
Active Uncertainty Calibration in Bayesian ODE Solvers

Hans Kersting and Philipp Hennig

Max-Planck-Institute for Intelligent Systems
Spemannstraße, 72076 Tübingen, Germany
[hkersting|phennig]@tue.mpg.de

Abstract

There is resurging interest, in statistics and machine learning, in solvers for ordinary differential equations (ODEs) that return probability measures instead of point estimates. Recently, Conrad et al. introduced a sampling-based class of methods that are ‘well-calibrated’ in a specific sense. But the computational cost of these methods is significantly above that of classic methods. On the other hand, Schober et al. pointed out a precise connection between classic Runge-Kutta ODE solvers and Gaussian filters, which gives only a rough probabilistic calibration, but at negligible cost overhead. By formulating the solution of ODEs as approximate inference in linear Gaussian SDEs, we investigate a range of probabilistic ODE solvers, that bridge the trade-off between computational cost and probabilistic calibration, and identify the inaccurate gradient measurement as the crucial source of uncertainty. We propose the novel filtering-based method Bayesian Quadrature filtering (BQF) which uses Bayesian quadrature to actively learn the imprecision in the gradient measurement by collecting multiple gradient evaluations.

1 INTRODUCTION

The numerical solution of an initial value problem (IVP) based on an *ordinary differential equation* (ODE)

$$u^{(n)} = f(t, u(t), u'(t), \dots, u^{(n-1)}(t)), u(0) = u_0 \in \mathbb{R}^D \quad (1)$$

of order $n \in \mathbb{N}$, with $u : \mathbb{R} \rightarrow \mathbb{R}^D$, $f : [0, T] \times \mathbb{R}^{nD} \rightarrow \mathbb{R}^D$, $T > 0$, is an essential topic of numerical mathematics, because ODEs are the standard model for dynamical systems. Solving ODEs with initial values is an exceedingly well-studied problem (see Hairer et al., 1987, for a comprehensive presentation) and modern solvers are designed

very efficiently. Usually, the original ODE (1) of order n is reduced to a system of n ODEs of first order

$$u'(t) = f(t, u(t)), u(0) = u_0 \in \mathbb{R}^D, \quad (2)$$

which are solved individually. The most popular solvers in practice are based on some form of Runge-Kutta (RK) method (as first introduced in Runge (1895) and Kutta (1901)) which employ a weighted sum of a fixed amount of gradients in order to iteratively extrapolate a discretized solution. That is, these methods iteratively collect ‘observations’ of approximate gradients of the solved ODE, by evaluating the dynamics f at an estimated solution, which is a linear combination of previously collected ‘observations’:

$$y_i = f \left(t_i, u_0 + \sum_{j < i} w_{ij} y_j \right). \quad (3)$$

The weights of s -stage RK methods of p -th order are carefully chosen so that the numerical approximation \hat{u} and the Taylor series of the true solution u coincide up to the p -th term, thereby yielding a local truncation error of high order, $\|u(t_0 + h) - \hat{u}(t_0 + h)\| = \mathcal{O}(h^{p+1})$, for $h \rightarrow 0$. One can prove that $s \geq p$ in general, but for $p \leq 4$ there are RK methods with $p = s$. Hence, allowing for more function evaluations can drastically improve the speed of convergence to the true solution.

The polynomial convergence is impressive and helpful; but it does not actually quantify the true epistemic uncertainty about the accuracy of the approximate solution \hat{u} for a concrete non-vanishing step-size h . One reason one may be concerned about this in machine learning is that ODEs are often one link of a chain of algorithms performing some statistical analysis. When employing classic ODE solvers and just plugging in the solution of the numerical methods in subsequent steps, the resulting uncertainty of the whole computation is ill-founded, resulting in overconfidence in a possibly wrong solution. It is thus desirable to model the epistemic uncertainty. Probability theory provides the framework to do so. Meaningful probability measures of the uncertainty about the result of determin-

istic computations (such as ODE solvers) can then be combined with probability measures modeling other sources of uncertainty, including ‘real’ aleatoric randomness (from e.g. sampling). Apart from quantifying our certainty over a computation, pinning down the main sources of uncertainty could furthermore improve the numerical solution and facilitate a more efficient allocation of the limited computational budget.

A closed framework to measure uncertainty over numerical computations was proposed by Skilling (1991) who pointed out that numerical methods can be recast as statistical inference of the latent true solution based on the observable results of tractable computations. Applying this idea to ordinary differential equations, Hennig & Hauberg (2014) phrased this notion more formally, as Gaussian process (GP) regression. Their algorithm class, however, could not guarantee the high polynomial convergence orders of Runge-Kutta methods. In parallel development, Chkrebtii et al. (2013) also introduced a probabilistic ODE solver of similar structure (i.e. based on a GP model), but using a Monte Carlo updating scheme. These authors showed a linear convergence rate of their solver, but again not the high-order convergence of classic solvers.

Recently, Schober et al. (2014) solved this problem by finding prior covariance functions which produce GP ODE solvers whose posterior means *exactly* match those of the optimal Runge-Kutta families of first, second and third order. While producing only a slight computational overhead compared to classic Runge-Kutta, this algorithm — as any GP-based algorithm — only returns Gaussian measures over the solution space.

In contrast, Conrad et al. (2015) recently provided a novel sampling-based class of ODE solvers which returns asymptotically non-parametric measures over the solution space, but creates significant computational overhead by running the whole classic ODE solvers multiple times over the whole time interval $[0, T]$ in order to obtain meaningful approximations for the desired non-parametric measure.

For practitioners, there is a trade-off between the desire for quantified uncertainty on the one hand, and low computational cost on the other. The currently available probabilistic solvers for ODEs either provide only a roughly calibrated uncertainty (Schober et al., 2014) at negligible overhead or a more fine-grained uncertainty supported by theoretical analysis (Conrad et al., 2015), at a computational cost increase so high that it rules out most practical applications. In an attempt to remedy this problem, we propose an algorithm enhancing the method of Schober et al. (2014) by improving the gradient measurement using modern probabilistic integration methods. By modeling the uncertainty where it arises, i.e. the imprecise prediction of where to evaluate f , we hope to gain better knowledge of the propagated uncertainty and arrive at well-calibrated posterior variances as uncertainty measures.

2 BACKGROUND

2.1 SAMPLING-BASED ODE SOLVERS

The probabilistic ODE solver by Conrad et al. (2015) modifies a classic deterministic one-step numerical integrator Ψ_h (e.g. Runge-Kutta or Taylor methods, cf. Hairer et al. (1987)) and models the discretization error of Ψ_h by adding suitably scaled i.i.d. Gaussian random variables $\{\xi_k\}_{k=0,\dots,K}$ after every step. Hence, it returns a discrete solution $\{U_k\}_{k=0,\dots,K}$ on a mesh $\{t_k = kh\}_{k=0,\dots,K}$ according to the rule

$$U_{k+1} = \Psi_h(U_k) + \xi_k. \quad (4)$$

This discrete solution can be extended into a continuous time approximation of the ODE, which is random by construction and can therefore be interpreted as a draw from a non-parametric probability measure Q_h on the solution space $C([0, T], \mathbb{R}^n)$, the Banach space of continuous functions. This probability measure can then be interpreted as a notion of epistemic uncertainty about the solution. This is correct in so far as, under suitable assumptions, including a bound on the variance of the Gaussian noise, the method converges to the true solution, in the sense that Q_h contracts to the Dirac measure on the true solution δ_u with the *same* convergence rate as the original numerical integrator Ψ_h , for $h \rightarrow 0$. This is a significant step towards a well-founded notion of uncertainty calibration for ODE solvers: It provides a probabilistic extension to classic method which does not break the convergence rate of these methods.

In practice, however, the precise shape of Q_h is *not* known and Q_h can only be interrogated by sampling, i.e. repeatedly running the entire probabilistic solver. After S samples, Q_h can be approximated by an empirical measure $Q_h(S)$. In particular, the estimated solution and uncertainty can only be expressed in terms of statistics of $Q_h(S)$, e.g. by the usual choices of the empirical mean and empirical variance respectively or alternatively by confidence intervals. For $S \rightarrow \infty$, $Q_h(S)$ converges in distribution to Q_h which again converges in distribution to δ_u for $h \rightarrow 0$:

$$Q_h(S) \xrightarrow{S \rightarrow \infty} Q_h \xrightarrow{h \rightarrow 0} \delta_u. \quad (5)$$

The theoretical analysis in Conrad et al. (2015) only concerns the convergence of the latent probability measures $\{Q_h\}_{h>0}$. Only the empirical measures $\{Q_h(S)\}_{S \in \mathbb{N}}$, however, can be observed. Consequently, it remains unclear whether the empirical mean of $Q_h(S)$ for a fixed step-size $h > 0$ converges to the true solution as $S \rightarrow \infty$ and whether the empirical variance of $Q_h(S)$ is directly related, in an analytical sense, to the approximation error. In order to extend the given convergence results to the practically observable measures $\{Q_h(S)\}_{S \in \mathbb{N}}$ an analysis of the first convergence in (5) remains missing. The deterministic algorithm proposed below avoids this problem, by instead

constructing a (locally parametric) measure from prior assumptions.

The computational cost of this method also seems to mainly depend on the rate of convergence of $Q_h(S) \rightarrow Q_h$ which determines how many (possibly expensive) runs of the numerical integrator Ψ_h over $[0, T]$ have to be computed and how many samples have to be stored for a sufficient approximation of Q_h . Furthermore, we expect that in practice the mean of Q_h , as approximated by $Q_h(S)$ might not be the best possible approximation, since in one step the random perturbation of the predicted solution by Gaussian noise ξ_k worsens our solution estimate with a probability of more than $1/2$, since - due to the full support of Gaussian distributions - the numerical sample solution is as likely to be perturbed away from as towards the true solution and - due to the tails of Gaussian distributions - it can also be perturbed way past the true solution with positive probability.

2.2 A FRAMEWORK FOR GAUSSIAN FILTERING FOR ODES

Describing the solution of ODEs as inference in a joint Gaussian model leverages state-space structure to achieve efficient inference. Therefore, we employ a Gauss-Markov prior on the state-space: *A priori* we model the solution function and $(q-1)$ derivatives $(u, \dot{u}, u^{(2)}, \dots, u^{(q-1)}) : [0, T] \rightarrow \mathbb{R}^{qD}$ as a draw from a q -times integrated Wiener process $X = (X_t)_{t \in [0, T]} = (X_t^{(1)}, \dots, X_t^{(q)})_{t \in [0, T]}$, i.e. the dynamics of X_t are given by the linear Itô stochastic differential equation (Karatzas & Shreve, 1991; Øksendal, 2003):

$$\begin{aligned} dX_t &= F \cdot X_t dt + Q \cdot dW_t, \\ X_0 &= \xi, \quad \xi \sim \mathcal{N}(m(0), P(0)), \end{aligned} \quad (6)$$

with constant drift $F \in \mathbb{R}^{q \times q}$ and diffusion $Q \in \mathbb{R}^q$ given by

$$F = \begin{pmatrix} 0 & f_1 & 0 & \dots & 0 \\ 0 & 0 & f_2 & \dots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & f_{q-1} & \\ 0 & \dots & 0 & 0 & \end{pmatrix}, Q = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \sigma^2 \end{pmatrix} \quad (8)$$

for all $t \in [0, T]$ and some $f_1, \dots, f_{q-1} \in \mathbb{R}$, where W_t denotes a q -dimensional Wiener process ($q \geq n$). Hence, we are a priori expecting that $u^{(q)}$ behaves like a Brownian motion with variance σ^2 and that $u^{(i)}$ is modeled by $(q-1-i)$ -times integrating this Brownian motion. The fact that the $(i+1)$ -th component is the derivative of the i -th component in our state space is captured by a drift matrix with non-zero entries only on the first off-diagonal. The entries f_1, \dots, f_{q-1} are damping factors. A standard choice is e.g. $f_i = i$. Without additional information, it seems natural to put white noise on the q -th derivative as the first

derivative which is not captured in the state space. This gives rise to Brownian noise on the $(q-1)$ -th derivative which is encoded in the diffusion matrix scaled by variance σ^2 . Hence, we consider the integrated Wiener process a natural prior. For notational simplicity, only the case of scalar-valued functions, i.e. $D = 1$, is presented in the following. The framework can be extended to $D \geq 2$ in a straightforward way by modeling the output dimensions of f as independent stochastic processes.

Since X is the strong solution of a linear equation (6) with normally distributed initial value X_0 , it follows from the theory of linear SDEs (Karatzas & Shreve, 1991) that X is a uniquely-determined Gauss-Markov process. This enables Bayesian inference in a highly efficient way by Gaussian filtering (Saatici, 2011). For time invariant linear SDEs like (6), the fixed matrices for Gaussian filtering can be precomputed analytically (Särkkä, 2006).

In addition, Schober et al. (2014) showed that for $q \leq 3$ inference in this linear SDE yields Runge-Kutta steps.

Equipped with this advantageous prior we can perform Bayesian inference. The linearity and time-invariance of the underlying SDE permits to formulate the computation of the posterior as a Kalman filter (KF) (cf. (Särkkä, 2013) for a comprehensive introduction) with step size $h > 0$. The prediction step of the KF is given by

$$m_{t+h}^- = A(h)m_t, \quad (9)$$

$$P_{t+h}^- = A(h)P_t A(h)^T + Q(h), \quad (10)$$

with matrices $A(h), Q(h) \in \mathbb{R}^{q \times q}$ with entries

$$A(h)_{i,j} = \exp(h \cdot F)_{i,j} = \chi_{j \geq i} \frac{h^{j-i}}{(j-i)!} \left(\prod_{k=0}^{j-i-1} f_{i+k} \right), \quad (11)$$

$$Q(h)_{i,j} = \sigma^2 \cdot \left(\prod_{k_1=0}^{q-1+i} f_{i+k_1} \right) \cdot \left(\prod_{k_2=0}^{q-1+j} f_{j+k_2} \right) \cdot \frac{h^{2q+1-i-j}}{(q-i)!(q-j)!(2q+1-i-j)}, \quad (12)$$

and followed by the update step

$$z = y - H m_{t+h}^-, \quad (14)$$

$$S = H P_{t+h}^- H^T + R, \quad (15)$$

$$K = P_{t+h}^- H^T S^{-1}, \quad (16)$$

$$m_{t+h} = m_{t+h}^- + K z, \quad (17)$$

$$P_{t+h} = P_{t+h}^- - K H P_{t+h}^-, \quad (18)$$

where $H = e_n^T \in \mathbb{R}^{1 \times q}$ is the n -th unit vector.

Between the prediction and update step the n -th derivative of the true solution $\frac{\partial^n u}{\partial x^n}$ at time $t+h$ as a measurement for the n -th derivative and the noise of this measurement

are estimated by the variable y and R . In order to derive precise values of y and R from the Gaussian prediction $\mathcal{N}(m_{t+h}, P_{t+h})$, we would have to compute the integrals

$$y = \int f(t+h, m_{t+h}^- + x) \cdot \mathcal{N}(x; 0, P_{t+h}^-) dx \quad (19)$$

and

$$R = \int f(t+h, m_{t+h}^- + x) \cdot f(t+h, m_{t+h}^- + x)^T \cdot \mathcal{N}(x; 0, P_{t+h}^-) dx - yy^T, \quad (20)$$

which are intractable for most choices of f . Below we investigate different ways to address the challenge of computing these integrals precisely while not creating too much computational overhead.

2.3 MEASUREMENT GENERATION OPTIONS FOR GAUSSIAN FILTERING

Schober et al. (2014) as, to the best of our knowledge, the first ones to point out the connection between Gaussian filtering and probabilistic ODE solvers, presents an algorithm which simply evaluates the gradient at the predicted mean, which is equivalent to setting y to be equal to its maximum likelihood estimator:

$$y = f(t+h, m_{t+h}^-), \quad R = 0. \quad (21)$$

While ensuring maximum speed, this is clearly not an ideal measurement. In our atomless predicted probability measure $\mathcal{N}(m_{t+h}^-, P_{t+h}^-)$ the mean predictor m_{t+h}^- is different from its true value $(u^{(0)}(t+h), \dots, u^{(n)}(t+h))^T$ almost surely. Hence, for a non-constant f the estimate will be inaccurate most of the times. In particular this method deals poorly with ‘skewed’ gradient fields (a problem that leads to a phenomenon known as ‘Lady Windermere’s fan’ (Hairer et al., 1987)). To get a better estimate of the true value of y , more evaluations of f seem necessary.

Therefore, we want to find numerical integration methods which capture y and R with sufficient precision, while using a minimal number of evaluations of f . Possible choices are:

(i) *Monte Carlo integration by sampling:*

$$y = \frac{1}{N} \sum_{i=1}^N f(t+h, x_i), \quad (22)$$

$$R = \frac{1}{N} \sum_{i=1}^N f(t+h, x_i) \cdot f(t+h, x_i)^T - y \cdot y^T, \quad (23)$$

$$x_i \sim \mathcal{N}(m_{t+h}^-, P_{t+h}^-), \quad (24)$$

(which is *not* the same as the sampling over the whole time axis in (Conrad et al., 2015)).

(ii) *Approximation by a first-order Taylor series expansion:*

$$\begin{aligned} f(t+h, m_{t+h}^- + x) \\ \simeq f(t+h, m_{t+h}^-) + \nabla f(t+h, m_{t+h}^-) \cdot x \end{aligned} \quad (25)$$

and thereby deriving moments of the linear transform of Gaussian distributions:

$$y = f(t+h, m_{t+h}^-), \quad (26)$$

$$R = \nabla f(t+h, m_{t+h}^-) P_{t+h}^- \nabla f(t+h, m_{t+h}^-)^T. \quad (27)$$

(iii) *Integration by Bayesian quadrature with Gaussian weight function:*

$$y = \alpha^T \cdot K^{-1} \cdot (f(x_1), \dots, f(x_n))^T, \quad (28)$$

$$R = \int \int k(x, x') w(x) w(x') dx dx' - \alpha^T K^{-1} \alpha. \quad (29)$$

with $w(x) = \mathcal{N}(x; m_{t+h}^-, P_{t+h}^-)$, kernel matrix $K \in \mathbb{R}^{N \times N}$ with $K_{i,j} = k(x_i, x_j)$ and $\alpha = (\alpha(1), \dots, \alpha(N))^T \in \mathbb{R}^N$ with $\alpha(i) = \int k(x, x_i) w(x) dx$ for a predefined covariance function k and evaluation points $(x_i)_{i=1, \dots, N}$ (cf. section 2.4).

Our experiments, presented in Section 3, suggest that BQ is the most useful option.

Monte Carlo integration by sampling behaves poorly if the trajectory of the numerical solution passes through domain areas (as e.g. in the spikes of oscillators governed by non-stiff ODEs) where f takes highly volatile values since the random spread of samples from the domain are likely to return a skewed spread of values resulting in bad predictions of y with huge uncertainty R . Hence, the posterior variance explodes and the mean drifts back to its zero prior mean, i.e. $m_t \rightarrow 0$ and $\|P_t\| \rightarrow \infty$, for $t \rightarrow \infty$. Thus, we consider this method practically useless.

One may consider it a serious downside of Taylor-approximation based methods that the gradient only approximates the shape of f and thereby its mapping of the error on an ‘infinitesimally small neighborhood’ of m_{t+h}^- . Hence, it might ignore the true value of y completely, if the mean prediction is far off. However, for a highly regular f (e.g. Lipschitz-continuous in the space variable) this gradient approximation is very good.

Moreover, the approximation by a first-order Taylor series expansion needs an approximation of the gradient, which explicit ODE solvers usually do not receive as an input. However, in many numerical algorithms (e.g. optimization) the gradient is provided anyway. Therefore the gradient

might already be known in real-world applications. While we find this method promising when the gradient is known or can be efficiently computed, we exclude it from our experiments because the necessity of a gradient estimate breaks the usual framework of ODE solvers.

In contrast, Bayesian quadrature avoids the risk of a skewed distortion of the samples for Monte Carlo integration by actively spreading a grid of deterministic sigma-points. It does not need the gradient of f and still can encode prior knowledge over f by the choice of the covariance function if more is known (Briol et al., 2015). The potential of using Bayesian quadrature as a part of a filter was further explored by Prüher & Simandl (2015), however in the less structured setting of nonlinear filtering where additional inaccuracy from the linear approximation in the prediction step arises. Moreover, Särkkä et al. (2015) recently pointed out that GPQ can be seen as sigma-point methods and gave covariance functions and evaluation points which reproduce numerical integration methods known for their favorable behavior (for example Gauss-Hermite quadrature, which is used for a Gaussian weight function).

Due to these advantages, we propose a new class of BQ-based probabilistic ODE filters named BQ Filtering.

2.4 BAYESIAN QUADRATURE FILTERING (BQF)

The crucial source of error for filtering-based ODE solvers is the calculation of the gradient measurement y and its variance R (c.f. Section 2.2). We propose the novel approach to use BQ to account for the uncertainty of the input and thereby estimate y and R . This gives rise a novel class of filtering-based solvers named *BQ Filter* (BQF). As a filtering-based method, one BQF-step consists of the KF prediction step (9) – (10), the calculation of y and R by BQ and the KF update step (14) – (18).

The KF prediction step outputs a Gaussian belief $\mathcal{N}(m_{t+h}^-, P_{t+h}^-)$ over the true solution $u(t+h)$. This input value is propagated through f yielding a distribution over the gradient at time $t+h$. In other words, our belief over $\nabla f(t+h, u(t+h))$ is equal to the distribution of $Y := f(t, X)$, with uncertain input $X \sim \mathcal{N}(m_{t+h}^-, P_{t+h}^-)$. For general f the distribution of Y will be neither Gaussian nor unimodal (as e.g. in Figure 1). But it is possible to compute the moments of this distribution under Gaussian assumptions on the input and the uncertainty over f (see for example Deisenroth (2009)). The equivalent formulation of prediction under uncertainty clarifies as numerical integration clarifies the connection to sigma-point methods, i.e. quadrature rules (Särkkä et al., 2015). Quadrature is as extensively studied and well-understood as the solution of ODEs. A basic overview can be found in Press et al. (2007). Marginalizing over X yields an integral with Gaus-

sian weight function

$$\mathbb{E}[Y] = \int f(t+h, x) \cdot \underbrace{\mathcal{N}(x; m_{t+h}^-, P_{t+h}^-)}_{=:w(x)} dx, \quad (30)$$

which is classically solved by quadrature, i.e. evaluating f at a number of evaluation points $(x_i)_{i=1, \dots, N}$ and calculating a weighted sum of these evaluations. BQ can be interpreted as a probabilistic extension of these quadrature rules in the sense that their posterior mean estimate of the integral coincides with classic quadrature rules, while adding a posterior variance estimate at low cost (Särkkä et al., 2015).

By choosing a kernel k over the input space of f and evaluation points $(x_i)_{i=1, \dots, N}$, the function f is approximated by a GP regression (Rasmussen & Williams, 2006) with respect to the function evaluations $(f(x_i))_{i=1, \dots, N}$, yielding a GP posterior over f with mean m_f and covariance k_f denoted by $\mathcal{GP}(f)$. The integral is then approximated by integrating the GP approximation, yielding the predictive distribution for $\mathcal{I}[f]$:

$$\mathcal{I}[f] \sim \int \mathcal{GP}(f)(x) \cdot \mathcal{N}(x; m_{t+h}^-, P_{t+h}^-) dx. \quad (31)$$

The uncertainty arising from the probability measure over the input is now split up in two parts: the uncertainty over the input value $x \sim \mathcal{N}(0, I)$ and the uncertainty over the precise value at this uncertain input, which can only be approximately inferred by its covariance with the evaluation points $(x_i)_{i=1, \dots, N}$, i.e. by $\mathcal{GP}(f)$. These two kinds of uncertainty are depicted in Figure 1. From the predictive distribution in (31), we can now compute a posterior mean and variance of $\mathcal{I}[f]$ which results in a weighted sum for the mean

$$y := \mathbb{E}[\mathcal{I}[f]] = \alpha^T K^{-1} (f(x_1), \dots, f(x_n))^T \quad (32)$$

with

$$\alpha(i) = \int k(x, x_i) \mathcal{N}(x; 0, I) dx \quad (33)$$

and variance

$$R := \text{Var}[\mathcal{I}(f)] \quad (34)$$

$$= \int \int k(x, x') w(x) w(x') dx dx' - \alpha^T K^{-1} \alpha, \quad (35)$$

where $K \in \mathbb{R}^{N \times N}$ denotes the kernel matrix, i.e. $K_{i,j} = k(x_i, x_j)$.

The measurement generation in BQF is hence completely defined by the two free choices of BQ: the kernel k and the evaluation points $(x_i)_{i=1, \dots, n}$. By these choices, BQ and thereby the measurement generation in BQF is completely defined. For the squared exponential kernel (Rasmussen & Williams, 2006)

$$k(x, x') = \theta^2 \exp\left(-\frac{1}{2\lambda^2} \|x - x'\|^2\right), \quad (36)$$




Figure 1: Prediction of function $f(x) = 8 \sin(x) + x^2$ (red) under uncertain input $x \sim \mathcal{N}(x; 1, 1)$ (density in blue). $\mathcal{GP}(f)$ (black) derived from Gaussian grid evaluation points with $N = 3$ (blue crosses) as mean ± 2 standard deviation. True distribution of prediction in blue. Gaussian fit to true distribution in yellow and predicted distribution by BQ in green with crosses at means.

with lengthscale $\lambda > 0$ and output variance $\theta^2 > 0$, it turns out that y and R can be computed in closed form and that many classic quadrature methods which are known for their favorable properties can be computed in closed form (Särkkä et al., 2015), significantly speeding up computations. For the scalar case $nD = 1$, we obtain for (33) by straightforward computations:

$$\alpha(i) = \frac{\lambda \theta^2}{\sqrt{\lambda^2 + \sigma^2}} \exp\left(-\frac{(x_i - \mu)^2}{2(\lambda^2 + \sigma^2)}\right), \quad (37)$$

and

$$\int \int k(x, x') w(x) w(x') dx dx' = \frac{\theta^2}{\sqrt{1 + 2\sigma^2/\lambda^2}} \quad (38)$$

Hence, our BQ estimate for y is given by the sigma-point rule

$$y \approx \sum_{i=1}^N W_i f(t+h, x_i) \quad (39)$$

with easily computable weights

$$W_i = [\alpha \cdot K]_i. \quad (40)$$

Also the variance R takes a convenient shape

$$R = \frac{\theta^2}{\sqrt{1 + 2\sigma^2/\lambda^2}} - \alpha^T K^{-1} \alpha. \quad (41)$$

For $nD > 1$, we get slightly more complicated formulas which are given in Deisenroth (2009).

The other free choice in BQ, the evaluation points $(x_i)_{i=1, \dots, n}$, can also be chosen freely in every step of BQ. Usually, the nodes of BQ chosen are chosen so that the variance of the integral estimate is minimized (cf. Briol et al. (2015)). For this algorithm, the uncertainty has to be measured, not minimized though. Hence, we propose just to take a uniform grid scaled by $\mathcal{N}(m_{t+h}^-, P_{t+h}^-)$ to measure the uncertainty in a comprehensive way.

Another promising choice is given by the roots of the physicists' version of the Hermite polynomials, since they yield Gauss-Hermite quadrature (GHQ), the standard numerical quadrature against Gaussian measures, as a posterior mean for a suitable covariance function (Särkkä et al., 2015). For GHQ, efficient algorithms to compute the roots and the weights are readily available (Press et al., 2007).

2.5 COMPUTATIONAL COST

All of the presented algorithms buy their probabilistic extension to classic ODE solvers by adding computational cost, sometimes more sometimes less. In most cases, evaluation of the dynamic function f forms the computational bottleneck, so we will focus on it here. Of course, the internal computations of the solver adds cost as well. Since all the models discussed here have linear inference cost, though, this additional overhead is manageable.

The ML-algorithm by Schober et al. (2014) is the fastest algorithm. By simply recasting a Runge-Kutta step as Gaussian filtering, rough probabilistic uncertainty is achieved with negligible computational overhead.

For the sampling method, the calculation of one individual sample of Q_h amounts to running the entire underlying ODE solver once, hence the overall cost is S times the original cost.

In contrast, the BQ-algorithm only has to run through $[0, T]$ once, but has to invert a $ND \times ND$ covariance matrix to perform Bayesian quadrature with N evaluation points. Usually, N will be small, since BQ performs well for a relatively small number of function evaluations (as e.g. illustrated by the experiments below). However, if the output dimension D is very large, Bayesian quadrature—like all quadrature methods—is not practical. BQ thus tends to be faster for small D , while MC tends to be faster for large D .

When considering these computational overheads, there is a nuanced point to be made about the value-to-cost trade-off of constructing a posterior uncertainty measure. If a classic numerical solver of order p is allotted a budget of M times its original one, it can use it to reduce its step-size by a factor of M , and thus reduce its approximation error by an order M^p . It may thus seem pointless to invest even such a linear cost increase into constructing an uncertainty measure around the classic estimate. But, in some practical settings, it may be more helpful to have a notion of uncertainty on a slightly less precise estimate than to produce a

more precise estimate without a notion of error. In addition, classic solvers are by nature sequential algorithms, while the probabilistic extensions (both the sampling-based and Gaussian-filtering based ones) can be easily parallelized. Where parallel hardware is available, the effective time cost of probabilistic functionality may thus be quite limited (although we do not investigate this possibility in our present experiments).

With regards to memory requirements, the MC-method needs significantly more storage, since it requires saving all sample paths, in order to statistically approximate the entire non-parametric measure Q_h on $C([0, T], \mathbb{R})$. The BQ-algorithm only has to save the posterior GP, i.e. a mean and a covariance function, which is arguably the minimal amount to provide a notion of uncertainty. If MC reduces the approximation of Q_h to its mean and variance, it only requires this minimal storage as well.

3 EXPERIMENTS

This section explores applications of the probabilistic ODE solvers discussed in Section 2. The sampling-based algorithm by (Conrad et al., 2015) will be abbreviated as MC, the maximum-likelihood Gaussian filter ((Schober et al., 2014)) as ML and our novel BQ-based filter (BQF) as BQ. In particular, we assess how the performance of the purely deterministic class of Gaussian filtering based solvers compares to the inherently random class of sampling-based solvers.

We experiment on the Van der Pol oscillator (Hairer et al., 1987), a non-conservative oscillator with non-linear damping, which is a standard example for a non-stiff dynamical system. It is governed by the equation

$$\frac{\partial^2 u}{\partial t^2} = \mu(1 - u^2) \frac{\partial u}{\partial t} - u, \quad (42)$$

where the parameter $\mu \in \mathbb{R}$ indicates the non-linearity and the strength of the damping. We set $\mu = 5$ on a time axis $[10, 60]$, with initial values $(u(10), \dot{u}(10)) = (2, 10)$.

All compared methods use a model of order $q = 3$, and a step size $h = 0.01$. This induces a state-space model given by a twice-integrated Wiener process prior (cf. (6)) which yields a version of ML close to second-order Runge-Kutta (Schober et al., 2014). The same solver is used as the underlying numerical solver Ψ_h in MC. For the noise parameter, which scales the deviation of the evaluation point of f from the numerical extrapolation (i.e. the variance of the driving Wiener process for ML and BQ, and the variance of ξ_k for MC), we choose $\sigma^2 = 0.1$. The drift matrix F of the underlying integrated Wiener process is set to the default values $f_i = i$ for $i = 1, \dots, q - 1$. The covariance function used in BQ is the widely popular squared exponential (36), with lengthscale $\lambda = 1$ and output variance $\theta^2 = 1$. (Since all methods use the same model, this tuning does not favor

one algorithm over the other. In practice all these parameters should of course be set by statistical estimation.)

For a fair comparison in all experiments, we allow MC and BQ to make the same amount of function evaluations per time step. If MC draws N samples, BQ uses N evaluation points. The first experiment presents the solutions of the presented algorithms on the van der Pol oscillator (42) on the whole time axis in one plot, when we allow BQ and MC to make five function evaluations. Then, we examine more closely how the error of each methods changes as a function of the number of evaluations of f in Figure 3 and Figure 4.

3.1 SOLUTION MEASURES ON VAN DER POL OSCILLATOR

Figure 3 shows the solution estimates constructed by the three solvers across the time domain. In all cases, the mean estimates roughly follow the true solution (which e.g. Gaussian filtering with Monte Carlo integration by sampling (22) – (24) does not achieve). A fundamental difference between the filtering-based methods (ML and BQ) and the sampling-based MC algorithm is evident in both the mean and the uncertainty estimate.

While the filtering-based methods output a trajectory quite similar to the true solution with a small time lag, the MC algorithm produces a trajectory of a more varying shape. Characteristic points of the MC mean estimate (such as local extrema) are placed further away from the true value than for filtering-based methods.

The uncertainty estimation of MC appears more flexible as well. ML and BQ produce an uncertainty estimate which runs parallel to the mean estimate and appears to be strictly increasing. It appears to increase slightly in every step, resulting in an uncertainty estimate, which only changes very slowly. The solver accordingly appears overconfident in the spikes and underconfident in the valleys of the trajectory. The uncertainty of MC varies more, scaling up at the steep parts of the oscillator and decreasing again at the flat parts, which is a desirable feature.

Among the class of filtering-based solvers, the more refined BQ method outputs a better mean estimate with more confidence than ML.

3.2 QUALITY OF ESTIMATE AS A FUNCTION OF ALLOWED EVALUATIONS

Figure 3 and Figure 4 depict the value of the error of the mean approximation as a function of the allowed function evaluations N (i.e. N evaluation points for BQ and N samples for MC) at time points $t_1 = 18$ and $t_2 = 54$. Since the desired solution measure Q_h for MC can only be statistically approximated by the N samples, the mean estimate of MC is random. For comparison, the average of five MC-




Figure 2: Solution estimates constructed on the Van der Pol oscillator (42). True solution in red. Mean estimates of ML, MC and BQ in black, green, blue, respectively. Uncertainty measures (drawn at two times standard deviation) as thin lines of the same color.




Figure 3: Plot of errors of the mean estimates at $t = 18$ of the methods MC (green) and BQ (blue) as a function of the allowed function evaluations. Maximum likelihood error in black. Single runs of the probabilistic MC solver as green crosses. Average over all runs as green line.

runs is computed.

At the early time point $t_1 = 18$, all trajectories are still close together and the methods perform roughly the same, as we allow more evaluations. There is a slight improvement for BQ with more evaluations, but the error remains above the one of ML error.

At the later time $t_2 = 54$, BQ improves drastically when at least five evaluations are allowed, dropping much below the ML error.




Figure 4: Plot of errors of the mean estimates at $t = 54$ of the methods MC (green) and BQ (blue) as a function of the allowed function evaluations. Maximum likelihood error in black. Single runs of the probabilistic MC solver as green crosses. Average over all runs as green line.

The average error by MC appears to be not affected by the number of samples. The ML error is constant, because it always evaluates only once.

4 DISCUSSION

The conducted experiments provide an interesting basis to discuss the differences between filtering-based methods (ML and BQ) and the sampling-based MC algorithm. We make the following observations:

- (i) *Additional samples do not improve the random mean estimate of MC in expectation:*

Since the samples of MC are independent and identically distributed, the expectation of the random mean estimate of MC is the same, regardless of the amounts of samples. This property is reflected in Figure 3 and Figure 4, by the constant green line (up to random fluctuation). Additional samples are therefore only useful to improve the uncertainty calibration.

- (ii) *The uncertainty calibration of MC appears more adaptive than of ML and BQ:*

Figure 3 suggests that MC captures the uncertainty more flexibly: It appropriately scales up in the steep parts of the oscillator, while expressing high confidence in the flat parts of the oscillator. The true trajectory is inside the interval between mean ± 2 standard deviations, which is not the case for BQ and ML. Moreover, MC produces a more versatile measure. The filtering-based methods appear to produce a strictly increasing uncertainty measure by adding to the posterior uncertainty in every step. MC avoids this problem by sampling multiple time over the whole time interval. We deem the resulting flexibility a highly desirable feature. BQ also outputs a meaningful uncertainty measure and we expect that adding Bayesian smoothing (Särkkä, 2013) would enable filtering-based methods to produce more adaptive measures as well.

- (iii) *The expected error of MC-samples (and their mean) is higher than the error of ML:*

In the experiments, MC produced a higher error for the mean estimate, compared to both ML and BQ. We expect that this happens on all dynamical systems *by construction*: Given U_k , the next value U_{k+1} of a MC-sample is calculated by adding Gaussian noise ψ_k to the ML-extrapolation starting in U_k (cf. equation (4)). Due to the symmetry and full support of Gaussian distributions, the perturbed solution has a higher error than the unperturbed prediction, which coincides with the ML solution. Hence, every MC-sample accumulates with every step a positive expected error increment compared to the ML estimate. By the linearity of the average, the mean over all samples inherits the same higher error than the ML mean (and thereby also than the error of the more refined BQ mean).

Summing up, we argue that — at their current state — filtering-based methods appear to produce a ‘better’ mean estimate, while sampling-based methods produce in some sense a ‘better’ uncertainty estimate. Many applications might put emphasis on a good mean estimate, while needing a still well-calibrated uncertainty quantification. Our method BQF provides a way of combining a precise

mean estimate with a meaningful uncertainty calibration. Sampling-based methods might not be able to provide this due to their less accurate mean estimate. For future work (which is beyond the scope of this paper), it could be possible to combine the advantages of both approaches in a unified method.

5 CONCLUSION

We have presented theory and methods for the probabilistic solution of ODEs which provide uncertainty measures over the solution of the ODE, contrasting the classes of (deterministic) filtering-based and (random) sampling-based solvers. We have provided a theoretical framework for Gaussian filtering as state space inference in linear Gaussian SDEs, highlighting the prediction of the gradient as the primary source of uncertainty. Of all investigated approximations of the gradient, Bayesian Quadrature (BQ) produces the best results, by actively learning the shape of the dynamic function f through deterministic evaluations. Hence, we propose a novel filtering-based method named *Bayesian Quadrature Filtering* (BQF), which employs BQ for the gradient measurement.

For the same amount of allowed gradient evaluations, the mean estimate of BQF appears to outperform the mean estimate of state-of-the-art sampling-based solvers on the Van der Pol oscillator, while outputting a better calibrated uncertainty than other filtering-based methods.

References

- Briol, F-X., Oates, C. J., Girolami, M., Osborne, M. A., and Sejdinovic, D. Probabilistic integration: A role for statisticians in numerical analysis? *arXiv:1512.00933 [stat.ML]*, December 2015.
- Chkrebtii, O., Campbell, D.A., Girolami, M.A., and Calderhead, B. Bayesian Uncertainty Quantification for Differential Equations. *arXiv prePrint 1306.2365*, 2013.
- Conrad, P.R., Girolami, M., Särkkä, S., Stuart, A.M., and Zygalakis, K.C. Probability measures for numerical solution of differential equations. *arXiv:1506.04592 [stat.ME]*, June 2015.
- Deisenroth, M.P. *Efficient Reinforcement Learning Using Gaussian Processes*. PhD thesis, Karlsruhe Institute of Technology, 2009.
- Hairer, E., Nørsett, S.P., and Wanner, G. *Solving Ordinary Differential Equations I – Nonstiff Problems*. Springer, 1987.
- Hennig, P. and Hauberg, S. Probabilistic Solutions to Differential Equations and their Application to Riemannian Statistics. In *Proc. of the 17th int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, volume 33. JMLR, W&CP, 2014.

- Karatzas, I. and Shreve, S.E. *Brownian Motion and Stochastic Calculus*. Springer, 1991.
- Kutta, W. Beitrag zur näherungsweise Integration totaler Differentialgleichungen. *Zeitschrift für Mathematik und Physik*, 46:435–453, 1901.
- Øksendal, B. *Stochastic Differential Equations: An Introduction with Applications*. Springer, 6 edition, 2003.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.P. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 3 edition, 2007. ISBN 0521880688, 9780521880688.
- Prüher, J. and Simandl, M. Bayesian quadrature in nonlinear filtering. In *Informatics in Control, Automation and Robotics (ICINCO), 2015 12th International Conference on*, volume 01, pp. 380–387, 2015.
- Rasmussen, C.E. and Williams, C.K.I. *Gaussian Processes for Machine Learning*. MIT, 2006.
- Runge, C. Über die numerische Auflösung von Differentialgleichungen. *Mathematische Annalen*, 46:167–178, 1895.
- Saatci, Y. *Scalable Inference for Structured Gaussian Process Models*. PhD thesis, University of Cambridge, 2011.
- Särkkä, S. *Recursive Bayesian Inference on Stochastic Differential Equations*. PhD thesis, Helsinki University of Technology, 2006.
- Särkkä, S. *Bayesian filtering and smoothing*. Cambridge University Press, 2013.
- Särkkä, S., Hartikainen, J., Svensson, L., and Sandblom, F. On the relation between gaussian process quadratures and sigma-point methods. *arXiv:1504.05994 [stat.ME]*, April 2015.
- Schober, M., Duvenaud, D., and Hennig, P. Probabilistic ODE Solvers with Runge-Kutta Means. *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Skilling, J. Bayesian solution of ordinary differential equations. *Maximum Entropy and Bayesian Methods, Seattle*, 1991.