
\text{maximize}_{\substack{\mathbf{x}_0 \in \mathcal{X}_0, \\ \{t_1, t_2, \dots, t_N\} \in \tau, \\ N \in \{1, \dots, N_{max}\}}} \alpha(\mathbf{x}_0, t_1, \dots, t_N) = \text{maximize}_{\substack{\mathbf{x}_0 \in \mathcal{X}_0, \\ \{t_1, t_2, \dots, t_N\} \in \tau, \\ N \in \{\lceil N_{max}/2 \rceil, \dots, N_{max}\}}} \alpha(\mathbf{x}_0, t_1, \dots, t_N).
\end{aligned}$$

where HVI stands for *hypervolume improvement* based on Pareto frontier \mathcal{F}^* , see Definition 2 of Daulton et al. (2020). Appendix C.1 provides the proof and shows that the search space reduction property holds for generic acquisition functions that are monotonic w.r.t. set inclusions. Consequently, we define \mathcal{N}_{opt} as $[\lceil \frac{N_{max}}{2} \rceil, N_{max}]$ thereafter.

Optimization Framework is outlined in Algorithm. 2 in Appendix. C.3. We refer to the Appendix. C.1 for details of how SANODEP is utilized within qEHVI, together with acquisition function optimizer described in Appendix. C.1. Finally, we highlight that this Bayesian Optimization framework is agnostic to the model choice of $p_\theta(\mathbf{X}^\mathbb{T} | \mathbb{C}, \mathbf{T}^\mathbb{T}, \mathbf{x}_0)$, as we will compare in Section. 6, both NP and GP (as a non-meta-learn model) can be used.

5 Related Works

Meta-learning of Dynamical Systems: Our work is built upon a meta-learned continuous time-based model for state prediction in dynamical systems. Singh et al. (2019) extended Neural Processes (Garnelo et al., 2018b) by incorporating temporal information from the perspective of the state-space model. However, their approach is limited to discrete-time systems. Following the NODEP framework, Dang et al. (2023) extended the model for forecasting purposes in multi-modal scenarios. Recently, Jiang et al. (2023) used meta-learning to tackle high-dimensional time series forecasting problems for generic sequential latent variable models, they leverage spatial-temporal convolution to extract the \mathbf{u}_{sys} , which only works on regularly spaced time steps (e.g., image frames in their use cases). However, SANODEP extracts system dynamic's feature from multiple irregularly sampled multivariate time series using the newly proposed encoder structure. Song & Jeong (2023) explored Hamiltonian representations, which are flexible enough for cross-domain generalizations. Beyond the Bayesian meta-learning paradigm, Li et al. (2023) used classical gradient-based meta-learning (Finn et al., 2017) to meta-learn dynamics with shared properties. Auzina et al. (2024) investigated the separation and modeling of both dynamic variables that influence state evolution and static variables that correspond to invariant properties, enhancing model performance.

Process-Constrained Bayesian Optimization: Vellanki et al. (2017) address process-constrained batch optimization where each batch shares identical constrained variables within a subspace. Our problem similarly treats initial conditions as constrained variables in batch optimization but differs by using a joint expected utility method instead of a greedy strategy. Additionally, our approach involves on-the-fly optimization with system evolution, leading to a novel time-constrained problem addressed by an adaptive batch BO algorithm with search space reduction. Folch et al. (2022; 2024) also consider optimization under movement (initial condition) constraints.

Few-Shot Bayesian Optimization A significant amount of recent work has focused on meta-learning stochastic processes, to be used as alternatives to GP for BO. Within the (Conditional) Neural Process framework ((C)NP) (Garnelo et al., 2018a;b), Transformer Neural Processes (TNP) (Nguyen & Grover, 2022) have demonstrated robust uncertainty quantification capabilities for BO. Dutordoir et al. (2023) introduced a diffusion model-based Neural Process, providing a novel framework for modeling stochastic processes which enables joint sampling and inverse design functionality. Similar to TNP, Müller et al. (2021) approached

the meta-learning problem as Bayesian inference, subsequently integrating a BO framework into this model (Müller et al., 2023). We note that while a majority of the aforementioned approaches lead to promising empirical performance on BO for regular black-box functions, their backend models (if there are any) are mainly motivated by a drop-in replacement of GP or non-temporal stochastic process. A meta-learned BO method specifically for dynamical systems is, to our best of knowledge, not revealed yet.

6 Experiments

This section conducts experiments on meta-learning and few-shot BO for dynamical systems. All models are implemented using **Flax** (Heek et al., 2023) and are open source, available in: <https://github.com/Anonymous7>. The optimization framework is based on **Trieste** (Picheny et al., 2023).

6.1 Modeling Comparison

We evaluate model performance comparisons on meta-learning diverse ODE system distributions with a twofold purpose. First, we investigate whether leveraging the system information can help modeling. We also investigate whether the inductive bias could bring performance benefits.

Baseline: To see if adding system information helps with the modeling, we compare SANODEP against NODEP under its original testing scenario (interpolation). For the second purpose, we compare SANODEP with Neural Processes (NP) (Garnelo et al., 2018b). As mentioned in Section 5, we were not able to compare with the meta ODE models of Jiang et al. (2023) as it is not been able to be applied on irregular time series. We also provide GP results for reference. For models that do not explicitly incorporate the ODE in their structure, we use the trajectory’s initial condition augmented with time as its input.

Training Following Norcliffe et al. (2021), we treat \mathcal{F} as a parametric function of a specific *kinetic model* with stochasticity induced by model parameter distributions P . We refer to Appendix. D.1 for the model formulations, its parameter distribution P ’s form, and the training epochs. Interpolation-only experiments use SANODEP trained using Eq. 6, setting $\lambda = 0$, and NODEP trained using Eq. 5. For the second experimental purpose, we train all models using Eq. 6 setting $\lambda = 0$ for interpolation only training and $\lambda = 0.5$ for bi-scenario training, with 5 different random seeds for model initialization.

Evaluation: We evaluate the model performance on target data generated from systems sampled from the same \mathcal{F} , which is known as *in-distribution* generalization evaluation. Excluding GP, each model was evaluated on 10^4 random systems, each consisting of 100 trajectories to predict in a minibatch fashion. As it is computationally infeasible to evaluate GPs on the same scale, we used a random subset of the test set, 5,000 trajectories. Appendix. B.2 provides full data generation settings. Following Norcliffe et al. (2021), we assess model performance in terms of Mean-Squared Error by conducting a sequential evaluation with varying numbers of context trajectories, as illustrated in Fig. 3 (and Fig. 9 in Appendix D.2 for Negative-log-Likelihood). We also provide the model prediction illustration in Appendix D.3 and the averaged-context-trajectory-size performance results in Table 2 and Table 3 in Appendix D.2.

Table 1: Model comparison on **interpolation** tasks (mean-squared-error $\times 10^{-2}$) for a range of dynamical systems.

Models	Lotka-Volterra (2d)	Brusselator (2d)	Selkov (2d)	SIR (3d)	Lotka-Volterra (3d)	SIRD (4d)
NODEP- $\lambda = 0$	38.9 ± 0.52	15.5 ± 0.379	7.12 ± 0.56	367.5 ± 58.0	26.8 ± 1.53	188.1 ± 21.8
SANODEP- $\lambda = 0$	32.1 ± 0.55	8.7 ± 1.11	0.52 ± 0.03	254.3 ± 60.3	25.1 ± 1.85	137.2 ± 12.8

Does System Awareness Improve Predictive Performance? Table. 1 demonstrates that SANODEP is better than NODEP in all problems, indicating the benefit of capturing system information.

Does Incorporating Temporal Information Enhance Model Performance? Illustrated in Fig. 3, comparing with NP, except for the Lotka-Volterra problem where the SANODEP trained with mixing probability $\lambda = 0.5$ shows larger predictive variance, and Lotka-Volterra (3d) where NP demonstrates better performance when context trajectory numbers are small, for the rest of the problems, SANODEP is either on

⁷We provide the code in supplementary and promise to release our code upon acceptance.

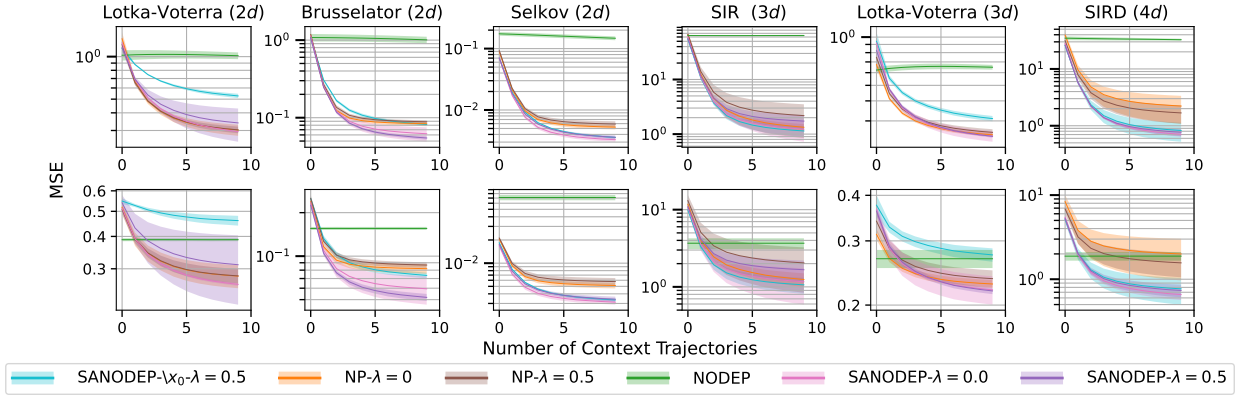


Figure 3: Model evaluation performance comparison on a different number of context trajectories. The first row corresponds to forecasting performance (prediction with only the initial condition known), and the second row represents an interpolating setting. It can be seen that SANODEP is either on-par or marginally better than NP for most problems, and the initial condition augmented encoder-based model variant provides more robust performance across problems.

par or noticeably better than NP irrespective with the mixing probability. Indicating that the incorporation of right inductive bias has the potential to help modeling.

Does Augmenting Initial Conditions Improve Modeling? We further evaluate the performance of SANODEP by comparing models with and without initial condition augmentation, as depicted in Fig. 3 (labeled as SANODEP- $\setminus x_0$ for without augmentation). The results intriguingly show that even for first order systems, models augmented with initial conditions exhibit greater robustness across various systems. This demonstrates the empirical benefits of our context aggregation approach.

How Do Meta-Learned Models Compare with Non-Meta-Learned Models? We provide the sequential model evaluation plot in Fig. 8 including GPs in Appendix D.2. As expected, the meta-learned models show superior accuracy when the number of trajectories is small (as we have also illustrated in Fig. 1), aligning with the problem setting of this study. We also remark that GP can quickly overtake when the context trajectory number is abundant, since the Neural Process typed model is known to suffer from *under fit* when the data is abundant.

6.2 Few-Shot Bayesian Optimization




Figure 4: Comparison of few-shot BO performances in terms of mean \pm standard deviation. Except for the Lotka-Voterra (3d) problem, SANODEP-based few-shot BO demonstrates competitive performance on all the rest of the problems.

Baselines: We conduct few-shot BO using the meta-learned models (SANODEP, NODEP, NP) trained in the previous section using our optimization framework (Algorithm 2). For reference, we provide GP-based BO (GP-BO) using the same framework.

Optimization: We optimize dynamical system realizations sampled from \mathcal{F} . The parameter settings, the objective function definitions and performance indicator calculation are provided in Table 4. For simplicity, in each optimization problem, we assume the maximum number of observations per trajectory is $\Delta t = \frac{t_{max}-t_0}{10}$. We start with one randomly sampled trajectory with uniformly spaced observations and have an additional budget to query 10 trajectories. Each optimization is repeated 10 times with different random seeds.

Results are reported in terms of the scaled experimental time versus the *hypervolume indicator*, with the reference point also provided in Table 4 (we also report the results of PI-SANODEP that will be introduced in next section.).

Does Meta-learning Helps Few-Shot Optimization in Dynamical Systems We highlight the merits of few-shot BO, as it consistently demonstrates significant convergence speed improvements during the early stages of BO.

Does Dynamical System Informed Model Behaves Better than Standard Model for Optimization Except for Lotka-Volterra (3d), SANODEP demonstrate better performance compared with Neural Processes. Again validate the potential of a deliberate consideration of temporal information.

Does Weight Training Objective help Optimization: For SANODEP, except in the case of the Brusselator, combining training with bi-scenario losses shows either comparable or slightly better performance than training solely in the interpolation setting. This suggests that careful design of the loss function for optimization purposes may offer modest benefits.

Finally, we provide the time profile of different models in the Appendix. D.5, in general, as the involvement of simulation in inference process, SANODEP and NODEP takes more time than NP and GP when perform optimization.

7 On the strength of the prior information, \mathcal{F}

Previously, we assumed that the prior information, the kinetic model, was known to formulate the task distribution. In practice, the level of prior information can vary across problems. In order to better understand how the different levels of prior information can impact the model performance in meta-learning of dynamical systems. We conduct some preliminary investigation on both strengthened and weakened priors in this section.

7.1 Strong Prior information: Physics-Informed SANODEP

When \mathbf{f} is a realization of \mathcal{F} and we know the parametric form of the kinetic model, we can benefit from explicitly incorporating the dynamical system within the model, instead of relying on a latent ODE model structure that is agnostic to the dynamical system, we can easily integrate this kinetic form in SANODEP’s model structure. We call this model variant Physical Informed (PI)-SANODEP.

As a concrete example, consider \mathbf{f} as a realization of the Lotka-Volterra (LV) problem with unknown parameters. Under the SANODEP framework, we can replace \mathbf{f}_{nn} in Eq. 3 with the following physics-informed form:

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} \alpha_\theta x_1 - \beta_\theta x_1 x_2 \\ \delta_\theta x_1 x_2 - \gamma_\theta x_2 \end{bmatrix} \quad (11)$$

In this form, $\mathbf{u}_{sys} = [\alpha_\theta, \beta_\theta, \delta_\theta, \gamma_\theta]$, and $\mathbf{l}(t_0) = \mathbf{x}_0$ and $\mathbf{l}(t) = \mathbf{x}(t)$ ⁸. Since the parameters of the LV system are strictly positive. We use log-normal distributions for the conditional prior and variational posterior for \mathbf{u}_{sys} in the encoder. By additionally incorporating a likelihood regularization term in Eq. 8 (Appendix. E.1), as we have spotlighted in Fig. 5 and exhaustively demonstrated in Fig. 15, we found PI-SANODEP can significantly enhance the model performance.

Meta-learn for Concurrent Prediction and Parameter Estimation: Benefit from the amortized variational inference through the encoder, as we have additionally outlined in Fig. 5, we note that PI-

⁸In PI-SANODEP, since the ODE is explicitly defined as the kinetic form (e.g., Eq. 11) the decoder is only used to predict variance.

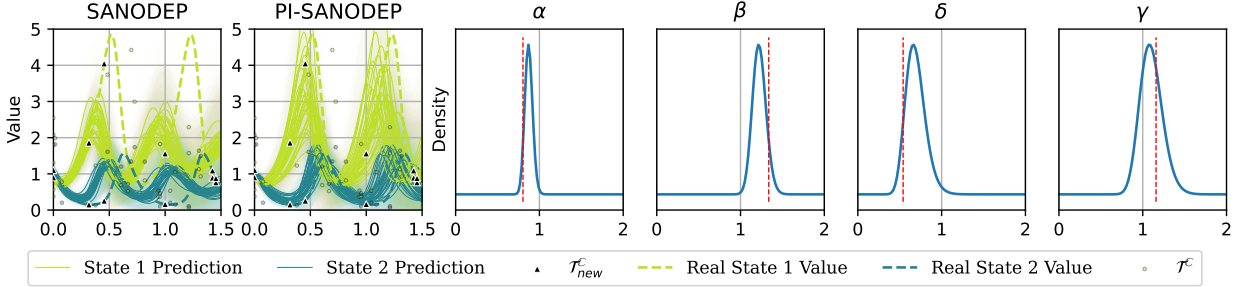


Figure 5: Model Comparison of SANODEP vs Physics Informed (PI) SANODEP, PI-SANODEP provides remarkable better prediction accuracy. In addition, PI-SANODEP provides a reasonable estimation (blue curve) on system parameters (dashed red line).

SANODEP also demonstrates the potential of concurrent few-shot prediction together with parameter estimation of dynamical systems. Compared with Bayesian inference-based parameter estimation strategies (e.g. Wu & Lysy (2024)), which may take minutes to perform estimation, PI-SANODEP shares the same advantage as *simulation-based inference* (Cranmer et al., 2020) and can estimate almost instantly since it only requires a forward pass of the encoder. We report the parameter estimation on the rest of the problems in Fig. 15 in Appendix. E.1 and leave an extensive investigation of meta-learning-based parameter estimation as future work.

7.2 Weaker Prior of \mathcal{F} : Cross-domain generalization possibility

What if we know nothing about \mathcal{F} ? From meta-learning perspective, to still frame the problem as in-distribution evaluation, this demands an extremely flexible prior for \mathcal{F} to be utilizable in the target dynamical system. To our best of knowledge, meta-learning dynamical systems to be leveraged in cross-domain scenarios is still an open direction. Certain approaches have resorted to incorporating additional dynamical system properties (e.g., energy conserving system Song & Jeong (2023)) which limit the scope of system types. We propose and leverage a vector valued GP prior for \mathcal{F} as a flexible task distribution. We show in Appendix. E.2 that it is able generate diverse dynamical systems. We also demonstrate in Fig. 16,17 that SANODEP can also be trained on this task distribution albeit at detriment to model fitting capability.

8 Discussion

Bayesian Optimization (BO) of unknown dynamical systems is an underexplored area, with existing approaches often relying on unsuitable surrogate models. Addressing this gap, while minimising the number of costly optimization steps, we have developed a novel few-shot BO framework. Our approach extends the Neural ODE Processes (NODEP) into the System-Aware Neural ODE Processes (SANODEP), which meta-learns the system’s prior information to enhance the optimization process. Through extensive benchmark experiments, we demonstrated SANODEP’s potential over NODEP. Our results show that SANODEP, equipped with an optimization-driven loss function, offers competitive performance compared to other non-dynamic-aware meta-learning models, and can offer additional possibilities including parameter estimation, which can’t be said with other non-dynamic-aware meta-learning models.

Limitations: A key limitation of SANODEP, inherited from Neural Processes, is the issue of underfitting. This necessitates a trade-off between few-shot functionality and model fitting capability. Additionally, while strong prior information enhances performance in scenarios with well-defined priors, the model performance is limited by a lack of prior information about the dynamical system. Although our proposed task distribution shows promise for cross-domain generalization, SANODEP still faces a trade-off between model fitting and prior flexibility, indicating that it cannot be used out of the box as a generic probabilistic model to optimize arbitrary dynamical system without knowing its kinetic form.

Future investigations: Future work will aim to develop more adaptable conditional distributions to better

capture a wider range of dynamical systems. This will address the trade-off between model performance and the flexibility needed to handle diverse and complex dynamical systems without relying heavily on strong priors.

References

- Ilze Amanda Auzina, Çağatay Yıldız, Sara Magliacane, Matthias Bethge, and Efstratios Gavves. Modulated neural ODEs. *NeurIPS*, 36, 2024.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *NeurIPS*, volume 31, 2018.
- Yuqi Chen, Kan Ren, Yansen Wang, Yuchen Fang, Weiwei Sun, and Dongsheng Li. ContiFormer: Continuous-time transformer for irregular time series modeling. In *NeurIPS*, volume 36, 2024.
- Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 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.
- Ting Dang, Jing Han, Tong Xia, Erika Bondareva, Chloë Siegle-Brown, Jagmohan Chauhan, Andreas Grammenos, Dimitris Spathis, Pietro Cicuta, and Cecilia Mascolo. Conditional neural ODE processes for individual disease progression forecasting: A case study on COVID-19. In *29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3914–3925, 2023.
- Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. In *NeurIPS*, volume 33, pp. 9851–9864, 2020.
- Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural ODEs. In *NeurIPS*, volume 32, 2019.
- Vincent Dutordoir, Alan Saul, Zoubin Ghahramani, and Fergus Simpson. Neural diffusion processes. In *ICML*, pp. 8990–9012, 2023.
- Katharina Ensinger, Nicholas Tagliapietra, Sebastian Ziesche, and Sebastian Trimpe. Exact inference for continuous-time Gaussian process dynamics. In *AAAI*, volume 38, pp. 11883–11891, 2024.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pp. 1126–1135, 2017.
- Jose Pablo Folch, Shiqiang Zhang, Robert Lee, Behrang Shafei, David Walz, Calvin Tsay, Mark van der Wilk, and Ruth Misener. SnAKE: Bayesian optimization with pathwise exploration. In *NeurIPS*, volume 35, pp. 35226–35239, 2022.
- Jose Pablo Folch, Calvin Tsay, Robert M Lee, Behrang Shafei, Weronika Ormaniec, Andreas Krause, Mark van der Wilk, Ruth Misener, and Mojmír Mutný. Transition constrained Bayesian optimization via Markov decision processes. In *NeurIPS*, volume 37, 2024.
- Andrew Foong, Wessel Bruinsma, Jonathan Gordon, Yann Dubois, James Requeima, and Richard Turner. Meta-learning stationary stochastic process prediction with convolutional neural processes. In *NeurIPS*, volume 33, pp. 8284–8295, 2020.
- Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *ICML*, pp. 1704–1713, 2018a.

- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018b.
- Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.
- Javier González, Michael Osborne, and Neil Lawrence. GLASSES: Relieving the myopia of Bayesian optimisation. In *AISTATS*, pp. 790–799, 2016.
- Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, Andreas Steiner, and Marc van Zee. Flax: A neural network library and ecosystem for JAX, 2023. URL <http://github.com/google/flax>.
- Pashupati Hegde, Çağatay Yıldız, Harri Lähdesmäki, Samuel Kaski, and Markus Heinonen. Variational multiple shooting for Bayesian ODEs with Gaussian processes. In *UAI*, pp. 790–799, 2022.
- Markus Heinonen, Cagatay Yildiz, Henrik Mannerström, Jukka Intosalmi, and Harri Lähdesmäki. Learning unknown ODE models with Gaussian processes. In *ICML*, pp. 1959–1968, 2018.
- Shali Jiang, Henry Chai, Javier Gonzalez, and Roman Garnett. Binoculars for efficient, nonmyopic sequential experimental design. In *ICML*, pp. 4794–4803, 2020.
- Xiajun Jiang, Ryan Missel, Zhiyuan Li, and Linwei Wang. Sequential latent variable models for few-shot high-dimensional time-series forecasting. In *ICLR*, volume 11, 2023.
- Patrick Kidger. On neural differential equations. *arXiv preprint arXiv:2202.02435*, 2022.
- Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *ICLR*, 2019.
- Qiaofeng Li, Tianyi Wang, Vwani Roychowdhury, and Mohammad Khalid Jawed. Metalearning generalizable dynamics from trajectories. *Physical Review Letters*, 131(6):067301, 2023.
- Ernest 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 (3):454–457, 1894.
- Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Transformers can do Bayesian inference. In *ICLR*, 2021.
- Samuel Müller, Matthias Feurer, Noah Hollmann, and Frank Hutter. PFNs4BO: In-context learning for Bayesian optimization. In *ICML*, pp. 25444–25470, 2023.
- Tung Nguyen and Aditya Grover. Transformer neural processes: Uncertainty-aware meta learning via sequence modeling. In *ICML*, pp. 16569–16594, 2022.
- Alexander Norcliffe, Cristian Bodnar, Ben Day, Nikola Simidjievski, and Pietro Liò. On second order behaviour in augmented neural ODEs. In *NeurIPS*, volume 33, pp. 5911–5921, 2020.
- Alexander Norcliffe, Cristian Bodnar, Ben Day, Jacob Moss, and Pietro Liò. Neural ODE processes. In *ICLR*, 2021.
- Derek Onken and Lars Ruthotto. Discretize-optimize vs. optimize-discretize for time-series regression and continuous normalizing flows. *arXiv preprint arXiv:2005.13420*, 2020.
- Katharina Ott, Michael Tiemann, and Philipp Hennig. Uncertainty and structure in neural ordinary differential equations. *arXiv preprint arXiv:2305.13290*, 2023.
- Joel A Paulson and Calvin Tsay. Bayesian optimization as a flexible and efficient design framework for sustainable process systems. *arXiv preprint arXiv:2401.16373*, 2024.

- Victor Picheny, Joel Berkeley, Henry B Moss, Hrvoje Stojic, Uri Granta, Sebastian W Ober, Artem Artemev, Khurram Ghani, Alexander Goodall, Andrei Paleyes, et al. Trieste: Efficiently exploring the depths of black-box functions with TensorFlow. *arXiv preprint arXiv:2302.08436*, 2023.
- Thomas Pinder and Daniel Dodd. GPJax: A Gaussian process framework in JAX. *Journal of Open Source Software*, 7(75):4455, 2022. doi: 10.21105/joss.04455. URL <https://doi.org/10.21105/joss.04455>.
- Ilya Prigogine and René Lefever. Symmetry breaking instabilities in dissipative systems. ii. *The Journal of Chemical Physics*, 48(4):1695–1700, 1968.
- Jixiang Qing, Tom Dhaene, and Ivo Couckuyt. Spectral representation of robustness measures for optimization under input uncertainty. In *ICML*, volume 38, 2022.
- Jixiang Qing, Henry B Moss, Tom Dhaene, and Ivo Couckuyt. $\{PF\}^2$ es: Parallel feasible pareto frontier entropy search for multi-objective Bayesian optimization. In *AISTATS*, 2023.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *NIPS*, 20, 2007.
- Oliver Schilter, Daniel Pacheco Gutierrez, Linnea M Folkmann, Alessandro Castrogiovanni, Alberto García-Durán, Federico Zipoli, Loïc M Roch, and Teodoro Laino. Combining Bayesian optimization and automation to simultaneously optimize reaction conditions and routes. *Chemical Science*, 2024.
- Mona Schirmer, Mazin Eltayeb, Stefan Lessmann, and Maja Rudolph. Modeling irregular time series with continuous recurrent units. In *ICML*, pp. 19388–19405, 2022.
- Alexandre Schoepfer, Jan Weinreich, Ruben Laplaza, Jerome Waser, and Clemence Corminboeuf. Cost-informed Bayesian reaction optimization. 2024. doi: 10.26434/chemrxiv-2024-44ft2.
- Satya Narayan Shukla and Benjamin M Marlin. Multi-time attention networks for irregularly sampled time series. In *ICLR*, 2021.
- Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential neural processes. In *NeurIPS*, volume 32, 2019.
- Yeongwoo Song and Hawoong Jeong. Towards cross domain generalization of Hamiltonian representation via meta learning. In *ICLR*, volume 12, 2023.
- Connor J Taylor, Alexander Pomberger, Kobi C Felton, Rachel Grainger, Magda Barecka, Thomas W Chamberlain, Richard A Bourne, Christopher N Johnson, and Alexei A Lapkin. A brief introduction to chemical reaction optimization. *Chemical Reviews*, 123(6):3089–3126, 2023.
- Alexander Thebelt, Johannes Wiebe, Jan Kronqvist, Calvin Tsay, and Ruth Misener. Maximizing information from chemical engineering data sets: Applications to machine learning. *Chemical Engineering Science*, 252: 117469, 2022.
- Pratibha Vellanki, Santu Rana, Sunil Gupta, David Rubin, Alessandra Sutti, Thomas Dorin, Murray Height, Paul Sanders, and Svetha Venkatesh. Process-constrained batch Bayesian optimisation. *NeurIPS*, 30, 2017.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, volume 29, 2016.
- Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Mohan Wu and Martin Lysy. Data-adaptive probabilistic likelihood approximation for ordinary differential equations. In *AISTATS*, pp. 1018–1026, 2024.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In *NeurIPS*, volume 30, 2017.

Appendix

Table of Contents

A Problem Setting and Model Structures	15
A.1 Problem Setting Illustrations	15
A.2 Model Structure	15
A.3 On the Differentiability of the Encoder and Decoder	17
B ELBO Derivation and Training Details	18
B.1 SANODEP ELBO derivation	18
B.2 Learning and Prediction Process of SANOEP	18
C Time Delay Constraint Optimization	19
C.1 Acquisition Functions	19
C.2 Search Space Reduction	20
C.3 Optimization Framework	20
D Experimental Details and Additional Results	21
D.1 Meta Training Data Definition	21
D.2 Additional Experimental Results	22
D.3 Model Prediction Visual Comparison	22
D.4 Optimization Problem Definition	22
D.5 Inference Time	23
E Investigation of Prior Strength on Model Performance	26
E.1 Parameter estimation via PI-SANODEP	26
E.2 Vector Valued GP Prior as Task Distribution	27
F Notations	29

A Problem Setting and Model Structures

A.1 Problem Setting Illustrations

A.2 Model Structure

The model structure and hyperparameters of SANODEP are summarized as follows:

- **Initial State Encoder:** $\mathbf{r}_{init} = \phi_{init}([t_0, \mathbf{x}_0])$, $\phi_{init} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^r$ consists of three dense layers and activation functions in between.
- **Augmented State Encoder:** $\mathbf{r}_i = \phi_r([t_i, \mathbf{x}_0, \mathbf{x}_i])$, $\phi_r := \mathbb{R}^{2d+1} \rightarrow \mathbb{R}^r$ consists of three dense layers and activation functions in between.
- **System Context Aggregation:** $\mathbf{r}_{sys} = \frac{1}{N} \sum_{i=1}^N \mathbf{r}_i$.
- **Context to Hidden Representation :** $\mathbf{h}_{sys} = \phi_{sys}(\mathbf{r}_{sys})$, $\mathbf{h}_{init} = \phi_{init}(\mathbf{r}_{init})$, $\phi_{sys}(\cdot)$ and $\phi_{init}(\cdot)$ are one dense layer followed by a context encoder activation function, $\mathbf{h}_{sys} \in \mathbb{R}^h$, $\mathbf{h}_{init} \in \mathbb{R}^h$.




Figure 6: Illustration of the initial condition optimization problem within one state variable ($d = 1$) ODE system. The optimization involves a forecasting scenario: having observed the context set $\mathbb{C} := (\cup_{i=1}^M \mathcal{T}_i^{\mathbb{C}}) \cup \mathcal{T}_{new}^{\mathbb{C}}$ (including the *new* trajectory’s initial condition, illustrated as \bullet), one tries to predict the target state values $\mathbf{X}_{new}^{\mathbb{T}} = [\mathbf{x}_{new1}^{\mathbb{T}}, \dots, \mathbf{x}_{newN}^{\mathbb{T}}]$ at target times $\mathbf{T}_{new}^{\mathbb{T}} = [t_{new1}^{\mathbb{T}}, \dots, t_{newN}^{\mathbb{T}}]$ (illustrated as \star).

- **Hidden to Variational posterior of \mathbf{u}_{sys} :** $q(\mathbf{u}_{sys}|\cdot) = \mathcal{N}(\phi_{\mu_{sys}}(\mathbf{h}_{sys}), \text{diag}(\phi_{\sigma_{sys}}(\mathbf{h}_{sys})^2))$, $\phi_{\mu_{sys}}: \mathbb{R}^h \rightarrow \mathbb{R}^{d_{sys}}$ is a dense layer, $\phi_{\sigma_{sys}}: \mathbb{R}^h \rightarrow \mathbb{R}^{d_{sys}}$ is defined as $\sigma_{lb} + 0.9 * \text{softplus}(\text{Dense}(\cdot))$, σ_{lb} is a hyperparameter.
- **Hidden to Variational posterior of \mathbf{L}_0 :** $q(\mathbf{L}_{0new}^{\mathbb{T}}|\cdot) = \mathcal{N}(\phi_{\mu_{init}}(\mathbf{h}_{init}), \phi_{\sigma_{init}}(\mathbf{h}_{init})^2)$. $\phi_{\mu_{init}}: \mathbb{R}^h \rightarrow \mathbb{R}^l$ and $\phi_{\sigma_{init}}: \mathbb{R}^h \rightarrow \mathbb{R}^l$ has the same structure as above, but with independent weights.
- **ODE block:** $\phi_{ODE}[\mathbf{l}, \mathbf{u}_{sys}, t] \rightarrow \mathbf{l}'$, $\phi_{ODE}: \mathbb{R}^{d_{sys}+l+1} \rightarrow \mathbb{R}^l$, ϕ_{ODE} is consisted of three dense layers with nonlinear activation functions in between.
- **Decoder:** $\phi_{dec}[\mathbf{l}(t), \mathbf{u}_{sys}, t] \rightarrow \mathbf{x}(t)$, $\mathbf{l} = \mathcal{N}(\phi_{\mu_{dec}}[\mathbf{l}(t), \mathbf{u}_{sys}, t], \phi_{\sigma_{dec}}^2[\mathbf{l}(t), \mathbf{u}_{sys}, t])$, where $\phi_{\mu_{dec}}$ and $\phi_{\sigma_{dec}}$ has the same structure as variational posterior, but with independent weights.

Model Hyperparameters

- Encoder ϕ_r output dimension r : 50
- Context encoder $(\phi_r, \phi_{sys}, \phi_{init})$ activation function: SiLU
- Encoder hidden dimension h : 50
- ODE layer (ϕ_{ODE}) activation function: Tanh
- ODE layer hidden dimension: 50
- Decoder hidden dimension: 50
- Latent ODE state dimension l : 10
- Context to latent dynamics $(\phi_{\mu_{sys}}, \phi_{\sigma_{sys}})$ activation function: SiLU
- Context to latent initial $(\phi_{\mu_{init}}, \phi_{\sigma_{init}})$ condition activation function: SiLU
- Latent dynamics dimension d_{sys} : 45
- Variational posterior $q(\mathbf{L}_{0new}^{\mathbb{T}}|\cdot)$ variance lower bound: $\sigma_{lb} = 0.1$
- Variational posterior $q(\mathbf{u}_{sys}|\cdot)$ variance lower bound: $\sigma_{lb} = 0.1$
- Decoder $(\phi_{dec}, \phi_{\mu_{dec}}, \phi_{\sigma_{dec}})$ activation function: SiLU

- ODE solver: Dopri5 with `rtol` = $1e-5$ and `atol` = $1e-5$

We use SiLU instead of ReLU as in the original NODEP is to enforce the differentiability, which we elaborate on in the following subsection.

A.3 On the Differentiability of the Encoder and Decoder

Since the optimization takes the initial condition together with the time as decision variables, we need the model output to be differentiable w.r.t. these quantities. NODEP overlooked this part and utilized ReLU activation functions originally, we elaborate that the differentiability requirement practically affects our choice of activation functions.

Time derivative Without loss of generality, we assume we optimize a function g that takes the output of the decoder $\phi_{\mu_{dec}}[\mathbf{l}(t), \mathbf{u}_{sys}, t]$ (defined in Appendix. A.2) as its input ⁹. The derivative of g w.r.t t is:

$$\frac{dg(\phi_{\mu_{dec}}[\mathbf{l}(t), \mathbf{u}_{sys}, t], t)}{dt} = \nabla_{\phi_{\mu_{dec}}} g \left(\frac{\partial \phi_{\mu_{dec}}}{\partial \mathbf{l}(t)} \cdot \phi_{ODEs}[\mathbf{l}(t), \mathbf{u}_{sys}, t] + \frac{\partial \phi_{\mu_{dec}}}{\partial t} \right) + \frac{\partial g}{\partial t} \quad (12)$$

where $\phi_{ODEs} = \cdot$ represents the neural network structure modeling the latent ODEs. It is clear that the differentiability of model output w.r.t time t needs to be enforced through $\frac{\partial \phi_{\mu_{dec}}}{\partial t}$, as long as the time t is taken into account in the decoder part.

Initial Condition Derivative The derivative of the model output w.r.t the initial condition \mathbf{x}_0 can be represented as:

$$\begin{aligned} & \frac{dg(\phi_{\mu_{dec}}[\mathbf{l}(t), \mathbf{u}_{sys}, t], t)}{d\mathbf{x}(t_0)} \\ &= \nabla_{\phi_{\mu_{dec}}} g \left(\frac{\partial \phi_{\mu_{dec}}}{\partial \mathbf{l}(t)} \cdot \frac{\partial \mathbf{l}(t)}{\partial \mathbf{x}(t_0)} + \frac{\partial \phi_{\mu_{dec}}}{\partial \mathbf{u}_{sys}} \cdot \frac{\partial \mathbf{u}_{sys}}{\partial \mathbf{x}(t_0)} \right) \\ &= \nabla_{\phi_{\mu_{dec}}} g \left[\frac{\partial \phi_{\mu_{dec}}}{\partial \mathbf{l}(t)} \cdot \left(\frac{\partial \mathbf{l}(t_0)}{\partial \mathbf{x}(t_0)} + \cdot \right) + \frac{\partial \phi_{\mu_{dec}}}{\partial \mathbf{u}_{sys}} \cdot \frac{\partial \left(\phi_{\mu_{sys}}(\mathbf{h}_{sys}) + \text{diag}(\phi_{\sigma_{sys}}(\mathbf{h}_{sys})\epsilon) \right)}{\partial \mathbf{x}(t_0)} \right. \\ & \quad \left. \dots \cdot \frac{\partial \mathbf{h}_{sys}}{\partial \mathbf{r}_{sys}} \cdot \frac{\partial \mathbf{r}_{sys}}{\partial \mathbf{x}_0} \right] \end{aligned} \quad (13)$$

where \mathbf{r}_{sys} and \mathbf{h}_{sys} are defined in Appendix. A.2, $\phi_{\mu_{sys}}(\mathbf{h}_{sys}) + \text{diag}(\phi_{\sigma_{sys}}(\mathbf{h}_{sys})\epsilon)$ represent the reparameterization based sampling of the latent variable \mathbf{u}_{sys} , where ϵ is a random sample from Gaussian prior. Note that the derivative of the state value w.r.t to initial condition $\frac{\partial \mathbf{l}(t)}{\partial \mathbf{x}(t_0)}$ can be obtained through the adjoint method, a derivation has been provided in Appendix C of Chen et al. (2018), without a full specification of such cumbersome terms, it is sufficient to indicate the sensibility of having differentiable activation functions in context encoding $\left(\frac{\partial \mathbf{r}_{sys}}{\partial \mathbf{x}_0} \right)$, context to latent representations $\left(\frac{\partial \mathbf{h}_{sys}}{\partial \mathbf{r}_{sys}} \right)$, as well as the decoder $\left(\frac{\partial \phi_{\mu_{dec}}}{\partial \mathbf{l}(t)} \right)$. For a practical calculation of the gradient of the model w.r.t decision variables in both training and optimization, we simply leverage the standard backpropagation through the ODE solver (also known as the "discretize-and-optimize" (e.g. Onken & Ruthotto (2020); Kidger (2022))) to obtain the gradient due to its speed advantage and minimum implementation effort.

In Figure. 7, we also empirically show the necessity of these differentiable choices by inspecting the time and initial condition derivative with respect to $g(\cdot)$.

⁹we note that this is sufficient to show that the requirements of time differentiability irrespective to the consideration of $\phi_{\sigma_{dec}}[\mathbf{l}(t), \mathbf{u}_{sys}, t]$




Figure 7: Initial condition gradient and time derivative w.r.t different activation function choices, the comparison is conducted with the same model structure but different activation functions.

B ELBO Derivation and Training Details

B.1 SANODEP ELBO derivation

Note that since the forecasting scenario is a special case of interpolating scenario, below we derive through a unified representation:

$$\begin{aligned}
& \log p(\mathbf{X}_{new}^T | \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^T) \\
&= \log \int p(\mathbf{X}_{new}^T | \mathbf{T}_{new}^T, \mathbf{u}_{sys}, \mathbf{L}_{0_{new}}^T) p(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C) p(\mathbf{L}_{0_{new}}^T) d\{\mathbf{u}_{sys}, \mathbf{L}_{0_{new}}^T\} \\
&= \left(\log \int p(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C) p(\mathbf{L}_{0_{new}}^T) \frac{q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C \cup \mathcal{T}_{new}^T)}{q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C \cup \mathcal{T}_{new}^T)} \right. \\
&\quad \left. \frac{q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C))}{q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C))} \cdot p(\mathbf{X}_{new}^T | \mathbf{T}_{new}^T, \mathbf{u}_{sys}, \mathbf{L}_{0_{new}}^T) d\{\mathbf{u}_{sys}, \mathbf{L}_{0_{new}}^T\} \right) \\
&\geq \mathbb{E}_{q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C \cup \mathcal{T}_{new}^T) q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C))} \log \left(\frac{p(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C)}{q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C \cup \mathcal{T}_{new}^T)} \right. \\
&\quad \left. \cdot \frac{p(\mathbf{L}_{0_{new}}^T)}{q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C))} \cdot p(\mathbf{X}_{new}^T | \mathbf{T}_{new}^T, \mathbf{u}_{sys}, \mathbf{L}_{0_{new}}^T) \right) \\
&\approx \mathbb{E}_{q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C \cup \mathcal{T}_{new}^T) q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C))} \log p(\mathbf{X}_{new}^T | \mathbf{T}_{new}^T, \mathbf{u}_{sys}, \mathbf{L}_{0_{new}}^T) \\
&\quad - \text{KL}[q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C \cup \mathcal{T}_{new}^T) || q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C)] - \text{KL}[q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C)) || p(\mathbf{L}_{0_{new}}^T)]
\end{aligned} \tag{14}$$

We note the variational posterior $q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C)$, $q(\mathbf{L}_{0_{new}}^T | (t_0^C, \mathbf{x}_{0_{new}}^C))$ has been obtained through the encoder in an amortized approach similar in Garnelo et al. (2018b).

B.2 Learning and Prediction Process of SANOEP

The practical learning process of SANODEP is summarized in Algorithm. 1, with parameters listed in section. F as well. In all of our subsequent experiments, we use $M_{min} = 0$, $M_{max} = 10$, $N_{x_0} = 100$, $N_{sys} = 20$, $N_{grid} = 100$, $m_{min} = 1$, $m_{max} = 10$, $n_{min} = 0$, $n_{max} = 45$. We note that all the loops have been practically implemented in batch, and all the subsampling approach has been implemented practically through `masking` in Jax.

Algorithm 1 Learning and Inference in System Aware Neural ODE Processes (SANODEP)

Require: ODE system inducing distributions P , known trajectory range $[M_{min}, M_{max}]$, batch size of Monte Carlo approximation of initial condition sample N_{x_0} , batch size of dynamical systems sample N_{sys} , prespecified time grid $\mathbf{T}_{grid} := \text{linspace}(t_0, t_{max}, N_{grid})$, minimum and maximum context points within a trajectory m_{min}, m_{max} , minimum and maximum extra target points n_{min}, n_{max} , initial condition space \mathcal{X}_0 .

- 1: Initialize SANODEP model parameters θ with random seeds.
- 2: **for** step in training_steps **do**
- 3: # Training data generation
- 4: **for** $j = 1$ to N_{sys} **do**
- 5: Sample ODE system $\mathbf{f} \sim \mathcal{F}$ and N_{x_0} different initial conditions from \mathcal{X}_0 , solve N_{x_0} odes at time grids \mathbf{T}_{grid} to obtain the dataset $\{\mathcal{T}_1, \dots, \mathcal{T}_{N_{x_0}}\}$.
- 6: Sample known trajectory number M from $\text{Uniform}(M_{min}, M_{max})$.
- 7: **for** $l = 1$ to M **do**
- 8: # Sample context and target elements within each trajectory
- 9: Sample m_l from $\text{Uniform}(m_{min}, m_{max})$, sample n_l from $\text{Uniform}(n_{min}, n_{max})$.
- 10: Randomly subsample from \mathcal{T}_l to extract the context dataset \mathcal{T}_l^C and target dataset \mathcal{T}_l^T , where $|\mathcal{T}_l^C| = m_l$ and $|\mathcal{T}_l^T| = m_l + n_l$ and $\mathcal{T}_l^C \subseteq \mathcal{T}_l^T$.
- 11: **end for**
- 12: Concatenate M trajectories context set $\mathcal{T}^C = \{\mathcal{T}_1^C, \dots, \mathcal{T}_M^C\}$ and target set $\mathcal{T}^T = \{\mathcal{T}_1^T, \dots, \mathcal{T}_M^T\}$.
- 13: **end for**
- 14: # Model Prediction
- 15: **for** $j = 1$ to N_{sys} **do**
- 16: **for** $k = 1$ to N_{x_0} **do**
- 17: sample $\mathbb{I}_{forecast}$ from $\text{Bernoulli}(\lambda)$ # determine forecasting or interpolating
- 18: **if** $\mathbb{I}_{forecast} = 1$ **then**
- 19: $\mathcal{T}_{new}^C = \{(t_{0k}, \mathbf{x}_{0k})\}$ where $(t_{0k}, \mathbf{x}_{0k}) \in \mathcal{T}_k$
- 20: **else if** $\mathbb{I}_{forecast} = 0$ and $k > M$ **then**
- 21: Randomly subsample from \mathcal{T}_k to extract the context dataset \mathcal{T}_k^C and target dataset \mathcal{T}_k^T , where $|\mathcal{T}_k^C| = m$, $|\mathcal{T}_k^T| = m + n$ and $\mathcal{T}_k^C \subseteq \mathcal{T}_k^T$.
- 22: **else**
- 23: $\mathcal{T}_{new}^C = \mathcal{T}_k^C$, $\mathcal{T}^C = \mathcal{T}^C \cup \mathcal{T}_k^C$.
- 24: **end if**
- 25: Augmented the context trajectories and target trajectories with the new trajectory $\mathcal{T}^C = \mathcal{T}^C \cup \mathcal{T}_{new}^C$, $\mathcal{T}^T = \mathcal{T}^T \cup \mathcal{T}_{new}^T$.
- 26: Compute variational posterior $q(\mathbf{u}_{sys}|\mathcal{T}^T)$, $q(\mathbf{L}_{0_{new}}^T|\mathcal{T}^T)$, $q(\mathbf{u}_{sys}|\mathcal{T}^C)$, $q(\mathbf{L}_{0_{new}}^T|\mathcal{T}^C)$ through the encoder block.
- 27: Sample $\mathbf{l}(t_0), \mathbf{u}_{sys}$ from $q(\mathbf{l}(t_0)|\mathcal{T}^T)$, $q(\mathbf{u}_{sys}|\mathcal{T}^T)$.
- 28: Solve latent odes as in Eq. 3 for all times and decode to obtain the model prediction.
- 29: Calculate the trajectory wise loss \mathcal{L}_{ELBO_k} based on Eq. 8.
- 30: **end for**
- 31: **end for**
- 32: Average the trajectory wise loss $\mathcal{L}_{ELBO} = \frac{1}{N_{x_0}} \sum_{o=1}^{N_{x_0}} \mathcal{L}_{ELBO_o}$
- 33: Update the model parameter through optimizer $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}_{ELBO}$
- 34: **end for**

C Time Delay Constraint Optimization

C.1 Acquisition Functions

We advocate that the framework is agnostic to the acquisition function forms as long as it can be written as expected utility form and satisfy the property mentioned in Lemma. C.1, we however mention the acquisition function form we specifically utilized for performing few-shot BO in Section 6.2 as:

$$\begin{aligned}
& \alpha(\mathbf{x}_0, t_1, \dots, t_N, p_\theta(\mathbf{X}^\mathbb{T} | \mathbb{C}, \mathbf{T}^\mathbb{T}, \mathbf{x}_0)) \\
& \approx \frac{1}{N_{MC}} \sum_{i=1}^{N_{MC}} \text{HVI}(\phi_{\mu_{dec}}[l_i(\mathbf{T}^\mathbb{T}), \mathbf{u}_{sys_i}, \mathbf{T}^\mathbb{T}], \mathbf{T}^\mathbb{T}, \mathcal{F}^*)
\end{aligned} \tag{15}$$

In practice, we infer the current Pareto frontier \mathcal{F}^* from SANODEP as well from the M context trajectories $\mathcal{T}^\mathbb{C}$. We use $N_{MC} = 32$ as the MC sample size to approximate the expected hypervolume improvement in all our experiments.

Acquisition Function Optimization We note the acquisition optimizer of both initial condition (Eq. 9), as well as optimal time scheduling (Eq. 10) involves constraint optimization. We utilize trust region-based constraint optimization available in `Scipy`¹⁰ to optimize acquisition function in both optimization problems starting with a uniform time scheduling (i.e., `linspace`(t_0, t_1, N)). Note that for the initial stage optimization, we additionally leverage a (10 instances) multi-start procedure on the initial condition decision variable only to boost the optimization performance.

C.2 Search Space Reduction

Proof of the search space reduction We prove that the search space can be reduced, without eliminating the global maximum in acquisition function optimization processes (Eq. 9).

Lemma C.1. *For batch acquisition function α defined as expected utility $\alpha = \mathbb{E}_{p(\mathbf{x})}(u(\mathbf{x}))$ where the utility function $u()$ is monotonic w.r.t set inclusion ($S \subseteq T \Rightarrow u(T) \geq u(S)$), the acquisition function is also monotonic w.r.t set inclusion: $\alpha(T) \geq \alpha(S)$.*

Proof. Let $T = S \cup T/S$, then $\alpha(T) - \alpha(S) = \mathbb{E}_{p(S)}\mathbb{E}_{p(T/S|S)}(u(T)) - \mathbb{E}_{p(S)}\mathbb{E}_{p(T/S|S)}(u(S)) = \mathbb{E}_{p(S)}\mathbb{E}_{p(T/S|S)}(u(T) - u(S)) \geq 0$, hence prove complete. \square

Corollary C.1.1. *Assume that the utility function $u(\cdot)$ is monotonic with respect to the inclusion of the set in each set separately: Let $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$ and $\mathbf{Y} = \{Y_1, Y_2, \dots, Y_n\}$ be two collections of sets, if $X_i \subseteq Y_i \forall i \in \{1, \dots, n\}$ then $u(\mathbf{X}) \leq u(\mathbf{Y})$. Then, the batch acquisition function defined as $\alpha = \mathbb{E}_{p(\mathbf{X})}(u(\mathbf{X}))$ satisfies $\alpha(\mathbf{Y}) \geq \alpha(\mathbf{X})$.*

Lemma C.2. *qEHVI is a monotonic acquisition function with respect to set inclusion*

Proof. Using Corollary. C.1.1, Lemma C.2 and the well-acknowledged fact that the hypervolume indicator is monotonic with respect to set inclusion on each output dimensionality, the proof is straightforward. We note that the above property is asymptotically hold when qEHVI is monte carlo approximated. \square

Finally, we start the proof of Theorem. 1:

Proof. when $N \in [1, \lceil \frac{N_{max}}{2} \rceil]$, $\exists i \in [1, \lceil \frac{N_{max}}{2} \rceil]$ s.t. $t_i - t_{i-1} \geq 2\Delta t$ meaning that one can insert a new batch point in between, forming an augmented set T as input, with Lemma C.2 and Corollary. C.1.1, the proof hence complete. \square

We finally remark that, with Lemma. C.1, a large series of batch improvement-based acquisition functions (e.g., parallel Expected Improvement, parallel Probability of Improvement, and parallel information-theoretic acquisition functions that can be written as expected utility form and also preserve the set inclusion monotonicity property (e.g., Qing et al. (2023))) also embrace the same search space reduction benefit.

C.3 Optimization Framework

¹⁰<https://docs.scipy.org/doc/scipy/reference/optimize.minimize-trustconstr.html>

Algorithm 2 Model Assisted Ordinary Differential Equation Optimization Framework

```

1: Input: maximum number of experiment trajectories  $N_{exp}$ , minimum spacing  $\Delta t$ , maximum observations
   per trajectory  $N_{max}$ , initial evaluated context datasets  $\mathbb{C}$ , acquisition function  $\alpha$ , evaluated trajectory
   number  $n = 1$ , initial time lower bound  $t_{lb} = t_0$ , time range  $\tau = [t_{lb}, t_{max}]$ , observer  $\mathbf{f}_{evolve}$  for
   measurement.
2: while  $n \leq N_{exp}$  do
3:    $\mathcal{T} = \phi$ 
4:   Construct the probabilistic model for  $p(\mathbf{x}|t, \mathbb{C}, \mathbf{x}_0)$ 
5:   # Initial Condition Optimization
6:    $\mathbf{x}_0^*, N^* = \arg \max_{\mathbf{x}_0 \in \mathcal{X}, \{t_1=t_0, t_2, \dots, t_N \in \tau\}, N \in [\lceil \frac{N_{max}}{2} \rceil, N_{max}]} \alpha(\mathbf{x}_0, t_1, \dots, t_N)$ 
       s. t.  $\forall i \in \{1, \dots, N\} : t_i - t_{i-1} \geq \Delta t$ 
        $t_{lb} = t_0 + \Delta t, \tau := [t_{lb}, t_{max}], \mathcal{T} = \mathcal{T} \cup \{t_0, \mathbf{x}_0^*\}, \mathbb{C} = \mathbb{C} \cup \mathcal{T}$ 
7:    $N_{traj} = N_{max} - 1$ 
8:   # Within trajectory Next Measurement Time Scheduling
9:   while  $N_{traj} >= 1$  do
10:     $N^*, t_1^* = \arg \max_{\{t_1, t_2, \dots, t_N \in \tau\}, N \in [\lceil \frac{N_{traj}}{2} \rceil, N_{traj}]} \alpha(t_1, \dots, t_N)$ 
        s. t.  $\forall i \in \{1, \dots, N\} : t_i - t_{i-1} \geq \Delta t$ 
11:     $\mathbf{x}_1^* = \mathbf{f}_{evolve}(\mathbf{x}_0^*, t_1^*)$ 
12:     $t_{lb} = t_1^* + \Delta t, \tau := [t_{lb}, t_{ub}], \mathbb{C} = (\mathbb{C} \setminus \mathcal{T}) \cup (\mathcal{T} \cup \{t_1, \mathbf{x}_1^*\}), \mathcal{T} = \mathcal{T} \cup \{t_1, \mathbf{x}_1^*\}$ 
13:     $N_{traj} = \lfloor \frac{t_{max} - t_{lb}}{\Delta t} \rfloor$ 
14:  end while
15:   $n = n + 1$ 
16: end while

```

D Experimental Details and Additional Results**D.1 Meta Training Data Definition****Lotka-Voterra****2D Cases**

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} \alpha x_1 - \beta x_1 x_2 \\ \delta x_1 x_2 - \gamma x_2 \end{bmatrix} \quad (16)$$

The initial condition is sample from $\mathbf{x}(t_0) \sim U(0.1, 3)^2$, the dynamics \mathcal{F} are sampled from $\alpha \sim U(\frac{1}{3}, 1)$, $\beta \sim U(1, 2)$, $\delta \sim U(0.5, 1.5)$, $\gamma \sim U(0.5, 1.5)$, $t_{max} = 15$. We run the meta-learn on this system distribution with 300 epochs.

3D cases

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} \alpha x_1 - \beta x_1 x_2 - \epsilon x_1 x_3 \\ \delta x_2 x_1 - \gamma x_2 - \zeta x_2 x_3 \\ \eta x_3 x_2 - \theta x_3 \end{bmatrix} \quad (17)$$

The initial condition is sampled from $\mathbf{x}(t_0) \sim U(0.1, 3)^3$. The dynamics \mathcal{F} are sampled from $\alpha \sim U(\frac{1}{3}, 1)$, $\beta \sim U(1, 2)$, $\delta \sim U(0.5, 1.5)$, $\gamma \sim U(0.5, 1.5)$, $\epsilon \sim U(0.5, 1.5)$, $\zeta \sim U(0.5, 1.5)$, $\eta \sim U(0.5, 1.5)$, and $\theta \sim U(0.5, 1.5)$, with $t_{max} = 15$. We run the meta-learning algorithm on this system distribution for 300 epochs.

Brusselator Prigogine & Lefever (1968)

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} A + x_1^2 x_2 - (B + 1)x_1 \\ Bx_1 - x_1^2 x_2 \end{bmatrix} \quad (18)$$

we leverage $A \sim U(0, 1)$, $B \sim U(0.1, 3)$, $\mathbf{x}_0 \sim U(0.1, 2.0)^2$, $t_{max} = 15$. We run the meta-learning algorithm on this system distribution for 300 epochs.

SIR model

$$\frac{dS}{dt} = -\beta SI, \quad \frac{dI}{dt} = \beta SI - \gamma I, \quad \frac{dR}{dt} = \gamma I \quad (19)$$

The initial condition is sampled from $S(t_0) \sim U(1.0, 3.0)$, $I(t_0) = 0.01$, $R(t_0) = 0$, and $\beta \sim U(0.1, 2)$, $\gamma \sim U(0.1, 10)$. $t_{max} = 1$. We run the meta-learning algorithm on this system distribution for 300 epochs.

SIRD model

$$\frac{dS}{dt} = -\beta SI, \quad \frac{dI}{dt} = \beta SI - \gamma I - \mu I, \quad \frac{dR}{dt} = \gamma I, \quad \frac{dD}{dt} = \mu I \quad (20)$$

where $S(t_0) \sim U(10, 30)$, $I(t_0) = 0.01$, $R(t_0) = D(t_0) = 0$. $t_{max} = 1$, $\beta \in [0.5, 2.0]$, $\gamma \in [0.1, 10]$, and $\mu \in [0.1, 5.0]$. We run the meta-learn on this system distribution with 300 epochs.

GP Vector Field

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}) \quad (21)$$

We assume a vector-valued GP as the non-parametric prior for the vector field $\mathbf{f} \sim \mathcal{GP}(\mathbf{0}, K(\mathbf{x}, \mathbf{x}'))$, with no correlations between each output as the most generic cases. To sample the vector field, we leverage the parametric approximation of GPs through the random Fourier feature approximation of the RBF kernel Rahimi & Recht (2007) as a Bayesian linear model, which is a common approach (e.g. Qing et al. (2022)) to obtain differentiable GP samples, the vector fields in practice are generated from kernel lengthscale 0.8 and signal variance 1. The implementation also utilizes the parametric sampling approach of the GPJax (Pinder & Dodd, 2022) library.

D.2 Additional Experimental Results

We report all model’s performance on interpolating tasks and forecasting tasks, averaged over known trajectory ranges M , in Table. 2 and Table. 3.

Table 2: Model comparison (mean-squared-error $\times 10^{-2}$) for a range of dynamical systems.

Models	Lotka-Volterra (2d)	Brusselator (2d)	Selkov (2d)	SIR (3d)	Lotka-Volterra (3d)	SIRD (4d)
Forecasting						
GP	87.2 \pm 31.9	39.3 \pm 61.0	31.8 \pm 6.72	1769.93 \pm 980.6	41.4 \pm 15.8	1322.845 \pm 489.4
NODEP- $\lambda = 0$	102.9 \pm 6.85	105.2 \pm 10.8	16.1 \pm 1.14	6367.784 \pm 161.8	64.8 \pm 2.56	3368.685 \pm 189.5
NP- $\lambda = 0.0$	48.6 \pm 3.1	21.6 \pm 0.56	1.6 \pm 0.0916	952.3 \pm 73.6	32.8 \pm 0.85	724.2 \pm 168.8
NP- $\lambda = 0.5$	47.3 \pm 1.34	22.4 \pm 1.47	1.7 \pm 0.0882	1057.981 \pm 207.2	34.8 \pm 1.53	573.0 \pm 88.3
SANODEP- $\lambda = 0$	48.4 \pm 2.09	20.1 \pm 0.816	1.22 \pm 0.0376	861.6 \pm 89.0	36.2 \pm 1.17	426.4 \pm 44.0
SANODEP- $\lambda = 0.5$	50.6 \pm 6.4	18.8 \pm 0.467	1.28 \pm 0.0459	873.1 \pm 82.4	35.5 \pm 0.824	410.5 \pm 27.3
Interpolating						
GP	52.2 \pm 20.2	9.62 \pm 19.0	11.7 \pm 2.76	173.7 \pm 296.8	16.8 \pm 13.7	280.7 \pm 137.3
NODEP- $\lambda = 0$	38.9 \pm 0.517	15.6 \pm 0.379	7.12 \pm 0.559	367.5 \pm 58.0	26.8 \pm 1.53	188.1 \pm 21.8
NP- $\lambda = 0.0$	33.2 \pm 2.59	10.4 \pm 0.523	0.742 \pm 0.0436	286.2 \pm 29.7	24.7 \pm 0.557	302.0 \pm 110.9
NP- $\lambda = 0.5$	33.2 \pm 1.48	11.0 \pm 0.477	0.807 \pm 0.0638	378.3 \pm 143.0	25.9 \pm 1.39	244.6 \pm 54.3
SANODEP- $\lambda = 0$	32.1 \pm 0.556	8.74 \pm 1.09	0.53 \pm 0.03	257.5 \pm 60.3	25.2 \pm 1.84	139.5 \pm 13.5
SANODEP- $\lambda = 0.5$	36.6 \pm 7.73	8.06 \pm 0.265	0.573 \pm 0.0247	307.8 \pm 53.5	25.2 \pm 0.464	143.8 \pm 14.0

D.3 Model Prediction Visual Comparison

We provide the prediction comparison of different models on a specific realization of ODE systems in Fig. 10, 13, 11, 12, 14.

D.4 Optimization Problem Definition

The definitions of the optimization problem are provided in Table. 4.

Table 3: Model comparison (negative log-likelihood $\times 10^2$) for a range of dynamical systems.

Models	Lotka-Volterra (2d)	Brusselator (2d)	Selkov (2d)	SIR (3d)	Lotka-Volterra (3d)	SIRD (4d)
Forecasting						
GP	18.2 ± 7.71	30.2 ± 25.7	17.1 ± 4.95	11143.352 ± 31169.575	19.3 ± 5.75	27321.033 ± 57173.003
NODEP- $\lambda = 0$	119.3 ± 21.7	79.1 ± 13.9	136.4 ± 32.1	1086.616 ± 331.3	625.5 ± 174.2	447.7 ± 86.0
NP- $\lambda = 0.0$	98.9 ± 52.2	17.7 ± 4.41	4.28 ± 1.34	162.6 ± 78.6	491.2 ± 122.1	197.5 ± 151.7
NP- $\lambda = 0.5$	200.9 ± 88.0	37.7 ± 18.5	5.93 ± 0.836	279.7 ± 198.2	961.1 ± 454.1	75.7 ± 42.8
SANODEP- $\lambda = 0$	15.5 ± 6.9	22.6 ± 2.98	7.98 ± 1.39	206.9 ± 79.4	277.5 ± 77.9	48.7 ± 17.4
SANODEP- $\lambda = 0.5$	27.0 ± 13.0	23.7 ± 4.55	8.55 ± 1.42	219.4 ± 139.7	329.0 ± 103.8	48.3 ± 11.0
Interpolating						
GP	3.27 ± 1.24	1.88 ± 2.0	5.27 ± 1.87	718.0 ± 2426.002	2.69 ± 1.61	1045.453 ± 3397.646
NODEP- $\lambda = 0$	32.6 ± 6.73	12.7 ± 2.89	13.9 ± 2.94	66.4 ± 18.2	67.1 ± 22.9	22.9 ± 8.44
NP- $\lambda = 0.0$	18.7 ± 7.46	0.254 ± 0.61	-2.29 ± 0.215	22.4 ± 10.1	65.8 ± 11.2	30.1 ± 23.1
NP- $\lambda = 0.5$	34.8 ± 14.0	2.88 ± 2.01	-1.97 ± 0.147	40.1 ± 28.4	122.1 ± 49.6	12.5 ± 6.12
SANODEP- $\lambda = 0$	2.88 ± 1.59	1.24 ± 0.568	-1.74 ± 0.308	28.1 ± 10.5	41.3 ± 17.7	7.17 ± 2.39
SANODEP- $\lambda = 0.5$	6.08 ± 3.63	1.13 ± 0.596	-1.5 ± 0.284	31.3 ± 19.3	47.3 ± 14.6	7.4 ± 1.97




Figure 8: Sequential model evaluation performance including GPs. The first row corresponds to forecasting performance (prediction with only the initial condition known), and the second row represents an interpolating setting.

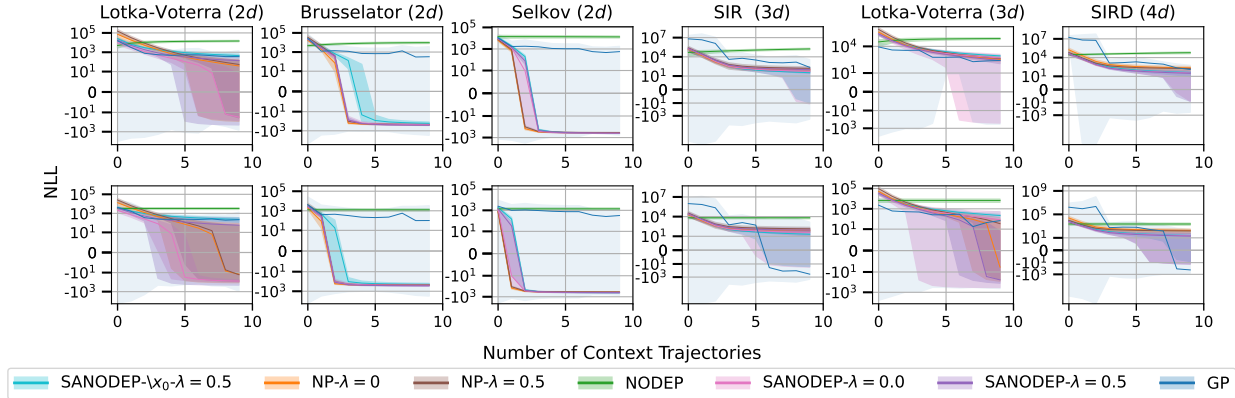


Figure 9: Sequential model evaluation performance in terms of Negative-Log-Likelihood (NLL) including GPs. The first row corresponds to forecasting performance (prediction with only the initial condition known), and the second row represents an interpolating setting.

D.5 Inference Time

We provide the parameter counts, training time and prediction time comparison of SANODEP vs. baselines in Table. 5, we note that all the models have the same latent dimensionality for a fair performance comparison,

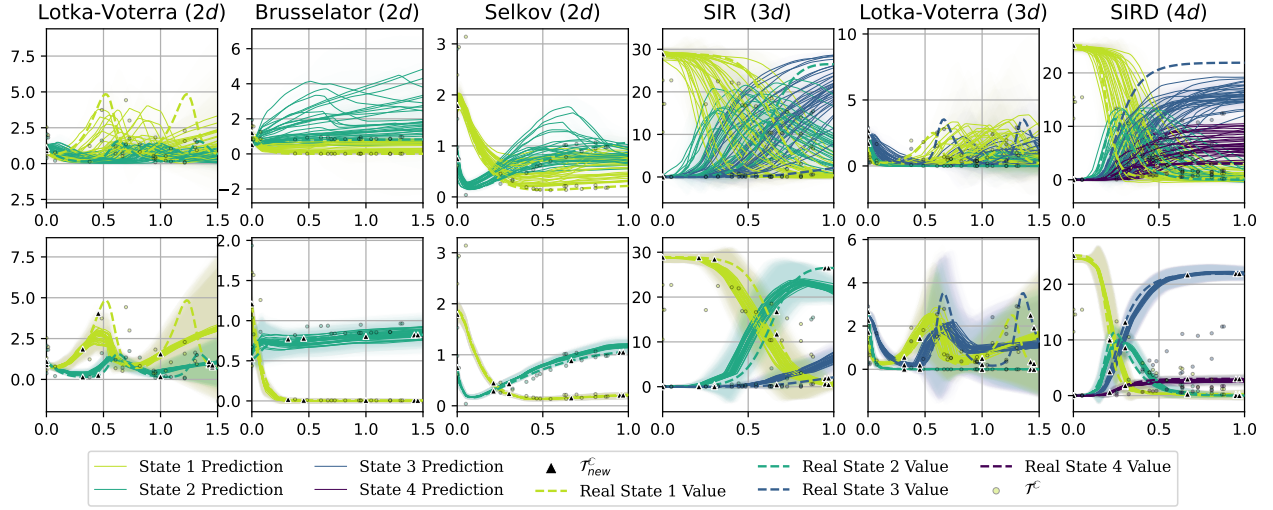


Figure 10: Neural Processes ($\lambda = 0.5$) model performance on test system in different meta-learning ODE problems.

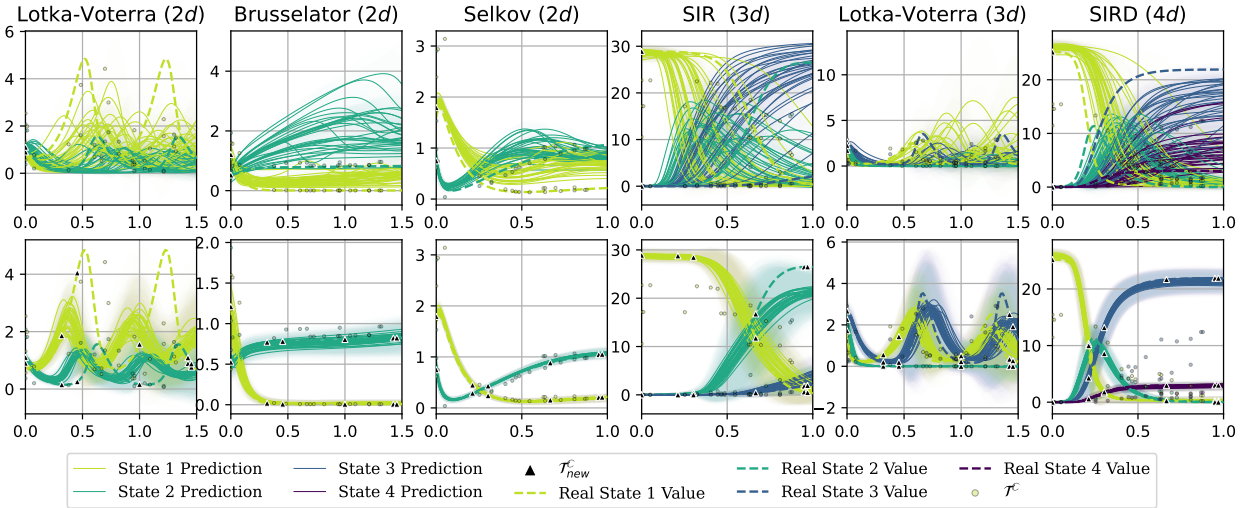


Figure 11: SANODEP- ($\lambda = 0.5$) model performance on test system in different meta-learning ODE problems.

Table 4: Optimization problem formulations

Problem	Lotka-Volterra (2d)	Brusselator (2d)	Selkov (2d)	SIR model (3d)	Lotka-Volterra (3d)	SIRD (4d)
g formulation	x_1	x_1	x_2	$x_1 / \sum_{i=1}^3 x_i - 0.05 \sum_{i=1}^3 x_i$	x_1	$x_1 / \sum_{i=1}^3 x_i - 0.05 \sum_{i=1}^3 x_i$
Design space	$\tau \in [0, 15]$, $\mathbf{x}_{dec} \in [0.1, 2.0]^2$ $\mathbf{x}_0 = \mathbf{x}_{dec}$	$\tau \in [0, 15]$, $\mathbf{x}_{dec} \in [0.1, 2]^2$ $\mathbf{x}_0 = \mathbf{x}_{dec}$	$\tau \in [0, 10]$, $\mathbf{x}_{dec} \in [0.1, 0.5]^2$ $\mathbf{x}_0 = \mathbf{x}_{dec}$	$\tau \in [0, 1]$, $x_{dec} \in [10, 30]$ $\mathbf{x}_0 = [x_{dec}, 0.1 \times x_{dec}, 0]$	$\tau \in [0, 15]$, $\mathbf{x}_{dec} \in [0.0, 2]^3$ $\mathbf{x}_0 = \mathbf{x}_{dec}$	$\tau \in [0, 1]$, $x_{dec} \in [10, 30]$ $\mathbf{x}_0 = [x_{dec}, 0.1 \times x_{dec}, 0, 0]$
Δt	1.5	1.5	1	0.1	1.5	0.1
Optimization ODE definition	$\alpha = 0.5$, $\beta = 1.2$, $\delta = 1.0$, $\gamma = 1.5$	$A = 0.8$, $B = 1.5$	$a = 0.25$, $b = 0.45$	$\beta = 1.5$, $\gamma = 5$	$\alpha = 0.5$, $\beta = 1.2$, $\delta = 1.0$, $\gamma = 1.5$, $\epsilon = 0.5$, $\zeta = 1.2$, $\eta = 1.0$, $\theta = 1.5$	$\beta = 1$, $\gamma = 0.5$, $\mu = 1$
Reference point	$[-1.771, 12.686]$	$[-1.467, 3.887]$	$[-0.474, 5.440]$	$[0.51151, 0.79646]$	$[-1.7557, 13.1687]$	$[0.52198, 1.04]$

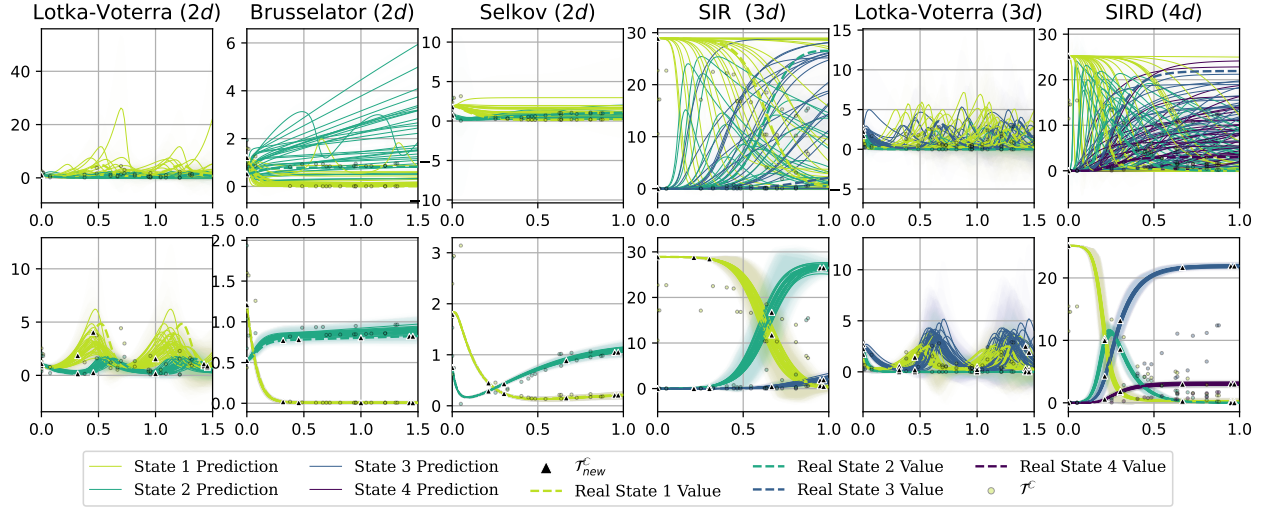


Figure 12: PI-SANODEP- $(\lambda = 0.5)$ model performance on test system in different meta-learning ODE problems.

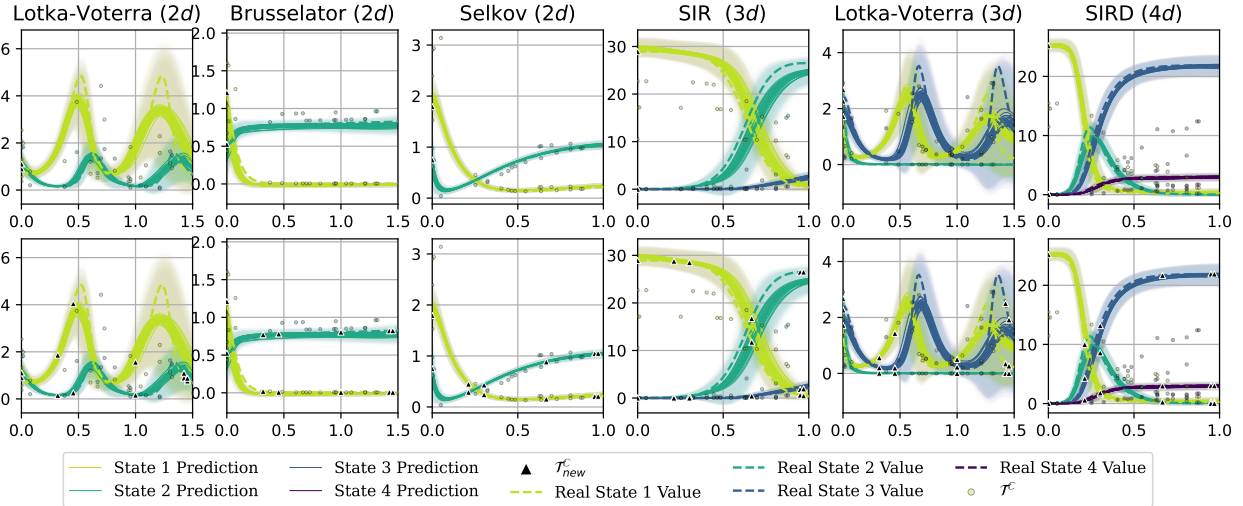


Figure 13: NODEP $(\lambda = 0)$ model performance on the test system in different meta-learning ODE problems.

which is regarded as the key bottleneck of the latent variable typed meta-learning model (Kim et al., 2019). The parameter increase of SANODEP compared to NODEP is the additional encoder ϕ_r to handle the augmented state variable.

Table 5: Comparison of the number of parameters and prediction time of different models. We measure run time on the Lotka-Voterra ($d = 2$) problem using an NVIDIA A40 GPU, We note training time and prediction time is per mini-batch, the optimization query time is per querying one observations.

Model	Number of Parameters	Training Time (s)	Prediction Time (s)	Optimization Query Time (s)
GP	NA	NA	NA	23.81 ± 43.94
NP	21814	0.3016 ± 0.0791	0.2551 ± 0.2337	18.04 ± 18.36
NODEP	30286	0.3792 ± 0.1892	0.3421 ± 0.2412	34.71 ± 20.62
SANODEP	35514	0.3813 ± 0.2423	0.3482 ± 0.2731	39.47 ± 37.30

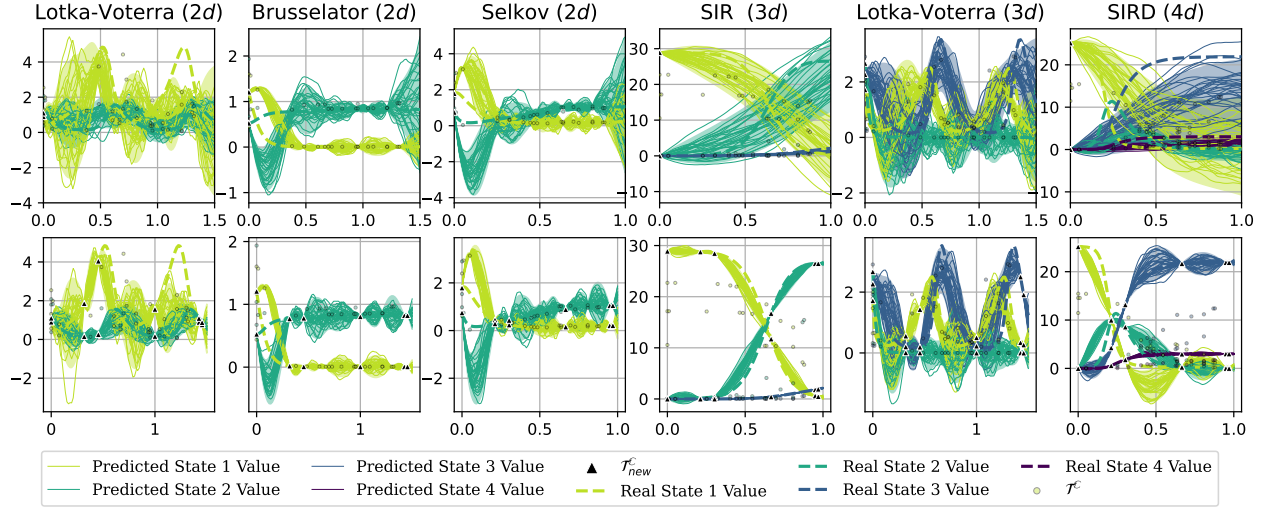


Figure 14: Performance of the GP model on the test system in different meta-learning ODE problems.

E Investigation of Prior Strength on Model Performance

E.1 Parameter estimation via PI-SANODEP

One advantage of SANODEP as a meta-learned dynamical model is its ability to estimate parameters. When the parametric form of the dynamical system is known, SANODEP can directly integrate this form into its governing ODE (Eq. 3), yielding a variant we call Physics-Informed SANODEP (PI-SANODEP).

For PI-SANODEP’s training loss function \mathcal{L}'_θ , we denote the system parameters for each meta-training task as \mathbf{u}_f , since these parameters are known during meta-training, we augment our loss (Eq. 8) by incorporating a likelihood term for these parameters, leading to the following expression:

$$\begin{aligned} \log p(\mathbf{X}_{new}^\top, \mathbf{u}_f | \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^\top) \\ = \log p(\mathbf{X}_{new}^\top | \mathbf{u}_f, \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^\top) + \log p(\mathbf{u}_f | \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^\top) \end{aligned} \quad (22)$$

Since \mathbf{u}_f is known during meta-training, we use the encoder (described in Appendix. A.2) as an inference network to calculate its likelihood $p(\mathbf{u}_f | \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^\top) = \mathcal{N}(\mathbf{u}_f; \phi_{\mu_{sys}}(\mathbf{h}_{sys}), \text{diag}(\phi_{\sigma_{sys}}(\mathbf{h}_{sys})^2))$. Regarding $p(\mathbf{X}_{new}^\top | \mathbf{u}_f, \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^\top)$, we however omit the dependence of \mathbf{u}_f and still make use of Eq. 8 to calculate. Thus, the final loss function for PI-SANODEP extends the original SANODEP loss with an added likelihood term $\log p(\mathbf{u}_f | \mathcal{T}^C \cup \mathcal{T}_{new}^C, \mathbf{T}_{new}^\top)$, capturing both the dynamics and parameter estimation. For the model structure of PI-SANODEP, since $\mathbf{l}(t) = \mathbf{x}(t)$, hence $\phi_{\mu_{dec}}[\mathbf{x}(t), \mathbf{u}_{sys}, t] = \mathbf{x}(t)$. the rest of the model structure is the same as SANODEP only except that the latent ODE \mathbf{f}_{nn} becomes explicit.

We note that, while PI-SANODEP can improve itself toward the right estimation of parameters along the training process, as a non-local parameter estimation strategy, we empirically noticed that at the initial training stage of some dynamical systems, the inferred parameter distribution can be far from the correct ones and hence either trigger an exhaustive discretization step or have divergent trajectories. To mitigate such *non-stable initialization* which causes numerical issues, given that we know the (interval) support \mathcal{S} of our training system parameters P , we clip the sample parameter from amortized variation posterior $q(\mathbf{u}_{sys} | \mathcal{T}^C \cup \mathcal{T}_{new}^C(\mathbb{I}_{forecast}) \cup \mathcal{T}_{new}^\top)$ by the support bound (provided in the Appendix.D.1), which significantly improves the stability of PI-SANODEP training.

We report the PI-SANODEP model performance in comparison with SANODEP, together with its parameter estimation on the rest of the problem in Fig. 15. For consistency, we use log-normal for all parameters in different problems. We conclude by remarking that the parameter estimation functionality of PI-SANODEP is standalone besides the optimization application.

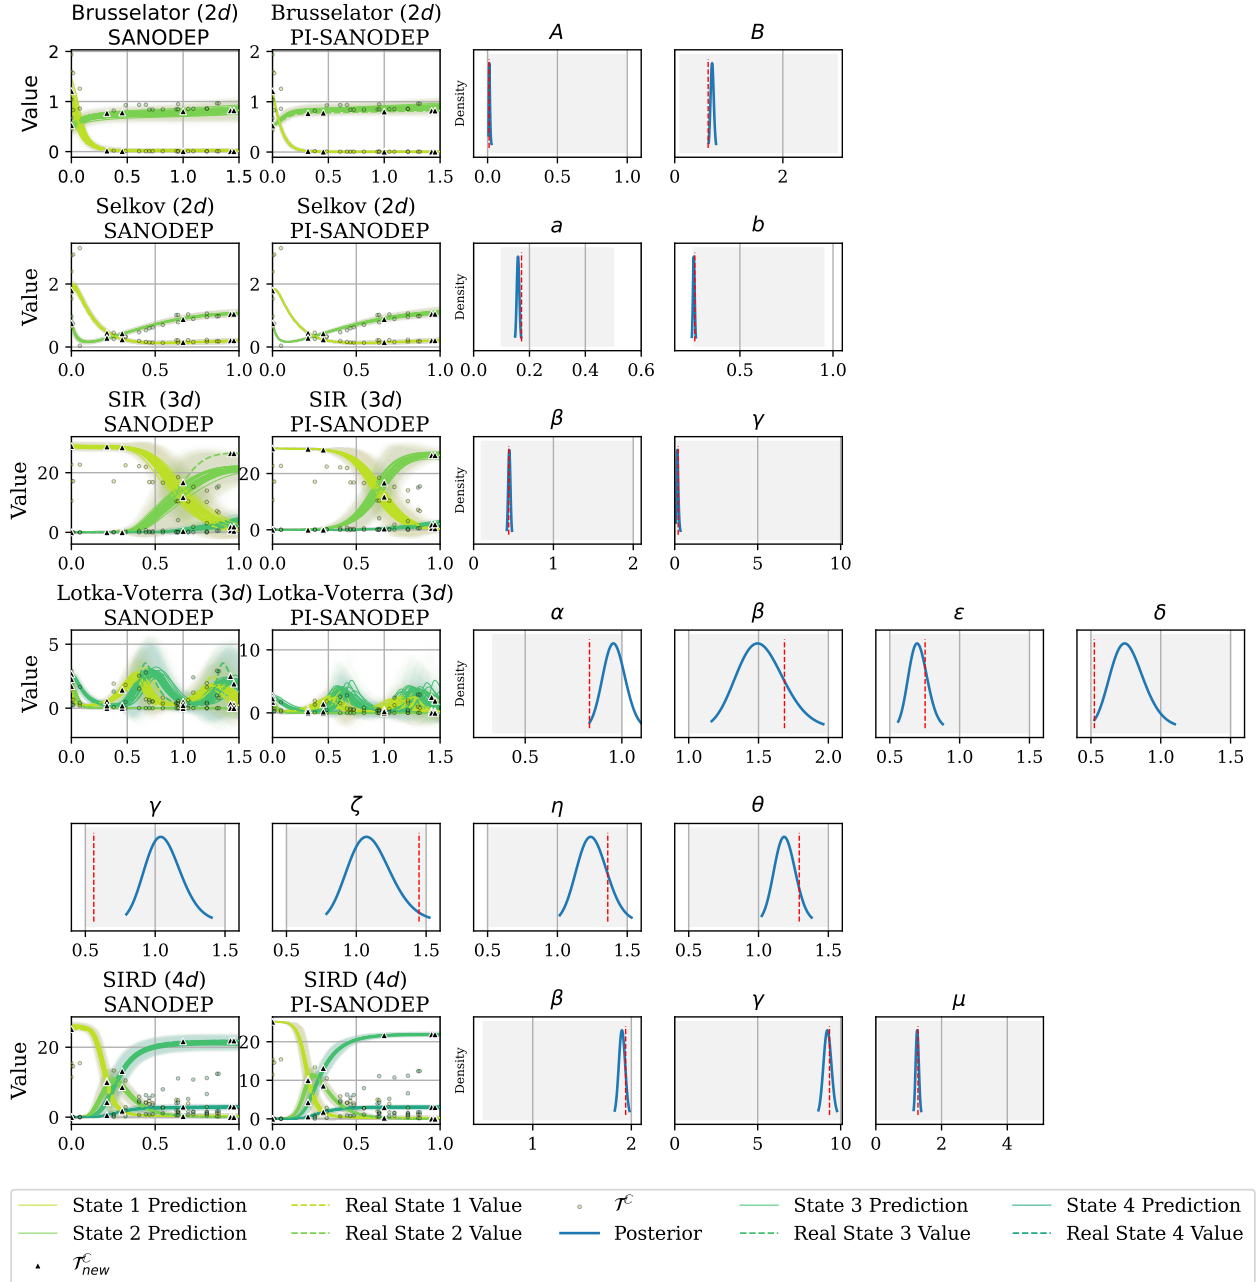


Figure 15: Additional experiments for meta-parameter estimation based on PI-SANODEP, the shaded area represents the meta-training parameter support $\mathcal{S} = \text{supp}(P)$. It can be seen that PI-SANODEP not only improves upon SANODEP over prediction (e.g., in the SIR example) but also provides a reasonable parameter estimation of the test system. Parameter estimation tends to become difficult when the number of parameters is high as can be seen from the 3-dimensional Lotka-Voterra system.

E.2 Vector Valued GP Prior as Task Distribution

As an initial investigation with SANODEP, inspired by the recent approach of using Gaussian Processes (GPs) to model vector fields (e.g., Heinonen et al. (2018)), we propose to use vector-valued GP as the non-parametric prior for the vector field, denoted by $\mathbf{f} \sim \mathcal{GP}(\mathbf{0}, K(\mathbf{x}, \mathbf{x}'))$, as task distributions for meta-learning

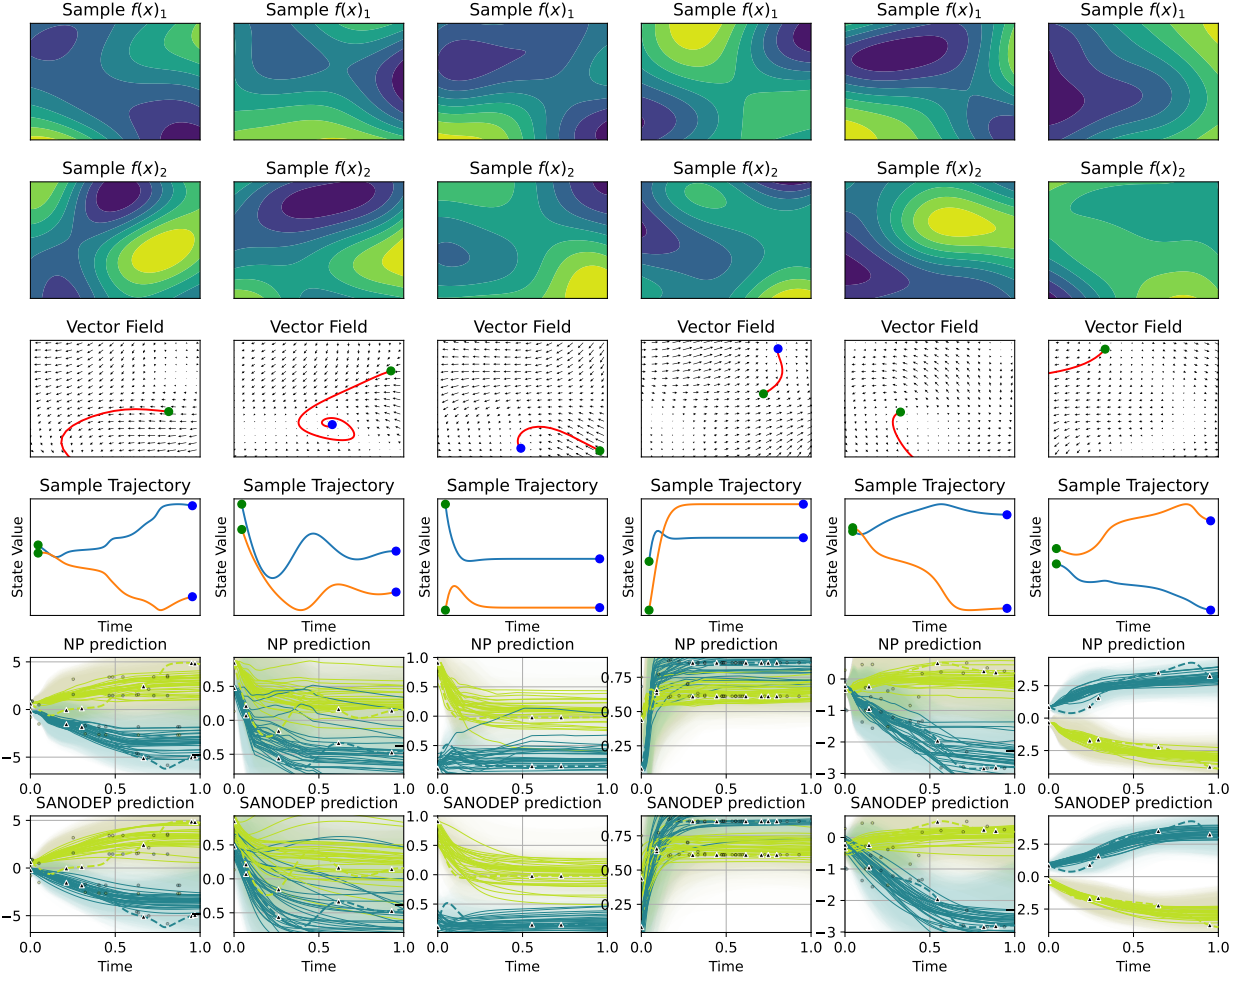


Figure 16: Task distribution generated by vector-valued Gaussian Process (GP) priors. The first two rows display samples from parametric GP priors. The third and fourth rows show the sampled vector fields and the resulting trajectories starting from sampled initial conditions. The last two rows depict the meta-model predictions on these trajectories. It can be seen that while the GP-based vector field has introduced a very flexible task distribution, both meta-learn models tend to underfit on these trajectories.

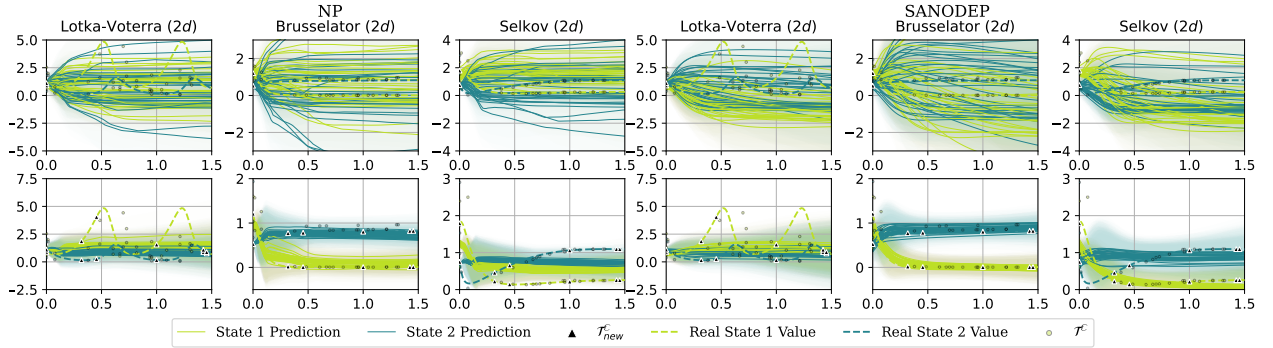


Figure 17: Cross domain generalization exploration: SANODEP and NP meta trained on GP-ODE seems to be able to capture the trajectories from Bruseelator system.

for dynamical systems. Further details on the data used for meta-training are provided in Appendix D.1. The dynamical systems sampled from this vector-valued GP prior, as illustrated in Fig. 16, demonstrate its ability to represent a broad distribution of dynamical systems. Additionally, we offer a visual comparison of SANODEP and NP on the sampled trajectories. This comparison reveals a trade-off between fitting accuracy and trajectory flexibility for both models. Specifically, both models tend to significantly underfit trajectories that involve oscillations, indicating the noticable complicity of this meta-learning problem.

We also explore whether this flexible task distribution facilitates cross-domain generalization. To this end, we employ the trained model on different dynamical systems that share the same state variable dimensionality, as shown in Fig. 17. Both meta-learned models demonstrate some potential for cross-domain generalization particularly from the Brusselator system. However, generalization to systems involving oscillatory behaviors proves to be especially challenging.

F Notations

Table 6: Nomenclature Table

Notation	Meaning
\mathbf{x}	State variable vector
$d_{\mathbf{x}}$	Dimension of the state space
t_0	Initial time for the dynamical system
t_i	Trajectory time samples, sampled irregularly
t	Termination time for a trajectory
\mathbb{C}	Context Sets
\mathbf{f}	Vector field governing the dynamics
$g(\cdot)$	Aggregation function used in optimization
L_0	Latent initial condition
L_D	Dynamics of the system
ϕ_r	Context encoding
\mathbf{r}	Context representation vector
\mathbf{l}_0	Realizations of latent initial conditions
$\mathbf{l}(t)$	Latent states at time t
\mathbf{u}	Realizations of the control term representing latent dynamics
\mathbf{f}_{nn}	Neural network parameterized vector field
θ_{ode}	Parameterisation of the ODE system
\mathbf{f}_{evolve}	Numerical solver implementation of the vector field
h	Instantaneous time variable in continuous dynamics
d_l	Dimension of the latent state space
\mathbb{T}	Target Sets
$\mathbf{T}^{\mathbb{T}}$	Target times for predictions
$\mathbf{X}^{\mathbb{T}}$	Target state values
\mathcal{F}	Distribution of dynamical systems
P	Parametric Task distribution, often stochastic
$\mathcal{X}_{\mathbf{x}}$	State variable space
\mathcal{X}_0	Initial condition space
τ	Time domain for the system evolution
M	Number of context trajectories
N_l	Number of context elements in trajectory l
$\mathcal{T}^{\mathbb{C}}$	Context observations from M trajectories
$\mathcal{T}_l^{\mathbb{C}}$	Context observations from the l th trajectory
\mathcal{T}_{new}	A new trajectory which augments the context set \mathbb{C}
$\mathcal{T}_{new}^{\mathbb{C}}$	Context observations from the new trajectory
$\mathcal{T}_{new}^{\mathbb{T}}$	Target observations from the new trajectory

Table 6: Nomenclature Table (Continued)

Notation	Meaning
D_{sys}	Latent variable conditioned on $M + 1$ trajectories
\mathbf{u}_{sys}	Realization of system-aware latent dynamics
\mathbf{r}_{sys}	System context representation vector
$\phi_{r_{sys}}$	Context encoding augmented with the initial state
$\mathbf{X}_{new}^{\mathbb{T}}$	Target state values for the new trajectory
$\mathbf{T}_{new}^{\mathbb{T}}$	Target times for the new trajectory
$\mathbf{x}_{new_i}^{\mathbb{C}}$	State vector samples from the new trajectory
$t_{new_i}^{\mathbb{C}}$	Time samples from the new trajectory
K	Number of observations in the new trajectory
\mathcal{L}_{θ}	SANODEP’s multi scenario loss function parametrised by θ
p_{θ}	SANODEP prediction parametrised by θ
$\mathbb{1}_{forecast}$	Bernoulli Indicator
λ	Parameter for Bernoulli distribution in decision-making process
\mathcal{N}_{opt}	Search space of the number of observations in the trajectory
Δt	Minimum delay between observations
t_{max}	Maximum time observation
N_{max}	Maximum number of observations per trajectory
$\alpha(\cdot)$	Batch acquisition function
HVI	Hypervolume improvement based on the Pareto frontier
\mathcal{F}^*	Pareto frontier
M_{min}	Minimum number of context trajectories observed in a system
M_{max}	Maximum number of context trajectories observed in a system
$N_{\mathbf{x}_0}$	Sampled number of initial conditions within a system
N_{sys}	Number of dynamical system samples per training iteration
N_{grid}	Number of points in the time grid for trajectory evaluation
m_{min}, m_{max}	Min and max context points within a trajectory
n_{min}, n_{max}	Min and max target points beyond context in a trajectory
\mathcal{S}	The support set of the stochastic parameters for meta-task distribution
\mathbf{T}_{grid}	Time grid used for trajectory evaluations