Differentiable Likelihoods for Fast Inversion of 'Likelihood-Free' Dynamical Systems

Hans Kersting [*1 2 3] **Nicholas Krämer** [*1]
Martin Schiegg [3] **Christian Daniel** [3] **Michael Tiemann** [3] **Philipp Hennig** [1 2]

Abstract

Likelihood-free (a.k.a. simulation-based) inference problems are inverse problems with expensive, or intractable, forward models. ODE inverse problems are commonly treated as likelihood-free, as their forward map has to be numerically approximated by an ODE solver. This, however, is not a fundamental constraint but just a lack of functionality in classic ODE solvers, which do not return a likelihood but a point estimate. To address this shortcoming, we employ Gaussian ODE filtering (a probabilistic numerical method for ODEs) to construct a local Gaussian approximation to the likelihood. This approximation yields tractable estimators for the gradient and Hessian of the (log-) likelihood. Insertion of these estimators into existing gradient-based optimization and sampling methods engenders new solvers for ODE inverse problems. We demonstrate that these methods outperform standard likelihood-free approaches on three benchmark-systems.

1. Introduction

Inferring the parameters of dynamical systems that are defined by ordinary differential equations (ODEs) is of importance in almost all areas of science and engineering. Despite the wide range of available ODE inverse problem solvers, simple random-walk Metropolis methods remain the go-to solution; see e.g. Tarantola (2005, Section 2.4). That is to say that ODE inverse problems are routinely treated as if their forward problems were black boxes. The reason usually cited for this generic approach is that ODE forward solutions are highly non-linear and numerically intractable for all but the most trivial cases. Therefore, it is common to consider ODE inverse problems as 'likelihood-free' inference (read: intractable likelihood)—a.k.a. simulation-based inference or, in the Bayesian case, Approximate Bayesian Computation (ABC); see Cranmer et al. (2020) for an up-to-date examination of these closely-related areas.

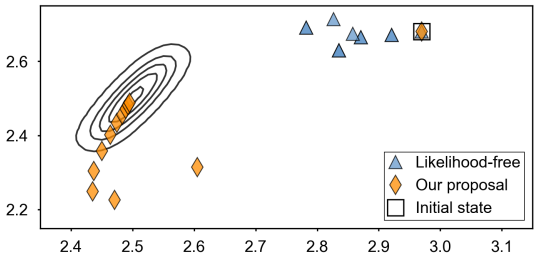

Figure 1. Inference on the logistic ODE. First twelve sampled parameters of likelihood-free inference and our proposed method. Details in text.

We here argue that, at least for ODEs, this approach is mistaken. If a dynamical system is accurately described by an ODE, its explicit mathematical definition should be exploited to design efficient algorithms—not ignored and treated as a black-box, likelihood-free inference problem.

To this end, we construct a local Gaussian approximation of the likelihood by Gaussian ODE Filtering, a probabilistic numerical method (PNM) for ODE forward problems. (Supplement A provides a concise introduction to Gaussian ODE filtering; Tronarp et al. (2019) offer a more detailed presentation. See Hennig et al. (2015) or Oates & Sullivan (2019) for a broad introduction to PNMs.) The key insight of our work is that there *is* a likelihood in simulations of ODEs, and in fact it can be approximated cheaply, and analytically: The mean estimate $\boldsymbol{m}_\theta$ of the forward solution computed by Gaussian ODE filters can be linearized in the parameter θ, so that gradient, Hessian, etc. of the approximated log-likelihood can—via a cheap estimator J of the Jacobian of the map $\theta \mapsto \boldsymbol{m}_\theta$—be computed in closed form (Section 5). In this way, the probabilistic information from Gaussian ODE filtering yields a tractable, twice-differentiable like-

[*]Primary author [1]University of Tübingen, Tübingen, Germany [2]Max Planck Institute for Intelligent Systems, Tübingen, Germany [3]Bosch Center for Artificial Intelligence, Renningen, Germany. Correspondence to: Hans Kersting <hans.kersting@uni-tuebingen.de>, Nicholas Krämer <nicholas.kraemer@uni-tuebingen.de>.

Proceedings of the 37^{th} *International Conference on Machine Learning*, Online, PMLR 119, 2020.

lihood for 'likelihood-free' ODE inverse problems. This enables the use of first and second-order optimization or sampling methods (see Figure 1).

Much thought has been devoted to improving the slow runtimes of ODE inverse inference—which is due to the laborious explicit numerical integration per parameter. In machine learning, e.g., authors have proposed to reduce the amount of necessary parameters by active learning with Gaussian process (GP) surrogate likelihoods (Meeds & Welling, 2014), or even to avoid numerical integration altogether by gradient matching (Calderhead et al., 2008). This paper adds a new way to reduce the amount of parameters by employing gradient (and Hessian) estimates of the log-likelihood.

Contributions The main contributions are twofold: *Firstly*, we introduce tractable estimators for the gradients and Hessian matrices of the log-likelihood of ODE inverse problems by Gaussian ODE filtering. To derive these estimators, we construct a new estimator J for the Jacobian of the forward map. We theoretically support the use of J by a decomposition of the true Jacobian into J and a sensitivity term S (see Theorem 1), as well as an upper bound on its approximation error (see Theorem 2). *Secondly*, we propose a range of new solvers which require gradients and/or Hessians, by inserting these estimators into first and second-order optimization and sampling methods. The utility of these algorithms is demonstrated by experiments on three benchmark ODEs where they outperform their gradient-free counterparts.

2. Problem Setting

We consider a dynamical system defined by the ODE

$$\dot{x}(t) = f\left(x(t), \theta\right), \qquad x(0) = x_0 \in \mathbb{R}^d, \tag{1}$$

on the finite time domain $t \in [0, T]$ for some $T > 0$, with parametrized vector field $f : \mathbb{R}^d \times \mathbb{R}^n \to \mathbb{R}^d$. We restrict our attention to choices of f satisfying the following

Assumption 1. *$f(x, \theta) = \sum_{i=1}^{n} \theta_i f_i(x)$, for some continuously differentiable $f_i : \mathbb{R}^d \to \mathbb{R}^d$, for all $i = 1, \ldots, n$.*

The necessity for this assumption will become evident in Section 3.1.1. It is not very restrictive: e.g. the corresponding assumption in Gorbach et al. (2017, eq. (10)) is stronger. In fact, most standard ODEs collected in Hull et al. (1972, Appendix I), a standard set of ODE benchmarking problems, satisfy Assumption 1 either immediately or after reparametrization. Otherwise, we can still transform a non-conforming ODE into a system that obeys Assumption 1, as exemplified for the protein signalling transduction pathway in Section 7.2.2. While this adds an additional layer of imprecision, the experiments appear to be equally good—which suggests a wider applicability of our methods than Assumption 1.

If the initial value x_0 is unknown too (as is often the case in practice), it can be treated as a parameter by defining a new parameter vector $\left(x_0^\intercal, \theta^\intercal\right)^\intercal \in \mathbb{R}^{d+n}$; see eq. (10). Solving eq. (1), for a given θ, with a numerical method is known as the *forward problem*.

For the *inverse problem*, we assume the dynamical system described by eq. (1) with *unknown* true parameter θ^*. The true trajectory $x = x_{\theta^*}$ is observed under additive, zero-mean Gaussian noise at M discrete times $0 \leq t_1 < \cdots < t_M \leq T$:

$$z(t_i) \coloneqq x(t_i) + \varepsilon_i \in \mathbb{R}^d, \quad \varepsilon_i \sim \mathcal{N}(0, \Sigma_i), \tag{2}$$

for all $i \in \{1, \ldots, M\}$. Below we assume, w.l.o.g., that $\Sigma_i = \Sigma$, for all $i \in \{1, \ldots, M\}$. We define the stacked data across M time points and d dimensions as

$$\boldsymbol{z} \coloneqq \left[z_1(t_1), \ldots, z_1(t_M), \ldots, z_d(t_1), \ldots, z_d(t_M)\right]^\intercal,$$

and analogously, for all $\theta \in \Theta$, the true solution at these points as $\boldsymbol{x}_\theta$. The inverse problem consists of inferring the parameter θ^* that generated the data through eq. (2). For the sake of readability, we will assume w.l.o.g. that $d = 1$; this restriction is purely notational as can be seen from the multi-dimensional experiments below. Under these conventions, eq. (2) is equivalent to

$$p(\boldsymbol{z} \mid \boldsymbol{x}) = \mathcal{N}\left(\boldsymbol{z}; \boldsymbol{x}, \sigma^2 I_M\right) \tag{3}$$

for some $\sigma^2 > 0$, where I_M is the $M \times M$ identity matrix. Heteroscedastic noise can be modelled by replacing $\sigma^2 I_M$ with a diagonal matrix with varying diagonal entries.

3. Likelihoods by Gaussian ODE Filtering

The prevailing view on the uncertainty in inverse problems only considers the aleatoric uncertainty Σ_i from eq. (2) and ignores the epistemic uncertainty over the quality of the employed numerical approximation $\hat{x}_\theta$ of x_θ. In other words, the likelihood of the forward problem, $p(\boldsymbol{x}_\theta \mid \theta)$, is commonly treated as a Dirac distribution $\delta(\boldsymbol{x}_\theta - \hat{\boldsymbol{x}}_\theta)$ which yields the *uncertainty-unaware likelihood*

$$p(\boldsymbol{z} \mid \theta) = \int p(\boldsymbol{z} \mid \boldsymbol{x}_\theta) p(\boldsymbol{x}_\theta \mid \theta) \, \mathrm{d}\boldsymbol{x}_\theta \tag{4}$$

$$= \int p(\boldsymbol{z} \mid \boldsymbol{x}_\theta) \delta(\boldsymbol{x}_\theta - \hat{\boldsymbol{x}}_\theta) \, \mathrm{d}\boldsymbol{x}_\theta \tag{5}$$

$$\overset{\text{eq. (3)}}{=} \mathcal{N}\left(\boldsymbol{z}; \hat{\boldsymbol{x}}_\theta, \sigma^2 I_M\right). \tag{6}$$

as the 'true' intractable likelihood. This, however, ignores the epistemic uncertainty over the accuracy $\hat{x}_\theta$ which leads to overconfidence. This uncertainty is due to the discretization error of the numerical solver used to compute $\hat{x}_\theta$, and can only be avoided for the most trivial ODEs. This problem has previously been recognized in, e.g., Conrad et al. (2017, Section 3.2) and Abdulle & Garegnani (2020,

Section 8) who, as a remedy, construct a 'cloud' of possible solutions by running a classical solver multiple times with a prespecified accuracy. This, unfortunately, requires the computational invest of several forward solves for the same θ, which could instead be used for additional θ, or higher accuracy.

To obtain such uncertainty quantification more cheaply, we employ Gaussian ODE filtering with a once-integrated Brownian motion (IBM) prior on x; see Supplement A.2 for a short introduction. This amounts—e.g. in the notation of Tronarp et al. (2019)—to setting $q = 1$. Gaussian ODE filtering has the advantage over other numerical solvers, probabilistic or classical, that we can compute gradients of the likelihood, as demonstrated below. For a given θ, the Gaussian ODE filter computes a multivariate normal distribution over x_θ at a set of $N = T/h$, for notational simplicity, equidistant time points $\{0, h, \dots, Nh\}$ with step size $h > 0$. This set is, w.l.o.g., assumed to contain the data time points $\{t_1, \dots, t_M\}$ from eq. (2), i.e. we assume the existence of a set of integers $\{l_1, \dots, l_M\}$ such that $t_i = l_i h$. (The w.l.o.g. assumption can otherwise be satisfied by interpolating along the dynamic model; see eq. (A.1) in Supplement A.)

3.1. The Filtering Distribution

The Gaussian ODE filter returns the so-called (posterior) filtering distribution over the ODE solution $\boldsymbol{x}_\theta$, given by

$$p(\boldsymbol{x}_\theta \mid \theta) = \mathcal{N}(\boldsymbol{x}_\theta; \boldsymbol{m}_\theta, \boldsymbol{P}), \tag{7}$$

with $\boldsymbol{m}_\theta \in \mathbb{R}^M$ and $\boldsymbol{P} \in \mathbb{R}^{M \times M}$ given below by eq. (10) and eq. (18), respectively. This probabilistic likelihood yields the new *uncertainty-aware likelihood*

$$p(\boldsymbol{z} \mid \theta) = \int p(\boldsymbol{z} \mid \boldsymbol{x}_\theta) \mathcal{N}(\boldsymbol{x}_\theta; \boldsymbol{m}_\theta, \boldsymbol{P}) \, \mathrm{d}\boldsymbol{x}_\theta \tag{8}$$

$$\overset{eq.\ (3)}{=} \mathcal{N}(\boldsymbol{z}; \boldsymbol{m}_\theta, \boldsymbol{P} + \sigma^2 I_M) \tag{9}$$

which has two advantages over the uncertainty-unaware likelihood from eq. (6):

1. The filtering mean $\boldsymbol{m}_\theta$ can be linearized in θ, as specified below in eq. (10). This yields an estimate J of the Jacobian matrix of $\theta \mapsto \boldsymbol{m}_\theta$ which implies estimators of gradients and Hessian matrices of the likelihood; see eqs. (26) and (27). These estimators are useful to guide samples of θ into regions of high likelihood by the gradient-based sampling and methods defined in Section 6 below.

2. The variance $\boldsymbol{P}$ captures the average-case squared (epistemic) error $\|\boldsymbol{m}_\theta - \boldsymbol{x}_\theta\|^2$, and can be added to the (aleatoric) variance Σ_i; see eq. (9). Unless $\boldsymbol{P} \ll \sigma^2 I_M$, this prevents over-confidence, as visualized in Figure 2.

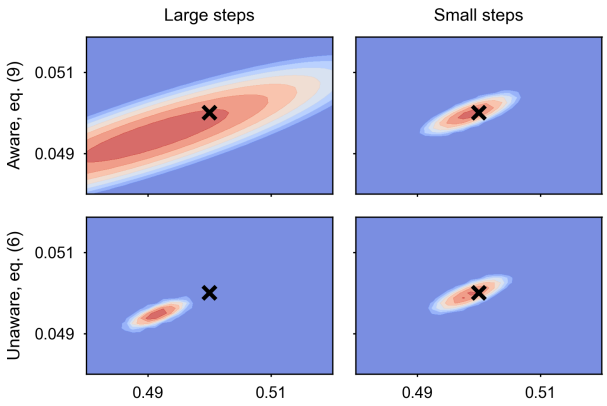

Figure 2. Uncertainty-(un)aware likelihoods, eqs. (6) and (9) w.r.t. (θ_1, θ_2) of Lotka-Volterra ODE, eq. (30), with fixed $(\theta_3, \theta_4) = (0.05, 0.5)$. θ_1 on x and θ_2 on y-axis. Black cross is true parameter. The unaware likelihood is overconfident for the large step size ($h = 0.2$), i.e. for large $\boldsymbol{P}$, while the aware likelihood has calibrated uncertainty. For the small step size ($h = 0.025$) this effect is less pronounced as $\boldsymbol{P}$ is small.

In the following two subsections, we provide explicit formulas for $\boldsymbol{m}_\theta$ and $\boldsymbol{P}$. A detailed derivation of these formulas is given in Supplement B.

3.1.1. The Filtering Mean

Under Assumption 1, the filtering mean $\boldsymbol{m}_\theta = [m_\theta(t_1), \dots, m_\theta(t_M)]^\intercal$ is given by

$$\boldsymbol{m}_\theta = \begin{bmatrix} 1_M & J \end{bmatrix} \begin{bmatrix} x_0 \\ \theta \end{bmatrix} = x_0 \cdot 1_M + J\theta \ \in \mathbb{R}^M, \tag{10}$$

where $1_M = [1, \dots, 1]^\intercal$ denotes a vector of M ones. Hence, $\boldsymbol{m}_\theta$ is linear in θ as well as in the extended parameter vector $[x_0, \theta^\intercal]^\intercal$. (A more detailed derivation of eq. (10) is provided in Supplement B.3.) Here,

$$J \coloneqq KY \ \in \mathbb{R}^{M \times n} \tag{11}$$

is an estimator of the Jacobian matrix of the map $\theta \mapsto \boldsymbol{m}_\theta$, as we show in Theorem 1 below. This estimator is equal to the product of the *kernel prefactor* K and the *evaluation factor* Y. The kernel prefactor K is given by

$$K \coloneqq \begin{bmatrix} \kappa_1, \dots, \kappa_M \end{bmatrix}^\intercal \ \in \mathbb{R}^{M \times N}, \tag{12}$$

whose i-th row is

$$\kappa_i \coloneqq \begin{bmatrix} \tilde{\kappa}_i^\intercal, 0, \dots, 0 \end{bmatrix}^\intercal \ \in \mathbb{R}^N, \tag{13}$$

which is defined by

$$\tilde{\kappa}_i \coloneqq \left[{}^{\partial}K^{\partial}(h : t_i) + R \cdot I_{l_i} \right]^{-1} k^{\partial}(h : t_i, t_i) \ \in \mathbb{R}^{l_i}, \tag{14}$$

for some measurement variance $R \geq 0$. Here, $k^{\partial} = \partial k(t,t')/\partial t'$ and ${}^{\partial}k^{\partial} = \partial^2 k(t,t')/\partial t \partial t'$ are derivatives of the IBM kernel k, and, analogously, the cross-covariance w.r.t. the kernel ${}^{\partial}k$ and the kernel Gram matrix w.r.t. the kernel ${}^{\partial}k^{\partial}$ up to time t_i are denoted by

$$k^{\partial}(h : t_i, t_i) \coloneqq \left[k^{\partial}(t_i, h), \ldots, k^{\partial}(t_i, t_i)\right]^{\intercal}, \text{ and} \quad (15)$$

$${}^{\partial}K^{\partial}(h : t_i) \coloneqq \begin{bmatrix} {}^{\partial}k^{\partial}(h,h) & \cdots & {}^{\partial}k^{\partial}(l_i h, l_i h) \\ \vdots & \ddots & \vdots \\ {}^{\partial}k^{\partial}(l_i h, h) & \cdots & {}^{\partial}k^{\partial}(l_i h, l_i h) \end{bmatrix}. \quad (16)$$

Now, recall Assumption 1. For a given θ, the entries of the evaluation factor $Y \in \mathbb{R}^{N \times n}$ are

$$y_{ij} \coloneqq f_j(m_\theta^-(ih)) - f_j(x_0), \quad (17)$$

for all $i = 1, \ldots, N$ and $j = 1, \ldots, n$, where $m_\theta^-(ih)$ is the predictive mean of the ODE Filter at $t = ih$. Note that the Gaussian ODE Filter computes the $f_j(m_\theta^-(ih))$ and $f_j(x_0)$ for every forward solve as intermediate quantities, to evaluate the right-hand side of eq. (1). Hence, Y is freely accessible with every filtering distribution, eq. (7). However, as an estimate of $x_\theta(ih)$, $m_\theta^-(ih)$ depends on θ in a nonlinear and potentially sensitive way. By ignoring this dependence in the above notation, we, strictly speaking, also omit the dependence of Y and, thereby, J on θ (more in Supplement B.3). For this reason, J is not the true Jacobian of $\theta \mapsto \boldsymbol{m}_\theta$ but only an estimator (see Section 3.2).

3.1.2. The Filtering Covariance

The entries of the covariance matrix $\boldsymbol{P} \coloneqq \operatorname{diag}(P(t_1), \ldots, P(t_M)) \in \mathbb{R}^{M \times M}$ of the filtering distribution from eq. (7) coincide with the GP-posterior variances, i.e.

$$P(t_i) = \begin{bmatrix} k(h,h) & \cdots & k(l_i h, l_i h) \\ \vdots & \ddots & \vdots \\ k(l_i h, h) & \cdots & k(l_i h, l_i h) \end{bmatrix} - k^{\partial}(h : t_i, t_i)^{\intercal} \times \left[{}^{\partial}K^{\partial}(h : t_i) + R \cdot I_l\right]^{-1} k^{\partial}(h : t_i, t_i), \quad (18)$$

and are hence independent of θ. (See Supplement B.2 for a detailed derivation of eq. (18).)

3.2. Decomposition of the True Jacobian

Next, we give an explicit decomposition of the true Jacobian into the estimator J, the kernel prefactor K and a sensitivity term S.

Theorem 1. *Under Assumption 1, the true Jacobian $D\boldsymbol{m}_\theta \in \mathbb{R}^{M \times n}$ of $\theta \mapsto \boldsymbol{m}_\theta$ has the analytic form*

$$D\boldsymbol{m}_\theta \coloneqq [\nabla_\theta m(t_1), \ldots, \nabla_\theta m(t_M)]^{\intercal} = J + KS, \quad (19)$$

where the sensitivity term S *is defined by*

$$S \coloneqq \left[\Lambda_1^{\intercal}\theta, \ldots, \Lambda_N^{\intercal}\theta\right]^{\intercal} \in \mathbb{R}^{N \times n}. \quad (20)$$

Here, $\Lambda_j = \left[\lambda_{kl}(jh)\right]_{kl}$ is the $n \times n$ matrix with entries

$$\lambda_{kl}(jh) \coloneqq \frac{\mathrm{d}}{\mathrm{d}x} f_l(m_\theta^-(jh)) \cdot \frac{\partial}{\partial \theta_k} m_\theta^-(jh). \quad (21)$$

Proof. See Supplement C. □

Thus, KS is the exact approximation error of J.

4. Bound on Approximation Error of J

In this section, we provide a bound on the approximation error of J under the following assumptions.

Assumption 2. *The first-order partial derivatives of f_i, $1 \leq i \leq N$, are bounded and globally L-Lipschitz, for $L > 0$.*

Assumption 2 is required to bound the global error of the ODE forward solution by Kersting et al. (2019, Thm. 6.7).

Assumption 3. *For the computation of J we only use a maximum of $\bar{N} \leq N$ time points, for some finite $\bar{N} \in \mathbb{N}$.*

Assumption 3 precludes the condition number of the K and S from growing arbitrarily large, thereby preventing numerical instability. While this restriction is necessary for Theorem 2, it is not relevant in practice because we are computing with a non-zero step size $h > 0$ anyway so that many different parameters θ can be simulated.

Theorem 2. *If $\Theta \subset \mathbb{R}^n$ is compact and $R > 0$, then it holds true, under Assumptions 1 to 3, that*

$$\|J - D\boldsymbol{m}_\theta\| \leq C(T)\left(\|\nabla_\theta x_\theta\| + h\right) \quad (22)$$

for sufficiently small $h > 0$, where $C(T) > 0$ is a constant that depends on T.

Proof. See Supplement D. □

Intuitively, this upper bound can be thought of as a decomposition of the approximation error of the 'sensitivity-unaware' estimator J into a summand proportional to the ignored sensitivity $\|\nabla_\theta x_\theta\|$ and the global integration error of the ODE filter, which is bounded by $C(T)h$ (Kersting et al., 2019, Thm. 6.7).

5. Gradient and Hessian Estimators

We observe that the uncertainty-aware likelihood, eq. (9), can be written in the form

$$p(\boldsymbol{z} \mid \theta) = \frac{e^{-E(\boldsymbol{z})}}{Z}, \quad (23)$$

with evidence $Z > 0$ and negative log-likelihood

$$E(\boldsymbol{z}) \coloneqq \frac{1}{2}\left[\boldsymbol{z} - \boldsymbol{m}_\theta\right]^\intercal \left[\boldsymbol{P} + \sigma^2 I_M\right]^{-1} \left[\boldsymbol{z} - \boldsymbol{m}_\theta\right] \quad (24)$$

$$\overset{eq.\ (10)}{=} \frac{1}{2}\left[\boldsymbol{z} - x_0 \cdot 1_M - J\theta\right]^\intercal \left[\boldsymbol{P} + \sigma^2 I_M\right]^{-1} \times \left[\boldsymbol{z} - x_0 \cdot 1_M - J\theta\right]. \quad (25)$$

For a given value of the Jacobian estimator J, the thereby-implied gradient and Hessian estimators are, by application of the chain rule,

$$\hat{\nabla}_\theta E(\boldsymbol{z}) \coloneqq -J^\intercal \left[\boldsymbol{P} + \sigma^2 I_M\right]^{-1} \left[\boldsymbol{z} - \boldsymbol{m}_\theta\right], \quad \text{and} \quad (26)$$

$$\hat{\nabla}^2_\theta E(\boldsymbol{z}) \coloneqq J^\intercal \left[\boldsymbol{P} + \sigma^2 I_M\right]^{-1} J. \quad (27)$$

(See Figure 3 for a visualization of these estimators.) Supplement E provides versions of these estimators for Bayesian inference of θ.

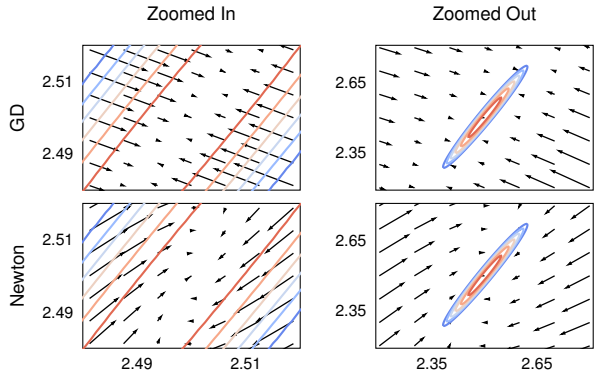

Figure 3. Directions of gradient descent (GD) and Newton using eqs. (26) and (27); around mode (left) and globally (right) of the likelihood, based on the logistic ODE. Globally, GD points more directly to the high-probability region. Within this region, however, Newton is better directed to the mode.

6. New Gradient-Based Methods

By deriving gradient and Hessian estimators of the negative log-likelihood, we have removed the need for 'likelihood-free' inference. This enables the use of two classes of inference methods for θ which could not otherwise be applied: gradient-based *optimization* and gradient-based *sampling*.

6.1. Gradient-Based Optimization

In principle, all first and second-order optimization algorithms (e.g. Bottou et al. (2018)), are now applicable by eqs. (26) and (27)—such as (stochastic) gradient descent (GD), (stochastic) Newton (NWT), Gauss-Newton and natural Gradient descent. This application of the estimators (26) and (27) unlocks fast computation of single parameter estimates by maximum-likelihood estimation, as we demonstrate in the experiments (see Section 7).

6.2. Gradient-Based Sampling

Likewise, all gradient-based MCMC schemes are now available. Classical gradient-based samplers include Langevin Monte Carlo (LMC) (Roberts & Tweedie, 1996) and Hamiltonian Monte Carlo (HMC) (Betancourt, 2017). They are known to be more efficient than gradient-free samplers in finding and covering regions of high probability (MacKay, 2003, Section 30.1). While their standard form only makes use of gradients, more sophisticated versions include second-order information as well: When the likelihood is ill-conditioned (i.e. it varies much more quickly in some directions than others), it is advantageous to precondition the proposal distribution with a suitable matrix (Girolami & Calderhead, 2011). A popular choice for the preconditioner is the Hessian (Qi & Minka, 2002). Hence, we can precondition LMC and HMC that use eq. (26) as a gradient with the Hessian estimator from eq. (27). For LMC, this leads to the proposal distribution

$$\pi(\theta^{i+1} \mid \theta^i) = \theta^i - \rho[\hat{\nabla}^2_\theta E_{\theta^i}(\boldsymbol{z}))]^{-1}\hat{\nabla}_\theta E_{\theta^i}(\boldsymbol{z})) + \xi^i, \quad (28)$$

$$\xi^i \sim \mathcal{N}(0, 2\rho[\hat{\nabla}^2_\theta E_{\theta^i}(\boldsymbol{z}))]^{-1}), \quad (29)$$

where ρ is the proposal width. (Analogous formulas hold for HMC.) Below, we refer to the so-preconditioned versions of LMC and HMC as PLMC and PHMC. In Section 7, we show that the gradient-based versions more aptly explore regions of high likelihood than their gradient-free counterparts.

6.3. Algorithm

The generic method that we propose is outlined in Algorithm 1. It includes all above-mentioned classical optimization and sampling methods (by a corresponding choice in Line 7). The only difference, compared to all of these existing gradient-based methods, are the additional Lines 5 and 6 where we compute our gradient and Hessian estimators from eqs. (26) and (27).

Algorithm 1 Gradient-based sampling/optimization

1: Precompute K and $(P + \sigma^2 I_M)^{-1}$ (see eqs. (12), (6))
2: Initialize $\theta = \theta^0$
3: **repeat**
4: Solve ODE with θ (this generates Y; see eq. (17))
5: Compute $J = KY$ (see eq. (11))
6: Compute $[\hat{\nabla}_\theta E, \hat{\nabla}^2_\theta E]$ (see eqs. (26), (27))
7: Update θ with gradient-based sampler/optimizer
8: **until** convergence/mixing

6.4. Computational Cost

The additional computational cost—on top of the employed classical optimization/sampling methods—is equal to the

cost of computing the inserted gradient (and Hessian) estimators: precomputation of K (Line 1 in Algorithm 1) requires the inversion of the M kernel Gram matrices $\{^{\partial}K^{\partial}(h : t_i),\ i = 1, \ldots, M\}$, which can have a maximum dimension of $(N-1) \times (N-1)$. This inversion can, however, be executed in linear time since $^{\partial}k^{\partial}$ is a Markov kernel (Hartikainen & Särkkä, 2010). Hence, K is in $\mathcal{O}(MN)$ and, as $M \leq N$, in $\mathcal{O}(N^2)$. The cost of inverting the $M \times M$ matrix $[\boldsymbol{P} + \sigma^2 I_M]$ is in $\mathcal{O}(N^3)$, as $M \leq N$. Since K and $\boldsymbol{P}$ are independent of θ, this $\mathcal{O}(N^3)$ cost is only required once. The Jacobian estimator $J = KY$ (Line 5 in Algorithm 1) is, by eq. (11), a matrix product of the precomputed kernel prefactor K and the evaluation factor Y. Y is almost free, as it is by eq. (17) only composed of terms that the Gaussian ODE filter computes anyway; see eq. (A.10) in Supplement A.2. Given J and $[\boldsymbol{P} + \sigma^2 I_M]^{-1}$, computing the gradient and Hessian estimators (Line 6 in Algorithm 1) is of the same complexity as computing J. Thus, the additional computational cost is in $\mathcal{O}(N^3)$ w.r.t. the number of time steps $N = T/h$ executed once and otherwise linear (but almost negligible) w.r.t. the number of simulated parameters θ. As a large number of θ is usually required, the overall overhead is small.

6.5. Choice of Hyperparameters

Recall that the parameters σ and R stem from the data and the accuracy of the ODE model (Kersting et al., 2019, Section 2.3), and that we only consider once-integrated Brownian motion priors in this paper. Therefore, the only remaining hyperparameter is the diffusion scale σ_{dif} which controls the width of the variance $\boldsymbol{P}$; see Supplements B.1 and B.2. There are two ways to set it: either as a local (**?**, eq. (46)) or as a global (Tronarp et al., 2019, eq. (41)) maximum-likelihood estimate, which can both be computed from intermediate quantities of the forward solves.

7. Experiments

To test the hypothesis that the gradient and Hessian estimators $[\hat{\nabla}_\theta E(\boldsymbol{z}), \hat{\nabla}^2_\theta E(\boldsymbol{z})]$ of the log-likelihood are useful despite their approximate nature, we compare the new optimization and sampling methods from Section 6—which use these estimators as if exact—with the standard 'likelihood-free' approach, i.e. with random search (RS) optimization and random-walk Metropolis (RWM) sampling.

7.1. Setup and Methods

As benchmark systems, we choose the popular Lotka–Volterra (LV) predator-prey model and the more challenging biochemical dynamics of glucose uptake in yeast (GUiY). For more generality, we add the chemical protein signalling transduction (PST) dynamics which violate Assumption 1 and have to be linearized. We consider our hypothesis validated if the new gradient-based algorithms outperform the conventional 'likelihood-free' methods (RS, RWM) on these three systems. All datasets are, as in eq. (3), generated by adding Gaussian noise to the solution x_{θ^*} for some true parameter θ^*.

Out of the new family of gradient-based optimizers and samplers introduced in Section 6, we evaluate only the most basic ones: gradient descent (GD) and Newton's method (NWT) for optimization, as well as PLMC and PHMC for sampling. This isolates the impact of the gradient and Hessian estimators more clearly. The required gradient and Hessian estimators are computed as detailed above. We employ the original fixed step-size RS by Rastrigin (1963), and the RWM version from MacKay (2003, Chapter 29). For all optimizers, we picked the best the step size and, for all samplers, the best proposal width within the interval $[10^{-16}, 10^{0}]$ which is wide enough to contain all plausible values. To make these experiments an ablation study for the gradient and Hessian estimators, we use Gaussian ODE filtering as a forward solver in all methods—which is similar to classical solvers anyway (**?**, Section 3). Since in all below experiments $\boldsymbol{P} \gg \sigma^2 I_M$, the gradient and Hessian estimates are scale-invariant w.r.t. hyperparameter σ^2_{dif}, as can be seen from eqs. (26) and (27): In this regime, $\boldsymbol{P}$ simply scales the step-size of the gradient, and $\boldsymbol{P}$ cancels out of the Hessian, making it invariant to this scale. The same applies in the regime $\boldsymbol{P} \ll \sigma^2 I_M$; adaptation of their relative scale, by choosing σ^2_{dif} as in Section 6.5, only matters when both error-sources are of comparable scale.

7.2. Results

We evaluate the performance of these methods over the first few iterations (steps), comparing the values of the negative log-likelihood E as well as the relative error in the parameter space, $\left\|\theta^i - \theta^*\right\| / \|\theta^*\|$. For optimizers, low values in both metrics indicate success and, in fact, both are important: ODE inverse problems are inherently ill-posed and can have parameters with high likelihood and large inference error that fit the data as well as the true parameter. Finding these parameters would not be a failure of the algorithms, but a success, as they are a mode of the true posterior.

Samplers, on the other hand, try to identify and explore regions of high probability (the typical set); see e.g. Betancourt (2017, Section 2). We opt for plotting the relative error in the parameter space additionally to the negative log-likelihood values to emphasize that, once a sampler creates samples near the typical set, MCMC methods keep exploring suitable values instead of relying on a single estimate with high likelihood. Despite maintaining a low near-constant negative log-likelihood, the error in the parameter space of a sampler may have (some) variation.

The details and results for each benchmark systems are

presented next, in ascending order of complexity.

7.2.1. LOTKA–VOLTERRA

First, we study the Lotka–Volterra (LV) ODE (Lotka, 1978)

$$\dot{x}_1 = \theta_1 x_1 - \theta_2 x_1 x_2, \quad \dot{x}_2 = -\theta_3 x_2 + \theta_4 x_1 x_2, \quad (30)$$

the standard model for predator-prey dynamics. We used this ODE with initial value $x_0 = [20, 20]$, time interval $[0, 5]$ and true parameter $\theta^* = [1, 0.1, 0.1, 1]$. To generate data by eq. (3), we added Gaussian noise with variance $\sigma^2 = 0.01$ to the corresponding solution at time points $[0.5, 1, 1.5, 2, 2.5, 3., 3.5, 4., 4.5]$. The optimizers and samplers were initialized at $\theta^0 = [0.8, 0.2, 0.05, 1.1]$, and the forward solutions for all likelihood evaluations were computed with step size $h = 0.05$. In order to turn this θ^0 into a useful initialization for the Markov chains, we accepted the first 45 states generated by PHMC and PLMC—the same would be counterproductive for RWM since a proposed sample may be further away from the region of nonzero probability. The results for optimization and sampling are depicted in Figure 4. In the case of optimizers, NWT outperforms GD which, in turn, outperforms RS. After roughly 25 samples, NWT generates iterations with relative error of less than 10^{-3}. While PLMC and PHMC quickly reach and explore regions of high probability, RWM does not find likelihood values within the first 250 samples. Thus, the gradient and Hessian estimators indeed appear to work well on LV.

7.2.2. PROTEIN SIGNALLING TRANSDUCTION

Next, we consider the protein signalling transduction (PST) pathway. It is governed by a combination of mass-action and Michaelis–Menten kinetics:

$$\begin{aligned}
\dot{S} &= -\theta_1 \times S - \theta_2 \times S \times R + \theta_3 \times RS, \\
\dot{dS} &= \theta_1 \times S, \\
\dot{R} &= -\theta_2 \times S \times R + \theta_3 \times RS + V \times \frac{Rpp}{K_m + Rpp}, \\
\dot{RS} &= \theta_2 \times S \times R - \theta_3 \times RS - \theta_4 \times RS, \\
\dot{Rpp} &= \theta_4 \times RS - \theta_5 \times \frac{Rpp}{K_m + Rpp}.
\end{aligned}$$

For more details, see Vyshemirsky & Girolami (2008). Due to the ratio $\frac{Rpp}{K_m + Rpp}$, Assumption 1 is violated. As a remedy, we follow Gorbach et al. (2017) in defining the latent variables $[x_1, x_2, x_3, x_4, x_5] := [S, dS, R, RS, \frac{Rpp}{K_m + Rpp}]$. This gives rise to the new linearized ODE

$$\dot{x}_1 = -\theta_1 x_1 - \theta_2 x_1 x_3 + \theta_3 x_4, \quad (31)$$
$$\dot{x}_2 = \theta_1 x_1, \quad (32)$$
$$\dot{x}_3 = -\theta_2 x_1 x_3 + \theta_3 x_4 + \theta_5 x_5, \quad (33)$$
$$\dot{x}_4 = \theta_2 x_1 x_3 - \theta_3 x_4 - \theta_4 x_4, \quad (34)$$
$$\dot{x}_5 = \theta_4 x_4 - \theta_5 x_5, \quad (35)$$

which is an approximation of the original ODE, since eq. (35) ignores the factor $(K_m + R_{pp})^{-1}$. We used this ODE with initial value $x_0 = [1, 0, 1, 0, 0]$ on time interval $[0, 100]$. We set the true parameter to $\theta^* = [0.07, 0.6, 0.05, 0.3, 0.017]$. To generate the data by eq. (3), we added Gaussian noise with variance $\sigma^2 = 10^{-8}$ to the corresponding solution at time points $[1., 2., 4., 5., 7., 10., 15., 20., 30., 40., 50., 60., 80., 100.]$. The optimizers and samplers were initialized at $\theta^0 = [0.24, 1.8, 0.15, 0.9, 0.05]$, and the forward solutions for all likelihood evaluations were computed with step size $h = 0.05$. We use the same burn-in procedure as on the Lotka-Volterra example, accepting the first 100 samples.
The results for optimization and sampling are depicted in Figure 5.

Again, the new methods outperform the conventional ones in both optimization and sampling. For optimization, NWT converges particularly fast. The final estimate that is returned by NWT is, rounded to two digits, $\theta^{200} = (0.07, 0.60, 0.05, 0.30, 0.02)$, and hence recovers four out of five parameters exactly. For sampling, both gradient-based samplers (after a fairly steep initial improvement) steadily stay in regions of high likelihood, while RWM only increases the likelihood in a much slower pace. Hence, the gradient and Hessian estimators are beneficial on PST as well—although we had to linearize the ODE first.

7.2.3. GLUCOSE UPTAKE IN YEAST

Last, we examine the challenging biochemical dynamics of glucose uptake in yeast (GUiY), as seen in Schillings et al. (2015). This ODE is 9-dimensional, has 10 parameters, and satisfies Assumption 1; see Supplement F for a complete mathematical definition and parameter choices. The results for optimization and sampling are depicted in Figure 6.

GD outperforms RS, and NWT converges even much faster than GD. Remarkably, NWT already finds parameters that are exact up to two relative digits after only five iterations which would take RS extremely long on this 10 dimensional domain. The gradient-based samplers (PLMC, PHMC), again, stay steadily within the region of significant likelihood, while RWM has difficulties sampling from this high dimensional problem in an efficient manner. Thus, this benchmark system also reaffirms the utility of the gradient

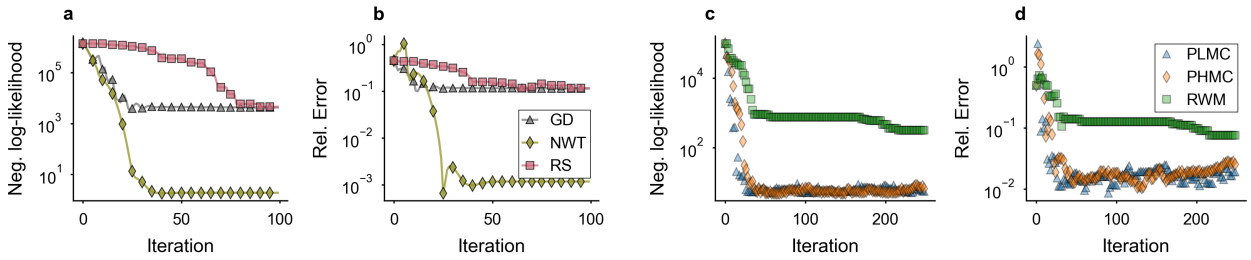

Figure 4. **Results for optimization (a, b) and sampling (c, d) on Lotka-Volterra**. Comparison of negative log-likelihood $E(\boldsymbol{z}) = E_{\theta^i}(\boldsymbol{z})$ (a and c, resp.) and relative error $\left\|\theta^i - \theta^*\right\| / \|\theta^*\|$ (b and d, resp.). 100 iterations of optimization (only every fifth iteration has a marker) and 250 Metropolis-Hastings samples (only every other sample has a marker).

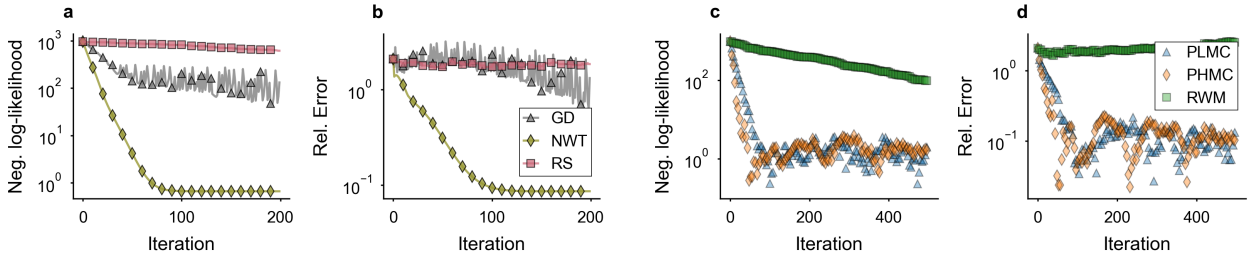

Figure 5. **Results for optimization (a, b) and sampling (c, d) on PST**. Comparison of negative log-likelihood $E(\boldsymbol{z}) = E_{\theta^i}(\boldsymbol{z})$ (a and c, resp.) and relative error $\left\|\theta^i - \theta^*\right\| / \|\theta^*\|$ (b and d, resp.). 200 iterations of optimization (only every tenth iteration has a marker) and 500 Metropolis-Hastings samples (only every fourth sample has a marker).

and Hessian estimators.

7.3. Summary of Experiments

On all three benchmark ODEs, the Jacobian and Hessian estimator proved useful to speed up both sampling and optimization. In the case of optimization, the new gradient-based methods consistently outperformed the classical random search. Notably, the second-order optimization was always significantly more sample-efficient than plain gradient descent—which indicates that not only the gradient but also the Hessian estimator is accurate enough to be useful. In the case of sampling, the gradient-based sampling methods, which were preconditioned by the Hessian, consistently outperformed the classical approach as well: PLMC and PHMC steadily explored regions of elevated likelihood, while the conventional random-walk Metropolis methods hardly ever reached regions of nonzero probability and wasted computational budget on less likely parameters. Overall, we consider these experiments first evidence for the hypothesis that the proposed gradient-based methods require drastically fewer samples than the standard 'likelihood-free' approach.

8. Related and Future Work

The following research areas are particularly closely related to this paper.

Probabilistic numerical methods (PNMs) There are two lines of work on PNMs for ODE forward problems: sampling- and filtering-based solvers; an up-to-date comparative discussion of these two approaches is given in Kersting et al. (2019, Section 1.2.). While this paper is the first to use filtering-based PNMs for inverse problems, there are previous methods—starting with Chkrebtii et al. (2016)—that use sampling-based solvers to integrate a non-Gaussian uncertainty-aware likelihood (cf. the Gaussian eq. (9)) into a pseudo-marginal MCMC framework; see Conrad et al. (2017), Teymur et al. (2018), Lie et al. (2019), and Abdulle & Garegnani (2020). Notably, Matsuda & Miyatake (2019) recently proposed to model the numerical errors as random variables without explicitly employing PNMs. On a related note, there are also first PNMs for PDE inverse problems; see Cockayne et al. (2017) and Oates et al. (2019).

GP-surrogate methods Modelling expensive likelihoods by GP regression is a common approach in statistics; see e.g. Sacks et al. (1989) and O'Hagan (2006). Notably, Meeds & Welling (2014) incorporated this approach into an ABC framework, and Perdikaris & Karniadakis (2016),

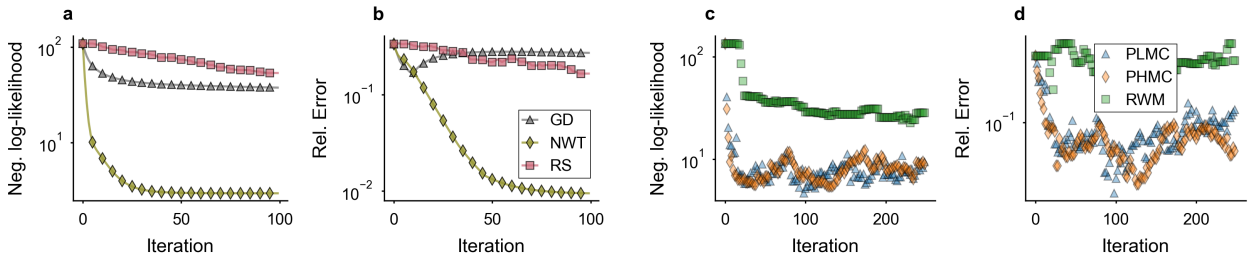

Figure 6. **Results for optimization (a, b) and sampling (c, d) on GUiY**. Comparison of negative log-likelihood $E(\boldsymbol{z}) = E_{\theta^i}(\boldsymbol{z})$ (a and c, resp.) and relative error $\left\|\theta^i - \theta^*\right\| / \|\theta^*\|$ (b and d, resp.). 100 iterations of optimization (only every fifth iteration has a marker) and 250 Metropolis-Hastings samples (only every other sample has a marker).

on the other hand, into a non-Bayesian setting by efficient global optimization. While these methods also compute a GP approximation to the likelihood, they are fundamentally different as they globally model the likelihood with a GP (instead of constructing a local Gaussian approximation (see eq. (9)), and do not exploit the shape of the ODE at all.

Gradient Matching This approach fits a joint GP model of the solution and its derivatives by conditioning on the ODE. Since introduced by Calderhead et al. (2008), it has received much attention in machine learning; see Macdonald & Husmeier (2015) for a detailed review, Wenk et al. (2019, Section 1) for an up-to-date overview, and Gorbach et al. (2017) for a paper that uses a slightly stronger version of our Assumption 1. As it avoids explicit numerical integration altogether, gradient matching is fundamentally different from our method (and PNMs in general).

Sensitivity analysis This field studies the derivatives of ODE solutions with respect to parameters; see, e.g., Rackauckas et al. (2018) for an overview spanning continuous (adjoint) sensitivity analysis and automatic differentiation. Therefore, the Jacobian estimator J of the map $\theta \mapsto \boldsymbol{m}_\theta \approx \boldsymbol{x}_\theta$ from eq. (11) can be interpreted as fast, approximate sensitivity analysis. This link is particularly interesting for modern machine learning, as sensitivity analysis is the mathematical corner stone of the recent advances by, e.g., Chen et al. (2018) in training neural networks as ODEs. It should be possible to use J for neural ODEs—as well as for all other applications of sensitivity analysis.

Future Work

We hope that this is the beginning of a new line of work on ODE inverse problems by ODE filtering. Here, we only used Gaussian ODE filtering with once-integrated Brownian motion prior. Future work could not only examine different priors (Kersting et al., 2019, Section 2.1), but also draw from the wide range of additional ODE filters (EKF, UKF, particle filter, etc.) that were unlocked by Tronarp et al. (2019). Notably, particle ODE filtering represents the belief over the ODE solution by a set of samples (particles), and could, therefore, be integrated in the above-mentioned existing framework for sampling-based PNMs.

The utility of the Jacobian estimator J is, however, not limited to inverse problems. As it constitutes fast, approximate sensitivity analysis, it should be compared with established methods, such as automatic differentiation and continuous sensitivity analysis (Rackauckas et al., 2018). If S (eq. (20)) could also be estimated with low overhead, it is in light of eq. (19) conceivable that the approximation error of J could be further reduced.

Either way, future work should examine which optimization and sampling methods are optimal—given that they received the (approximate) gradient and Hessian estimators $[\hat{\nabla}_\theta E(\boldsymbol{z}), \hat{\nabla}^2_\theta E(\boldsymbol{z})]$. For instance, the approximation error on these estimators might—according to Bottou et al. (2018, Section 3.3)—warrant optimization by stochastic methods such as SGD. On a related note, it should be examined whether classical theorems on limit behavior of the employed optimization and MCMC methods remain true when using these estimators, and whether our approach is indeed applicable to ODEs that violate Assumption 1—as the results from Section 7.2.2 suggest. Finally, this work should be, by the methods of lines (Schiesser & Griffiths, 2009), extendable to PDEs and, by John et al. (2019), to boundary value problems.

9. Concluding Remarks

We introduced a novel Jacobian estimator for ODE solutions w.r.t. their parameters which implies approximate estimators of the gradient and Hessian of the log-likelihood. Using these estimators, we proposed new first and second-order optimization and sampling methods for ODE inverse problems which outperformed standard 'likelihood-free' approaches—namely random search optimization and random-walk Metropolis MCMC—in all conducted experiments. Moreover, the employed Jacobian estimator constitutes a new method for fast, approximate sensitivity analysis.

Acknowledgements

We thank the anonymous reviewers for their careful, constructive comments. We thank Filip Tronarp and Katharina Ott for helpful feedback.

HK, NK and PH gratefully acknowledge financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence Machine Learning - New Perspectives for Science, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg. HK, NK and PH also gratefully acknowledge financial support by the German Federal Ministry of Education and Research (BMBF) through Project ADIMEM (FKZ 01IS18052B).

References

Abdulle, A. and Garegnani, G. Random time step probabilistic methods for uncertainty quantification in chaotic and geometric numerical integration. *Statistics and Computing*, pp. 1–26, 2020.

Betancourt, M. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv:1701.02434 [stat.ME]*, 2017.

Bottou, L., Curtis, F., and Nocedal, J. Optimization methods for large-scale machine learning. *SIAM Review*, 2018.

Calderhead, B., Girolami, M., and Lawrence, N. Accelerating Bayesian inference over nonlinear differential equations with Gaussian processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2008.

Chen, R., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Chkrebtii, O. A., Campbell, D. A., Calderhead, B., and Girolami, M. A. Bayesian solution uncertainty quantification for differential equations. *Bayesian Analysis*, 11 (4):1239–1267, 2016.

Cockayne, J., Oates, C., Sullivan, T., and Girolami, M. Probabilistic numerical methods for PDE-constrained Bayesian inverse problems. In *AIP Conference Proceedings*, volume 1853, pp. 060001. AIP Publishing LLC, 2017.

Conrad, P. R., Girolami, M., Särkkä, S., Stuart, A., and Zygalakis, K. Statistical analysis of differential equations: introducing probability measures on numerical solutions. *Statistics and Computing*, 2017.

Cranmer, K., Brehmer, J., and Louppe, G. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 2020.

Girolami, M. and Calderhead, B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 2011.

Gorbach, N. S., Bauer, S., and Buhmann, J. M. Scalable variational inference for dynamical systems. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Hartikainen, J. and Särkkä, S. Kalman filtering and smoothing solutions to temporal Gaussian process regression models. In *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2010.

Hennig, P., Osborne, M. A., and Girolami, M. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 471(2179), 2015.

Hull, T., Enright, W., Fellen, B., and Sedgwick, A. Comparing numerical methods for ordinary differential equations. *SIAM Journal on Numerical Analysis*, 9(4):603–637, 1972.

John, D., Heuveline, V., and Schober, M. GOODE: A Gaussian off-the-shelf ordinary differential equation solver. In *International Conference on Machine Learning (ICML)*, 2019.

Kersting, H., Sullivan, T. J., and Hennig, P. Convergence rates of Gaussian ODE filters. *arXiv:1807.09737v2 [math.NA]*, 2019.

Lie, H. C., Stuart, A. M., and Sullivan, T. J. Strong convergence rates of probabilistic integrators for ordinary differential equations. *Statistics and Computing*, 29(6): 1265–1283, 2019.

Lotka, A. The growth of mixed populations: two species competing for a common food supply. *The Golden Age of Theoretical Ecology: 19231940*, 1978.

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, 2015.

MacKay, D. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

Matsuda, T. and Miyatake, Y. Estimation of ordinary differential equation models with discretization error quantification. *arXiv:1907.10565 [stat.ME]*, 2019.

Meeds, E. and Welling, M. GPS-ABC: Gaussian process surrogate approximate Bayesian computation. In *Uncertainty in Artificial Intelligence (UAI)*, 2014.

Oates, C., Cockayne, J., Aykroyd, R., and Girolami, M. Bayesian probabilistic numerical methods in time-dependent state estimation for industrial hydrocyclone equipment. *Journal of the American Statistical Association*, 2019.

Oates, C. J. and Sullivan, T. J. A modern retrospective on probabilistic numerics. *Statistics and Computing*, 29(6): 1335–1351, 2019.

O'Hagan, A. Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering & System Safety*, 2006.

Perdikaris, P. and Karniadakis, G. E. Model inversion via multi-fidelity Bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. *Journal of the Royal Society Interface*, 13, 2016.

Qi, Y. and Minka, T. Hessian-based Markov chain Monte-Carlo algorithms. *Workshop on Monte Carlo Methods*, 2002.

Rackauckas, C., Ma, Y., Dixit, V., Guo, X., Innes, M., Revels, J., Nyberg, J., and Ivaturi, V. A comparison of automatic differentiation and continuous sensitivity analysis for derivatives of differential equation solutions. *arXiv:1812.01892 [math.NA]*, 2018.

Rastrigin, L. A. The convergence of the random search method in the extremal control of a many parameter system. *Automation and Remote Control*, 24, 1963.

Roberts, G. and Tweedie, R. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.

Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. Design and analysis of computer experiments. *Statistical Science*, 1989.

Schiesser, W. E. and Griffiths, G. W. *A Compendium of Partial Differential Equation Models: Method of Lines Analysis with Matlab*. Cambridge University Press, 1st edition, 2009.

Schillings, C., Sunnaker, M., Stelling, J., and Schwab, C. Efficient characterization of parametric uncertainty of complex (bio)chemical networks. *PLOS Computational Biology*, 11, 2015.

Schober, M., Särkkä, S., and Hennig, P. A probabilistic model for the numerical solution of initial value problems. *Statistics and Computing*, 29(1):99–122, 2019.

Tarantola, A. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2005.

Teymur, O., Lie, H. C., Sullivan, T. J., and Calderhead, B. Implicit probabilistic integrators for ODEs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Tronarp, F., Kersting, H., Särkkä, S., and Hennig, P. Probabilistic solutions to ordinary differential equations as nonlinear Bayesian filtering: a new perspective. *Statistics and Computing*, 29(6):1297–1315, 2019.

Vyshemirsky, V. and Girolami, M. A. Bayesian ranking of biochemical system models. *Bioinformatics*, 24(6): 833–839, 2008.

Wenk, P., Abbati, G., Bauer, S., Osborne, M. A., Krause, A., and Schölkopf, B. ODIN: ODE-informed regression for parameter and state inference in time-continuous dynamical systems. In *AAAI Conference on Artificial Intelligence*, 2019.