
Constraining the Dynamics of Deep Probabilistic Models

Marco Lorenzi¹ Maurizio Filippone²

Abstract

We introduce a novel generative formulation of deep probabilistic models implementing “soft” constraints on their function dynamics. In particular, we develop a flexible methodological framework where the modeled functions and derivatives of a given order are subject to inequality or equality constraints. We then characterize the posterior distribution over model and constraint parameters through stochastic variational inference. As a result, the proposed approach allows for accurate and scalable uncertainty quantification on the predictions and on all parameters. We demonstrate the application of equality constraints in the challenging problem of parameter inference in ordinary differential equation models, while we showcase the application of inequality constraints on the problem of monotonic regression of count data. The proposed approach is extensively tested in several experimental settings, leading to highly competitive results in challenging modeling applications, while offering high expressiveness, flexibility and scalability.

1. Introduction

Modern machine learning methods have demonstrated state-of-art performance in representing complex functions in a variety of applications. Nevertheless, the translation of complex learning methods in natural sciences and in the clinical domain is still challenged by the need of interpretable solutions. To this end, several approaches have been proposed in order to constrain the solution dynamics to plausible forms such as boundedness (Da Veiga & Marrel, 2012), monotonicity (Riihimäki & Vehtari, 2010), or mechanistic behaviors (Alvarez et al., 2013). This is a crucial require-

ment to provide a more precise and realistic description of natural phenomena. For example, monotonicity of the interpolating function is a common assumption when modeling disease progression in neurodegenerative diseases (Lorenzi et al., 2017; Donohue et al., 2014), while bio-physical or mechanistic models are necessary when analyzing and simulating experimental data in bio-engineering (Vyshemirsky & Girolami, 2007; Konukoglu et al., 2011).

However, accounting for the complex properties of biological systems in data-driven modeling approaches poses important challenges. For example, functions are often non-smooth and characterized by nonstationaries which are difficult to encode in “shallow” models. Complex cases can arise already in classical ODE systems for certain configurations of the parameters, where functions can exhibit sudden temporal changes (Goel et al., 1971; FitzHugh, 1955). Within this context, approaches based on stationary models, even when relaxing the smoothness assumptions, may lead to suboptimal results for both data modeling (interpolation), and estimation of dynamics parameters. To provide insightful illustrations of this problem we anticipate the results of Section 4.4.1 and Figure 5. Moreover, the application to real data requires to account for the uncertainty of measurements and underlying model parameters, as well as for the – often large – dimensionality characterizing the experimental data. Within this context, deep probabilistic approaches may represent a promising modeling tool, as they combine the flexibility of deep models with a systematic way to reason about uncertainty in model parameters and predictions. The flexibility of these approaches stems from the fact that deep models implement compositions of functions, which considerably extend the complexity of signals that can be represented with “shallow” models (LeCun et al., 2015). Meanwhile, their probabilistic formulation introduces a principled approach to quantify uncertainty in parameters estimation and predictions, as well as to model selection problems (Neal, 1996; Ghahramani, 2015).

In this work, we aim at extending deep probabilistic models to account for constraints on their dynamics. In particular, we focus on a general and flexible formulation capable of imposing a rich set of constraints on functions and derivatives of any order. We focus on: i) *equality constraints* on the function and its derivatives, required when the model should satisfy given physical laws implemented through

^{*}Equal contribution ¹University of Cote d’Azur, INRIA Sophia Antipolis, EPIONE research group, France ²EURECOM, Department of Data Science, Sophia Antipolis, France. Correspondence to: Marco Lorenzi <marco.lorenzi@inria.fr>, Maurizio Filippone <maurizio.filippone@eurecom.fr>.

a mechanistic description of a system of interest; and ii) *inequality constraints*, arising in problems where the class of suitable functions is characterized by specific properties, such as monotonicity or convexity/concavity (Riihimäki & Vehtari, 2010).

In case of equality constraints, we tackle the challenge of parameters inference in Ordinary Differential Equations (ODE). Exact parameter inference of ODE models is computationally expensive due to the need for repeatedly solving ODEs within the Bayesian setting. To this end, previous works attempted to recover tractability by introducing approximate solutions of ODEs (see, e.g., Macdonald & Husmeier (2015) for a review). Following these ideas, we introduce “soft” constraints through a probabilistic formulation that penalizes functions violating the ODE on a set of virtual inputs. Note that this is in contrast to previous approaches, such as the ones proposed with probabilistic ODE solvers (Wheeler et al., 2014; Schober et al., 2014), where a given dynamics is strictly enforced to the model posterior. By deriving a lower bound on the model evidence, we enable the use of stochastic variational inference to achieve end-to-end posterior inference over model and constraint parameters.

In what follows we shall focus on a class of deep probabilistic models implementing a composition of Gaussian processes (GPs) (Rasmussen & Williams, 2006) into Deep Gaussian Processes (DGPs) (Damianou & Lawrence, 2013). More generally, our formulation can be straightforwardly extended to probabilistic Deep Neural Networks (DNNs) (Neal, 1996). On the practical side, our formulation allows us to take advantage of automatic differentiation tools, leading to flexible and easy-to-implement methods for inference in constrained deep probabilistic models. As a result, our method scales linearly with the number of observations and constraints. Furthermore, in the case of mean-field variational inference, it also scales linearly with the number of parameters in the constraints. Finally, it can easily be parallelized/distributed and exploit GPU computing.

Through an in-depth series of experiments, we demonstrate that our proposal achieves state-of-the-art performance in a number of constrained modeling problems while being characterized by attractive scalability properties. The paper is organized as follows: Section 2 reports on related work, whereas the core of the methodology is presented in Section 3. Section 4 contains an in-depth validation of the proposed model against the state-of-the-art. We demonstrate the application of equality constraints in the challenging problem of parameter inference in ODE, while we showcase the application of inequality constraints in the monotonic regression of count data. Additional insights and conclusions are given in Section 5. Results that we could not fit in the manuscript are deferred to the supplementary material.

2. Related Work

Equality constraints where functions are enforced to model the solution of ODE systems have been considered in a variety of problems, particularly in the challenging task of accelerated inference of ODE parameters. Previous approaches to accelerate ODE parameter optimization involving interpolation date back to Varah (1982). This idea has been developed in several ways, including splines, GPs, and Reproducing Kernel Hilbert spaces. Works that employ GPs as interpolants have been proposed in Ramsay et al. (2007), Liang & Wu (2008), Calderhead et al. (2009), and Campbell & Steele (2012). Such approaches have been extended to introduce a novel formulation to regularize the interpolant based on the ODE system, notably Dondelinger et al. (2013); Barber & Wang (2014). An in-depth analysis of the model in Barber & Wang (2014) is provided by Macdonald et al. (2015). Recently, Gorbach et al. (2017) extended previous works by proposing mean-field variational inference to obtain an approximate posterior over ODE parameters. Our work improves previous approaches by considering a more general class of interpolants than “shallow” GPs, and proposes a scalable framework for inferring the family of interpolating functions jointly with the parameters of the constraint, namely ODE parameters.

Another line of research that builds on gradient matching approaches uses a Reproducing Kernel Hilbert space formulation. For example, González et al. (2014) proposes to exploit the linear part of ODEs to accelerate the interpolation, while Niu et al. (2016) exploits the quadratic dependency of the objective with respect to the parameters of the interpolant to improve the computational efficiency of the ODE regularization. Interestingly, inspired by Calandra et al. (2016), the latter approach was extended to handle nonstationarity in the interpolation through warping (Niu et al., 2018). The underlying idea is to estimate a transformation of the input domain to account for nonstationarity of the signal, in order to improve the fitting of stationary GP interpolants. A key limitation of this approach is the lack of a probabilistic formulation, which prevents one from approximating the posterior over ODE parameters. Moreover, the warping approach is tailored to periodic functions, thus limiting the generalization to more complex signals. In our work, we considerably improve on these aspects by effectively modeling the warping through GPs/DGPs that we infer jointly with ODE parameters.

Inequality constraints on the function derivatives have been considered in several works such as in Meyer (2008); Groeneboom & Jongbloed (2014); Mašić et al. (2017); Riihimäki & Vehtari (2010); Da Veiga & Marrel (2012); Salzmann & Urtasun (2010). In particular, the GP setting provides a solid and elegant theoretical background for tackling this problem; thanks to the linearity of differentiation,

both mean and covariance functions of high-order derivatives of GPs can be expressed in closed form, leading to exact formulations for linearly-constrained GPs (Da Veiga & Marrel, 2012). In case of inequality constraints on the derivatives, instead, this introduces non-conjugacy between the likelihood imposing the derivative constraint and the GP prior, thus requiring approximations (Riihimäki & Vehtari, 2010). Although this problem can be tackled through sampling schemes or variational inference methods, such as Expectation Propagation (Minka, 2001), scalability to large dimensions and sample size represents a critical limitation. In this work, we extend these methods by considering a more general class of functions based on DGPs, and develop scalable inference that makes our method applicable to large data and dimensions.

3. Methods

3.1. Equality constraints in probabilistic modeling

In this section we provide a derivation of the posterior distribution of our model when we introduce equality constraints in the dynamics. Let Y be a set of n observed multivariate variables $\mathbf{y}_i \in \mathbb{R}^s$ associated with measuring times t collected into \mathbf{t} ; the extension where the n variables are measured at different times is notationally heavier but straightforward. Let $\mathbf{f}(t)$ be a multivariate interpolating function with associated noise parameters $\boldsymbol{\psi}$, and define F similarly to Y to be the realization of \mathbf{f} at \mathbf{t} . In this work, $\mathbf{f}(t)$ will be either modeled using a GP, or deep probabilistic models based on DGPs. We introduce functional constraints on the dynamics of the components of $\mathbf{f}(t)$ by specifying a family of admissible functions whose derivatives of order h evaluated at the inputs \mathbf{t} satisfy some given constraint

$$C_{hi} = \left\{ \mathbf{f}(t) \left| \frac{d^h f_i(t)}{dt^h} = \mathcal{H}_{hi} \left(t, \mathbf{f}, \frac{d\mathbf{f}}{dt}, \dots, \frac{d^q \mathbf{f}}{dt^q}, \boldsymbol{\theta} \right) \right|_{\mathbf{t}} \right\}.$$

Here the constraint is expressed as a function of the input, the function itself, and high-order derivatives up to order q . The constraint also includes $\boldsymbol{\theta}$ as dynamics parameters that should be inferred. We are going to consider the intersection of all the constraints for a set of indices \mathcal{I} comprising pairs (h, i) of interest

$$C = \bigcap_{(h,i) \in \mathcal{I}} C_{hi}$$

To keep the notation uncluttered, and without loss of generality, in the following we will assume that all the terms are evaluated at \mathbf{t} ; we can easily relax this by allowing for the constraints to be evaluated at different sampling points than \mathbf{t} . As a concrete example, consider the constraints induced by the Lotka-Volterra ODE system (more details in the experiments section); for this system, $\boldsymbol{\theta} = \{\alpha, \beta, \gamma, \delta\}$,

and the family of functions is identified by the conditions

$$\begin{aligned} \left. \frac{dg_1(t)}{dt} \right|_{\mathbf{t}} &= \mathcal{H}_{11}(\mathbf{f}(t)) \Big|_{\mathbf{t}} = \alpha f_1(\mathbf{t}) - \beta f_1(\mathbf{t}) f_2(\mathbf{t}), \\ \left. \frac{dg_2(t)}{dt} \right|_{\mathbf{t}} &= \mathcal{H}_{12}(\mathbf{f}(t)) \Big|_{\mathbf{t}} = -\gamma f_2(\mathbf{t}) + \delta f_1(\mathbf{t}) f_2(\mathbf{t}), \end{aligned}$$

where the products $f_1(\mathbf{t})f_2(\mathbf{t})$ are element-wise.

Denote by $\tilde{F} = \{\mathbf{f}_{hi}\}$ the set of realizations of \mathbf{f} and of its derivatives at any required order h evaluated at timed \mathbf{t} . We define the constrained regression problem through two complementary likelihood-based elements: a data attachment term $p(Y|F, \boldsymbol{\psi})$, and a term quantifying the constraint on the dynamics, $p(C|\tilde{F}, \boldsymbol{\theta}, \boldsymbol{\psi}_D)$, where $\boldsymbol{\psi}_D$ is the associated noise parameter. To solve the inference problem, we shall determine a lower bound for the marginal

$$p(Y, C|\mathbf{t}, \boldsymbol{\psi}, \boldsymbol{\psi}_D) = \quad (1)$$

$$\int p(Y|F, \boldsymbol{\psi}) p(C|\tilde{F}, \boldsymbol{\theta}, \boldsymbol{\psi}_D) p(F, \tilde{F}|\mathbf{t}, \boldsymbol{\psi}) p(\boldsymbol{\theta}) dF d\tilde{F} d\boldsymbol{\theta},$$

where

$$p(F, \tilde{F}|\mathbf{t}, \boldsymbol{\psi}) = p(\tilde{F}|F) p(F|\mathbf{t}, \boldsymbol{\psi}).$$

Note that \tilde{F} is in fact completely identified by F .

Equation (1) requires specifying suitable models for both likelihood and functional constraints. This problem thus implies the definition of noise models for both observations and model dynamics. In the case of continuous observations, the likelihood can be assumed to be Gaussian:

$$p(Y|F, \boldsymbol{\psi}) = \mathcal{N}(Y|F, \Sigma(\boldsymbol{\psi})), \quad (2)$$

where $\Sigma(\boldsymbol{\psi})$ is a suitable multivariate covariance. Extensions to other likelihood functions are possible, and in the experiments we show an application to regression on counts where the likelihood is Poisson with rates equal to the exponential of the elements of F .

Concerning the noise model for the derivative observations, we assume independence across the constraints C_{hi} so that

$$p(C|\tilde{F}, \boldsymbol{\theta}, \boldsymbol{\psi}_D) = \prod_{(h,i) \in \mathcal{I}} p(C_{hi}|\tilde{F}, \boldsymbol{\theta}, \boldsymbol{\psi}_D). \quad (3)$$

We can again assume a Gaussian likelihood:

$$p(C_{hi}|\tilde{F}, \boldsymbol{\theta}, \boldsymbol{\psi}_D) = \prod_t \mathcal{N}(\mathbf{f}_{hi}(t) | \mathcal{H}_{hi}(t, \tilde{F}, \boldsymbol{\theta}), \boldsymbol{\psi}_D), \quad (4)$$

or, in order to account for potentially heavy-tailed error terms on the derivative constraints, we can assume a Student-t distribution:

$$p(C_{hi}|\tilde{F}, \boldsymbol{\theta}, \boldsymbol{\psi}_D) = \prod_t \mathcal{T}(\mathbf{f}_{hi}(t) | \mathcal{H}_{hi}(t, \tilde{F}, \boldsymbol{\theta}), \boldsymbol{\psi}_D, \nu), \quad (5)$$

where $\mathcal{T}(z|\mu, \lambda, \nu) \propto \frac{1}{\lambda} [1 + \frac{(z-\mu)^2}{\nu\lambda^2}]^{-(\nu+1)/2}$. We test these two noise models for \tilde{F} in the experiments.

3.2. Inequality constraints in probabilistic modeling

In the case of inequality constraints we can proceed analogously as in the previous section. In particular, we are interested in the class of functions satisfying the following conditions:

$$\mathcal{C}_{hi} = \left\{ \mathbf{f}(t) \left| \frac{d^h f_i(t)}{dt^h} > \mathcal{H}_{hi} \left(t, \mathbf{f}, \frac{d\mathbf{f}}{dt}, \dots, \frac{d^q \mathbf{f}}{dt^q}, \boldsymbol{\theta} \right) \right|_t \right\}.$$

For example, a monotonic univariate regression problem can be obtained with a constraint of the form $\frac{df}{dt} > 0$.

In this case, the model dynamics can be enforced by a logistic function:

$$p(\mathcal{C}_{hi} | \tilde{F}, \psi_D) = \prod_{j=1}^n \frac{1}{1 + \exp(-\psi_D \frac{df}{dt}(t_j))}, \quad (6)$$

where the parameter ψ_D controls the strength of the monotonicity constraint.

3.3. Optimization and inference in constrained regression with DGPs

After recalling the necessary methodological background, in this section we derive an efficient inference scheme for the model posterior introduced in Section 3.1.

To recover tractability, our scheme leverages on recent advances in modeling and inference in DGPs through approximation via random feature expansions (Rahimi & Recht, 2008; Cutajar et al., 2017). Denoting with $F^{(l)}$ the GP random variables at layer l , an (approximate) DGP is obtained by composing GPs approximated by Bayesian linear models, $F^{(l)} \approx \Phi^{(l)} W^{(l)}$. The so-called random features $\Phi^{(l)}$ are obtained by multiplying the layer input by a random matrix $\Omega^{(l)}$ and by applying a nonlinear transformation $h(\cdot)$. For example, in case of the standard RBF covariance, the elements in $\Omega^{(l)}$ are Gaussian distributed with covariance function parameterized through the length-scale of the RBF covariance. The nonlinearity is obtained through trigonometric functions, $\mathbf{h}(\cdot) = (\cos(\cdot), \sin(\cdot))$, while the prior over the elements of $W^{(l)}$ is standard normal. As a result, the interpolant becomes a Bayesian Deep Neural Network (DNN), where for each layer we have weights $\Omega^{(l)}$ and $W^{(l)}$, and activation functions $\mathbf{h}(\cdot)$ applied to the input of each layer multiplied by the weights $\Omega^{(l)}$.

3.3.1. DERIVATIVES IN DGPS WITH RANDOM FEATURE EXPANSIONS

To account for function derivatives consistently with the theory developed in Cutajar et al. (2017), we need to extend the random feature expansion formulation of DGPs to high-order derivatives. Fortunately, this is possible thanks to the chain rule and to the closure under linear operations of

the approximated GPs. More precisely, the derivatives of a ‘‘shallow’’ GP model with form $F = \mathbf{h}(t\Omega)W$ can still be expressed through linear composition of matrix-valued operators depending on W and Ω only: $\frac{dF}{dt} = \frac{d\mathbf{h}(t\Omega)}{dt} W$. The computational tractability is thus preserved and the GP function and derivatives are identified by the same sets of weights Ω and W . The same principle clearly extends to DGP architectures where the derivatives at each layer can be combined following the chain rule to obtain the derivatives of the output function with respect to the input.

3.3.2. VARIATIONAL LOWER BOUND

In the constrained DGP setting, we are interested in carrying out inference of the functions $F^{(l)}$ and of the associated covariance parameters at all layers. Moreover, we may want to infer any dynamics parameters $\boldsymbol{\theta}$ that parameterize the constraint on the derivatives. Within this setting, the inference of the latent variables $F^{(l)}$ in the marginal (1) is generally not tractable. Nevertheless, the Bayesian DNN structure provided by the random feature approximation allows the efficient estimation of its parameters, and the tractability of the inference is thus recovered.

In particular, let Ω , \mathbf{W} , and $\boldsymbol{\psi}$ be the collections of all $\Omega^{(l)}$, $W^{(l)}$, and covariance and likelihood parameters, respectively. Recalling that we can obtain random features at each layer by sampling the elements in Ω from a given prior distribution, we propose to tackle the inference problem through *variational inference* of the parameters \mathbf{W} and $\boldsymbol{\theta}$. We could also attempt to infer Ω , although in this work we are going to assume them sampled from the prior with fixed randomness, which allows us to optimize covariance parameters using the reparameterization trick (option PRIOR-FIXED in Cutajar et al. (2017)). We also note that we could infer, rather than optimize, $\boldsymbol{\psi}$; we leave this for future work.

Using Jensen’s inequality, the variational approach allows us to obtain a lower bound on the log-marginal likelihood $\mathcal{L} := \log [p(Y, \mathcal{C} | t, \Omega, \boldsymbol{\psi}, \psi_D)]$ of equation (1), as follows:

$$\begin{aligned} \mathcal{L} &\geq E_{q(\mathbf{W})} (\log [p(Y | \Omega, \mathbf{W}, \boldsymbol{\psi})]) \\ &\quad + E_{q(\mathbf{W})q(\boldsymbol{\theta})} (\log [p(\mathcal{C} | \Omega, \mathbf{W}, \boldsymbol{\psi}_D, \boldsymbol{\theta})]) \\ &\quad - \text{DKL}(q(\mathbf{W}) || p(\mathbf{W})) - \text{DKL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta})). \end{aligned} \quad (7)$$

The distribution $q(\mathbf{W})$ acts as a variational approximation and is assumed to be Gaussian, factorizing completely across weights and layers (l):

$$q(\mathbf{W}) = \prod_{j,k,l} p(W_{jk}^{(l)}) = \prod_{j,k,l} \mathcal{N} \left(m_{jk}^{(l)}, (s^2)_{jk}^{(l)} \right). \quad (8)$$

Extensions to approximations where we relax the factorization assumption are possible. Similarly, we are going to assume $q(\boldsymbol{\theta})$ to be Gaussian, and will assume no factorization, so that $q(\boldsymbol{\theta}) = \mathcal{N}(\mu_\theta, \Sigma_\theta)$.

4. Experiments

This section reports an in-depth validation of the proposed method on a variety of benchmarks. We are going to study the proposed variational framework for constrained dynamics in DGP models for ODE parameter estimates using equality constraints, and compare it against state-of-the-art methods. We will then consider the application of inequality constraints for a regression problem on counts, which was previously considered in the literature of monotonic GPs.

4.1. Settings for the proposed constrained DGP

We report here the configuration that we used across all benchmarks for the proposed method. Due to the generally low sample size n used across experiments (in most cases $n < 50$), unless specified otherwise the tests were performed with a two-layer DGP $\mathbf{f}(t) = \mathbf{f}^{(2)} \circ \mathbf{f}^{(1)}(t)$, with dimension of the “hidden” GP layer $\mathbf{f}^{(1)}(t)$ equal to 2, and RBF kernels. The length-scale of the RBF covariances was initialized to $\lambda_0 = \log(t_{\max} - t_{\min})$, while the marginal standard deviation to $\alpha_0 = \log(y_{\max} - y_{\min})$; the initial likelihood noise was set to $\sigma_0^2 = \alpha_0/10^5$. Finally, the initial ODE parameters were set to the value of 0.1. The optimization was carried out through stochastic gradient descent with Adaptive moment Estimation (Adam) (Kingma & Ba, 2017), through the alternate optimization of i) the approximated posterior over \mathbf{W} and likelihood/covariance parameters ($q(\mathbf{W})$ and ψ), and ii) likelihood parameters of ODE constraints and the approximate posterior over ODE parameters (ψ_D and $q(\theta)$). We note that the optimization of the ODE constraints parameters (the noise and scale parameters for Gaussian and Student-t likelihoods, respectively) is aimed at identifying in a fully data-driven manner the optimal trade-off between data attachment (likelihood term) and regularity (constraints on the dynamics). In what follows, DGP-t and DGP-G respectively denote the model tested with Student-t and Gaussian noise models on the ODE constraints.

4.2. Equality constraints from ODE systems

The proposed framework was tested on a set of ODE systems extensively studied in previous works: Lotka-Volterra (Goel et al., 1971), FitzHugh-Nagumo (FitzHugh, 1955), and protein biopathways from Vyshemirsky & Girolami (2007). For each experiment, we used the experimental setting proposed in previous studies (Niu et al., 2016; Macdonald & Husmeier, 2015). In particular, for each test, we identified two experimental configurations with increasing modeling difficulty (e.g. less samples, lower signal-to-noise ratio, . . .). A detailed description of the models and testing parameters is provided in the supplementary material. The experimental results are reported for parameter inference and model estimation performed on 5 different realizations of the noise.

4.2.1. BENCHMARK

We tested the proposed method against several reference approaches from the state-of-art to infer parameters of ODE systems.

RKG3: We tested the method presented in Niu et al. (2016) using the implementation in the R package `KGode`. This method implements gradient matching, where the interpolant is modeled using functions in Reproducing Kernel Hilbert spaces. This approach, for which ODE parameters are estimated and not inferred, was shown to achieve state-of-the-art performance on a variety of ODE estimation problems. We used values ranging from 10^{-4} to 1 for the parameter λ that the method optimizes using cross-validation.

Warp: In the R package `KGode` there is also an implementation of the warping approach presented in Niu et al. (2018). This method extends gradient matching techniques by attempting to construct a warping of the input where smooth Reproducing Kernel Hilbert spaces-based interpolants can effectively model nonstationary observations. The warping attempts to transform the original signal via assumptions on periodicity and regularity conditions. We used the default parameters and initialized the optimization of the warping function from a period equal to the interval where observations are available. Similarly to RKG3, ODE parameters are estimated and not inferred.

AGM: We report results on the Approximate Gradient Matching (AGM) approach in Dondelinger et al. (2013), implemented in the recently released R package `deGradInfer`. AGM implements a population Markov chain Monte Carlo approach tempering from the prior to the approximate posterior of ODE parameters based on an interpolation with GPs. In the experiments we use 10 parallel chains and we run them for 10^4 iterations. In the implementation of AGM, the variance of the noise on the observations is assumed known and it is fixed; we expect this to give a slight advantage to this method.

MCMC: In the R package `deGradInfer` there is also an implementation of a population Markov chain Monte Carlo sampler where the ODE is solved explicitly. In this case too we use 10 parallel chains that we run for 10^4 iterations. In contrast to AGM, in this implementation, the variance of the noise on the observations is learned together with ODE parameters.

4.2.2. RESULTS

Figure 4.2.2 shows the distribution of the root mean squared error (RMSE) across folds for each experimental setting (see supplement for details). We note that the proposed method consistently leads to better RMSE values compared to the reference approaches (except some folds in one of the Fitz-Hugh-Nagumo experiments, according to a Mann-

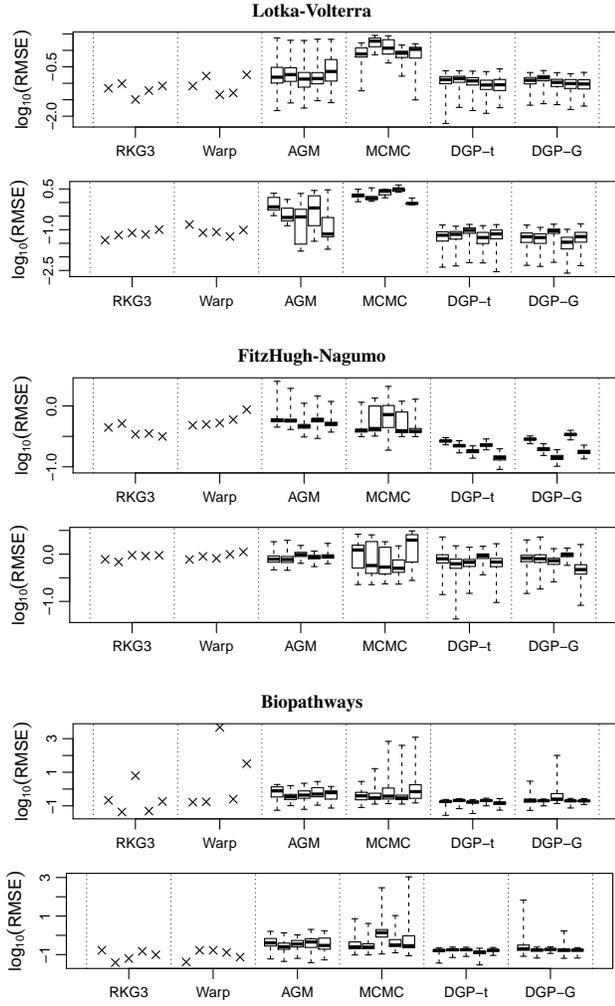


Figure 1. Boxplot of the RMSE on ODE parameters for the three ODE systems considered and for the two experimental settings. We report 5 bars for each method in the plots, corresponding to five different instantiations of the noise.

Whitney nonparametric test), and that DGP-t provides more consistent parameter estimates than DGP-G. This latter result may indicate a lower sensitivity to outliers derivatives involved in the functional constraint term. This is a crucial aspect due to the generally noisy derivative terms of nonparametric regression models. The distribution of the parameters for all the datasets tested in this study, which we report in the supplementary material, reveals that, unlike the nonprobabilistic methods RKG3 and WARP, our approach is capable of inferring ODE parameters yielding meaningful uncertainty estimation.

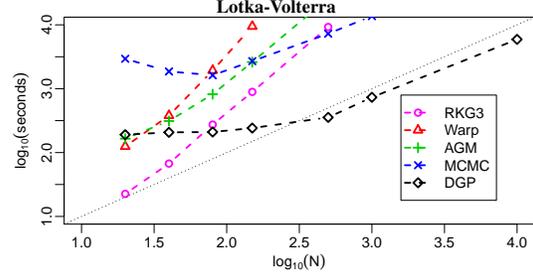


Figure 2. Execution time vs sample size – Lotka-Volterra ODE.

4.3. Scalability test - large n

We tested the scalability of the proposed method with respect to sample size. To this end, we repeated the test on the Lotka-Volterra system with $n = 20, 40, 80, 150, 500, 10^3$, and 10^4 observations. For each instance of the model, the execution time was recorded and compared with the competing methods. All the experiments were performed on a 1.3GHz Intel Core i5 MacBook. The proposed method scales linearly with n (Figure 2), while it has an almost constant execution when $n < 500$; we attribute this effect to overheads in the framework we used to code our method. For small n , the running time of our method is comparable with competing methods, and it is considerably faster in case of large n .

4.4. Scalability test - large s

In order to assess the ability of the framework to scale to a large number of ODEs, we tested our method on the Lorenz96 system with increasing number of equations, $s = 125$ to $s = 1000$ (Lorenz & Emanuel, 1998). To the best of our knowledge, the solution of this challenging problem via gradient matching approaches has only been previously attempted in Gorbach et al. (2017). We could not find an implementation of their method to carry out a direct comparison, so we are going to refer to the results reported in their paper. The system consists of a set of drift states functions $(f_1(\mathbf{x}(t), \theta), f_2(\mathbf{x}(t), \theta), \dots, f_s(\mathbf{x}(t), \theta))$ recursively linked by the relationship:

$$f_i(\mathbf{x}(t), \theta) = (x_{i+1}(t) - x_{i-2}(t))x_{i-1}(t) - x_i(t) + \theta,$$

where $\theta \in \mathbb{R}$ is the drift parameter. Consistently with the setting proposed in Gorbach et al. (2017); Vrettas et al. (2015), we set $\theta = 8$ and generated 32 equally spaced observations over the interval $[0, 4]$ seconds, with additive Gaussian noise $\sigma^2 = 1$. We performed two tests by training (i) on all the states, and (ii) by keeping one third of the states as unobserved, and by applying our method to identify model dynamics on both observed and unobserved states.

Figure 3 shows the average RMSE in the different experimental settings. As expected, the modeling accuracy is

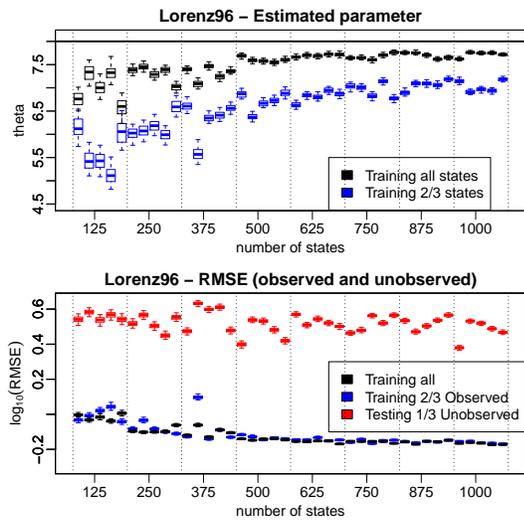


Figure 3. Top: parameter estimates in the different folds when training on all (black) or only on 2/3 (blue) of the states. The ground truth is indicated by the top horizontal bar ($\theta = 8$). Bottom: RMSE on the ODE curves fitting when training on all the states (black), and on the observed (blue) and unobserved (red) states when training on 2/3 of the states only.

sensibly higher when trained on the full set of equations. Moreover, the RMSE is lower on observed states compared to unobserved ones. This is confirmed by visual inspection of the modeling results for sample training and testing states (Figure 4). The observed states are generally associated with lower uncertainty in the predictions and by an accurate fitting of the solutions (Figure 4, top). The model still provides remarkable modeling results on unobserved states (Figure 4, bottom), although with decreased accuracy and higher uncertainty. We are investigating the reasons for the posterior distribution over θ not covering the true value of the parameter across different experimental conditions.

4.4.1. DEEP VS SHALLOW

We explore here the capability of a DGP to accommodate for the data nonstationarity typical of ODE systems. In particular, the tests are performed in two different settings with large and small sample size n . By using the same experimental setting of Section 4.1, we sampled 80 and 1000 points, respectively, from the FitzHugh-Nagumo equations. The data is modeled with DGPs composed by one (“shallow” GP), two and three layers, all with RBF covariances.

Figure 5 shows the modeling results obtained on the two configurations. We note that the shallow GP consistently underfits the complex dynamics producing smooth interpolants. On the contrary, DGPs provide a better representation of the nonstationarity. As expected, the three-layer DGP leads to sub-optimal results in the low-sample size

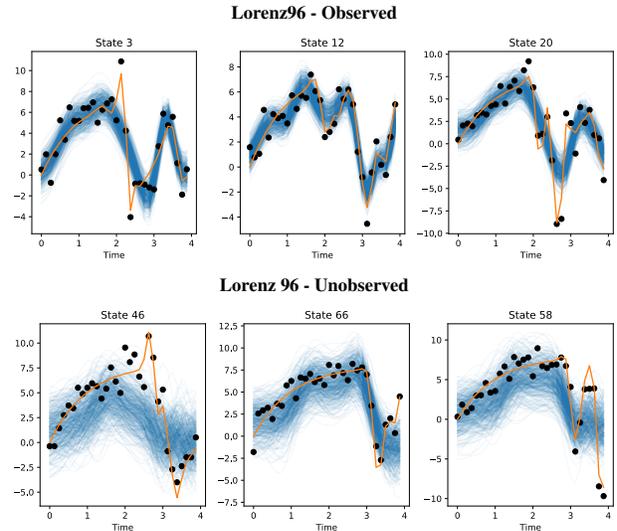


Figure 4. Model fit in Lorenz96. Randomly sampled observed (top) vs unobserved (bottom) states for $s = 125$ ODEs. Orange lines and black dots represent respectively the ground truth dynamics and noisy sample points. The blue lines are realizations of the DGP.

setting. Furthermore, in order to motivate the importance of nonstationarity, which we implement through DGPs, we further compared against shallow GPs with lower degrees of smoothness through the use of Matérn covariances with degrees $\nu = 1/2, 1, 3/2, 5/2$.

The overall performance in parameter estimation and data fit is reported in Table 1. According to the results, a two-layer DGP provides the best solution overall in terms of modeling accuracy and complexity. Interestingly, the Matérn covariance, with an appropriate degree of smoothness, achieves superior performance in parameter estimation in case of low sample size. However, the nonstationarity implemented by the DGP outperforms the stationary Matérn in the data fit, as well as in the parameter estimation when the sample size is large. For an illustration of the data fit of the Matérn GP we refer the reader to the supplementary material. Crucially, these results indicate that our approach provides a practical and scalable way to learn nonstationarity within the framework of variational inference for deep probabilistic models.

4.5. Inequality constraints

We conclude our experimental validation by applying monotonic regression on counts as an illustration of the proposed framework for inequality constraints in DGP models dynamics. We applied our approach to the mortality dataset from (Broffitt, 1988), with a two-layer DGP initialized with an analogous setting to the one proposed in Section 4.1. In

Table 1. Shallow and deep GP models under different experimental conditions in FitzHugh-Nagumo equations. Best results are highlighted in bold.

AVERAGE RMSE ACROSS PARAMETERS							
n	RBF LAYERS			MATÉRN ν			
	1	2	3	1/2	1	3/2	5/2
80	0.86	0.85	2.16	0.97	0.72	0.79	0.77
1000	0.66	0.52	0.53	0.87	0.63	0.70	0.74
DATA FIT RMSE							
80	0.23	0.19	0.42	0.23	0.20	0.23	0.25
1000	0.22	0.17	0.19	0.21	0.21	0.19	0.24

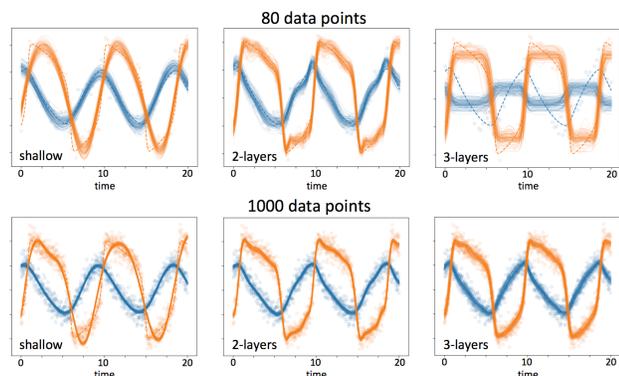


Figure 5. Modeling FitzHugh-Nagumo equations with GP and DGP. A deep model provides a more accurate description of data nonstationarity and associated dynamics (Table 1). Training points are denoted with circles; the ground truth trajectory is represented by the dashed line. Top: $N = 80$; Bottom: $N = 1000$. From left to right: Shallow GP, two-layers and three-layers DGP.

particular, the sample rates were modeled with with a Poisson likelihood of the form $p(y_i|\mu_i) = \frac{\exp(-\mu_i)\mu_i^{y_i}}{y_i!}$, and link function $\mu_i = \exp(f(t_i))$. Monotonicity on the solution was strictly enforced by setting $\psi_D = 5$. Figure 6 shows the regression results without (top) and with (bottom) monotonicity constraint. The effect of the constraint on the dynamics can be appreciated by looking at the distribution of the derivatives (right panel). In the monotonic case the GP derivatives lie on the positive part of the plane. This experiment leads to results compatible with those obtained with the monotonic GP proposed in (Riihimäki & Vehtari, 2010), and implemented in the GPstuff toolbox (Vanhatalo et al., 2013). However, our approach is characterized by appealing scalability properties and can implement monotonic constraints on DGPs, which offer a more general class of functions than GPs.

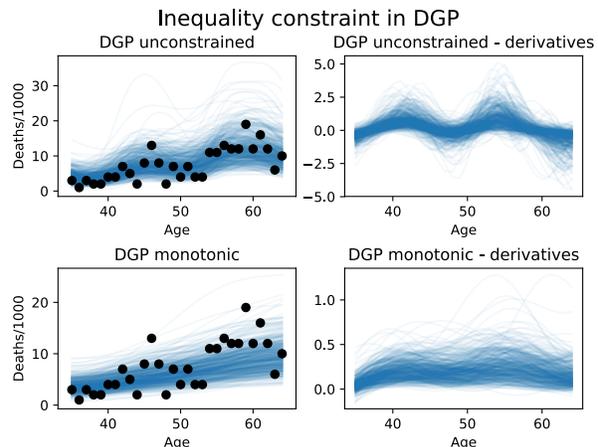


Figure 6. GP with Poisson likelihood: unconstrained (top) and monotonic (bottom). Black dots: observations from (Broffitt, 1988). Blue lines: GP realizations.

5. Conclusions

We introduced a novel generative formulation of deep probabilistic models implementing “soft” constraints on functions dynamics. The proposed approach was extensively tested in several experimental settings, leading to highly competitive results in challenging modeling applications, and favorably comparing with the state-of-the-art in terms of modeling accuracy and scalability. Furthermore, the proposed variational formulation allows for a meaningful uncertainty quantification of both model parameters and predictions. This is an important aspect intimately related to the application of our proposal in real scenarios, such as in biology and epidemiology, where data is often noisy and scarce.

Although in this study we essentially focused on the problem of ODE parameters inference and monotonic regression, the generality of our approach enables several other applications that will be subject of future investigations. We will focus on the extension to manifold valued data, such as spatio-temporal observations represented by graphs, meshes, and 3D volumes, occurring for example in medical imaging and system biology.

Acknowledgements

This work has been supported by the French government, through the UCAJEDI Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-15-IDEX-01 (project Meta-ImaGen). MF gratefully acknowledges support from the AXA Research Fund.

This work is dedicated to Mattia Filippone.

References

- Alvarez, M. A., Luengo, D., and Lawrence, N. D. Linear latent force models using Gaussian processes. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2693–2705, 2013.
- Barber, D. and Wang, Y. Gaussian Processes for Bayesian Estimation in Ordinary Differential Equations. In Xing, E. P. and Jebara, T. (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 1485–1493, Beijing, China, June 2014. PMLR.
- Broffitt, J. D. Increasing and increasing convex Bayesian graduation. *Transactions of the Society of Actuaries*, 40(1):115–48, 1988.
- Calandra, R., Peters, J., Rasmussen, C. E., and Deisenroth, M. P. Manifold Gaussian Processes for regression. In *2016 International Joint Conference on Neural Networks, IJCNN 2016, Vancouver, BC, Canada, July 24-29, 2016*, pp. 3338–3345, 2016.
- Calderhead, B., Girolami, M., and Lawrence, N. D. Accelerating Bayesian Inference over Nonlinear Differential Equations with Gaussian Processes. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems 21*, pp. 217–224. Curran Associates, Inc., 2009.
- Campbell, D. and Steele, R. J. Smooth functional tempering for nonlinear differential equation models. *Statistics and Computing*, 22(2):429–443, March 2012.
- Cutajar, K., Bonilla, E. V., Michiardi, P., and Filippone, M. Random feature expansions for deep Gaussian processes. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 884–893, International Convention Centre, Sydney, Australia, August 2017. PMLR.
- Da Veiga, S. and Marrel, A. Gaussian process modeling with inequality constraints. *Annales de la faculté des sciences de Toulouse Mathématiques*, 21(3):529–555, April 2012.
- Damianou, A. C. and Lawrence, N. D. Deep Gaussian Processes. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2013, Scottsdale, AZ, USA, April 29 - May 1, 2013*, volume 31 of *JMLR Proceedings*, pp. 207–215. JMLR.org, 2013.
- Dondelinger, F., Husmeier, D., Rogers, S., and Filippone, M. Ode parameter inference using adaptive gradient matching with gaussian processes. In Carvalho, C. M. and Ravikumar, P. (eds.), *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31 of *Proceedings of Machine Learning Research*, pp. 216–228, Scottsdale, Arizona, USA, 29 Apr–01 May 2013. PMLR.
- Donohue, M. C., Jacqmin-Gadda, H., Le Goff, M., Thomas, R. G., Raman, R., Gamst, A. C., Beckett, L. A., Jack, C. R., Weiner, M. W., Dartigues, J.-F., et al. Estimating long-term multivariate progression from short-term data. *Alzheimer’s & dementia: the journal of the Alzheimer’s Association*, 10(5):S400–S410, 2014.
- FitzHugh, R. Mathematical models of threshold phenomena in the nerve membrane. *The bulletin of mathematical biophysics*, 17(4):257–278, 1955.
- Ghahramani, Z. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459, May 2015.
- Goel, N. S., Maitra, S. C., and Montroll, E. W. On the volterra and other nonlinear models of interacting populations. *Reviews of modern physics*, 43(2):231, 1971.
- González, J., Vujačić, I., and Wit, E. Reproducing kernel Hilbert space based estimation of systems of ordinary differential equations. *Pattern Recognition Letters*, 45: 26–32, 2014.
- Gorbach, N. S., Bauer, S., and Buhmann, J. M. Scalable Variational Inference for Dynamical Systems. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 4809–4818. Curran Associates, Inc., 2017.
- Groeneboom, P. and Jongbloed, G. *Nonparametric Estimation under Shape Constraints: Estimators, Algorithms and Asymptotics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2014.
- Kingma, D. P. and Ba, J. Adam: A Method for Stochastic Optimization, January 2017. arXiv:1412.6980.
- Konukoglu, E., Relan, J., Cilingir, U., Menze, B. H., Chinchapatnam, P., Jadidi, A., Cochet, H., Hocini, M., Delingette, H., Jaïs, P., et al. Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to eikonal-diffusion models in cardiac electrophysiology. *Progress in biophysics and molecular biology*, 107(1):134–146, 2011.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Liang, H. and Wu, H. Parameter Estimation for Differential Equation Models Using a Framework of Measurement

- Error in Regression Models. *Journal of the American Statistical Association*, 103(484):1570–1583, 2008. PMID: 19956350.
- Lorenz, E. N. and Emanuel, K. A. Optimal sites for supplementary weather observations: Simulation with a small model. *Journal of the Atmospheric Sciences*, 55(3):399–414, 1998.
- Lorenzi, M., Filippone, M., Frisoni, G. B., Alexander, D. C., and Ourselin, S. Probabilistic disease progression modeling to characterize diagnostic uncertainty: Application to staging and prediction in Alzheimer’s disease. *NeuroImage*, 2017.
- Macdonald, B. and Husmeier, D. Gradient Matching Methods for Computational Inference in Mechanistic Models for Systems Biology: A Review and Comparative Analysis. *Frontiers in Bioengineering and Biotechnology*, 3: 180, 2015.
- Macdonald, B., Higham, C., and Husmeier, D. Controversy in mechanistic modelling with Gaussian processes. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1539–1547, Lille, France, July 2015. PMLR.
- Mašić, A., Srinivasan, S., Billeter, J., Bonvin, D., and Villez, K. Shape constrained splines as transparent black-box models for bioprocess modeling. *Computers & Chemical Engineering*, 99:96–105, 2017.
- Meyer, M. C. Inference using shape-restricted regression splines. *The Annals of Applied Statistics*, 2(3):1013–1033, 2008.
- Minka, T. P. Expectation Propagation for approximate Bayesian inference. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, UAI ’01, pp. 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- Neal, R. M. *Bayesian Learning for Neural Networks (Lecture Notes in Statistics)*. Springer, 1 edition, August 1996.
- Niu, M., Rogers, S., Filippone, M., and Husmeier, D. Fast Parameter Inference in Nonlinear Dynamical Systems using Iterative Gradient Matching. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1699–1707, New York, New York, USA, June 2016. PMLR.
- Niu, M., Macdonald, B., Rogers, S., Filippone, M., and Husmeier, D. Statistical inference in mechanistic models: time warping for improved gradient matching. *Computational Statistics*, 33(2):1091–1123, Jun 2018.
- Rahimi, A. and Recht, B. Random Features for Large-Scale Kernel Machines. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T. (eds.), *Advances in Neural Information Processing Systems 20*, pp. 1177–1184. Curran Associates, Inc., 2008.
- Ramsay, J. O., Hooker, G., Campbell, D., and Cao, J. Parameter estimation for differential equations: a generalized smoothing approach. *Journal of the Royal Statistical Society Series B*, 69(5):741–796, 2007.
- Rasmussen, C. E. and Williams, C. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Riihimäki, J. and Vehtari, A. Gaussian processes with monotonicity information. In Teh, Y. W. and Titterton, M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 645–652, Chia Laguna Resort, Sardinia, Italy, May 2010. PMLR.
- Salzmann, M. and Urtasun, R. Implicitly Constrained Gaussian Process Regression for Monocular Non-Rigid Pose Estimation. In Lafferty, J. D., Williams, C. K. I., Shawe-Taylor, J., Zemel, R. S., and Culotta, A. (eds.), *Advances in Neural Information Processing Systems 23*, pp. 2065–2073. Curran Associates, Inc., 2010.
- Schober, M., Duvenaud, D. K., and Hennig, P. Probabilistic ODE solvers with Runge-Kutta means. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 739–747. Curran Associates, Inc., 2014.
- Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., and Vehtari, A. GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research*, 14(Apr):1175–1179, 2013.
- Varah, J. M. A Spline Least Squares Method for Numerical Parameter Estimation in Differential Equations. *SIAM Journal on Scientific and Statistical Computing*, 3(1):28–46, 1982.
- Vrettas, M. D., Opper, M., and Cornford, D. Variational mean-field algorithm for efficient inference in large systems of stochastic differential equations. *Physical Review E*, 91(1):012148, 2015.
- Vysheirsky, V. and Girolami, M. A. Bayesian ranking of biochemical system models. *Bioinformatics*, 24(6): 833–839, 2007.
- Wheeler, M. W., Dunson, D. B., Pandalai, S. P., Baker, B. A., and Herring, A. H. Mechanistic hierarchical Gaussian processes. *Journal of the American Statistical Association*, 109(507):894–904, 2014.