
Learning unknown ODE models with Gaussian processes

Markus Heinonen^{*1,2} Çagatay Yıldız^{*1} Henrik Mannerström¹ Jukka Intosalmi¹ Harri Lähdesmäki¹

Abstract

In conventional ODE modelling coefficients of an equation driving the system state forward in time are estimated. However, for many complex systems it is practically impossible to determine the equations or interactions governing the underlying dynamics. In these settings, parametric ODE model cannot be formulated. Here, we overcome this issue by introducing a novel paradigm of non-parametric ODE modelling that can learn the underlying dynamics of arbitrary continuous-time systems without prior knowledge. We propose to learn non-linear, unknown differential functions from state observations using Gaussian process vector fields within the exact ODE formalism. We demonstrate the model’s capabilities to infer dynamics from sparse data and to simulate the system forward into future.

1. Introduction

Dynamical systems modelling is a cornerstone of experimental sciences. In biology, as well as in physics and chemistry, modelers attempt to capture the dynamical behavior of a given system or a phenomenon in order to improve its understanding and make predictions about its future state. Systems of coupled ordinary differential equations (ODEs) are undoubtedly the most widely used models in science. Even simple ODE functions can describe complex dynamical behaviours (Hirsch et al., 2004). Typically, the dynamics are firmly grounded in physics with only a few parameters to be estimated from data. However, equally ubiquitous are the cases where the governing dynamics are partially or completely unknown.

We consider the dynamics of a system governed by multi-

^{*}Equal contribution ¹Aalto University, Finland ²Helsinki Institute of Information Technology HIIT, Finland. Correspondence to: Markus Heinonen <markus.o.heinonen@aalto.fi>.

variate ordinary differential functions:

$$\dot{\mathbf{x}}(t) = \frac{d\mathbf{x}(t)}{dt} = \mathbf{f}(\mathbf{x}(t)) \quad (1)$$

where $\mathbf{x}(t) \in \mathcal{X} = \mathbb{R}^D$ is the state vector of a D -dimensional dynamical system at time t , and the $\dot{\mathbf{x}}(t) \in \dot{\mathcal{X}} = \mathbb{R}^D$ is the first order time derivative of $\mathbf{x}(t)$ that drives the state $\mathbf{x}(t)$ forward, and where $\mathbf{f} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is the vector-valued derivative function. The ODE solution is determined by

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_0^t \mathbf{f}(\mathbf{x}(\tau))d\tau, \quad (2)$$

where we integrate the system state from an initial state $\mathbf{x}(0) = \mathbf{x}_0$ for time t forward. We assume that $\mathbf{f}(\cdot)$ is *completely unknown* and we only observe one or several multivariate time series $Y = (\mathbf{y}_1, \dots, \mathbf{y}_N)^T \in \mathbb{R}^{N \times D}$ obtained from an additive noisy observation model at observation time points $T = (t_1, \dots, t_N) \in \mathbb{R}^N$,

$$\mathbf{y}(t) = \mathbf{x}(t) + \varepsilon_t, \quad (3)$$

where $\varepsilon_t \sim \mathcal{N}(\mathbf{0}, \Omega)$ follows a stationary zero-mean multivariate Gaussian distribution with diagonal noise variances $\Omega = \text{diag}(\omega_1^2, \dots, \omega_D^2)$. The observation time points do not need to be equally spaced. Our task is to learn the differential function $\mathbf{f}(\cdot)$ given observations Y , with no prior knowledge of the ODE system.

There is a vast literature on conventional ODEs (Butcher, 2016) where a parametric form for function $\mathbf{f}(\mathbf{x}; \boldsymbol{\theta}, t)$ is assumed to be known, and its parameters $\boldsymbol{\theta}$ are subsequently optimised with least squares or Bayesian approach, where the expensive forward solution $\mathbf{x}_{\boldsymbol{\theta}}(t_i) = \int_0^{t_i} \mathbf{f}(\mathbf{x}(\tau); \boldsymbol{\theta}, t)d\tau$ is required to evaluate the system responses $\mathbf{x}_{\boldsymbol{\theta}}(t_i)$ from parameters $\boldsymbol{\theta}$ against observations $\mathbf{y}(t_i)$. To overcome the computationally intensive forward solution, a family of methods denoted as gradient matching (Varah, 1982; Ellner et al., 2002; Ramsay et al., 2007) have proposed to replace the forward solution by matching $\mathbf{f}(\mathbf{y}_i) \approx \dot{\mathbf{y}}_i$ to empirical gradients $\dot{\mathbf{y}}_i$ of the data instead, which do not require the costly integration step. Recently several authors have proposed embedding a parametric differential function within a Bayesian or Gaussian process (GP) framework (Graepel, 2003; Calderhead et al., 2008;

Dondelinger et al., 2013; Wang and Barber, 2014; Macdonald, 2017) (see Macdonald et al. (2015) for a review). GPs have been successfully applied to model linear differential equations as they are analytically tractable (Gao et al., 2008; Raissi et al., 2017).

However, conventional ODE modelling can only proceed if a parametric form of the driving function $f(\cdot)$ is known. Recently, initial work to handle unknown or non-parametric ODE models have been proposed, although with various limiting approximations. Early works include spline-based smoothing and additive functions $\sum_j^D f_j(x_j)$ to infer gene regulatory networks (De Hoon et al., 2002; Henderson and Michailidis, 2014). Äijö and Lähdesmäki (2009) proposed estimating the unknown nonlinear function with GPs using either finite time differences, or analytically solving the derivative function as a function of only time, $\dot{\mathbf{x}}(t) = \mathbf{f}(t)$ (Äijö et al., 2013). In a seminal technical report of Heinonen and d’Alche Buc (2014) a full vector-valued kernel model $\mathbf{f}(\mathbf{x})$ was proposed, however using a gradient matching approximation. To our knowledge, there exists no model that can learn non-linear ODE functions $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t))$ over the state \mathbf{x} against the true forward solutions $\mathbf{x}(t_i)$.

In this work we propose NPODE¹: the first ODE model for learning arbitrary, and *a priori* completely unknown non-parametric, non-linear differential functions $\mathbf{f} : \mathcal{X} \rightarrow \dot{\mathcal{X}}$ from data in a Bayesian way. We do not use gradient matching or other approximative models, but instead propose to directly optimise the exact ODE system with the fully forward simulated responses against data. We parameterise our model as an augmented Gaussian process vector field with inducing points, while we propose sensitivity equations to efficiently compute the gradients of the system. Our model can forecast continuous-time systems arbitrary amounts to future, and we demonstrate the state-of-the-art performance in human motion datasets.

2. Nonparametric ODE Model

The differential function $\mathbf{f}(\mathbf{x})$ to be learned defines a *vector field*² \mathbf{f} , that is, an assignment of a gradient vector $\mathbf{f}(\mathbf{x}) \in \mathbb{R}^D$ to every state $\mathbf{x} \in \mathbb{R}^D$. We model the vector field as a vector-valued Gaussian process (Rasmussen and Williams, 2006)

$$\mathbf{f}(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, K(\mathbf{x}, \mathbf{x}')), \quad (4)$$

which defines *a priori* distribution over function values $\mathbf{f}(\mathbf{x})$ whose mean and covariances are

$$\mathbb{E}[\mathbf{f}(\mathbf{x})] = \mathbf{0} \quad (5)$$

$$\text{cov}[\mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{x}')] = K(\mathbf{x}, \mathbf{x}'), \quad (6)$$

¹The implementation is publicly available in <http://www.github.com/cagatayyildiz/npode>

²We use vector field and differential function interchangeably.

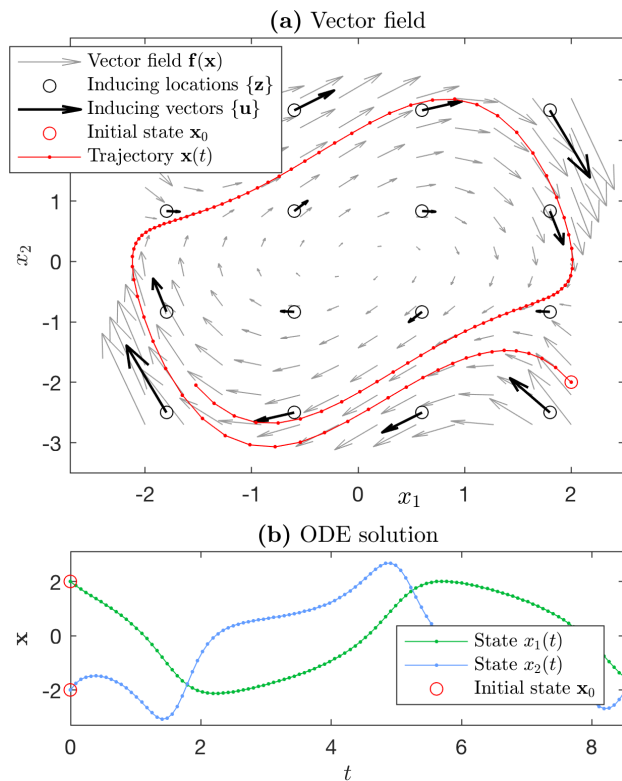


Figure 1. (a) Illustration of an ODE system vector field induced by the Gaussian process. The vector field $\mathbf{f}(\mathbf{x})$ (gray arrows) at arbitrary states \mathbf{x} is interpolated from the inducing points \mathbf{u}, \mathbf{z} (black arrows), with the trajectory $\mathbf{x}(t)$ (red points) following the differential system $\mathbf{f}(\mathbf{x})$ exactly. (b) The trajectory $\mathbf{x}(t)$ plotted over time t .

and where the kernel $K(\mathbf{x}, \mathbf{x}') \in \mathbb{R}^{D \times D}$ is matrix-valued. A GP prior defines that for any collection of states $X = (\mathbf{x}_1, \dots, \mathbf{x}_N)^T \in \mathbb{R}^{N \times D}$, the function values $F = (\mathbf{f}(\mathbf{x}_1), \dots, \mathbf{f}(\mathbf{x}_N))^T \in \mathbb{R}^{N \times D}$ follow a matrix-valued normal distribution,

$$p(F) = \mathcal{N}(\text{vec}(F) | \mathbf{0}, \mathbf{K}(X, X)), \quad (7)$$

where $\mathbf{K}(X, X) = (K(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^N \in \mathbb{R}^{ND \times ND}$ is a block matrix of matrix-valued kernels $K(\mathbf{x}_i, \mathbf{x}_j)$. The key property of Gaussian processes is that they encode functions where similar states \mathbf{x}, \mathbf{x}' induce similar differentials $\mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{x}')$, and where the state similarity is defined by the kernel $K(\mathbf{x}, \mathbf{x}')$.

In standard GP regression we would obtain the posterior of the vector field by conditioning the GP prior with the data (Rasmussen and Williams, 2006). In ODE models the conditional $\mathbf{f}(\mathbf{x})|Y$ of a vector field is intractable due to the integral mapping (2) between observed states $\mathbf{y}(t_i)$ and differentials $\mathbf{f}(\mathbf{x})$. Instead, we resort to augmenting the Gaussian process with a set of M inducing points $\mathbf{z} \in \mathcal{X}$ and $\mathbf{u} \in \dot{\mathcal{X}}$, such that $\mathbf{f}(\mathbf{z}) = \mathbf{u}$ (Quiñero-Candela and

Rasmussen, 2005). We choose to interpolate the differential function between the inducing points as (See Figure 1)

$$\mathbf{f}(\mathbf{x}) \triangleq \mathbf{K}_\theta(\mathbf{x}, Z)\mathbf{K}_\theta(Z, Z)^{-1}\text{vec}(U), \quad (8)$$

which supports the function $\mathbf{f}(\mathbf{x})$ with *inducing locations* $Z = (\mathbf{z}_1, \dots, \mathbf{z}_M)$, *inducing vectors* $U = (\mathbf{u}_1, \dots, \mathbf{u}_M)$, and θ are the kernel parameters. The function above corresponds to a vector-valued kernel function (Alvarez et al., 2012), or to a multi-task Gaussian process conditional mean without the variance term (Rasmussen and Williams, 2006). This definition is then compatible with the deterministic nature of the ODE formalism. Due to universality of several kernels and kernel functions (Shawe-Taylor and Cristianini, 2004), we can represent arbitrary vector fields with appropriate inducing point and kernel choices.

2.1. Operator-valued Kernels

The vector-valued kernel function (8) uses *operator-valued* kernels, which result in matrix-valued kernels $K_\theta(\mathbf{z}, \mathbf{z}') \in \mathbb{R}^{D \times D}$ for real valued states \mathbf{x}, \mathbf{z} , while the kernel matrix over data points becomes $\mathbf{K}_\theta = (K(\mathbf{z}_i, \mathbf{z}_j))_{i,j=1}^M \in \mathbb{R}^{MD \times MD}$ (See Alvarez et al. (2012) for a review). Most straightforward operator-valued kernel is the identity decomposable kernel $K_{dec}(\mathbf{z}, \mathbf{z}') = k(\mathbf{z}, \mathbf{z}') \cdot I_D$, where the scalar Gaussian kernel

$$K_\theta(\mathbf{z}, \mathbf{z}') = \sigma_f^2 \exp\left(-\frac{1}{2} \sum_{j=1}^D \frac{(z_j - z'_j)^2}{\ell_j^2}\right) \quad (9)$$

with differential variance σ_f^2 and dimension-specific lengthscales $\ell = (\ell_1, \dots, \ell_D)$ are expanded into a diagonal matrix of size $D \times D$. We collect the kernel parameters as $\theta = (\sigma_f, \ell)$.

We note that more complex kernels can also be considered given prior information of the underlying system characteristics. The divergence-free matrix-valued kernel induces vector fields that have zero divergence (Wahlström et al., 2013; Solin et al., 2015). Intuitively, these vector fields do not have sinks or sources, and every state always finally returns to itself after sufficient amount of time. Similarly, curl-free kernels induce curl-free vector fields that can contain sources or sinks, that is, trajectories can accelerate or decelerate. For theoretical treatment of vector field kernels, see (Narcowich and Ward, 1994; Bhatia et al., 2013; Fuselier and Wright, 2017). Non-stationary vector fields can be modeled with input-dependent lengthscales (Heinonen et al., 2016), while spectral kernels can represent stationary (Wilson et al., 2013) or non-stationary (Remes et al., 2017) recurring patterns in the differential function.

2.2. Joint Model

We assume a Gaussian likelihood over the observations \mathbf{y}_i and the corresponding simulated responses $\mathbf{x}(t_i)$ of Equation (2),

$$p(Y|\mathbf{x}_0, U, Z, \omega) = \prod_{i=1}^N \mathcal{N}(\mathbf{y}_i|\mathbf{x}(t_i), \Omega), \quad (10)$$

where $\mathbf{x}(t_i)$ are forward simulated responses using the integral Equation (2) and differential Equation (8), and $\Omega = \text{diag}(\omega_1^2, \dots, \omega_D^2)$ collects the dimension-specific noise variances.

The inducing vectors have a Gaussian process prior

$$p(U|Z, \theta) = \mathcal{N}(\text{vec}(U)|\mathbf{0}, \mathbf{K}_\theta(Z, Z)). \quad (11)$$

The model posterior is then

$$p(U, \mathbf{x}_0, \theta, \omega|Y) \propto p(Y|\mathbf{x}_0, U, \omega)p(U|\theta) = \mathcal{L}, \quad (12)$$

where we have for brevity omitted the dependency on the locations of the inducing points Z and also the parameter hyperpriors $p(\theta)$ and $p(\omega)$ since we assume them to be uniform, unless there is specific domain knowledge of the priors.

The model parameters are the initial state \mathbf{x}_0^3 , the inducing vectors U , the noise standard deviations $\omega = (\omega_1, \dots, \omega_D)$, and the kernel hyperparameters $\theta = (\sigma_f, \ell_1, \dots, \ell_D)$.

2.3. Noncentral Parameterisation

We apply a latent parameterisation using Cholesky decomposition $\mathbf{L}_\theta \mathbf{L}_\theta^T = \mathbf{K}_\theta(Z, Z)$, which maps the inducing vectors to whitened domain (Kuss and Rasmussen, 2005)

$$U = \mathbf{L}_\theta \tilde{U}, \quad \tilde{U} = \mathbf{L}_\theta^{-1} U. \quad (13)$$

The latent variables \tilde{U} are projected on the kernel manifold \mathbf{L}_θ to obtain the inducing vectors U . This non-centered parameterisation (NCP) transforms the hierarchical posterior \mathcal{L} of Equation (12) into a reparameterised form

$$p(\mathbf{x}_0, \tilde{U}, \theta, \omega|Y) \propto p(Y|\mathbf{x}_0, \tilde{U}, \omega, \theta)p(\tilde{U}), \quad (14)$$

where all variables to be optimised are decoupled, with the latent inducing vectors having a standard normal prior $\tilde{U} \sim \mathcal{N}(\mathbf{0}, I)$. Optimizing \tilde{U} and θ is now more efficient since they have independent contributions to the vector field via $U = \mathbf{L}_\theta \tilde{U}$. The gradients of the whitened posterior can be retrieved analytically as (Heinonen et al., 2016)

$$\nabla_{\tilde{U}} \log \mathcal{L} = \mathbf{L}_\theta^T \nabla_U \log \mathcal{L}. \quad (15)$$

³In case of multiple time-series, we will use one initial state for each time-series.

Finally, we find a maximum a posteriori (MAP) estimate for the initial state \mathbf{x}_0 , latent vector field \tilde{U} , kernel parameters θ and noise variances ω by gradient ascent,

$$\mathbf{x}_{0,\text{MAP}}, \tilde{U}_{\text{MAP}}, \theta_{\text{MAP}}, \omega_{\text{MAP}} = \arg \max_{\mathbf{x}_0, \tilde{U}, \theta, \omega} \log \mathcal{L}, \quad (16)$$

while keeping the inducing locations Z fixed on a sufficiently dense grid (See Figure 1). The partial derivatives of the posterior with respect to noise parameters ω can be found analytically, while the derivative with respect to σ_f is approximated with finite differences. We select the optimal lengthscales ℓ by cross-validation.

3. Sensitivity Equations

The key term to carry out the MAP gradient ascent optimization is the likelihood

$$\log p(Y|\mathbf{x}_0, \tilde{U}, \omega)$$

that requires forward integration and computing the partial derivatives with respect to the whitened inducing vectors \tilde{U} . Given Equation (15) we only need to compute the gradients with respect to the inducing vectors $\mathbf{u} = \text{vec}(U) \in \mathbb{R}^{MD}$,

$$\begin{aligned} & \frac{d \log p(Y|\mathbf{x}_0, \mathbf{u}, \omega)}{d\mathbf{u}} \\ &= \sum_{s=1}^N \frac{d \log \mathcal{N}(\mathbf{y}_s | \mathbf{x}(t_s, \mathbf{u}), \Omega)}{d\mathbf{x}} \frac{d\mathbf{x}(t_s, \mathbf{u})}{d\mathbf{u}}. \end{aligned} \quad (17)$$

This requires computing the derivatives of the simulated system response $\mathbf{x}(t, \mathbf{u})$ against the vector field parameters \mathbf{u} ,

$$\frac{d\mathbf{x}(t, \mathbf{u})}{d\mathbf{u}} \equiv S(t) \in \mathbb{R}^{D \times MD}, \quad (18)$$

which we denote by $S_{ij}(t) = \frac{\partial \mathbf{x}(t, \mathbf{u})_i}{\partial u_j}$, and expand the notation to make the dependency of \mathbf{x} on \mathbf{u} explicit. Approximating these with finite differences is possible in principle, but is highly inefficient and has been reported to cause instability (Raue et al., 2013). We instead turn to sensitivity equations for \mathbf{u} and \mathbf{x}_0 that provide computationally efficient, analytical gradients $S(t)$ (Kokotovic and Heller, 1967; Fröhlich et al., 2017).

The solution for $\frac{d\mathbf{x}(t, \mathbf{u})}{d\mathbf{u}}$ can be derived by differentiating the full nonparametric ODE system with respect to \mathbf{u} by

$$\frac{d}{d\mathbf{u}} \frac{d\mathbf{x}(t, \mathbf{u})}{dt} = \frac{d}{d\mathbf{u}} \mathbf{f}(\mathbf{x}(t, \mathbf{u})). \quad (19)$$

The sensitivity equation for the given system can be obtained by changing the order of differentiation on the left hand side and carrying out the differentiation on the right hand side.

The resulting sensitivity equation can then be expressed in the form

$$\frac{d}{dt} \frac{d\mathbf{x}(t, \mathbf{u})}{d\mathbf{u}} = \frac{J(t)}{\partial \mathbf{x}} \frac{S(t)}{d\mathbf{u}} + \frac{R(t)}{\partial \mathbf{u}}, \quad (20)$$

where $J(t) \in \mathbb{R}^{D \times D}$, $R(t), \dot{S}(t) \in \mathbb{R}^{D \times MD}$ (See Supplements for detailed specification). For our nonparametric ODE system the sensitivity equation is fully determined by

$$J(t) = \frac{\partial \mathbf{K}(\mathbf{x}, Z)}{\partial \mathbf{x}} \mathbf{K}(Z, Z)^{-1} \mathbf{u} \quad (21)$$

$$R(t) = \mathbf{K}(\mathbf{x}, Z) \mathbf{K}(Z, Z)^{-1}. \quad (22)$$

The sensitivity equation provides us with an additional ODE system which describes the time evolution of the derivatives with respect to the inducing vectors $S(t)$. The sensitivities are coupled with the actual ODE system and, thus both systems $\mathbf{x}(t)$ and $S(t)$ are concatenated as the new augmented state that is solved jointly by Equation (2) driven by the differentials $\dot{\mathbf{x}}(t)$ and $\dot{S}(t)$ (Leis and Kramer, 1988). The initial sensitivities are computed as $S(0) = \frac{d\mathbf{x}_0}{d\mathbf{u}}$. In our implementation, we merge \mathbf{x}_0 with \mathbf{u} for sensitivity analysis to obtain the partial derivatives with respect to the initial state which is estimated along with the other parameters. We use the CVODES solver from the SUNDIALS package (Hindmarsh et al., 2005) to solve the nonparametric ODE models and the corresponding gradients numerically. The sensitivity equation based approach is superior to the finite differences approximation because we have exact formulation for the gradients of state over inducing points, which can be solved up to the numerical accuracy of the ODE solver.

4. Simple Simulated Dynamics

As first illustration of the proposed nonparametric ODE method we consider three simulated differential systems: the Van der Pol (VDP), FitzHugh-Nagumo (FHN) and Lotka-Volterra (LV) oscillators of form

$$\begin{aligned} \text{VDP : } \quad & \dot{x}_1 = x_2 & \dot{x}_2 &= (1 - x_1^2)x_2 - x_1 \\ \text{FHN : } \quad & \dot{x}_1 = 3(x_1 - \frac{x_1^3}{3} + x_2) & \dot{x}_2 &= \frac{0.2 - 3x_1 - 0.2x_2}{3} \\ \text{LV : } \quad & \dot{x}_1 = 1.5x_1 - x_1x_2 & \dot{x}_2 &= -3x_2 + x_1x_2. \end{aligned}$$

In the conventional ODE case the coefficients of these equations can be inferred using standard statistical techniques if sufficient amount of time series data is available (Girolami, 2008; Raue et al., 2013). Our main goal is to infer *unknown* dynamics, that is, when these equations are unavailable and we instead represent the dynamics with a nonparametric

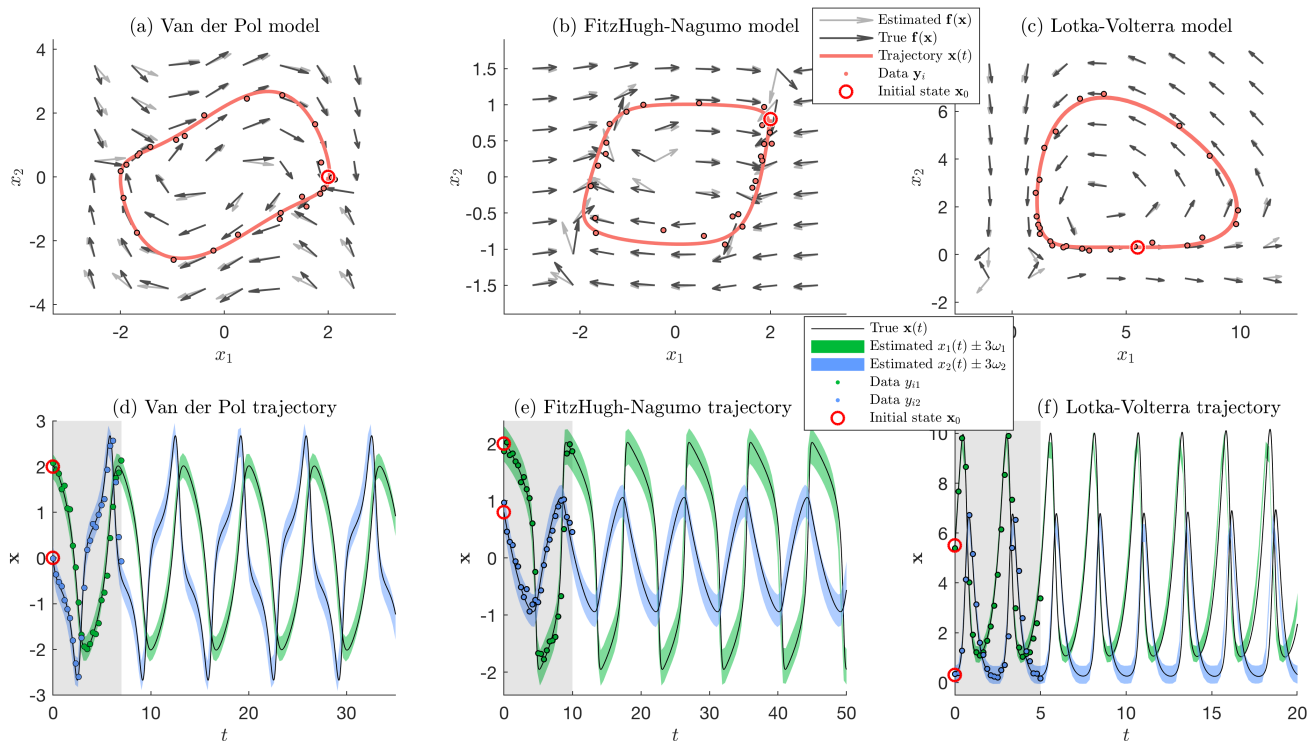


Figure 2. Estimated dynamics from Van der Pol, FitzHugh-Nagumo and Lotka-Volterra systems. The top part (a-c) shows the learned vector field (grey arrows) against the true vector field (black arrows). The bottom part (d-f) shows the training data (grey region points) and forecasted future cycle likelihoods with the learned model (shaded region) against the true trajectory (black line).

vector field of Equation (8). We use these simulated models to only illustrate our model behavior against the true dynamics.

We employ 25 data points from one cycle of noisy observation data from VDP and FHN models, and 25 data points from 1.7 cycles from the LV model with a noise variance of $\sigma_n^2 = 0.1^2$. We learn the nODE model with five training sequences using $M = 6^2$ inducing locations on a fixed grid, and forecast between 4 and 8 future cycles starting from true initial state x_0 at time 0. Training takes approximately 100 seconds per oscillator. Figure 2 (bottom) shows the training datasets (grey regions), initial states, true trajectories (black lines) and the forecasted trajectory likelihoods (colored regions). The model accurately learns the dynamics from less than two cycles of data and can reproduce them reliably into future.

Figure 2 (top) shows the corresponding true vector field (black arrows) and the estimated vector field (grey arrows). The vector field is a continuous function, which is plotted on a 8×8 grid for visualisation. In general the most difficult part of the system is learning the middle of the loop (as seen in the FHN model), and learning the most outermost regions (bottom left in the LV model). The model learns the

underlying differential $f(x)$ accurately close to observed points, while making only few errors in the border regions with no data.

5. Unknown System Estimation

Next, we illustrate how the model estimates realistic, unknown dynamics from noisy observations $y(t_1), \dots, y(t_N)$. As in Section 4, we make no assumptions on the structure or form of the underlying system, and capture the underlying dynamics with the nonparametric system alone. We employ no subjective priors, and assume no inputs, controls or other sources of information. The task is to infer the underlying dynamics $f(x)$, and interpolate or extrapolate the state trajectory outside the observed data.

We use a benchmark dataset of human motion capture data from the Carnegie Mellon University motion capture (CMU mocap) database. Our dataset contains 50-dimensional pose measurements $y(t_i)$ from humans walking, where each pose dimension records a measurement in different parts of the body during movement (Wang et al., 2008). We apply the preprocessing of Wang et al. (2008) by downsampling the datasets by a factor of four and centering the data. This resulted in a total of 4303 datapoints spread across 43 trajec-

tories with on average 100 frames per trajectory. In order to tackle the problem of dimensionality, we project the original dataset with PCA to a three dimensional latent space where the system is specified, following [Damianou et al. \(2011\)](#) and [Wang et al. \(2006\)](#). We place $M = 5^3$ inducing vectors on a fixed grid, and optimize our model starting from 100 different initial values, which we set by perturbing the projected empirical differences $\mathbf{y}(t_i) - \mathbf{y}(t_{i-1})$ to the inducing vectors. We use an L-BFGS optimizer in Matlab. The whole inference takes approximately few minutes per trajectory.

We evaluate the method with two types of experiments: imputing missing values and forecasting future cycles. For the forecasting the first half of the trajectory is reserved for model training, and the second half is to be forecasted. For imputation we remove roughly 20% of the frames from the middle of the trajectory, which are to be filled by the models. We perform model selection for lengthscales ℓ with cross-validation split of 80/20. We record the root mean square error (RMSE) over test points in the original feature space in both cases, where we reconstruct the original dimensions from the latent space trajectories.

Due to the current lack of ODE methods suitable for this non-parametric inference task, we instead compare our method to the state-of-the-art state-space models where such problems have been previously considered ([Wang et al., 2008](#)). In a state-space or dynamical model a transition function $\mathbf{x}(t_{k+1}) = \mathbf{g}(\mathbf{x}(t_k))$ moves the system forward in discrete steps. With sufficiently high sampling rate, such models can estimate and forecast finite approximations of smooth dynamics. In Gaussian process dynamical model ([Wang et al., 2006](#); [Frigola et al., 2014](#); [Svensson et al., 2016](#)) a GP transition function is inferred in a latent space, which can be inferred with a standard GPLVM ([Lawrence, 2004](#)) or with a dependent GPLVM ([Zhao and Sun, 2016](#)). In dynamical systems the transition function is replaced by a GP interpolation ([Damianou et al., 2011](#)). The discrete time state-space models emphasize inference of a low-dimensional manifold as an explanation of the high-dimensional measurement trajectories.

We compare our method to the dynamical model GPDM of [Wang et al. \(2006\)](#) and to the dynamical system VGPLVM of [Damianou et al. \(2011\)](#), where we directly apply the implementations provided by the authors at inverseprobability.com/vargplvm and dgp.toronto.edu/~jmwang/gpdm. Both methods optimize their latent spaces separately, and they are thus not directly comparable.

5.1. Forecasting

In the forecasting task we train all models with the first half of the trajectory, while forecasting the second half starting from the first frame. The models are trained and forecasted

Table 1. Means and standard deviations of RMSEs of 43 datasets in forecasting and filling experiments.

MODEL	FORECASTING	IMPUTATION
npODE	4.52 \pm 2.31	3.94 \pm 3.50
GPDM	4.94 \pm 3.3	5.31 \pm 3.39
VGPLVM	8.74 \pm 3.43	3.91 \pm 1.80

within a low-dimensional space, and subsequently projected back into the original space via inverting the PCA or with GPLVM mean predictions. As all methods optimize their latent spaces separately, they are not directly comparable. Thus, the mean errors are computed in the original high-dimensional space. Note that the low-dimensional representation necessarily causes some reconstruction errors.

Figure 3 illustrates the models on one of the trajectories 35_12.amc. The top part (a) shows the training data in the PCA space for npODE, and optimized training data representation for GPDM and VGPLVM (black points). The colored lines (npODE) and points (GPDM, VGPLVM) indicate the future forecast. The bottom part (b) shows the first 9 reconstructed original pose dimensions reconstructed from the latent forecasted trajectories. The training data is shown in grey background, while test data is shown with circles.

The VGPLVM has most trouble forecasting future points, and reverts quickly after training data to a value close to zero, failing to predict future points. The GPDM model produces more realistic trajectories, but fails to predict any of the poses accurately. Finally, npODE can accurately predict five poses, and still retains adequate performance on remaining poses, except for pose 2.

Furthermore, Table 1 indicates that npODE is also best performing method on average over the whole dataset in the forecasting.

5.2. Imputation

In the imputation task we remove approximately 20% of the training data from the middle of the trajectory. The goals are to learn a model with the remaining data and to forecast the missing values. Figure 4 highlights the performance of the three models on the trajectory 07_07.amc. The top part (a) shows the training data (black points) in the PCA space (npODE) or optimized training locations in the latent space (GPDM, VGPLVM). The middle part imputation is shown with colored points or lines. Interestingly both npODE and GPDM operate on cyclic representations, while VGPLVM is not cyclic.

The bottom panel (b) shows the first 9 reconstructed pose

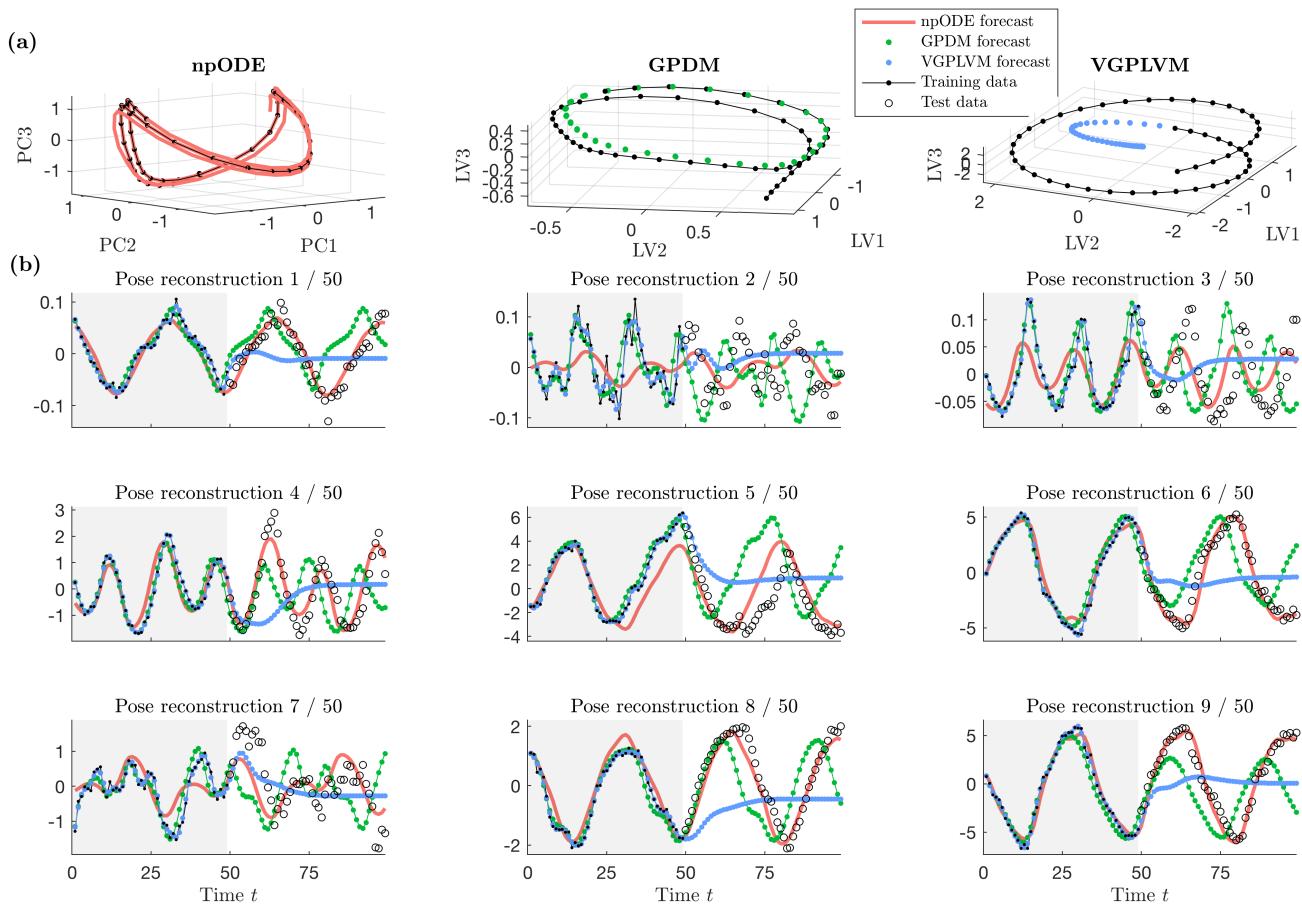


Figure 3. Forecasting 50 future frames after 49 frames of training data of human motion dataset 35_12.amc. (a) The estimated locations of the trajectory in a latent space (black points) and future forecast (colored lines). (b) The original features reconstructed from the latent predictions with grey region showing the training data.

dimensions from the three models. The missing values are shown in circles, while training points are shown with black dots. All models can accurately reproduce the overall trends, while npODE seems to fit slightly worse than the other methods. The PCA projection causes the seemingly perfect fit of the npODE prediction (at the top) to lead to slightly warped reconstructions (at the bottom). All methods mostly fit the missing parts as well. Table 1 shows that on average the npODE and VGPLVM have approximately equal top performance on the imputing missing values task.

6. Discussion

We proposed the framework of nonparametric ODE model that can accurately learn arbitrary, nonlinear continuous-time dynamics from purely observational data without making assumptions of the underlying system dynamics. We demonstrated that the model excels at learning dynamics that can be forecasted into the future. We consider this work as the

first in a line of studies of nonparametric ODE systems, and foresee several aspects as future work. Currently we do not handle non-stationary vector fields, that is time-dependent differentials $\mathbf{f}_t(\mathbf{x})$. Furthermore, an interesting future avenue is the study of various vector field kernels, such as divergence-free, curl-free or spectral kernels (Remes et al., 2017). Finally, including inputs or controls to the system would allow precise modelling in interactive settings, such as robotics.

The proposed nonparametric ODE model operates along a continuous-time trajectory, while dynamic models such as hidden Markov models or state-space models are restricted to discrete time steps. These models are unable to consider system state at arbitrary times, for instance, between two successive timepoints.

Conventional ODE models have also been considered from the stochastic perspective with stochastic differential equation (SDE) models that commonly model the deterministic

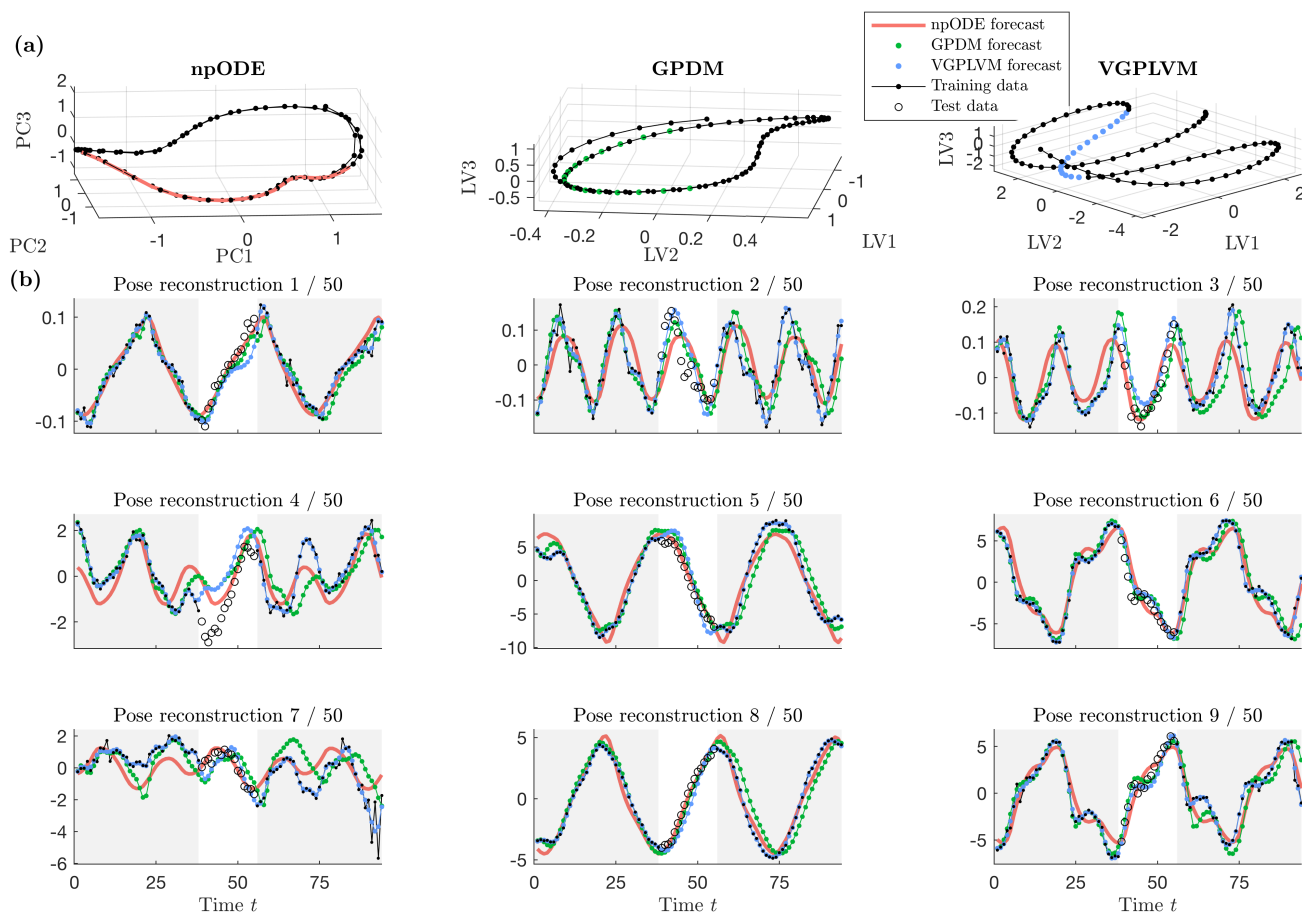


Figure 4. Imputation of 17 missing frames in the middle of a 94-length trajectory of human motion dataset 07_07 .amc (subsampled every fourth frame). (a) The estimated locations of the missing points in the latent space are colored. (b) The original features reconstructed from the latent trajectory.

system drift and diffusion processes separately leading to a distribution of trajectories $p(\mathbf{x}(t))$ (Archambeau et al., 2007; García et al., 2017). As future work we will consider stochastic extensions of our nonparametric ODE model, as well as MCMC sampling of the inducing point posterior $p(U|Y)$, leading to trajectory distribution as well.

Acknowledgements. The data used in this project was obtained from mocap.cs.cmu.edu. The database was created with funding from NSF EIA-0196217. This work has been supported by the Academy of Finland Center of Excellence in Systems Immunology and Physiology, the Academy of Finland grants no. 284597, 311584, 313271, 299915.

References

- Tarmo Äijö, Kirsi Granberg, and Harri Lähdesmäki. Sorrad: a systems biology approach to predict and modulate dynamic signaling pathway response from phosphoproteome time-course measurements. *Bioinformatics*, 29(10):1283–1291, 2013.
- M. Alvarez, L. Rosasco, and N. Lawrence. Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning*, 2012.
- C. Archambeau, D. Cornford, M. Opper, and J. Shawe-Taylor. Gaussian process approximations of stochastic differential equations. In *Gaussian Processes in Practice*, 2007.
- H. Bhatia, G. Norgard, V. Pascucci, and P. Bremer. The helmholtz-hodge decomposition — a survey. *IEEE Transactions on visualization and computer graphics*, 2013.
- J. Butcher. *Numerical methods for ordinary differential equations*. John Wiley & Sons, 2016.

- B. Calderhead, M. Girolami, and N. Lawrence. Accelerating bayesian inference over nonlinear differential equations with gaussian processes. *NIPS*, 2008.
- Andreas Damianou, Michalis K Titsias, and Neil D Lawrence. Variational gaussian process dynamical systems. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2011.
- Michiel JL De Hoon, Seiya Imoto, Kazuo Kobayashi, Naotake Ogasawara, and Satoru Miyano. Inferring gene regulatory networks from time-ordered gene expression data of bacillus subtilis using differential equations. In *Bio-computing 2003*, pages 17–28. World Scientific, 2002.
- F. Dondelinger, M. Filippone, S Rogers, and D. Husmeier. Ode parameter inference using adaptive gradient matching with gaussian processes. *JMLR*, 2013.
- S. P. Ellner, Y. Seifu, and R. H. Smith. Fitting population dynamic models to time-series data by gradient matching. *Ecology*, 83(8):2256–2270, 2002.
- Roger Frigola, Yutian Chen, and Carl Edward Rasmussen. Variational gaussian process state-space models. In *Advances in Neural Information Processing Systems*, pages 3680–3688, 2014.
- Fabian Fröhlich, Barbara Kaltenbacher, Fabian J. Theis, and Jan Hasenauer. Scalable parameter estimation for genome-scale biochemical reaction networks. *PLoS Computational Biology*, 13(1):1–18, 01 2017. doi: 10.1371/journal.pcbi.1005331.
- E. Fuselier and G. Wright. A radial basis function method for computing helmholtz–hodge decompositions. *IMA Journal of Numerical Analysis*, 2017.
- Pei Gao, Antti Honkela, Magnus Rattray, and Neil D Lawrence. Gaussian process modelling of latent chemical species: applications to inferring transcription factor activities. *Bioinformatics*, 24(16):i70–i75, 2008.
- C. García, A. Otero, P. Felix, J. Presedo, and D. Marquez. Nonparametric estimation of stochastic differential equations with sparse Gaussian processes. *Physical Review E*, 96(2):022104, 2017.
- Mark Girolami. Bayesian inference for differential equations. *Theor. Comput. Sci.*, 408(1):4–16, 2008.
- T. Graepel. Solving noisy linear operator equations by gaussian processes: Application to ordinary and partial differential equations. *ICML*, 2003.
- M. Heinonen and F. d’Alche Buc. Learning nonparametric differential equations with operator-valued kernels and gradient matching. arxiv, Telecom ParisTech, 2014.
- M. Heinonen, H. Mannerström, J. Rousu, S. Kaski, and H. Lähdesmäki. Non-stationary Gaussian process regression with Hamiltonian Monte Carlo. In *AISTATS*, volume 51, pages 732–740, 2016.
- J. Henderson and G. Michailidis. Network reconstruction using nonparametric additive ode models. *PLOS ONE*, 2014.
- Alan C Hindmarsh, Peter N Brown, Keith E Grant, Steven L Lee, Radu Serban, Dan E Shumaker, and Carol S Woodward. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Softw.*, 31(3):363–396, 2005.
- M. Hirsch, S. Smale, and Devaney. *Differential Equations, Dynamical Systems, and an Introduction to Chaos (Edition: 2)*. Elsevier Science & Technology Books, 2004.
- P Kokotovic and J Heller. Direct and adjoint sensitivity equations for parameter optimization. *IEEE Transactions on Automatic Control*, 12(5):609–610, 1967.
- Malte Kuss and Carl Edward Rasmussen. Assessing approximate inference for binary gaussian process classification. *Journal of machine learning research*, 6(Oct):1679–1704, 2005.
- Neil D Lawrence. Gaussian process latent variable models for visualisation of high dimensional data. In *Advances in neural information processing systems*, pages 329–336, 2004.
- Jorge R. Leis and Mark A. Kramer. The simultaneous solution and sensitivity analysis of systems described by ordinary differential equations. *ACM Trans. Math. Softw.*, 14(1), 1988.
- Benn Macdonald. *Statistical inference for ordinary differential equations using gradient matching*. PhD thesis, University of Glasgow, 2017.
- Benn Macdonald, Catherine Higham, and Dirk Husmeier. Controversy in mechanistic modelling with gaussian processes. In *International Conference on Machine Learning*, pages 1539–1547, 2015.
- Francis J Narcowich and Joseph D Ward. Generalized hermite interpolation via matrix-valued conditionally positive definite functions. *Mathematics of Computation*, 63(208):661–687, 1994.
- Joaquin Quiñonero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate gaussian process regression. *Journal of Machine Learning Research*, 6(Dec):1939–1959, 2005.

- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Inferring solutions of differential equations using noisy multi-fidelity data. *Journal of Computational Physics*, 335:736–746, 2017.
- J. Ramsay, G. Hooker, D. Campbell, and J. Cao. Parameter estimation for differential equations: a generalized smoothing approach. *Journal of the Royal Statistical Society: Series B*, 69:741–796, 2007.
- C.E. Rasmussen and K.I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.
- Andreas Raue, Marcel Schilling, Julie Bachmann, Andrew Matteson, Max Schelker, Daniel Kaschek, Sabine Hug, Clemens Kreutz, Brian D. Harms, Fabian J. Theis, Ursula Klingmüller, and Jens Timmer. Lessons learned from quantitative dynamical modeling in systems biology. *PLOS ONE*, 8(9):1–17, 2013.
- S. Remes, M. Heinonen, and S. Kaski. Non-stationary spectral kernels. *NIPS*, 2017.
- John Shawe-Taylor and Nello Cristianini. *Kernel methods for pattern analysis*. Cambridge university press, 2004.
- A. Solin, M. Kok, N. Wahlstrom, T. Schon, and S. Särkkä. Modeling and interpolation of the ambient magnetic field by gaussian processes. *arXiv*, 2015. arXiv:1509.04634.
- Andreas Svensson, Arno Solin, Simo Särkkä, and Thomas Schön. Computationally efficient bayesian learning of gaussian process state space models. In *Artificial Intelligence and Statistics*, pages 213–221, 2016.
- J. M. Varah. A spline least squares method for numerical parameter estimation in differential equations. *SIAM J.sci. Stat. Comput.*, 3(1):28–46, 1982.
- N. Wahlström, M. Kok, and T. Schön. Modeling magnetic fields using gaussian processes. *IEEE conf on Acoustics, Speech and Signal Processing (ICASSP)*, 2013.
- Jack Wang, Aaron Hertzmann, and David M Blei. Gaussian process dynamical models. In *Advances in neural information processing systems*, pages 1441–1448, 2006.
- Jack M Wang, David J Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *IEEE transactions on pattern analysis and machine intelligence*, 30(2):283–298, 2008.
- Y. Wang and D. Barber. Gaussian processes for bayesian estimation in ordinary differential equations. *ICML*, 2014.
- A. Wilson, E. Gilboa, A. Nehorai, and J. Cunningham. Fast multidimensional pattern extrapolation with gaussian processes. *AISTATS*, 2013.
- Jing Zhao and Shiliang Sun. Variational dependent multi-output gaussian process dynamical systems. *The Journal of Machine Learning Research*, 17(1):4134–4169, 2016.
- T Äijö and H. Lähdesmäki. Learning gene regulatory networks from gene expression measurements using non-parametric molecular kinetics. *Bioinformatics*, 25:2937–2944, 2009.