



Convergence rates of Gaussian ODE filters

Hans Kersting¹ · T. J. Sullivan^{2,3} · Philipp Hennig¹

Received: 20 March 2020 / Accepted: 27 August 2020 / Published online: 12 September 2020
© The Author(s) 2020

Abstract

A recently introduced class of probabilistic (uncertainty-aware) solvers for ordinary differential equations (ODEs) applies Gaussian (Kalman) filtering to initial value problems. These methods model the true solution x and its first q derivatives *a priori* as a Gauss–Markov process X , which is then iteratively conditioned on information about \dot{x} . This article establishes worst-case local convergence rates of order $q + 1$ for a wide range of versions of this Gaussian ODE filter, as well as global convergence rates of order q in the case of $q = 1$ and an integrated Brownian motion prior, and analyses how inaccurate information on \dot{x} coming from approximate evaluations of f affects these rates. Moreover, we show that, in the globally convergent case, the posterior credible intervals are well calibrated in the sense that they globally contract at the same rate as the truncation error. We illustrate these theoretical results by numerical experiments which might indicate their generalizability to $q \in \{2, 3, \dots\}$.

Keywords Probabilistic numerics · Ordinary differential equations · Initial value problems · Numerical analysis · Gaussian processes · Markov processes

Mathematics Subject Classification 65L20 · 37H10 · 68W20 · 93E11

1 Introduction

A solver of an initial value problem (IVP) outputs an approximate solution $\hat{x}: [0, T] \rightarrow \mathbb{R}^d$ of an ordinary differential equation (ODE) with initial condition:

$$\begin{aligned} x^{(1)}(t) &:= \frac{dx}{dt}(t) = f(x(t)), \quad \forall t \in [0, T], \\ x(0) &= x_0 \in \mathbb{R}^d. \end{aligned} \quad (1)$$

(Without loss of generality, we simplify the presentation by restricting attention to the autonomous case.) The numerical

solution \hat{x} is computed by iteratively collecting information on $x^{(1)}(t)$ by evaluating $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ at a numerical estimate $\hat{x}(t)$ of $x(t)$ and using these approximate evaluations of the time derivative to extrapolate along the time axis. In other words, the numerical solution (or *estimator*) \hat{x} of the exact solution (or *estimand*) x is calculated based on evaluations of the vector field f (or *data*). Accordingly, we treat \hat{x} itself as an estimator, i.e., a statistic that translates evaluations of f into a probability distribution over $C^1([0, T]; \mathbb{R}^d)$, the space of continuously differentiable functions from $[0, T]$ to \mathbb{R}^d .

This probabilistic interpretation of numerical computations of tractable from intractable quantities as statistical inference of latent from observable quantities applies to all numerical problems and has been repeatedly recommended in the past (Poincaré 1896; Diaconis 1988; Skilling 1991; O’Hagan 1992; Ritter 2000). It employs the language of probability theory to account for the epistemic uncertainty (i.e., limited knowledge) about the accuracy of intermediate and final numerical computations, thereby yielding algorithms which can be more aware of—as well as more robust against—uncertainty over intermediate computational results. Such algorithms can output probability measures, instead of point estimates, over the final quantity of inter-

✉ Hans Kersting
hans.kersting@uni-tuebingen.de

T. J. Sullivan
t.j.sullivan@warwick.ac.uk; sullivan@zib.de

Philipp Hennig
philipp.hennig@uni-tuebingen.de

¹ University of Tübingen and Max Planck Institute for Intelligent Systems, Maria-von-Linden-Straße 6, 72076 Tübingen, Germany

² University of Warwick, Coventry CV4 7AL, United Kingdom

³ Zuse Institute Berlin, Takustraße 7, 14195 Berlin, Germany

est. This approach, now called *probabilistic numerics (PN)* (Hennig et al. 2015; Oates and Sullivan 2019), has in recent years been spelled out for a wide range of numerical tasks, including linear algebra, optimization, integration, and differential equations, thereby working towards the long-term goal of a coherent framework to propagate uncertainty through chained computations, as desirable, e.g., in statistical machine learning.

In this paper, we determine the convergence rates of a recent family of PN methods (Schober et al. 2014; Kersting and Hennig 2016; Magnani et al. 2017; Schober et al. 2019; Tronarp et al. 2019) which recast an IVP as a *stochastic filtering problem* (Øksendal 2003, Chapter 6), an approach that has been studied in other settings (Jazwinski 1970), but has not been applied to IVPs before. These methods assume *a priori* that the solution x and its first $q \in \mathbb{N}$ derivatives follow a Gauss–Markov process X that solves a stochastic differential equation (SDE).

The evaluations of f at numerical estimates of the true solution can then be regarded as imperfect evaluations of \dot{x} , which can then be used for a Bayesian update of X . Such recursive updates along the time axis yield an algorithm whose structure resembles that of Gaussian (Kalman) filtering (Särkkä 2013, Chapter 4). These methods add only slight computational overhead compared to classical methods (Schober et al. 2019) and have been shown to inherit local convergence rates from equivalent classical methods in specific cases (Schober et al. 2014; Schober et al. 2019). These equivalences (i.e., the equality of the filtering posterior mean and the classical method) are only known to hold in the case of the integrated Brownian motion (IBM) prior and noiseless evaluations of f (in terms of our later notation, the case $R \equiv 0$), as well as under the following restrictions:

Firstly, for $q \in \{1, 2, 3\}$, and if the first step is divided into sub-steps resembling those of Runge–Kutta methods, an equivalence of the posterior mean of the first step of the filter and the explicit Runge–Kutta method of order q was established in Schober et al. (2014) (but for $q \in \{2, 3\}$ only in the limit as the initial time of the IBM tends to $-\infty$). Secondly, it was shown by Schober et al. (2019) that, for $q = 1$, the posterior mean after each step coincides with the trapezoidal rule if it takes an additional evaluation of f at the end of each step, known as P(EC)1. The same paper shows that, for $q = 2$, the filter coincides with a third-order Nordsieck method (Nordsieck 1962) if the filter is in the steady state, i.e., after the sequence of error covariance matrices has converged. These results neither cover filters with the integrated Ornstein–Uhlenbeck process (IOUP) prior (Magnani et al. 2017) nor nonzero noise models on evaluations of f .

In this paper, we directly prove convergence rates without first fitting the filter to existing methods, and thereby lift many of the above restrictions on the convergence rates. While the more-recent work by Tronarp et al. (2020) also

provide convergence rates of estimators of x in the Bayesian ODE filtering/smoothing paradigm, they concern the maximum a posteriori estimator (as computed by the iterated extended Kalman ODE smoother), and therefore differ from our convergence rates of the filtering mean (as computed by the Kalman ODE filter).

1.1 Contribution

Our main results—Theorems 8 and 14—provide local and global convergence rates of the ODE filter when the step size h goes to zero. Theorem 8 shows local convergence rates of h^{q+1} without the above-mentioned previous restrictions—i.e., for a generic Gaussian ODE filter for all $q \in \mathbb{N}$, both IBM and IOUP prior, flexible Gaussian initialization (see Assumptions 2 and 3), and arbitrary evaluation noise $R \geq 0$. As a first global convergence result, Theorem 14 establishes global convergence rates of h^q in the case of $q = 1$, the IBM prior and all fixed measurement uncertainty models R of order $p \in [1, \infty]$ (see Assumption 4). This global rate of the worst-case error is matched by the contraction rate of the posterior credible intervals, as we show in Theorem 15. Moreover, we also give closed-form expressions for the steady states in the global case and illustrate our results as well as their possible generalizability to $q \geq 2$ by experiments in Sect. 9.

1.2 Related work on probabilistic ODE solvers

The Gaussian ODE filter can be thought of as a self-consistent Bayesian decision agent that iteratively updates its prior belief X over $x: [0, T] \rightarrow \mathbb{R}^d$ (and its first q derivatives) with information on \dot{x} from evaluating f .¹ For Gauss–Markov priors, it performs exact Bayesian inference and optimally (with respect to the L^2 -loss) extrapolates along the time axis. Accordingly, all of its computations are deterministic and—due to its restriction to Gaussian distributions—only slightly more expensive than classical solvers. Experiments demonstrating competitive performance with classical methods are provided in Schober et al. (2019, Section 5).

Another line of work (comprising the methods from Chkrebti et al. (2016); Conrad et al. (2017); Teymur et al. (2016); Lie et al. (2019); Abdulle and Garegnani (2020); Teymur et al. (2018)) introduces probability measures to ODE solvers in a fundamentally different way—by representing

¹ Here, the word ‘Bayesian’ describes the algorithm in the sense that it employs a prior over the quantity of interest and updates it by Bayes rule according to a prespecified measurement model (as also used in Skilling (1991); Chkrebti et al. (2016); Kersting and Hennig (2016)). The ODE filter is not Bayesian in the stronger sense of Cockayne et al. (2019), and it remains an open problem to construct a Bayesian solver in this strong sense without restrictive assumptions, as discussed in Wang et al. (2018).

the distribution of all numerically possible trajectories with a set of sample paths. To compute these sample paths, Chkrebti et al. (2016) draws them from a (Bayesian) Gaussian process (GP) regression; Conrad et al. (2017); Teymur et al. (2016); Lie et al. (2019); Teymur et al. (2018) perturb classical estimates after an integration step with a suitably scaled Gaussian noise; and Abdulle and Garegnani (2020) perturbs the classical estimate instead by choosing a stochastic step-size. While Conrad et al. (2017); Teymur et al. (2016); Lie et al. (2019); Abdulle and Garegnani (2020); Teymur et al. (2018) can be thought of as (non-Bayesian) ‘stochastic wrappers’ around classical solvers, which produce samples with the same convergence rate, Chkrebti et al. (2016) employs—like the filter—GP regression to represent the belief on x . While the Gaussian ODE filter can convergence with polynomial order (see results in this paper), However, Chkrebti et al. (2016) only show first-order convergence rates and also construct a sample representation of numerical errors, from which samples are drawn iteratively. A conceptual and experimental comparison between the filter and Chkrebti et al. (2016) can be found in Schober et al. (2019). An additional numerical test against Conrad et al. (2017) was given by Kersting and Hennig (2016). Moreover, Tronarp et al. (2019) recently introduced a particle ODE filter, which combines a filtering-based solver with a sampling-based uncertainty quantification (UQ), and compared it numerically with Conrad et al. (2017) and Chkrebti et al. (2016).

All of the above sampling-based methods can hence represent more expressive, non-Gaussian posteriors (as, e.g., desirable for bifurcations), but multiply the computational cost of the underlying method by the number of samples. ODE filters are, in contrast, not a perturbation of known methods, but novel methods designed for computational speed and for a robust treatment of intermediate uncertain values (such as the evaluations of f at estimated points). Unless parallelization of the samples in the sampling-based solvers is possible and inexpensive, one can spend the computational budget for generating additional samples on dividing the step size h by the number of samples, and can thereby polynomially decrease the error. Its Gaussian UQ, however, should not be regarded as the true UQ—in particular for chaotic systems whose uncertainty can be better represented by sampling-based solvers, see, e.g., Conrad et al. (2017, Figure 1) and Abdulle and Garegnani (2020, Figure 2)—but as a rough inexpensive probabilistic treatment of intermediate values and final errors which is supposed to, on average, guide the posterior mean towards the true x . Therefore, it is in a way more similar to classical non-stochastic solvers than to sampling-based stochastic solvers and, unlike sampling-based solvers, puts emphasis on computational speed over statistical accuracy. Nevertheless, its Gaussian UQ is sufficient to make the forward models in ODE inverse problems

more ‘uncertainty-aware’; see Kersting et al. (2020, Section 3).

Accordingly, the convergence results in this paper concern the convergence rate of the posterior mean to the true solution, while the theoretical results from Teymur et al. (2016); Chkrebti et al. (2016); Conrad et al. (2017); Lie et al. (2019); Abdulle and Garegnani (2020); Teymur et al. (2018) provide convergence rates of the variance of the non-Gaussian empirical measure of samples (and not for an individual sample).

1.3 Relation to filtering theory

While Gaussian (Kalman) filtering was first applied to the solution of ODEs by Kersting and Hennig (2016) and Schober et al. (2019), it has previously been analyzed in the filtering, data assimilation as well as linear system theory community. The convergence results in this paper are concerned with its asymptotics when the step size h (aka time step between data points) goes to zero. In the classical filtering setting, where the data comes from an external sensor, this quantity is not treated as a variable, as it is considered a property of the data and not, like in our case, of the algorithm. Accordingly, the standard books lack such an analysis for $h \rightarrow 0$ —see Jazwinski (1970); Anderson and Moore (1979); Maybeck (1979) for filtering, Law et al. (2015); Reich and Cotter (2015) for data assimilation and Callier and Desoer (1991) for linear system theory—and we believe that our convergence results are completely novel. It is conceivable that, also for these communities, this paper may be of interest in settings where the data collection mechanism can be actively chosen, e.g., when the frequency of the data can be varied or sensors of different frequencies can be used.

1.4 Outline

The paper begins with a brief introduction to Gaussian ODE filtering in Sect. 2. Next, Sects. 3 and 5 provide auxiliary bounds on the flow map of the ODE and on intermediate quantities of the filter, respectively. With the help of these bounds, Sects. 6 and 7 establish local and global convergence rates of the filtering mean, respectively. In light of these rates, Sect. 8 analyses for which measurement noise models the posterior credible intervals are well calibrated. These theoretical results are experimentally confirmed and discussed in Sect. 9. Section 10 concludes with a high-level discussion.

1.5 Notation

We will use the notation $[n] := \{0, \dots, n - 1\}$. For vectors and matrices, we will use zero-based numbering, e.g., $x = (x_0, \dots, x_{d-1}) \in \mathbb{R}^d$. For a matrix $P \in \mathbb{R}^{n \times m}$ and $(i, j) \in [n] \times [m]$, we will write $P_{i,\cdot} \in \mathbb{R}^{1 \times m}$ for the i th row and $P_{\cdot,j}$

for the j th column of P . A fixed but arbitrary norm on \mathbb{R}^d will be denoted by $\|\cdot\|$. The minimum and maximum of two real numbers a and b will be denoted by $a \wedge b$ and $a \vee b$, respectively. Vectors that span all q modeled derivatives will be denoted by bold symbols, such as \mathbf{x} .

2 Gaussian ODE filtering

This section defines how a Gaussian filter can solve the IVP Eq. (1). In the various subsections, we first explain the choice of prior on x , then describe how the algorithm computes a posterior output from this prior (by defining a numerical integrator Ψ), and add explanations on the measurement noise of the derivative observations. To alternatively understand how this algorithm can be derived as an extension of generic Gaussian filtering in probabilistic state space models, see the concise presentation in (Kersting et al. 2020, Supplement A).

2.1 Prior on x

In PN, it is common (Hennig et al. 2015, Section 3(a)) to put a prior measure on the unknown solution x . Often, for fast Bayesian inference by linear algebra (Rasmussen and Williams 2006, Chapter 2), this prior is Gaussian. To enable GP inference in linear time by Kalman filtering (Särkkä 2013, Chapter 4.3), we further restrict the prior to Markov processes. As discussed in Särkkä and Solin (2019, Chapter 12.4), a wide class of such Gauss–Markov processes can be captured by a law of the (strong) solution (Øksendal 2003, Chapter 5.3) of a linear SDE with Gaussian initial condition. Here—as we, by Eq. (1), have information on at least one derivative of x —the prior also includes the first $q \in \mathbb{N}$ derivatives. Therefore, for all $j \in [d]$, we define the vector of time derivatives by $\mathbf{X}_j = (X_j^{(0)}, \dots, X_j^{(q)})^\top$. We define \mathbf{X}_j as a $(q + 1)$ -dimensional stochastic process via the SDE

$$\begin{aligned} d\mathbf{X}_j(t) &= (dX_j^{(0)}(t), \dots, dX_j^{(q-1)}(t), dX_j^{(q)}(t))^\top \\ &= \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & 0 & 1 & \\ c_0 & \dots & \dots & \dots & c_q \end{pmatrix} \begin{pmatrix} X_j^{(0)}(t) \\ \vdots \\ X_j^{(q-1)}(t) \\ X_j^{(q)}(t) \end{pmatrix} dt + \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \sigma_j \end{pmatrix} dB_j(t), \end{aligned} \tag{2}$$

driven by mutually independent one-dimensional Brownian motions $\{B_j; j \in [d]\}$ (independent of $\mathbf{X}(0)$) scaled by $\sigma_j > 0$, with initial condition $\mathbf{X}_j(0) \sim \mathcal{N}(m_j(0), P_j(0))$. We assume that $\{X_j(0); j \in [d]\}$ are independent. In other words, we model the unknown i th derivative of the j th dimension of the solution x of the IVP Eq. (1), denoted

by $x_j^{(i)}$, as a draw from a real-valued, one-dimensional GP $X_j^{(i)}$, for all $i \in [q + 1]$ and $j \in [d]$, such that $X_j^{(q)}$ is defined by (c_0, \dots, c_q) as well as the Brownian motion scale σ_j and $X_j^{(i-1)}$ is defined to be the integral of $X_j^{(i)}$. Note that, by the independence of the components of the d -dimensional Brownian motion, the components $\{\{X_j(t); 0 \leq t \leq T\}; j \in [d]\}$ of $\{\mathbf{X}(t); 0 \leq t \leq T\}$ are independent². The (strong) solution of Eq. (2) is a Gauss–Markov process with mean $m_j: [0, T] \rightarrow \mathbb{R}^{q+1}$ and covariance matrix $P_j: [0, T] \rightarrow \mathbb{R}^{(q+1) \times (q+1)}$ given by

$$m_j(t) = A(t)m_j(0), \tag{3}$$

$$P_j(t) = A(t)P_j(0)A(t)^\top + Q(t), \tag{4}$$

where the matrices $A(t), Q(t) \in \mathbb{R}^{(q+1) \times (q+1)}$ yielded by the SDE Eq. (2) are known in closed form Särkkä (2006, Theorem 2.9) (see Eq. (77)). The precise choice of the prior stochastic process \mathbf{X} depends on the choice of $(c_0, \dots, c_q) \in \mathbb{R}^{q+1}$ in Eq. (2). While the below algorithm works for all choices of c , we restrict our attention to the case of

$$(c_0, \dots, c_q) := (0, \dots, 0, -\theta), \quad \text{for some } \theta \geq 0, \tag{5}$$

where the q -times integrated Brownian motion (IBM) and the q -times integrated Ornstein–Uhlenbeck process (IOUP) with drift parameter θ is the unique solution of Eq. (2), in the case of $\theta = 0$ and $\theta > 0$, respectively, (Karatzas and Shreve 1991, Chapter 5: Example 6.8). In this case, the matrices A and Q from Eqs. (3) and (4) are given by

$$A(t)_{ij} = \begin{cases} \mathbb{I}_{i \leq j} \frac{t^{j-i}}{(j-i)!}, & \text{if } j \neq q, \\ \frac{t^{q-i}}{(q-i)!} - \theta \sum_{k=q+1-i}^{\infty} \frac{(-\theta)^{k+i-q-1} t^k}{k!}, & \text{if } j = q, \end{cases} \tag{6}$$

$$Q(t)_{ij} = \sigma^2 \frac{t^{2q+1-i-j}}{(2q+1-i-j)(q-i)!(q-j)!} + \Theta(t^{2q+2-i-j}). \tag{7}$$

(Derivations of Eqs. (6) and (7), as well as the precise form of Q without $\Theta(t^{2q+2-i-j})$, are presented in Appendix A.) Hence, for all $i \in [q + 1]$, the prediction of step size h of the

² More involved correlation models of $\{\{X_j(t); 0 \leq t \leq T\}; j \in [d]\}$ are straightforward to incorporate into the SDE Eq. (2), but seem complicated to analyze. Therefore, we restrict our attention to independent dimensions. See Appendix B for an explanation of this restriction. Note that one can also use a state space vector $\mathbf{X}(t)$ which models other features of $x(t)$ than the derivatives, as demonstrated with Fourier summands in Kersting and Mahserici (2020).

i th derivative from any state $u \in \mathbb{R}^{q+1}$ is given by

$$[A(t)u]_i = \sum_{k=i}^q \frac{t^{k-i}}{(k-i)!} u_k - \theta \left[\sum_{k=q+1-i}^{\infty} \frac{(-\theta)^{k+i-q-1}}{k!} t^k \right] u_q. \tag{8}$$

2.2 The algorithm

To avoid the introduction of additional indices, we will define the algorithm Ψ for $d = 1$; for statements on the general case of $d \in \mathbb{N}$ we will use the same symbols from Eq. (10)–(15) as vectors over the whole dimension—see, e.g., Eq. (31) for a statement about a general $r \in \mathbb{R}^d$. By the independence of the dimensions of X , due to Eq. (2), extension to $d \in \mathbb{N}$ amounts to applying Ψ to every dimension independently (recall Footnote 2). Accordingly, we may in many of the below proofs w.l.o.g. assume $d = 1$. Now, as previously spelled out in Kersting and Hennig (2016); Schober et al. (2019), Bayesian filtering of X —i.e., iteratively conditioning X on the information on $X^{(1)}$ from evaluations of f at the mean of the current conditioned $X^{(0)}$ —yields the following numerical method Ψ . Let $\mathbf{m}(t) = (m^{(0)}(t), \dots, m^{(q)}(t))^T \in \mathbb{R}^{q+1}$ be an arbitrary state at some point in time $t \in [0, T]$ (i.e., $m^{(i)}(t)$ is an estimate for $x^{(i)}(t)$), and let $P(t) \in \mathbb{R}^{(q+1) \times (q+1)}$ be the covariance matrix of $x^{(i)}(t)$. For $t \in [0, T]$, let the current estimate of $x(t)$ be a normal distribution $\mathcal{N}(\mathbf{m}(t), P(t))$, i.e., the mean $\mathbf{m}(t) \in \mathbb{R}^{q+1}$ represents the best numerical estimate (given data $\{y(h), \dots, y(t)\}$, see Eq. (12)) and the covariance matrix $P(t) \in \mathbb{R}^{(q+1) \times (q+1)}$ its uncertainty. For the time step $t \rightarrow t + h$ of size $h > 0$, the ODE filter first computes the prediction step consisting of *predictive mean*

$$\mathbf{m}^-(t+h) := A(h)\mathbf{m}(t) \in \mathbb{R}^{q+1}, \tag{9}$$

and *predictive covariance*

$$P^-(t+h) := A(h)P(t)A(h)^T + Q(h) \in \mathbb{R}^{(q+1) \times (q+1)}, \tag{10}$$

with A and Q generally defined by Eq. (77) and, in the considered particular case of Eq. (5), by Eqs. (6) and (7). In the subsequent step, the following quantities are computed first: the *Kalman gain*

$$\begin{aligned} \boldsymbol{\beta}(t+h) &= (\beta^{(0)}(t+h), \dots, \beta^{(q)}(t+h))^T \\ &:= \frac{P^-(t+h)_{:,1}}{(P^-(t+h))_{11} + R(t+h)} \in \mathbb{R}^{(q+1) \times 1}, \end{aligned} \tag{11}$$

the *measurement/data on \dot{x}*

$$y(t+h) := f\left(m^{-(0)}(t+h)\right) \in \mathbb{R}, \tag{12}$$

and *innovation/residual*

$$r(t+h) := y(t+h) - m^{-(1)}(t+h) \in \mathbb{R}. \tag{13}$$

Here, R denotes the variance of y (the ‘measurement noise’) and captures the squared difference between the data $y(t+h) = f(m^-(t+h))$ that the algorithm actually receives and the idealized data $\dot{x}(t+h) = f(x(t+h))$ that it ‘should’ receive (see Sect. 2.3). Finally, the mean and the covariance matrix are conditioned on this data, which yields the *updated mean*

$$\begin{aligned} \Psi_{P(t),h}(\mathbf{m}(t)) &:= \mathbf{m}(t+h) \\ &= \mathbf{m}^-(t+h) + \boldsymbol{\beta}(t+h)r(t+h), \end{aligned} \tag{14}$$

and the *updated covariance*

$$P(t+h) := P^-(t+h) - \frac{P^-(t+h)_{:,1}P^-(t+h)_{1,:}}{P^-(t+h)_{11} + R(t+h)}. \tag{15}$$

This concludes the step $t \rightarrow t + h$, with the Gaussian distribution $\mathcal{N}(\mathbf{m}(t+h), P(t+h))$ over $x(t+h)$. The algorithm is iterated by computing $\mathbf{m}(t+2h) := \Psi_{P(t+h),h}(\mathbf{m}(t+h))$ as well as repeating Eq. (10) and (15), with $P(t+h)$ instead of $P(t)$, to obtain $P(t+2h)$. In the following, to avoid notational clutter, the dependence of the above quantities on t, h and σ will be omitted if their values are unambiguous. Parameter adaptation reminiscent of classical methods (e.g., for σ s.t. the added variance per step coincide with standard error estimates) has been explored in Schober et al. (2019, Section 4).

This filter is essentially an iterative application of Bayes rule (see, e.g., Särkkä (2013, Chapter 4)) based on the prior X on x specified by Eq. (2) (entering the algorithm via A and Q) and the measurement model $y \sim \mathcal{N}(\dot{x}, R)$. Since the measurement model is a likelihood by another name and therefore forms a complete Bayesian model together with the prior X , it remains to detail the measurement model (recall Sect. 2.1 for the choice of prior). Concerning the data generation mechanism for y Eq. (12), we only consider the maximum-a-posteriori point estimate of $\dot{x}(t)$ given $\mathcal{N}(m^{-(0)}(t), P_{00}^-(t))$; a discussion of more involved statistical models for y as well as an algorithm box for the Gaussian ODE filter can be found in Schober et al. (2019, Subsection 2.2). Next, for lack of such a discussion for R , we will examine different choices of R —which have proved central to the UQ of the filter (Kersting and Hennig 2016) and will turn out to affect global convergence properties in Sect. 7.

2.3 Measurement noise R

Two sources of uncertainty add to $R(t)$: noise from imprecise knowledge of $x(t)$ and f . Given f , previous integration steps of the filter (as well as an imprecise initial value) inject uncertainty about how close $m^-(t)$ is to $x(t)$ and how close $y = f(m^-(t))$ is to $\dot{x}(t) = f(x(t))$. This uncertainty stems from the discretization error $\|m^{-(0)}(t) - x(t)\|$ and, hence, tends to increase with h . Additionally, there can be uncertainty from a misspecified f , e.g., when f has estimated parameters, or from numerically imprecise evaluations of f , which can be added to R —a functionality which classical solvers do not possess. In this paper, since R depends on h via the numerical uncertainty on $x(t)$, we analyze the influence of noise R of order $p \in [1, \infty]$ (see Assumption 4) on the quality of the solution to illuminate for which orders of noise we can trust the solution to which extent and when we should, instead of decreasing h , rather spend computational budget on specifying or evaluating f more precisely. The explicit dependence of the noise on its order p in h resembles, despite the fundamentally different role of R compared to additive noise in Conrad et al. (2017); Abdulle and Garegnani (2020), the variable p in Conrad et al. (2017, Assumption 1) and Abdulle and Garegnani (2020, Assumption 2.2) in the sense that the analysis highlights how uncertainty of this order can still be modeled without breaking the convergence rates. (Adaptive noise models are computationally feasible (Kersting and Hennig 2016) but lie outside the scope of our analysis.)

3 Regularity of flow

Before we proceed to the analysis of Ψ , we provide all regularity results necessary for arbitrary $q, d \in \mathbb{N}$ in this section.

Assumption 1 The vector field $f \in C^q(\mathbb{R}^d; \mathbb{R}^d)$ is globally Lipschitz, and all its derivatives of order up to q are uniformly bounded and globally Lipschitz, i.e., there exists some $L > 0$ such that $\|D^\alpha f\|_\infty \leq L$ for all multi-indices $\alpha \in \mathbb{N}_0^d$ with $1 \leq \sum_i \alpha_i \leq q$, and $\|D^\alpha f(a) - D^\alpha f(b)\| \leq L\|a - b\|$ for all multi-indices $\alpha \in \mathbb{N}_0^d$ with $0 \leq \sum_i \alpha_i \leq q$.

Assumption 1 and the Picard–Lindelöf theorem imply that the solution x is a well-defined element of $C^{q+1}([0, T]; \mathbb{R}^d)$. For $i \in [q + 1]$, we denote $\frac{d^i x}{dt^i}$ by $x^{(i)}$. Recall that, by a bold symbol, we denote the vector of these derivatives: $\mathbf{x} \equiv (x^{(0)}, \dots, x^{(q)})^\top$. In particular, the solution x of Eq. (1) is denoted by $x^{(0)}$. Analogously, we denote the flow of the ODE Eq. (1) by $\Phi^{(0)}$, i.e., $\Phi_t^{(0)}(x_0) \equiv x^{(0)}(t)$, and, for all $i \in [q + 1]$, its i th partial derivative with respect to t by $\Phi^{(i)}$, so that $\Phi_t^{(i)}(x_0) \equiv x^{(i)}(t)$.

Lemma 1 Under Assumption 1, for all $a \in \mathbb{R}^d$ and all $h > 0$,

$$\left\| \Phi_h^{(i)}(a) - \sum_{k=i}^q \frac{h^{k-i}}{(k-i)!} \Phi_0^{(k)}(a) \right\| \leq Kh^{q+1-i}. \tag{16}$$

Here, and in the sequel, $K > 0$ denotes a constant independent of h and θ which may change from line to line.

Proof By Assumption 1, $\Phi^{(q+1)}$ exists and is bounded by $\|\Phi^{(q+1)}\| \leq L$, which can be seen by applying the chain rule q times to both sides of Eq. (1). Now, applying $\|\Phi^{(q+1)}\| \leq L$ to the term $\Phi_\tau^{(q+1)}(a)$ (for some $\tau \in (0, h)$) in the Lagrange remainder of the $(q - i)$ th-order Taylor expansion of $\Phi_h^{(i)}(a)$ yields Eq. (16). \square

Lemma 2 Under Assumption 1 and for all sufficiently small $h > 0$,

$$\sup_{a \neq b \in \mathbb{R}^d} \frac{\|\Phi_h^{(0)}(a) - \Phi_h^{(0)}(b)\|}{\|a - b\|} \leq 1 + 2Lh. \tag{17}$$

Proof Immediate corollary of Teschl (2012, Theorem 2.8). \square

Global convergence (Sect. 7) will require the following generalization of Lemma 2.

Lemma 3 Let $q = 1$. Then, under Assumption 1 and for all sufficiently small $h > 0$,

$$\sup_{a \neq b \in \mathbb{R}^d} \frac{\|\|\Phi_h(a) - \Phi_h(b)\|\|_h}{\|a - b\|} \leq 1 + Kh, \tag{18}$$

where given the norm $\|\cdot\|$ on \mathbb{R}^d and $h > 0$, the new norm $\|\|\cdot\|\|_h$ on $\mathbb{R}^{(q+1) \times d}$ is defined by

$$\|\|a\|\|_h := \sum_{i=0}^q h^i \|a_{i,\cdot}\|. \tag{19}$$

Remark 1 The necessity of $\|\|\cdot\|\|_h$ stems from the fact that—unlike other ODE solvers—the ODE filter Ψ additionally estimates and uses the first q derivatives in its state $\mathbf{m} \in \mathbb{R}^{(q+1) \times d}$, whose development cannot be bounded in $\|\cdot\|$, but in $\|\|\cdot\|\|_h$. The norm $\|\|\cdot\|\|_h$ is used to make rigorous the intuition that the estimates of the solution’s time derivative are ‘one order of h worse per derivative.’

Proof We bound the second summand of

$$\begin{aligned} \|\|\Phi_h(a) - \Phi_h(b)\|\|_h &\stackrel{\text{eq. (19)}}{=} \\ &\underbrace{\left\| \Phi_h^{(0)}(a) - \Phi_h^{(0)}(b) \right\|}_{\leq (1+2Lh)\|a-b\|, \text{ by eq. (17)}} + h \left\| \underbrace{\Phi_h^{(1)}(a)}_{=f(\Phi_h^{(0)}(a))} - \underbrace{\Phi_h^{(1)}(b)}_{=f(\Phi_h^{(0)}(b))} \right\| \end{aligned} \tag{20}$$

by

$$\begin{aligned} \left\| f\left(\Phi_h^{(0)}(a)\right) - f\left(\Phi_h^{(0)}(b)\right) \right\| &\stackrel{\text{Ass. 1}}{\leq} \\ L \left\| \Phi_h^{(0)}(a) - \Phi_h^{(0)}(b) \right\| &\stackrel{\text{eq. (17)}}{\leq} L(1 + 2Lh) \|a - b\|. \end{aligned} \tag{21}$$

Inserting Eq. (21) into Eq. (20) concludes the proof. \square

4 The role of the state misalignments δ

In Gaussian ODE filtering, the interconnection between the estimates of the ODE solution $x(t) = x^{(0)}(t)$ and its first q derivatives $\{x^{(1)}(t), \dots, x^{(q)}(t)\}$ is intricate. From a purely analytical point of view, every possible estimate $m(t)$ of $x(t)$ comes with a fixed set of derivatives, which are implied by the ODE, for the following reason: Clearly, by Eq. (1), the estimate $m^{(1)}(t)$ of $x^{(1)}(t)$ ought to be $f(m(t))$. More generally (for $i \in [q + 1]$) the estimate $m^{(i)}(t)$ of $x^{(i)}(t)$ is determined by the ODE as well. To see this, let us first recursively define $f^{(i)}: \mathbb{R}^d \rightarrow \mathbb{R}^d$ by $f^{(0)}(a) := a$, $f^{(1)}(a) := f(a)$ and $f^{(i)}(a) := [\nabla_x f^{(i-1)} \cdot f](a)$. Now, differentiating the ODE, Eq. (1), $(i - 1)$ -times by the chain rule yields

$$x^{(i)}(t) = f^{(i-1)}(t) \left(x^{(0)}(t) \right), \tag{22}$$

which implies that $m^{(i)}(t)$ ought to be $f^{(i-1)}(t) (m^{(0)}(t))$. Since

$$\Phi_0^{(i)} \left(m^{(0)}(nh) \right) = f^{(i-1)} \left(m^{(0)}(nh) \right) \tag{23}$$

(which we prove in Appendix E), this amounts to requiring that

$$m^{(i)}(t) \stackrel{!}{=} \Phi_0^{(i)} \left(m^{(0)}(nh) \right). \tag{24}$$

Since $\Phi_0^{(i)}$ is (recall Sect. 3) the i th time derivative of the flow map $\Phi^{(0)}$ at $t = 0$, this simply means that $m^{(i)}(t)$ would be set to the ‘true’ derivatives in the case where the initial condition of the ODE, Eq. (1), is $x(0) = m^{(0)}(t)$ instead of $x(0) = x_0$ —or, more loosely speaking, that the derivative estimates $m^{(i)}(t)$ are forced to comply with $m^{(0)}(t)$, irrespective of our belief $x^{(i)}(t) \sim \mathcal{N}(m^{(i)}(t), P_{ii}(t))$. The Gaussian ODE filter, however, does not use this (intractable) analytical approach. Instead, it jointly models and infers $x^{(0)}(t)$ and its first q derivatives $\{x^{(1)}(t), \dots, x^{(q)}(t)\}$ in a state space X , as detailed in Sect. 2. The thus-computed filtering mean estimates $m^{(i)}(t)$ depend not only on the ODE but also on the statistical model—namely on the prior (SDE) and measurement noise R ; recall Sects. 2.1 and 2.3. In fact, the analytically desirable derivative estimate, Eq. (24), is,

for $i = 1$, only satisfied if $R = 0$ (which can be seen from Eq. (14)), and generally does not hold for $i \geq 2$ since both $f^{(i-1)}$ and $\Phi^{(i)}$ are inaccessible to the algorithm. The numerical example in Appendix C clarifies that $\delta^{(i)}$ is likely to be strictly positive, even after the first step $0 \rightarrow h$.

This inevitable mismatch, between exact analysis and approximate statistics, motivates the following definition of the i th state i th state misalignment at time t :

$$\delta^{(i)}(t) := \left\| m^{(i)}(t) - \Phi_0^{(i)} \left(m^{(0)}(t) \right) \right\| \geq 0. \tag{25}$$

Intuitively speaking, $\delta^{(i)}(t)$ quantifies how large this mismatch is for the i th derivative at time t . Note that $\delta^{(i)}(t) = 0$ if and only if Eq. (24) holds—i.e., for $i = 1$ iff $R = 0$ (which can be seen from Eq. (14)) and only by coincidence for $i \geq 2$ since both $f^{(i-1)}$ and $\Phi_0^{(i)}$ are inaccessible to the algorithm. (Since $\Phi_0^{(0)} = \text{Id}$, $\delta^{(0)}(t) = 0$ for all t .)

The possibility of $\delta^{(i)} > 0$, for $i \geq 1$, is inconvenient for the below worst-case analysis since (if Eq. (24) held true and $\delta^{(i)} \equiv 0$) the prediction step of the drift-less IBM prediction ($\theta = 0$) would coincide with a Taylor expansions of the flow map $\Phi_0^{(i)}$; see Eq. (8). But, because $\delta^{(i)} \neq 0$ in general, we have to additionally bound the influence of $\delta \geq 0$ which complicates the below proofs further.

Fortunately, we can *locally* bound the import of $\delta^{(i)}$ by the easy Lemma 7 and *globally* by the more complicated Lemma 11 (see Sect. 7.3). Intuitively, these bounds demonstrate that the order of the deviation from a Taylor expansion of the state $\mathbf{m} = [m^{(0)}, \dots, m^{(q)}]$ due to δ is not smaller than the remainder of the Taylor expansion. This means, more loosely speaking, that the import of the $\delta^{(i)}$ is swallowed by the Taylor remainder. This effect is locally captured by Lemma 4 and globally by Lemma 12. The global convergence rates of $\delta^{(i)}(T)$, as provided by Lemma 12, are experimentally demonstrated in Appendix D.

5 Auxiliary bounds on intermediate quantities

Recall from Eq. (5) that $\theta = 0$ and $\theta > 0$ denote the cases of IBM and IOUP prior with drift coefficient θ , respectively. The ODE filter Ψ iteratively computes the filtering mean $\mathbf{m}(nh) = (m^{(0)}(nh), \dots, m^{(q)}(nh))^T \in \mathbb{R}^{(q+1)}$ as well as error covariance matrices $P(nh) \in \mathbb{R}$ on the mesh $\{nh\}_{n=0}^{T/h}$. (Here and in the following, we assume w.l.o.g. that $T/h \in \mathbb{N}$.) Ideally, the truncation error over all derivatives

$$\boldsymbol{\varepsilon}(nh) := (\varepsilon^{(0)}(nh), \dots, \varepsilon^{(q)}(nh))^T := \mathbf{m}(nh) - \mathbf{x}(nh), \tag{26}$$

falls quickly as $h \rightarrow 0$ and is estimated by the standard deviation $\sqrt{P_{00}(nh)}$. Next, we present a classical worst-case convergence analysis over all f satisfying Assumption 1; see Sect. 10 for a discussion of the desirability and feasibility of an average-case analysis. To this end, we bound the added error of every step by intermediate values, defined in Eqs. (11) and (13),

$$\Delta^{(i)}((n+1)h) := \left\| \Psi_{P(nh),h}^{(i)}(\mathbf{m}(nh)) - \Phi_h^{(i)}(m^{(0)}(nh)) \right\| \quad (27)$$

$$\stackrel{\text{eq. (14)}}{\leq} \underbrace{\left\| (A(h)\mathbf{m}(nh))_i - \Phi_h^{(i)}(m^{(0)}(nh)) \right\|}_{=: \Delta^{-(i)}((n+1)h)} + \left\| \beta^{(i)}((n+1)h) \right\| \|r((n+1)h)\|, \quad (28)$$

and bound these quantities in the order $\Delta^{-(i)}$, r , $\beta^{(i)}$. These bounds will be needed for the local and global convergence analysis in Sects. 6 and 7, respectively. Note that, intuitively, $\Delta^{-(i)}((n+1)h)$ and $\Delta^{(i)}((n+1)h)$ denote the *additional* numerical error which is added in the $(n+1)$ th step to the i th derivative of the predictive mean $m^{-(i)}(t+h)$ and the updated mean $m^{(i)}(t+h)$, respectively.

Lemma 4 *Under Assumption 1, for all $i \in [q+1]$ and all $h > 0$,*

$$\Delta^{-(i)}((n+1)h) \leq K \left[1 + \theta \left\| m^{(q)}(nh) \right\| \right] h^{q+1-i} + \sum_{k=i}^q \frac{h^{k-i}}{(k-i)!} \delta^{(k)}(nh). \quad (29)$$

Proof We may assume, as explained in Sect. 2.2, without loss of generality that $d = 1$. We apply the triangle inequality to the definition of $\Delta^{-(i)}((n+1)h)$, as defined in Eq. (28), which, by Eq. (8), yields

$$\begin{aligned} \Delta^{-(i)}((n+1)h) &\leq \sum_{k=i}^q \frac{h^{k-i}}{(k-i)!} \delta^{(k)}(nh) + K\theta \left\| m^{(q)}(nh) \right\| h^{q+1-i} \\ &\quad + \underbrace{\left\| \sum_{l=i}^q \frac{h^{l-i}}{(l-i)!} \Phi_0^{(l)}(m^{(0)}(nh)) - \Phi_h^{(i)}(m^{(0)}(nh)) \right\|}_{\leq Kh^{q+1-i}, \text{ by eq. (16)}}. \end{aligned} \quad (30)$$

□

Lemma 5 *Under Assumption 1 and for all sufficiently small $h > 0$,*

$$\|r((n+1)h)\| \leq K \left[1 + \theta \left\| m^{(q)}(nh) \right\| \right] h^q + K \sum_{k=1}^q \frac{h^{k-1}}{(k-1)!} \delta^{(k)}(nh). \quad (31)$$

Proof See Appendix F. □

To bound the Kalman gains $\beta(nh)$, we first need to assume that the orders of the initial covariance matrices are sufficiently high (matching the latter required orders of the initialization error; see Assumption 3).

Assumption 2 The entries of the initial covariance matrix $P(0)$ satisfy, for all $k, l \in [q+1]$, $\|P(0)_{k,l}\| \leq K_0 h^{2q+1-k-l}$, where $K_0 > 0$ is a constant independent of h .

We make this assumption, as well as Assumption 3, explicit (instead of just making the stronger assumption of exact initializations with zero variance), because it highlights how statistical or numerical uncertainty on the initial value effects the accuracy of the output of the filter—a novel functionality of PN with the potential to facilitate a management of the computational budget across a computational chain with respect to the respective perturbations from different sources of uncertainty (Hennig et al. 2015, Section 3(d)).

Lemma 6 *Under Assumption 2, for all $i \in [q+1]$ and for all $h > 0$, $\|\beta^{(i)}(h)\| \leq Kh^{1-i}$.*

Proof Again, w.l.o.g. $d = 1$. Application of the orders of A and Q from Eqs. (6) and (7), the triangle inequality and Assumption 2 to the definition of P^- in Eq. (10) yields

$$\begin{aligned} |P^-(h)_{k,l}| &\stackrel{\text{eq. (10)}}{\leq} \left| [A(h)P(0)A(h)^\top]_{k,l} \right| + |Q(h)_{k,l}| \\ &\stackrel{\text{eqs. (6),(7)}}{\leq} K \left[\sum_{a=k}^q \sum_{b=l}^q |P(0)_{a,b}| h^{a+b-k-l} \right. \\ &\quad \left. + 2\theta \sum_{b=l}^{q-1} |P(0)_{q,b}| \right. \\ &\quad \left. + \theta^2 |P(0)_{q,q}| + h^{2q+1-k-l} \right] \\ &\stackrel{\text{Ass. 2}}{\leq} K[1 + \theta + \theta^2] h^{2q+1-k-l}. \end{aligned} \quad (32)$$

Recall that P and Q are (positive semi-definite) covariance matrices; hence, $P^-(h)_{1,1} \geq Kh^{2q-1}$. Inserting these orders into the definition of $\beta^{(i)}$ (Eq. (11)), recalling that $R \geq 0$, and removing the dependence on θ by reducing the fraction conclude the proof. □

6 Local convergence rates

With the above bounds on intermediate algorithmic quantities (involving state misalignments $\delta^{(i)}$) in place, we only need an additional assumption to proceed—via a bound on $\delta^{(i)}(0)$ —to our first main result on local convergence orders of Ψ .

Assumption 3 The initial errors on the initial estimate of the i th derivative $m^{(i)}(0)$ satisfy $\|\varepsilon^{(i)}(0)\| = \|m^{(i)}(0) - x^{(i)}(0)\| \leq K_0 h^{q+1-i}$. (This assumption is, like Assumption 2, weaker than the standard assumption of exact initializations.)

Lemma 7 Under Assumptions 1 and 3, for all $i \in [q + 1]$ and for all $h > 0$, $\delta^{(i)}(0) \leq K h^{q+1-i}$.

Proof The claim follows, using Assumptions 1 and 3, from

$$\delta^{(i)}(0) \leq \underbrace{\|m^{(i)}(0) - x^{(i)}(0)\|}_{=\|\varepsilon^{(i)}(0)\| \leq K_0 h^{q+1-i}} + \underbrace{\|f^{(i-1)}(x^{(0)}(0)) - f^{(i-1)}(m^{(0)}(0))\|}_{\leq L\|\varepsilon^{(0)}(0)\| \leq LK_0 h^{q+1}}. \tag{33}$$

□

Now, we can bound the local truncation error $\varepsilon^{(0)}(h)$ as defined in Eq. (26).

Theorem 8 (Local Truncation Error) Under the Assumptions 1 to 3 and for all sufficiently small $h > 0$,

$$\|\varepsilon^{(0)}(h)\| \leq \|\mathbf{e}(h)\|_h \leq K \left[1 + \theta \|m^{(q)}(0)\| \right] h^{q+1}. \tag{34}$$

Proof By the triangle inequality for $\|\cdot\|_h$ and subsequent application of Lemma 3 and Assumption 3 to the second summand of the resulting inequality, we obtain

$$\|\mathbf{e}(h)\|_h \leq \underbrace{\|\Psi_{P(0),h}(m(0)) - \Phi_h(x^{(0)}(0))\|_h}_{=\sum_{i=0}^q h^i \Delta^{(i)}(h), \text{ by eq. (27)}} + \underbrace{\|\Phi_h(x^{(0)}(0)) - \Phi_h(m^{(0)}(0))\|_h}_{\leq (1+Kh)\|\varepsilon^{(0)}(0)\| \leq Kh^{q+1}}. \tag{35}$$

The remaining bound on $\Delta^{(i)}(h)$, for all $i \in [q + 1]$ and sufficiently small $h > 0$, is obtained by insertion of the bounds from Lemmas 4 to 6 (in the case of $n = 0$), into Eq. (28):

$$\Delta^{(i)}(h) \leq K \left[1 + \theta \|m^{(q)}(0)\| \right] h^{q+1-i} + K \sum_{k=1}^q \frac{h^{k-1}}{(k-1)!} \delta^{(k)}(nh) \tag{36}$$

$$\stackrel{\text{Lemma 7}}{\leq} K \left[1 + \theta \|m^{(q)}(0)\| \right] h^{q+1-i}. \tag{37}$$

Insertion of Eq. (37) into Eq. (35) and $\|\varepsilon^{(0)}(h)\| \leq \|\mathbf{e}(h)\|_h$ (by Eq. (19)) concludes the proof. □

Remark 2 Theorem 8 establishes a bound of order h^{q+1} on the local truncation error $\varepsilon^{(0)}(h)$ on $x(h)$ after one step h .

Moreover, by the definition Eq. (19) of $\|\cdot\|_h$, this theorem also implies additional bounds of order h^{q+1-i} on the error $\varepsilon^{(i)}(h)$ on the i th derivative $x^{(i)}(h)$ for all $i \in [q + 1]$. Such derivative bounds are (to the best of our knowledge) not available for classical numerical solvers, since they do not explicitly model the derivatives in the first place. These bounds could be useful for subsequent computations based on the ODE trajectory (Hennig et al. 2015).

Unsurprisingly, as the mean prediction (recall Eq. (8)) deviates from a pure q th order Taylor expansion by $K\theta \|m^{(q)}(0)\| h^{q+1}$ for an IOUP prior (i.e., $\theta > 0$ in Eq. (5)), the constant in front of the local h^{q+1} convergence rate depends on both θ and $m^{(q)}(0)$ in the IOUP case. A global analysis for IOUP is therefore more complicated than for IBM: Recall from Eq. (8) that, for $q = 1$, the mean prediction for $x((n + 1)h)$ is

$$\begin{aligned} & \left(\begin{matrix} m^{-,(0)}((n + 1)h) \\ m^{-,(1)}((n + 1)h) \end{matrix} \right) \stackrel{\text{eq. (8)}}{=} \\ & \left(\begin{matrix} m^{(0)}(nh) + hm^{(1)}(nh) - \theta \left[\frac{h^2}{2!} + \mathcal{O}(h^3) \right] m^{(1)}(nh) \\ e^{-\theta h} m^{-,(1)}(nh) \end{matrix} \right), \end{aligned} \tag{38}$$

which pulls both $m^{-,(0)}$ and $m^{-,(1)}$ towards zero (or some other prior mean) compared to the prediction given by its Taylor expansion for $\theta = 0$. While this is useful for ODEs converging to zero, such as $\dot{x} = -x$, it is problematic for diverging ODEs, such as $\dot{x} = x$ (Magnani et al. 2017). As shown in Theorem 8, this effect is asymptotically negligible for local convergence, but it might matter globally and, therefore, might necessitate stronger assumptions on f than Assumption 1, such as a bound on $\|f\|_\infty$ which would globally bound $\{y(nh); n = 0, \dots, T/h\}$ and thereby $\{m^{(1)}(nh); n = 0, \dots, T/h\}$ in Eq. (38). It is furthermore conceivable that a global bound for IOUP would depend on the relation between θ and $\|f\|_\infty$ in a non-trivial way. The inclusion of IOUP ($\theta > 0$) would hence complicate the below proofs further. Therefore, we restrict the following first global analysis to IBM ($\theta = 0$).

7 Global analysis

As explained in Remark 2, we only consider the case of the IBM prior, i.e., $\theta = 0$, in this section. Moreover, we restrict our analysis to $q = 1$ in this first global analysis. Although we only have definite knowledge for $q = 1$, we believe that the convergence rates might also hold for higher $q \in \mathbb{N}$ —which we experimentally test in Sect. 9.1. Moreover, we believe that proofs analogous to the below proofs

might work out for higher $q \in \mathbb{N}$ and that deriving a generalized version of Proposition 10 for higher q is the bottleneck for such proofs (see Sect. 10 for a discussion of these restrictions).

While, for local convergence, all noise models R yield the same convergence rates in Theorem 8, it is unclear how the order of R in h (as described in Sect. 2.3) affects global convergence rates: e.g., for the limiting case $R \equiv Kh^0$, the steady-state Kalman gains β^∞ would converge to zero (see Eqs. (43) and 44 below) for $h \rightarrow 0$, and hence the evaluation of f would not be taken into account—yielding a filter Ψ which assumes that the evaluations of f are equally off, regardless of $h > 0$, and eventually just extrapolates along the prior without global convergence of the posterior mean m . For the opposite limiting case $R \equiv \lim_{p \rightarrow \infty} Kh^p \equiv 0$, it has already been shown in Schober et al. (2019, Proposition 1 and Theorem 1) that—in the steady state and for $q = 1, 2$ —the filter Ψ inherits global convergence rates from known multistep methods in Nordsieck form Nordsieck (1962). To explore a more general noise model, we assume a fixed noise model $R \equiv Kh^p$ with arbitrary order p .

In the following, we analyze how small p can be in order for Ψ to exhibit fast global convergence (cf. the similar role of the order p of perturbations in Conrad et al. (2017, Assumption 1) and Abdulle and Garegnani (2020, Assumption 2.2)). In light of Theorem 8, the highest possible global convergence rate is $\mathcal{O}(h)$ —which will indeed be obtained for all $p \in [1, \infty]$ in Theorem 14. Since every extrapolation step of Ψ from t to $t + h$ depends not only on the current state, but also on the covariance matrix $P(t)$ —which itself depends on all previous steps— Ψ is neither a single-step nor a multistep method. Contrary to Schober et al. (2019), we do not restrict our theoretical analysis to the steady-state case, but provide our results under the weaker Assumptions 2 and 3 that were already sufficient for local convergence in Theorem 8—which is made possible by the bounds Eqs. (48) and (49) in Proposition 10.

7.1 Outline of global convergence proof

The goal of the following sequence of proofs in Sect. 7 is Theorem 14. It is proved by a special version of the discrete Grönwall inequality (Clark 1987) whose prerequisite is provided in Lemma 13. This Lemma 13 follows from Lemma 3 (on the regularity of the flow map Φ_t) as well as Lemma 12 which provides a bound on the maximal increment of the numerical error stemming from local truncation errors. For the proof of Lemma 12, we first have to establish

- (i) global bounds on the Kalman gains $\beta^{(0)}$ and $\beta^{(1)}$ by the inequalities Eqs. (48) and (49) in Proposition 10, and

- (ii) a global bound on the state misalignment $\delta^{(1)}$ in Lemma 11.

In Sects. 7.2–7.4, we will collect these inequalities in the order of their numbering to subsequently prove global convergence in Sect. 7.5.

7.2 Global bounds on Kalman gains

Since we will analyze the sequence of covariance matrices and Kalman gains using contractions in Proposition 10, we first introduce the following generalization of Banach fixed-point theorem (BFT).

Lemma 9 *Let (\mathcal{X}, d) be a non-empty complete metric space, $T_n: \mathcal{X} \rightarrow \mathcal{X}, n \in \mathbb{N}$, a sequence of L_n -Lipschitz continuous contractions with $\sup_n L_n \leq \bar{L} < 1$. Let u_n be the fixed point of T_n , as given by BFT, and let $\lim_{n \rightarrow \infty} u_n = u^* \in \mathcal{X}$. Then, for all $x_0 \in \mathcal{X}$, the recursive sequence $x_n := T_n(x_{n-1})$ converges to u^* as $n \rightarrow \infty$.*

Proof See Appendix G. □

In the following, we will assume that T is a multiple of h .

Proposition 10 *For constant $R \equiv Kh^p$ with $p \in [0, \infty]$, the unique (attractive) steady states for the following quantities are*

$$P_{11}^{-,\infty} := \lim_{n \rightarrow \infty} P_{11}^-(nh) = \frac{1}{2} \left(\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} \right), \tag{39}$$

$$P_{11}^\infty := \lim_{n \rightarrow \infty} P_{11}(nh) = \frac{\left(\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} \right) R}{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 2R}, \tag{40}$$

$$P_{01}^{-,\infty} := \lim_{n \rightarrow \infty} P_{01}^-(nh) = \frac{\sigma^4 h^2 + (2R + \sigma^2 h)\sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 4R\sigma^2 h}{2(\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2})} h, \tag{41}$$

$$P_{01}^\infty := \lim_{n \rightarrow \infty} P_{01}(nh) = \frac{R\sqrt{4R\sigma^2 h + \sigma^4 h^2}}{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2}} h, \tag{42}$$

$$\beta^{\infty,(0)} := \lim_{n \rightarrow \infty} \beta^{(0)}(nh) = \frac{\sqrt{4R\sigma^2 h + \sigma^4 h^2}}{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2}} h, \quad \text{and} \tag{43}$$

$$\beta^{\infty,(1)} := \lim_{n \rightarrow \infty} \beta^{(1)}(nh) = \frac{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2}}{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 2R}. \tag{44}$$

If furthermore Assumption 2 holds, then, for all sufficiently small $h > 0$,

$$\max_{n \in [T/h+1]} P_{11}^-(nh) \leq Kh^{1 \wedge \frac{p+1}{2}}, \tag{45}$$

$$\max_{n \in [T/h+1]} P_{11}(nh) \leq Kh^{p \vee \frac{p+1}{2}}, \tag{46}$$

$$\max_{n \in [T/h+1]} \|P_{01}(nh)\| \leq Kh^{p+1}, \tag{47}$$

$$\max_{n \in [T/h+1]} \|\beta^{(0)}(nh)\| \leq Kh, \quad \text{and} \tag{48}$$

$$\max_{n \in [T/h+1]} \|1 - \beta^{(1)}(nh)\| \leq Kh^{(p-1) \vee 0}. \tag{49}$$

All of these bounds are sharp in the sense that they fail for any higher order in the exponent of h .

Remark 3 The recursions for $P(nh)$ and $P^-(nh)$ given by Eqs. (10) and (15) follow a discrete algebraic Riccati equation (DARE)—a topic studied in many related settings (Lancaster and Rodman 1995). While the asymptotic behavior Eq. (39) of the completely detectable state $X^{(1)}$ can also be obtained using classical filtering theory (Anderson and Moore 1979, Chapter 4.4), the remaining statements of Proposition 10 also concern the undetectable state $X^{(0)}$ and are, to the best of our knowledge, not directly obtainable from existing theory on DAREs or filtering (which makes the following proof necessary). Note that, in the special case of no measurement noise ($R \equiv 0$), Eqs. (43) and (44) yield the equivalence of the filter in the steady state with the P(EC)1 implementation of the trapezoidal rule, which was previously shown in Schober et al. (2019, Proposition 1). For future research, it would be interesting to examine whether insertion of positive choices of R into Eqs. (43) and (44) can reproduce known methods as well.

Proof See Appendix H. □

7.3 Global bounds on state misalignments

For the following estimates, we restrict the choice of p to be larger than $q = 1$.

Assumption 4 The noise model is chosen to be $R \equiv Kh^p$, for $p \in [q, \infty] = [1, \infty]$, where $Kh^\infty := 0$.

Before bounding the added deviation of Ψ from the flow Φ per step, a global bound on the state misalignments defined in Eq. (25) is necessary. The result of the following lemma is discussed in Appendix D.

Lemma 11 Under Assumptions 1 to 4 and for all sufficiently small $h > 0$,

$$\max_{n \in [T/h+1]} \delta^{(1)}(nh) \leq Kh. \tag{50}$$

Proof See Appendix I. □

See Lemma 11 for a experimental demonstration of Eq. (33).

7.4 Prerequisite for discrete Grönwall inequality

Equipped with the above bounds, we can now prove a bound on the maximal increment of the numerical error stemming from local truncation errors which is needed to prove Eq. (56), the prerequisite for the discrete Grönwall inequality.

Lemma 12 Under Assumptions 1 to 4 and for all sufficiently small $h > 0$,

$$\begin{aligned} \max_{n \in [T/h+1]} \left\| \Psi_{P(nh),h}(\mathbf{m}(nh)) - \Phi_h(m^{(0)}(nh)) \right\|_h \\ \leq Kh^2. \end{aligned} \tag{51}$$

Proof By Eq. (19), we have

$$\begin{aligned} \left\| \Psi_{P(nh),h}(\mathbf{m}(nh)) - \Phi_h(m^{(0)}(nh)) \right\|_h \\ = S_1(h) + hS_2(h), \end{aligned} \tag{52}$$

with $S_1(h)$ and $S_2(h)$ defined and bounded by

$$\begin{aligned} S_1(h) &:= \left\| \Psi_h^{(0)}(\mathbf{m}(nh)) - \Phi_h^{(0)}(m^{(0)}(nh)) \right\| \\ &\stackrel{\text{eq. (28)}}{\leq} \underbrace{\Delta^{-(0)}((n+1)h)}_{\stackrel{\text{eq. (29)}}{\leq} Kh^2 + \delta^{(0)}(nh) + h\delta^{(1)}(nh)} \\ &\quad + \underbrace{\left\| \beta^{(0)}((n+1)h) \right\|}_{\stackrel{\text{eq. (48)}}{\leq} Kh} \underbrace{\|r((n+1)h)\|}_{\stackrel{\text{eq. (31)}}{\leq} Kh + (1+Kh)\delta^{(1)}(nh)}, \end{aligned} \tag{53}$$

and, analogously,

$$\begin{aligned} S_2(h) &:= \left\| \Psi_h^{(1)}(\mathbf{m}(nh)) - \Phi_h^{(1)}(m^{(0)}(nh)) \right\| \\ &\stackrel{\text{eq. (28)}}{\leq} \underbrace{\Delta^{-(1)}((n+1)h)}_{\stackrel{\text{eq. (29)}}{\leq} Kh + \delta^{(1)}(nh)} \\ &\quad + \underbrace{\left\| \beta^{(1)}((n+1)h) \right\|}_{\stackrel{\text{eq. (11)}}{\leq} 1} \underbrace{\|r((n+1)h)\|}_{\stackrel{\text{eq. (31)}}{\leq} Kh + (1+Kh)\delta^{(1)}(nh)} \end{aligned} \tag{54}$$

Insertion of Eqs. (53) and (54) into Eq. (52) yields

$$\begin{aligned} \left\| \Psi_{P(nh),h}(\mathbf{m}(nh)) - \Phi_h(m^{(0)}(nh)) \right\|_h \\ \leq Kh^2 + \delta^{(0)}(nh) + Kh\delta^{(1)}(nh), \end{aligned} \tag{55}$$

which—after recalling $\delta^{(0)}(nh) = 0$ and applying Lemma 11 to $\delta^{(1)}(nh)$ —implies Eq. (51). \square

The previous lemma now implies a suitable prerequisite for a discrete Grönwall inequality.

Lemma 13 *Under Assumptions 1 to 4 and for all sufficiently small $h > 0$,*

$$\|\boldsymbol{\varepsilon}((n + 1)h)\|_h \leq Kh^2 + (1 + Kh) \|\boldsymbol{\varepsilon}^{(0)}(nh)\|. \tag{56}$$

Proof We observe, by the triangle inequality for the norm $\|\cdot\|_h$, that

$$\begin{aligned} &\|\boldsymbol{\varepsilon}((n + 1)h)\|_h \\ &= \left\| \Psi_{P(nh),h}(\mathbf{m}(nh)) - \Phi_h(x^{(0)}(nh)) \right\|_h \\ &\leq \left\| \Psi_{P(nh),h}(\mathbf{m}(nh)) - \Phi_h(m^{(0)}(nh)) \right\|_h \\ &\quad + \left\| \Phi_h(m^{(0)}(nh)) - \Phi_h(x^{(0)}(nh)) \right\|_h. \end{aligned} \tag{57}$$

The proof is concluded by applying Lemma 12 to the first and Lemma 3 to the second summand of this bound (as well as recalling from Eq. (26) that $\|\boldsymbol{\varepsilon}^{(0)}(nh)\| = \|m^{(0)}(nh) - x^{(0)}(nh)\|$). \square

7.5 Global convergence rates

With the above bounds in place, we can now prove global convergence rates.

Theorem 14 (Global truncation error) *Under Assumptions 1 to 4 and for all sufficiently small $h > 0$,*

$$\max_{n \in [T/h+1]} \|\boldsymbol{\varepsilon}^{(0)}(nh)\| \leq \max_{n \in [T/h+1]} \|\boldsymbol{\varepsilon}(nh)\|_h \leq K(T)h, \tag{58}$$

where $K(T) > 0$ is a constant that depends on T , but not on h .

Remark 4 Theorem 14 not only implies that the truncation error $\|\boldsymbol{\varepsilon}^{(0)}(nh)\|$ on the solution of Eq. (1) has global order h , but also (by Eq. (19)) that the truncation error $\|\boldsymbol{\varepsilon}^{(1)}(nh)\|$ on the derivative is uniformly bounded by a constant K independent of h . The convergence rate of this theorem is sharp in the sense that it cannot be improved over all f satisfying Assumption 1 since it is one order worse than the local convergence rate implied by Theorem 8.

Proof Using $\|\boldsymbol{\varepsilon}^{(0)}(nh)\| \leq \|\boldsymbol{\varepsilon}(nh)\|_h$ (due to Eq. (19)), the bound Eq. (56), a telescoping sum, and $\|\boldsymbol{\varepsilon}(0)\|_h \leq Kh^2$ (by

Assumption 3), we obtain, for all sufficiently small $h > 0$, that

$$\begin{aligned} &\|\boldsymbol{\varepsilon}((n + 1)h)\|_h - \|\boldsymbol{\varepsilon}(nh)\|_h \\ &\stackrel{\text{eq. (19)}}{\leq} \|\boldsymbol{\varepsilon}((n + 1)h)\|_h - \|\boldsymbol{\varepsilon}^{(0)}(nh)\| \\ &\stackrel{\text{eq. (56)}}{\leq} Kh^2 + Kh \|\boldsymbol{\varepsilon}^{(0)}(nh)\| \\ &\stackrel{\text{eq. (19)}}{\leq} Kh^2 + Kh \|\boldsymbol{\varepsilon}(nh)\|_h \\ &\stackrel{(\text{tel. sum})}{=} Kh^2 + \|\boldsymbol{\varepsilon}(0)\|_h \\ &\quad + Kh \sum_{l=0}^{n-1} (\|\boldsymbol{\varepsilon}((l + 1)h)\|_h - \|\boldsymbol{\varepsilon}(lh)\|_h) \\ &\stackrel{(\|\boldsymbol{\varepsilon}(0)\|_h \leq Kh^2)}{\leq} Kh^2 \\ &\quad + Kh \sum_{l=0}^{n-1} (\|\boldsymbol{\varepsilon}((l + 1)h)\|_h - \|\boldsymbol{\varepsilon}(lh)\|_h). \end{aligned} \tag{59}$$

Now, by a special version of the discrete Grönwall inequality (Clark 1987), if z_n and g_n are sequences of real numbers (with $g_n \geq 0$), $c \geq 0$ is a nonnegative constant, and if

$$z_n \leq c + \sum_{l=0}^{n-1} g_l z_l, \quad \text{for all } n \in \mathbb{N}, \tag{60}$$

then

$$z_n \leq c \prod_{l=0}^{n-1} (1 + g_l) \leq c \exp\left(\sum_{l=0}^{n-1} g_l\right), \quad \text{for all } n \in \mathbb{N}.$$

Application of this inequality to Eq. (59) with $z_n := \|\boldsymbol{\varepsilon}((n + 1)h)\|_h - \|\boldsymbol{\varepsilon}(nh)\|_h$, $g_n := Kh$, and $c := Kh^2$ yields

$$\|\boldsymbol{\varepsilon}((n + 1)h)\|_h - \|\boldsymbol{\varepsilon}(nh)\|_h \leq K(T)h^2 \exp(nKh) \tag{61}$$

$$\stackrel{n \leq T/h}{\leq} K(T)h^2. \tag{62}$$

By another telescoping sum argument and $\|\boldsymbol{\varepsilon}(0)\|_h \leq Kh^2$, we obtain

$$\begin{aligned} \|\boldsymbol{\varepsilon}(nh)\|_h &\stackrel{(\text{tel. sum})}{=} \sum_{l=0}^{n-1} (\|\boldsymbol{\varepsilon}((l + 1)h)\|_h - \|\boldsymbol{\varepsilon}(lh)\|_h) \\ &\quad + \|\boldsymbol{\varepsilon}(0)\|_h \end{aligned} \tag{63}$$

$$\stackrel{\text{eq. (62)}}{\leq} nK(T)h^2 + Kh^2 \tag{64}$$

$$\stackrel{n \leq T/h}{\leq} K(T)h + Kh^2 \tag{65}$$

$$\leq K(T)h + Kh^2, \tag{66}$$

for all sufficiently small $h > 0$. Recalling that $\|\varepsilon^{(0)}(nh)\| \leq \|e(nh)\|_h$, by Eq. (19), concludes the proof. \square

8 Calibration of credible intervals

In PN, one way to judge calibration of a Gaussian output $\mathcal{N}(m, V)$ is to check whether the implied 0.95 credible interval $[m - 2\sqrt{V}, m + 2\sqrt{V}]$ contracts at the same rate as the convergence rate of the posterior mean to the true quantity of interest. For the filter, this would mean that the rate of contraction of $\max_n \sqrt{P_{00}(nh)}$ should contract at the same rate as $\max_{n \in [T/h+1]} \|\varepsilon^{(0)}(nh)\|$ (recall its rates from Theorem 14). Otherwise, for a higher or lower rate of the interval it would eventually be under- or overconfident, as $h \rightarrow 0$. The following proposition shows—in light of the sharp bound Eq. (58) on the global error—that the credible intervals are well calibrated in this sense if $p \in [1, \infty]$.

Theorem 15 *Under Assumption 2 and for $R \equiv Kh^p$, $p \in [0, \infty]$, as well as sufficiently small $h > 0$,*

$$\max_{n \in [T/h+1]} P_{00}^-(nh) \leq K(T)h^{(p+1)\wedge 2}, \quad \text{and} \quad (67)$$

$$\max_{n \in [T/h+1]} P_{00}(nh) \leq K(T)h^{(p+1)\wedge 2}. \quad (68)$$

Proof See Appendix J. \square

9 Numerical experiments

In this section, we empirically assess the following hypotheses:

- (i) the worst-case convergence rates from Theorem 14 hold not only for $q = 1$ but also for $q \in \{2, 3\}$ (see Sect. 9.1),
- (ii) the convergence rates of the credible intervals from Theorem 15 hold true (see Sect. 9.2), and
- (iii) Assumption 4 is necessary to get these convergence rates (see Sect. 9.3).

The three hypotheses are all supported by the experiments. These experiments are subsequently discussed in Sect. 9.4. Appendix D contains an additional experiment illustrating the convergence rates for the state misalignment δ from Lemma 11.

9.1 Global convergence rates for $q \in \{1, 2, 3\}$

We consider the following three test IVPs: Firstly, the following linear ODE

$$\begin{aligned} \dot{x}(t) &= \Lambda x(t), \quad \forall t \in [0, 10], \\ \text{with } \Lambda &= \begin{pmatrix} 0 & -\pi \\ \pi & 0 \end{pmatrix} \text{ and } x(0) = (0, 1)^T, \end{aligned} \quad (69)$$

and has the harmonic oscillator

$$x(t) = e^{t\Lambda}x(0) = (-\sin(t\pi) \cos(t\pi))^T \quad (70)$$

as a solution. Secondly, the logistic equation

$$\begin{aligned} \dot{x}(t) &= \lambda_0 x(t) (1 - x(t)/\lambda_1), \quad \forall t \in [0, 1.5], \\ \text{with } (\lambda_0, \lambda_1) &= (3, 1) \text{ and } x(0) = 0.1, \end{aligned} \quad (71)$$

which has the logistic curve

$$x(t) = \frac{\lambda_1 \exp(\lambda_0 t)x(0)}{\lambda_1 + x(0)(\exp(\lambda_0 t) - 1)}. \quad (72)$$

And, thirdly, the FitzHugh–Nagumo model

$$\begin{pmatrix} \dot{x}_1(t) \\ \dot{x}_2(t) \end{pmatrix} = \begin{pmatrix} x_1(t) - \frac{x_1(t)}{3} - x_2(t) \\ \frac{1}{\tau}(x_1(t) + a - bx_2(t)) \end{pmatrix}, \quad \forall t \in [0, 10] \quad (73)$$

with $(a, b, c) = (0.08, 0.07, 1.25)$ and $x(0) = (1, 0)$ which does not have a closed-form solution. Its solution, which we approximate by Euler’s method with a step size of $h = 10^{-6}$ for the below experiments, is depicted in Fig. 1. We numerically solve these three IVPs with the Gaussian ODE filter for multiple step sizes $h > 0$ and with a q -times IBM prior (i.e., $\theta = 0$ in Eq. (5)) for $q \in \{1, 2, 3\}$ and scale $\sigma = 20$. As a measurement model, we employ the minimal $R \equiv 0$ and maximal measurement variance $R \equiv K_R h^q$ (for $h \leq 1$) which are permissible under Assumption 4 whose constant $K > 0$ is denoted explicitly by K_R in this section. The resulting convergence rates of global errors $\|m(T) - x(T)\|$ are depicted in a work-precision diagram in Fig. 2; cf. Hairer et al. (1987, Chapter II.1.4) for such diagrams for Runge–Kutta methods. Now, recall from Theorem 14 that, for $q = 1$, the global truncation error decreases at a rate of at least h^q in the worst case. Figure 2 shows that these convergence rates of q th order hold true in the considered examples for values of up to $q = 3$ if $R \equiv 0$ and, for values of up to $q = 3$. In the case of $R \equiv 0$, even $(q + 1)$ th order convergence rates appear to hold true for all three ODEs and $q \in \{1, 2, 3\}$. Note that it is more difficult to validate these convergence rates for $q = 4$, for all three test problems and small $h > 0$, since numerical instability can contaminate the analytical rates.




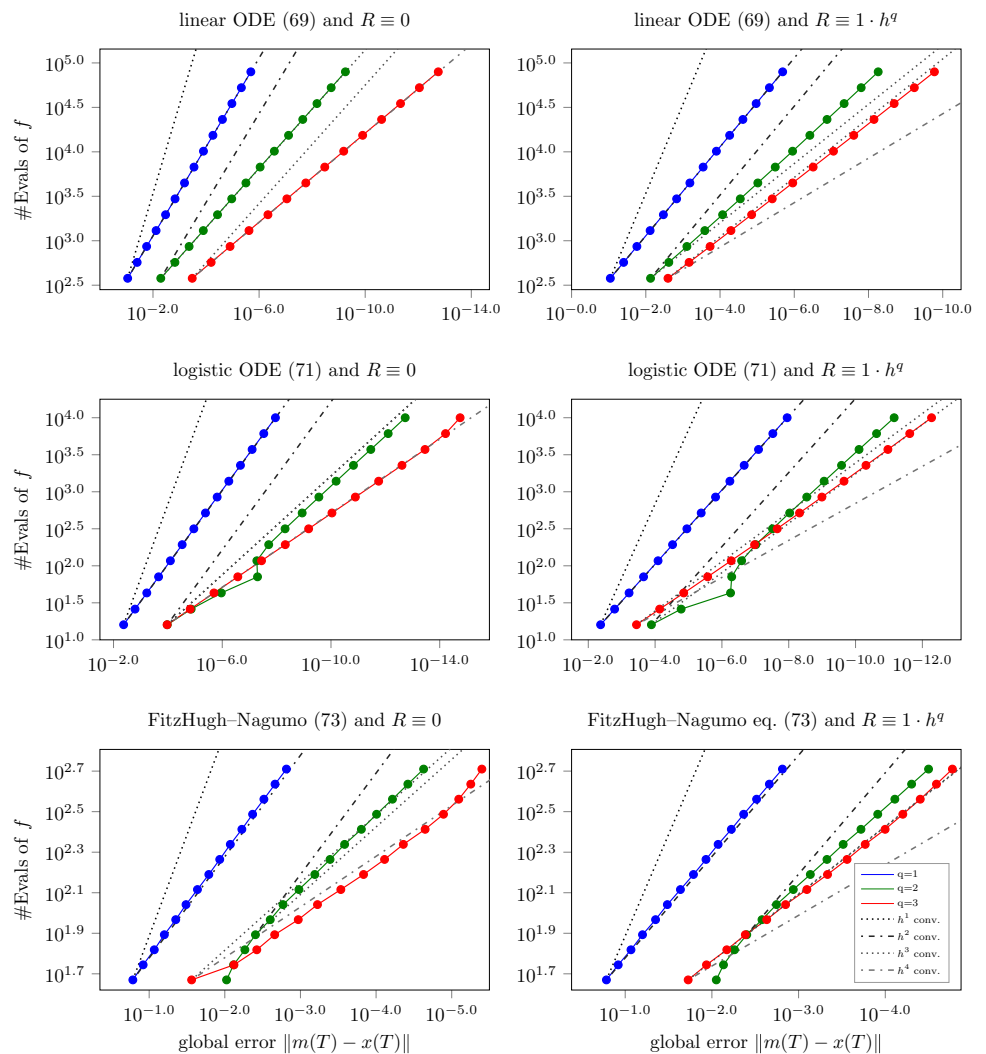
Fig. 1 True solution of the FitzHugh–Nagumo model, Eq. (73); x_1 in blue and x_2 in orange

9.2 Calibration of credible intervals

To demonstrate the convergence rates of the posterior credible intervals proved in Theorem 15, we now restrict our attention to the case of $q = 1$, that was considered therein. As in Sect. 9.1, we numerically solve the IVPs eqs. (69) and

(71) with the Gaussian ODE filter with a once IBM prior with fixed scale $\sigma = 1$. We again employ the minimal $R \equiv 0$ and maximal measurement variance $R \equiv K_R h^q$ (for $h \leq 1$) which are permissible under Assumption 4 as a measurement model. Figure 3 depicts the resulting convergence rates in work-precision diagrams. As the parallel standard deviation (std. dev.) and h^1 convergence curves show, the credible intervals asymptotically contract at the rate of h^1 guaranteed by Theorem 15. In all four diagrams of Fig. 3, the global error shrinks at a faster rate than the width of the credible intervals. This is unsurprising for $R \equiv 0$ as we have already observed convergence rates of h^{q+1} in this case. While this effect is less pronounced for $R \equiv K_R h^q$, it still results in underconfidence as $h \rightarrow 0$. Remarkably, the shrinking of the standard deviations seems to be ‘adaptive’ to the numerical error—by which we mean that, as long as the numerical error hardly decreases (up to $10^{1.75}$ evaluations of f), the standard deviation also stays almost constant, before adopting its h^1 convergence asymptotic (from $\approx 10^{2.00}$).

Fig. 2 Work-precision diagrams for the Gaussian ODE filter with q -times IBM prior, for $q \in \{1, 2, 3\}$, applied to the linear Eq. (71), logistic ODE Eq. (69) and the FitzHugh–Nagumo model. The number of function evaluations (# Evals of f), which is inversely proportional to the step size h , is plotted in color against the logarithmic global error at the final time T . The (dash-)dotted gray lines visualize idealized convergence rates of orders one to four. The left and right columns employ the minimal $R \equiv 0$ and maximal measurement variance $R \equiv K_R h^q$ ($K_R = 1$) which are permissible under Assumption 4



9.3 Necessity of Assumption 4

Having explored the asymptotic properties under Assumption 4 in Sects. 9.1 and 9.2, we now turn our attention to the question of whether this assumption is necessary to guarantee the convergence rates from Theorems 14 and 15. This question is of significance, because Assumption 4 is weaker than the $R \equiv 0$ assumption of the previous theoretical results (i.e., Proposition 1 and Theorem 1 in Schober et al. (2019)) and it is not self-evident that it cannot be further relaxed. To this end, we numerically solve the logistic ODE Eq. (71) with the Gaussian ODE filter with a once IBM prior with fixed scale $\sigma = 1$ and measurement variance $R \equiv K_R h^{1/2}$, which is impermissible under Assumption 4, for increasing choices of K_R from 0.00×10^0 to 1.00×10^7 . In the same way as in Fig. 3, the resulting work-precision diagrams are plotted in Fig. 4.

In contrast to the lower left diagram in Fig. 3, which presents the same experiment for $R \equiv K_R h^q$ (the maximal measurement variance permissible under Assumption 4), the rate of h^2 , that is again observed for $K_R = 0$ in the first diagram, is already missed for $K_R = 1.00 \times 10^0$ in the second diagram. With growing constants, the convergence rates of the actual errors as well as the expected errors (standard deviation) decrease from diagram to diagram. In the center diagram with $K_R = 3.73 \times 10^3$, the rates are already slightly worse than the h^1 convergence rates guaranteed by Theorems 14 and 15 under Assumption 4, whereas, for $K_R = 5.00 \times 10^3$, the convergence rates in the lower left plot of Fig. 3 were still significantly better than h^1 . For the greater constants up to $K_R = 1.00 \times 10^7$, the rates even become significantly lower. Notably, as in the lower right diagram of Fig. 3, the slope of the standard deviation curve matches the slope of the global error curve, as can be seen best in the lower right subfigure—thereby asymptotically exhibiting neither over- nor underconfidence. These experiments suggest that the convergence rates from Theorems 14 and 15 do not hold in general for $R \equiv K_R h^{1/2}$. Hence, it seems likely that Assumption 4 is indeed necessary for our results and cannot be further relaxed without lowering the implied worst-case convergence rates.

9.4 Discussion of experiments

Before proceeding to our overall conclusions, we close this section with a comprehensive discussion of the above experiments. First and foremost, the experiments in Sect. 9.1 suggest that Theorem 14, the main result of this paper, might be generalizable to $q \in \{2, 3\}$ and potentially even higher $q \in \mathbb{N}$ —although unresolved issues with numerical instability for small step sizes prevent us from confidently asserting that these theoretical results would hold in practice for $q \geq 4$. Moreover, we demonstrated the contraction

rates of the posterior credible intervals from Theorem 15 and evidence for the necessity of Assumption 4 in Sects. 9.3 and 9.2. The asymptotics revealed by these experiments can be divided by the employed measurement model into three cases: the zero-noise case $R \equiv 0$, the permissible nonzero case $R \leq K_R h^q$ (under Assumption 4) and the non-permissible case $R \not\leq K_R h^q$. First, if $R \equiv 0$, the diagrams in the left column of Fig. 2 reaffirm the h^{q+1} convergence reported for $q \in \{1, 2\}$ in Schober et al. (2019, Figure 4) and extend them to $q = 3$ (see Sect. 10 for a discussion on why we expect the above global convergence proofs to be extensible to $q \geq 2$)

The contraction rates of the credible intervals, for $q = 1$, appear to be asymptotically underconfident in this case as they contract faster than the error. This underconfidence is not surprising in so far as the posterior standard deviation is a worst-case bound for systems modeled by the prior, while the convergence proofs require smoothness of the solution of one order higher than sample paths from the prior. This is a typical result that highlights an aspect known to, but on the margins of classic analysis: The class of problems for which the algorithm converges is rougher than the class on which convergence order proofs operate. How to remedy such overly cautious UQ remains an open research question in PN as well as classical numerical analysis.

Secondly, in the case of $R > 0$, as permissible under Assumption 4, the convergence rates are slightly reduced compared to the case $R \equiv 0$, exhibiting convergence between h^q and h^{q+1} . The asymptotic underconfidence of the credible intervals, however, is either reduced or completely removed as depicted in the right column of Fig. 3. Thirdly, in the final case of an impermissibly large $R > 0$, the h^q convergence speed guaranteed by Theorem 14 indeed does not necessarily hold anymore—as depicted in Fig. 4. Note, however, that even then the convergence rate is only slightly worse than h^q . The asymptotic UQ matches the observed global error in this case, as the parallel standard deviation and the h^1 curves in all but the upper left $R \equiv 0$ diagram show.

Overall, the experiments suggest that, in absence of statistical noise on f , a zero-variance measurement model yields the best convergence rates of the posterior mean. Maybe this was expected as, in this case, R only models the inaccuracy from the truncation error, that ideally should be treated adaptively (Kersting and Hennig 2016, Section 2.2). The convergence rates of adaptive noise models should be assessed in future work. As the observed convergence rates in practice sometimes outperform the proved worst-case convergence rates, we believe that an average-case analysis of the filter in the spirit of Ritter (2000) may shed more light upon the expected practical performance. Furthermore, it appears that the UQ becomes asymptotically accurate as well as adaptive to the true numerical error as soon as the $R > 0$ is large

Fig. 3 Work-precision diagrams for the Gaussian ODE filter with q -times IBM prior, for $q = 1$, applied to the linear Eq. (69) and logistic ODE Eq. (71) in the upper and lower row, respectively. The number of function evaluations (# Evals of f), which is inversely proportional to the step size h , is plotted in color against the logarithmic global error at the final time T . The (dash)-dotted gray lines visualize idealized convergence rates of orders one and two. The dashed blue lines show the posterior standard deviations calculated by the filter. The left and right columns, respectively, employ the minimal $R \equiv 0$ and maximal measurement variance $R \equiv K_R h^q$ ($K_R = 5.00 \times 10^3$) which are permissible under Assumption 4


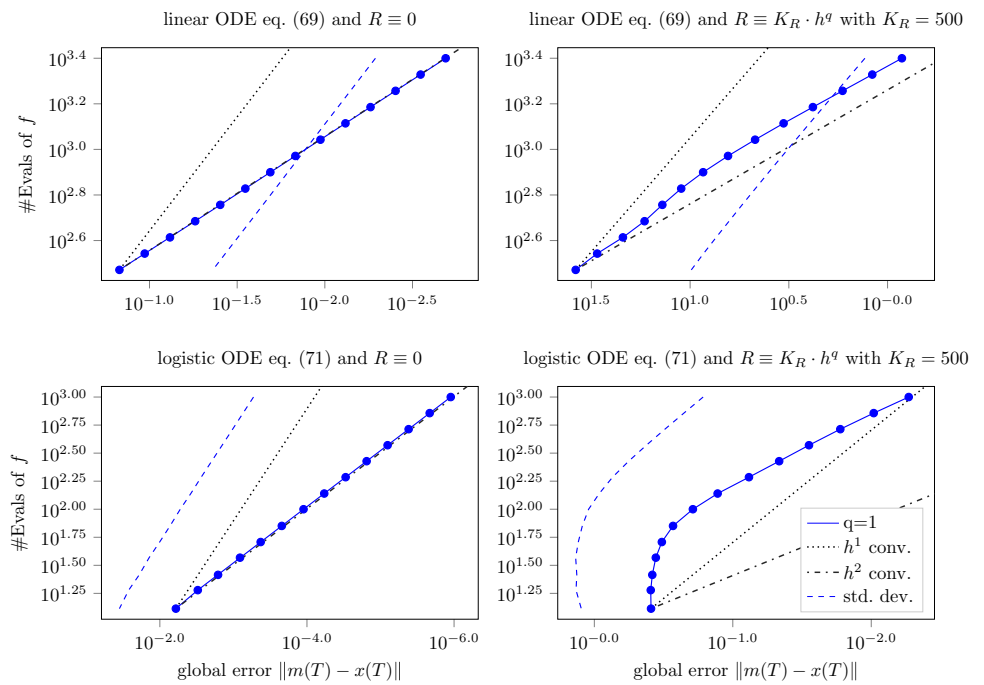


Fig. 4 Work-precision diagrams for the Gaussian ODE filter with $q = 1$ and $R \equiv K_R h^{1/2}$, applied to the logistic ODE Eq. (71) for increasing values of K_R . The number of function evaluations (# Evals of f), which is inversely proportional to the step

size h , is plotted in blue against the logarithmic global error at the final time T . The (dash)-dotted gray lines visualize idealized convergence rates of orders one and two. The dashed blue lines show the posterior standard deviations calculated by the filter

enough. This reinforces our hope that these algorithms will prove useful for IVPs when f is estimated itself (Hennig et al. 2015, Section 3(d)), thereby introducing a $R > 0$.

10 Conclusions

We presented a worst-case convergence rate analysis of the Gaussian ODE filter, comprising both local and global convergence rates. While local convergence rates of h^{q+1} were shown to hold for all $q \in \mathbb{N}$, IBM and IOUP prior as well as any noise model $R \geq 0$, our global convergence results is restricted to the case of $q = 1$, IBM prior and fixed noise model $R \equiv Kh^p$ with $p \in [1, \infty]$. While a restriction of the noise model seems inevitable, we believe that the other two restrictions can be lifted: In light of Theorem 8, global convergence rates for the IOUP prior might only require an additional assumption that ensures that all possible data sequences $\{y(nh); n = 1, \dots, T/h\}$ (and thereby all possible q th-state sequences $\{m^{(q)}(nh); n = 0, \dots, T/h\}$) remain uniformly bounded (see discussion in Remark 2). For the case of $q \geq 2$, it seems plausible that a proof analogous to the presented one would already yield global convergence rates of order h^q ,³ as suggested for $q \in \{2, 3\}$ by the experiments in Sect. 9.1.

The orders of the predictive credible intervals can also help to intuitively explain the threshold of $p = 1$ (or maybe more generally: $p = q$; see Fig. 2) below which the performance of the filter is not as good, due to Eqs. (45)–(49): According to Kersting and Hennig (2016, Equation (20)), the ‘true’ (push-forward) variance on $y(t)$ given the predictive distribution $\mathcal{N}(m^-(t), P^-(t))$ is equal to the integral of ff^T with respect to $\mathcal{N}(m^-(t), P^-(t))$, whose maximum over all time steps, by Eq. (67), has order $\mathcal{O}(h^{\frac{p+1}{2} \wedge 1})$ if ff^T is globally Lipschitz—since $P^-(t)$ enters the argument of the integrand ff^T , after a change of variable, only under a square root. Hence, the added ‘statistical’ noise R on the evaluation of f is of lower order than the accumulated ‘numerical’ variance $P^-(t)$ (thereby preventing numerical convergence) if and only if $p < 1$. Maybe this, in the spirit of Hennig et al. (2015, Subsection 3(d)), can serve as a criterion for vector fields f that are too roughly approximated for a numerical solver to output a trustworthy result, even as $h \rightarrow 0$.

Furthermore, the competitive practical performance of the filter, as numerically demonstrated in Schober et al. (2019, Section 5), might only be completely captured by an

average-case analysis in the sense of Ritter (2000), where the average error is computed with respect to some distribution $p(f)$, i.e., over a distribution of ODEs. To comprehend this idea, recall that the posterior filtering mean is the Bayes estimator with minimum mean squared error in linear dynamical systems with Gauss–Markov prior (as defined by the SDE Eq. (2)), i.e., when the data is not evaluations of f but real i.i.d. measurements, as well as in the special case of $\dot{x}(t) = f(t)$, when the IVP simplifies to a quadrature problem—see Solak et al. (2003) and O’Hagan (1991, Section 2.2), respectively. In fact, the entire purpose of the update step is to correct the prediction in the (on average) correct direction, while a worst-case analysis must assume that it corrects in the worst possible direction in every step—which we execute by the application of the triangle inequality in Eq. (28) resulting in a worst-case upper bound that is the sum of the worst-case errors from prediction and update step. An analysis of the probabilities of ‘good’ vs. ‘bad’ updates might therefore pave the way for such an average-case analysis in the setting of this paper. Since, in practice, truncation errors of ODE solvers tend to be significantly smaller than the worst case—as mirrored by the experiments in Section 9—such an analysis might be useful for applications.

Lastly, we hope that the presented convergence analysis can lay the foundations for similar results for the novel ODE filters (extended KF, unscented KF, particle filter) introduced in Tronarp et al. (2019), and can advance the research on uncertainty-aware likelihoods for inverse problems by ODE filtering (Kersting et al. 2020, Section 3).

Acknowledgements The authors are grateful to Han Cheng Lie for discussions and feedback to early versions of what is now Sect. 3 and 5 of this work, as well as Sect. 7.5. The authors also thank Michael Schober for valuable discussions and helpful comments on the manuscript. TJS’s work has been partially supported by the Freie Universität Berlin within the Excellence Initiative of the German Research Foundation (DFG), by the DFG through grant CRC 1114 ‘Scaling Cascades in Complex Systems,’ and by the National Science Foundation under grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute (SAMSI) and SAMSI’s QMC Working Group II ‘Probabilistic Numerics.’ HK and PH gratefully acknowledge financial support by the German Federal Ministry of Education and Research through BMBF grant 01IS18052B (ADIMEM). PH also gratefully acknowledges support through ERC StG Action 757275 / PANAMA. Any opinions, findings, and conclusions or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the above-named institutions and agencies.

Funding Open Access funding provided by Projekt DEAL.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

³ According to Loscalzo and Talbot (1967), the filter might, however, suffer from numerical instability for high choices of q . (See Schober et al. (2019, Section 3.1) for an explanation of how such results on spline-based methods concern the ODE filter.)

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Derivation of A and Q

As derived in Särkkä (2006, Section 2.2.6) the solution of the SDE Eq. (2), i.e.,

$$dX(t) = \begin{pmatrix} dX^{(0)}(t) \\ \vdots \\ dX^{(q-1)}(t) \\ dX^{(q)}(t) \end{pmatrix} = \underbrace{\begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ \vdots & & \ddots & \ddots & 1 \\ c_0 & \dots & \dots & \dots & c_q \end{pmatrix}}_{=:F} \underbrace{\begin{pmatrix} X^{(0)}(t) \\ \vdots \\ X^{(q-1)}(t) \\ X^{(q)}(t) \end{pmatrix}}_{=:X(t)} dt + \underbrace{\begin{pmatrix} 0 \\ \vdots \\ 0 \\ \sigma \end{pmatrix}}_{=:L} dB(t), \tag{74}$$

where we omitted the index j for simplicity, is a Gauss–Markov process with mean $m(t)$ and covariance matrix $P(t)$ given by

$$m(t) = A(t)m(0), \quad P(t) = A(t)P(0)A(t)^T + Q(t), \tag{75}$$

where the matrices $A, Q \in \mathbb{R}^{(q+1) \times (q+1)}$ are explicitly defined by

$$A(t) = \exp(tF), \tag{76}$$

$$Q(t) := \int_0^t \exp(F(t - \tau))LL^T \exp(F(t - \tau))^T d\tau. \tag{77}$$

Parts of the following calculation can be found in Magnani et al. (2017). If we choose $c_0, \dots, c_{q-1} = 0$ and $c_q = -\theta$ (for $\theta \geq 0$) in Eq. (74) the unique strong solution of the SDE is a q -times IOUP, if $\theta > 0$, and a q -times IBM, if $\theta = 0$; see, e.g., Karatzas and Shreve (1991, Chapter 5: Example 6.8). By Eq. (77) and

$$((tF)^k)_{i,j} = t^k \left[\mathbb{I}_{j-i=k} + (-\theta)^{k+i-q} \mathbb{I}_{\{j=q, i+k \geq q\}} \right], \tag{78}$$

it follows that

$$A(t)_{ij} = \left(\sum_{k=0}^{\infty} \frac{(tF)^k}{k!} \right)_{i,j} = \begin{cases} \mathbb{I}_{i \leq j} \frac{t^{j-i}}{(j-i)!}, & \text{if } j \neq q, \\ \frac{1}{(-\theta)^{q-i}} \sum_{k=q-i}^{\infty} \frac{(-\theta t)^k}{k!}, & \text{if } j = q, \end{cases} = \begin{cases} \mathbb{I}_{i \leq j} \frac{t^{j-i}}{(j-i)!}, & \text{if } j \neq q, \\ \frac{t^{q-i}}{(q-i)!} - \theta \sum_{k=q+1-i}^{\infty} \frac{(-\theta)^{k+i-q-1} t^k}{k!}, & \text{if } j = q. \end{cases} \tag{79}$$

Analogously, it follows that

$$\exp(F(t - \tau)) = \begin{cases} \mathbb{I}_{i \leq j} \frac{(t-\tau)^{j-i}}{(j-i)!}, & \text{if } j \neq q, \\ \frac{(t-\tau)^{q-i}}{(q-i)!} - \theta \sum_{k=q+1-i}^{\infty} \frac{(-\theta)^{k+i-q-1} (t-\tau)^k}{k!}, & \text{if } j = q, \end{cases} \tag{80}$$

If we insert Eq. (80) into Eq. (77), then we obtain, by the sparsity of L , that

$$Q(t)_{ij} = \frac{\sigma^2}{(-\theta)^{2q-i-j}} \int_0^t \left(\sum_{k=q-i}^{\infty} \frac{(-\theta \tau)^k}{k!} \right) \left(\sum_{l=q-j}^{\infty} \frac{(-\theta \tau)^l}{l!} \right) d\tau, \tag{81}$$

and the dominated convergence theorem (with dominating function $\tau \mapsto e^{2\theta\tau}$) yields

$$Q(t)_{ij} = \frac{\sigma^2}{(-\theta)^{2q-i-j}} \sum_{k=q-i}^{\infty} \sum_{l=q-j}^{\infty} \int_0^t \frac{(-\theta \tau)^{k+l}}{k!l!} d\tau = \frac{\sigma^2}{(-\theta)^{2q-i-j}} \sum_{k=q-i}^{\infty} \sum_{l=q-j}^{\infty} (-\theta)^{k+l} \frac{t^{k+l+1}}{(k+1+l)k!l!}. \tag{82}$$

Now, by extracting the first term and noticing that the rest of the series is in $\Theta(t^{2q+2-i-j})$, it follows that

$$Q(t)_{ij} = \sigma^2 \frac{t^{2q+1-i-j}}{(2q+1-i-j)(q-i)!(q-j)!} + \Theta \left(t^{2q+2-i-j} \right). \tag{83}$$

B Extension to x with dependent dimensions

The algorithm in Sect. 2.2 employs a prior X with independent dimensions $X_j = (X_j^{(0)}, \dots, X_j^{(q)})^T, j \in [d]$, by

Eq. (2). While this constitutes a loss of generality for our new theoretical results, which do not immediately carry over to the case of x with dependent dimensions, it is not a restriction to the class of models the algorithm can employ. To construct such a prior \mathbf{X} , we first stack its dimensions into the random vector $\mathbf{X} = (\mathbf{X}_0^\top, \dots, \mathbf{X}_{d-1}^\top)^\top$, choose symmetric positive semi-definite matrices $K_x, K_\varepsilon \in \mathbb{R}^{d \times d}$, and define, using the Kronecker product \otimes , its law according to the SDE

$$d\mathbf{X}(t) = [K_x \otimes F] \mathbf{X}(t) dt + [K_\varepsilon \otimes L] dB(t), \tag{84}$$

with initial condition $\mathbf{X}(0) \sim \mathcal{N}(m(0), P(0))$, mean $m(0) \in \mathbb{R}^{d(q+1)}$ and covariance matrix $P(0) \in \mathbb{R}^{d(q+1) \times d(q+1)}$, as well as an underlying d -dimensional Brownian motion B (independent of $\mathbf{X}(0)$). Now, insertion of $K_x \otimes F$ and $K_\varepsilon \otimes L$ for F and L into Eq. (77) yields new predictive matrices \tilde{A} and \tilde{Q} . If we now choose $K_x = I_d$ and $K_\varepsilon = I_d$, substitute \tilde{A} and \tilde{Q} for A and Q in Eqs. (9) and (10), and use the $d(q + 1)$ -dimensional GP \mathbf{X} from Eq. (84) with $m(0) \in \mathbb{R}^{d(q+1)}$ and $P(0) \in \mathbb{R}^{d(q+1) \times d(q+1)}$ as a prior, we have equivalently defined the version of Gaussian ODE filtering with independent dimensions from Sect. 2.2. If we, however, choose different symmetric positive semi-definite matrices for K_x and K_ε , we introduce, via \tilde{A} and \tilde{Q} , a correlation in the development of the solution dimensions $(x_0, \dots, x_{d-1})^\top$ as well as the error dimensions $(\varepsilon_0, \dots, \varepsilon_d)^\top$, respectively. Note that, while K_ε plays a similar role as C^h in Conrad et al. (2017, Assumption 1) in correlating the numerical errors, the matrix K_x additionally introduces a correlation of the numerical estimates, that is m , along the time axis. Even more flexible correlation models (over all modeled derivatives) can be employed by inserting arbitrary matrices (of the same dimensionality) for $K_x \otimes F$ and $K_\varepsilon \otimes L$ in Eq. (84), but such models seem hard to interpret. For future research, it would be interesting to examine whether such GP models with dependent dimensions are useful in practice. There are first publications (Xiaoyue et al. 2018; Gessner et al. 2019) on this topic for integrals, but not yet for ODEs.

C Illustrative example

To illustrate the algorithm defined in Sect. 2.2, we apply it to a special case of the Riccati equation (Davis 1962, p. 73)

$$\frac{dx}{dt}(t) = f(x(t)) = -\frac{(x(t))^3}{2}, \quad x(0) = 1, \tag{85}$$

$$\left(\text{solution: } x(t) = (t + 1)^{-1/2} \right), \tag{86}$$

with step size $h = 0.1$, measurement noise $R = 0.0$ (for simplicity) as well as prior hyperparameters $q = 1, \sigma^2 = 10.0$ and $c_i = 0$ for all $i \in [q + 1]$ (recall Eq. (2)), i.e., with

a 1-times integrated Brownian motion prior whose drift and diffusion matrices are, by Eq. (8), given by

$$A(h) = \begin{pmatrix} 1 & h \\ 0 & 1 \end{pmatrix}, \quad Q(h) = \begin{pmatrix} 1/300 & 1/20 \\ 1/20 & 1 \end{pmatrix}. \tag{87}$$

As the ODE Eq. (85) is one-dimensional (i.e., $d = 1$), the dimension index $j \in [d]$ is omitted in this section. Since the initial value and derivative are certain at $x(0) = 1$ and $\dot{x}(0) = f(x_0) = -1/2$, our prior GP is initialized with a Dirac distribution (i.e., $\mathbf{X}(0) = (X^{(0)}(0), X^{(1)}(0))^\top \sim \delta_{(x_0, f(x_0))} = \delta_{(1, -1/2)}$). Therefore, $\mathbf{m}(0) = (1, -1/2)^\top$ and $P(0) = 0 \in \mathbb{R}^{2 \times 2}$ for the initial filtering mean and covariance matrix. Now, the Gaussian ODE Filter computes the first integration step by executing the prediction step Eqs. (9) and (10)

$$\begin{aligned} \mathbf{m}^-(h) &= A(h)\mathbf{m}^-(0) \\ &= \left(m^{(0)}(0) + hm^{(1)}(0), m^{(1)}(0) \right)^\top \\ &= (19/20, -1/2)^\top, \quad \text{and} \end{aligned} \tag{88}$$

$$P^-(h) = 0 + Q(h) = \begin{pmatrix} 1/300 & 1/20 \\ 1/20 & 1 \end{pmatrix}. \tag{89}$$

Note that, for all $i \in [q + 1]$, $m^{-(i)}(h)$ is obtained by a $(q - i)$ th-order Taylor expansion of the state $\mathbf{m}(0) = (x_0, f(x_0))^\top \in \mathbb{R}^{q+1}$. Based on this prediction, the data is then generated by

$$\begin{aligned} y(h) &= f\left(m^{-(0)}(h)\right) \stackrel{\text{eq. (88)}}{=} f(19/20) \\ &\stackrel{\text{eq. (85)}}{=} -6859/16000 \end{aligned} \tag{90}$$

with variance $R = 0.0$. In the subsequent update step Eqs. (9) and (11) to (13), a Bayesian conditioning of the predictive distribution Eqs. (88) and (89) on this data is executed:

$$\begin{aligned} \boldsymbol{\beta}(h) &= \left(\beta^{(0)}(h), \beta^{(1)}(h) \right)^\top \\ &= \left(\frac{P^-(h)_{01}}{(P^-(h))_{11} + R}, \frac{P^-(h)_{11}}{(P^-(h))_{11} + R} \right)^\top \\ &\stackrel{\text{eq. (89)}}{=} \left(\frac{1}{20}, 1 \right)^\top, \end{aligned} \tag{91}$$

$$\begin{aligned} r(h) &= y(h) - m^{-(1)}(h) \\ &\stackrel{\text{eqs. (88), (90)}}{=} -6859/16000 + 1/2 \\ &= 1141/16000, \end{aligned} \tag{92}$$

$$\begin{aligned} \mathbf{m}(h) &\stackrel{\text{eq. (9)}}{=} \begin{pmatrix} m^{-(0)}(h) + \beta^{(0)}(h)r(h) \\ m^{-(1)}(h) + \beta^{(1)}(h)r(h) \end{pmatrix} \\ &\stackrel{\text{eqs. (88), (91), (92)}}{=} \begin{pmatrix} 305141/320000 \\ -6859/16000 \end{pmatrix}, \end{aligned} \tag{93}$$




Fig. 5 Work-precision diagram plotting the number of function evaluations (# Evals of f) against the final state misalignment $\delta^{(1)}(T)$ on the Riccati Eq. (85); cf. Fig. 2

which concludes the step from 0 to h . The next step $h \rightarrow 2h$ starts with computing $m^{-(i)}(2h)$ by a $(q - i)$ th-order Taylor expansion of the i th state $m^{(i)}(h)$, for all $i \in [q + 1]$. Note that, now, there is a nonzero *state misalignment* (recall Eq. (25)):

$$\delta^{(1)}(h) \stackrel{eq.(25)}{=} \left| m^{(1)}(h) - f\left(m^{(0)}(h)\right) \right| \tag{94}$$

$$= \left| -\frac{6859}{16000} - \frac{1}{2} \left(\frac{305141}{320000} \right)^3 \right| \tag{95}$$

$$\approx 0.00485 > 0 \tag{96}$$

which confirms the exposition on the possibility of $\delta^{(i)} > 0$ from Sect. 4. Note that δ tends to increase with R ; e.g., if $R = 1.0$ in the above example, then $\delta^{(1)}(h) \approx 0.03324$.

D Experiment: global convergence of state misalignments δ

Figure 5 depicts the global convergence of the state misalignment $\delta^{(1)}(T)$ in the above example Eq. (85), as detailed in Appendix C, for $q \in \{1, 2, 3\}$. The plotting is analogous to Fig. 2. The resulting convergence rates of h^{q+1} confirm Lemma 11 and suggest that it may also be generalizable to $q \in \{2, 3, \dots\}$.

E Proof of Eq. (23)

We prove the stronger statement

$$\Phi_t^{(i+1)}(a) = f^{(i)}\left(\Phi_t^{(0)}(a)\right), \tag{97}$$

from which Eq. (23) follows by inserting $t = 0$ and $\Phi_0^{(0)}(a) = a$. Hence, it remains to show Eq. (97).

Proof (of Eq. (97)) By induction over $i \in \{0, \dots, q\}$. The base case ($i = 0$) is obtained using the fundamental theorem of calculus and $f^{(1)} = f: \Phi_t^{(1)}(a) = f\left(\Phi_t^{(0)}(a)\right) = f^{(1)}\left(\Phi_t^{(0)}(a)\right)$. For the inductive step $(i - 1) \rightarrow i$, we conclude (using the inductive hypothesis (IH), the chain rule (CR), the base case (BC) and $f^{(i)} = \nabla_x f^{(i-1)} \cdot f$) that

$$\begin{aligned} \Phi_t^{(i+1)}(a) &= \frac{d}{dt} \Phi_t^{(i)}(a) \\ &\stackrel{(IH)}{=} \frac{d}{dt} f^{(i-1)}\left(\Phi_t^{(0)}(a)\right) \\ &\stackrel{(CR)}{=} \nabla_x f^{(i-1)}\left(\Phi_t^{(0)}(a)\right) \frac{d}{dt} \Phi_t^{(0)}(a) \\ &= \nabla_x f^{(i-1)}\left(\Phi_t^{(0)}(a)\right) \cdot f\left(\Phi_t^{(0)}(a)\right) \\ &= \left[\nabla_x f^{(i-1)} \cdot f \right]\left(\Phi_t^{(0)}(a)\right) \\ &\stackrel{(BC)}{=} f^{(i)}\left(\Phi_t^{(0)}(a)\right). \end{aligned} \tag{98}$$

□

F Proof of Lemma 5

Proof Again, w.l.o.g. $d = 1$. Recall that, by Eq. (13), r is implied by the values of $m^{-(0)}$ and $m^{-(1)}$. By insertion of

$$\begin{aligned} m^{-(i)}((n + 1)h) &= \sum_{k=i}^q \frac{h^{k-i}}{(k - i)!} m^{(k)}(nh) + K\theta \left| m^{(q)}(nh) \right| h^{q+1-i} \end{aligned} \tag{99}$$

(due to Eqs. (8) and (14)) into the definition Eq. (13) of $r((n + 1)h)$, we obtain the following equality which we then bound by repeated application of the triangle inequality:

$$\begin{aligned} &|r((n + 1)h)| \\ &= \left| f\left(\sum_{k=0}^q \frac{h^k}{k!} m^{(k)}(nh) + K\theta \left| m^{(q)}(nh) \right| h^{q+1}\right) \right. \\ &\quad \left. - \left(\sum_{k=1}^q \frac{h^{k-1}}{(k - 1)!} m^{(k)}(nh) + K\theta \left| m^{(q)}(nh) \right| h^q\right) \right| \\ &\leq \left| f\left(\sum_{k=0}^q \frac{h^k}{k!} m^{(k)}(nh) + K\theta \left| m^{(q)}(nh) \right| h^{q+1}\right) \right. \\ &\quad \left. - \left(\sum_{k=1}^q \frac{h^{k-1}}{(k - 1)!} m^{(k)}(nh)\right) \right| + K\theta \left| m^{(q)}(nh) \right| h^q \\ &\stackrel{eq.(25)}{\leq} I_1(h) + I_2(h) + I_3(h) \\ &\quad + \sum_{k=1}^q \frac{h^{k-1}}{(k - 1)!} \delta^{(k)}(nh) + K\theta \left| m^{(q)}(nh) \right| h^q, \end{aligned} \tag{100}$$

where I_1, I_2 , and I_3 are defined and bounded as follows, using Assumption 1 and Lemma 1:

$$\begin{aligned}
 I_1(h) &:= \left| f \left(\sum_{k=0}^q \frac{h^k}{k!} m^{(k)}(nh) + K\theta \left| m^{(q)}(nh) \right| h^{q+1} \right) \right. \\
 &\quad \left. - f \left(\sum_{k=0}^q \frac{h^k}{k!} \Phi_0^{(k)} \left(m^{(0)}(nh) \right) \right) \right| \\
 &\leq L \sum_{k=0}^q \frac{h^k}{k!} \delta^{(k)}(nh) + LK\theta \left| m^{(q)}(nh) \right| h^{q+1}, \tag{101}
 \end{aligned}$$

$$\begin{aligned}
 I_2(h) &:= \left| f \left(\sum_{k=0}^q \frac{h^k}{k!} \Phi_0^{(k)} \left(m^{(0)}(nh) \right) \right) - f \left(\Phi_h^{(0)} \left(m^{(0)}(nh) \right) \right) \right| \\
 &\leq L \left| \sum_{k=0}^q \frac{h^k}{k!} \Phi_0^{(k)} \left(m^{(0)}(nh) \right) - \Phi_h^{(0)} \left(m^{(0)}(nh) \right) \right| \\
 &\stackrel{\text{eq. (16)}}{\leq} Kh^{q+1}, \tag{102}
 \end{aligned}$$

and

$$\begin{aligned}
 I_3(h) &:= \left| \Phi_h^{(1)} \left(m^{(0)}(nh) \right) - \sum_{k=1}^q \frac{h^{k-1}}{(k-1)!} \Phi_0^{(k)} \left(m^{(0)}(nh) \right) \right| \\
 &\stackrel{\text{eq. (16)}}{\leq} Kh^q. \tag{103}
 \end{aligned}$$

Inserting Eqs. (101), (102), and (103) into Eq. (100) (and recalling $\delta^{(0)} = 0$) yields Eq. (31). \square

G Proof of Lemma 9

Proof Let $\tilde{u}_0 = u^*$ and $\tilde{u}_n = T_n(\tilde{u}_{n-1})$, for $n \in \mathbb{N}$. Then,

$$d(u^*, x_n) \leq \underbrace{d(u^*, u_n)}_{\rightarrow 0} + \underbrace{d(u_n, \tilde{u}_n)}_{=: a_n} + \underbrace{d(\tilde{u}_n, x_n)}_{\rightarrow 0}, \tag{104}$$

where the last summand goes to zero by

$$\begin{aligned}
 d(\tilde{u}_n, x_n) &= d \left((T_n \circ \dots \circ T_1)(u^*), (T_n \circ \dots \circ T_1)(x_0) \right) \\
 &\leq \bar{L}^n d(u^*, x_0) \rightarrow 0, \quad \text{as } n \rightarrow \infty.
 \end{aligned}$$

Hence, it remains to show that $\lim_{n \rightarrow \infty} a_n = 0$. The \bar{L} -Lipschitz continuity of T_n and the triangle inequality yield that

$$\begin{aligned}
 a_n &= d(T_n(u_n), T_n(\tilde{u}_{n-1})) \\
 &\leq \bar{L} [d(u_n, u_{n-1}) + d(u_{n-1}, \tilde{u}_{n-1})] \\
 &= \bar{L} a_{n-1} + b_{n-1}, \tag{105}
 \end{aligned}$$

where $b_n := \bar{L} d(u_{n+1}, u_n) \rightarrow 0$. Now, for all $m \in \mathbb{N}$, let $a_0^{(m)} := a_0$ and $a_n^{(m)} := \bar{L} a_{n-1}^{(m)} + b_m$. By BFT, $\lim_{n \rightarrow \infty} a_n^{(m)}$

$= b_m / (1 - \bar{L})$. Since, for all $m \in \mathbb{N}, a_n \leq a_n^{(m)}$ for sufficiently large n , it follows that

$$0 \leq \limsup_{n \rightarrow \infty} a_n \leq \lim_{n \rightarrow \infty} a_n^{(m)} = \frac{b_m}{1 - \bar{L}}, \quad \forall m \in \mathbb{N}. \tag{106}$$

Since the convergent sequence u_n is in particular a Cauchy sequence, $\lim_{m \rightarrow \infty} b_m = 0$ and, hence, $0 \leq \lim_{n \rightarrow \infty} a_n = \limsup_{n \rightarrow \infty} a_n \leq 0$. Hence, $\lim_{n \rightarrow \infty} a_n = 0$. \square

H Proof of Proposition 10

Proof Again, w.l.o.g. $d = 1$. We prove the claims in the following order: Eqs. (39), (45), (40), (46), (41), (43), (44), (42), (49), (48), (47). The sharpness of these bounds is shown, directly after they are proved. As a start, for Eq. (39), we show that $P_{11}^{-, \infty}$ is indeed the unique fixed point of the recursion for $\{P_{11}^-(nh)\}_n$ by checking that, if $P_{11}^-(nh) = \frac{1}{2} (\sigma^2 h + \sqrt{4\sigma^2 R h + \sigma^4 h^2})$, then also $P_{11}^-((n+1)h) = \frac{1}{2} (\sigma^2 h + \sqrt{4\sigma^2 R h + \sigma^4 h^2})$:

$$\begin{aligned}
 P_{11}^-((nh)) &\stackrel{\text{eq. (15)}}{=} P_{11}^-(nh) \left(1 - \frac{P_{11}^-(nh)}{P_{11}^-(nh) + R} \right) \\
 &= \frac{(\sigma^2 h + \sqrt{4\sigma^2 R h + \sigma^4 h^2}) R}{\sigma^2 h + \sqrt{4\sigma^2 R h + \sigma^4 h^2} + 2R}, \quad \text{and} \tag{107}
 \end{aligned}$$

$$\begin{aligned}
 P_{11}^-((n+1)h) &= P_{11}^-(nh) + \sigma^2 h \\
 &\stackrel{\text{eq. (107)}}{=} \frac{1}{2} (\sigma^2 h + \sqrt{4\sigma^2 R h + \sigma^4 h^2}) \\
 &= P_{11}^-(nh). \tag{108}
 \end{aligned}$$

After combining Eq. (107) and Eq. (108), the recursion for P_{11}^- is given by

$$\begin{aligned}
 P_{11}^-((n+1)h) &= \underbrace{\left(\frac{R}{P_{11}^-(nh) + R} \right)}_{=: \alpha(nh)} P_{11}^-(nh) + \sigma^2 h \tag{109} \\
 &=: \tilde{T} (P_{11}^-(nh)). \tag{110}
 \end{aligned}$$

Since R and $P_{11}^-(nh)$ are positive variances, we know that $\inf_{n \in [T/h+1]} P_{11}^-(nh) \geq \sigma^2 h$, and hence $\max_{n \in [T/h+1]} \alpha(nh) \leq R / (\sigma^2 h + R) < 1$. Hence, \tilde{T} is a contraction. By BFT, $P_{11}^{-, \infty}$ is the unique (attractive) fixed point of \tilde{T} , and the sequence $\{|P_{11}^-(nh) - P_{11}^{-, \infty}\}_n$ is strictly decreasing. Since, by Eqs. (15), (6) with $\theta = 0$ and Assumption 2,

$$P_{11}^-(h) = P_{11}(0) + \sigma^2 h \leq Kh, \tag{111}$$

we can, using the reverse triangle inequality and the (by BFT) strictly decreasing sequence $\{|P_{11}^-(nh) - P_{11}^{-,\infty}\}_n$, derive Eq. (45):

$$|P_{11}^-(nh)| \leq \underbrace{|P_{11}^-(nh) - P_{11}^{-,\infty}|}_{\leq |P_{11}^-(h) - P_{11}^{-,\infty}|} + |P_{11}^{-,\infty}| \tag{112}$$

$$\leq \underbrace{P_{11}^-(h)}_{\leq Kh} + \underbrace{2P_{11}^{-,\infty}}_{\leq Kh^{1 \wedge \frac{p+1}{2}}, \text{ by eq. (39)}} \tag{113}$$

$$\leq Kh^{1 \wedge \frac{p+1}{2}}, \tag{114}$$

which is sharp because it is estimated against the maximum of the initial P_{11}^- and the steady state that can both be attained. Recall that, by Eq. (107), $P_{11}(nh)$ depends continuously on $P_{11}^-(nh)$, and, hence, inserting Eq. (39) into Eq. (107) yields Eq. (40)—the necessary computation was already performed in Eq. (107). Since $P_{11}(nh)$ monotonically increases in $P_{11}^-(nh)$ (because the derivative of $P_{11}(nh)$ with respect to $P_{11}^-(nh)$ is non-negative for all $P_{11}^-(nh)$ due to $R \geq 0$; see Eq. (107)), we obtain Eq. (46):

$$P_{11}(nh) \stackrel{\text{eq. (107)}}{\leq} \frac{(\max_n P_{11}^-(nh)) R}{\max_n P_{11}^-(nh) + R} \tag{115}$$

$$\stackrel{R \sim h^p}{\leq} \frac{Kh^{1 \wedge \frac{p+1}{2}} Kh^p}{Kh^{1 \wedge \frac{p+1}{2}} + Kh^p} \tag{116}$$

$$\leq \frac{Kh^{(p+1) \wedge \frac{3p+1}{2}}}{Kh^{1 \wedge p}} \tag{117}$$

$$\leq \begin{cases} Kh^{\frac{p+1}{2}}, & \text{if } p \leq 1, \\ Kh^p, & \text{if } p \geq 1, \end{cases} \tag{118}$$

$$\leq Kh^{p \vee \frac{p+1}{2}}, \tag{119}$$

which is sharp because the steady state Eq. (45) has these rates. For Eq. (41), we again first construct the following recursion (from Eq. (10), Eqs. (15) and (6) with $\theta = 0$)

$$P_{01}^-((n+1)h) = \underbrace{\frac{R}{P_{11}^-(nh) + R}}_{=\alpha(nh)} P_{01}^-(nh) + \underbrace{\left(P_{11}(nh) + \frac{\sigma^2 h}{2}\right)}_{=:g(nh)} h \tag{120}$$

$$= T_n(P_{01}^-(nh)), \tag{121}$$

where the $\alpha(nh)$ -Lipschitz continuous contractions T_n satisfy the prerequisites of Lemma 9, since $\sup_n \alpha(nh) \leq R/(\sigma^2 h + R) < 1$ (due to $\inf_n P_{11}^-(nh) \geq \sigma^2 h$) and the sequence of fixed points $(1 - \alpha(nh))^{-1}g(nh)$ of T_n (defined

by BFT) converges. Both $\alpha(nh)$ and $g(nh)$ depend continuously on $P_{11}^-(nh)$. Hence, insertion of the limits Eqs. (39) and (40) yield

$$\lim_{n \rightarrow \infty} (1 - \alpha(nh))^{-1} = \frac{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 2R}{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2}}, \tag{122}$$

and

$$\begin{aligned} \lim_{n \rightarrow \infty} g(nh) &= \frac{(\sigma^4 h^2 + (2R + \sigma^2 h)\sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 4R\sigma^2 h)}{2(\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 2R)} h. \end{aligned} \tag{123}$$

Now, application of Lemma 9 implies convergence of the recursion Eq. (121) to the product of these two limits Eqs. (122) and (123), i.e., Eq. (41):

$$\begin{aligned} \lim_{n \rightarrow \infty} P_{01}^-(nh) &= \lim_{n \rightarrow \infty} (1 - \alpha(nh))^{-1} \times \lim_{n \rightarrow \infty} g(nh) \\ &= \frac{\sigma^4 h^2 + (2R + \sigma^2 h)\sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 4R\sigma^2 h}{2(\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 2R)} h. \end{aligned}$$

For Eqs. (43) and (44), we can simply insert Eqs. (39) and (41) for $P_{01}^-(nh)$ and $P_{11}^-(nh)$, respectively, into their definition Eq. (11):

$$\beta^{\infty,(0)} \stackrel{\text{eq. (11)}}{=} \frac{P_{01}^{-,\infty}}{P_{11}^{-,\infty} + R} \tag{124}$$

$$\stackrel{\text{eqs. (39) and (41)}}{=} \frac{\sqrt{4R\sigma^2 h + \sigma^4 h^2}}{\sigma^2 h + \sqrt{4R\sigma^2 h + \sigma^4 h^2}} h, \tag{125}$$

and

$$\beta^{\infty,(1)} \stackrel{\text{eqs. (11) and (39)}}{=} \frac{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2}}{\sigma^2 h + \sqrt{4\sigma^2 Rh + \sigma^4 h^2} + 2R}. \tag{126}$$

These steady states Eqs. (43) and (44) are again unique and attractive because $\beta^{(0)}(nh)$ and $\beta^{(1)}(nh)$ depend continuously on $P_{11}^-(nh)$ and $P_{01}^-(nh)$. Next, recall that

$$P_{01}(nh) \stackrel{\text{eq. (15)}}{=} \left(1 - \frac{P_{11}^-(nh)}{P_{11}^-(nh) + R}\right) P_{01}^-(nh) \tag{127}$$

$$= R \frac{P_{01}^-(nh)}{P_{11}^-(nh) + R} \stackrel{\text{eq. (11)}}{=} R\beta^{(0)}(nh), \tag{128}$$

which, since $P_{01}(nh)$ depends continuously on $\beta^{(0)}(nh)$, implies the unique (attractive) fixed point $P_{01}^\infty(nh) =$

$R\beta^{\infty,(0)}$, which yields Eq. (42). Now, exploiting Eq. (11) and $\inf_n P_{11}^-(nh) \geq \sigma^2 h$ yields Eq. (49):

$$|1 - \beta^{(1)}(nh)| = \frac{R}{P_{11}^-(nh) + R} \tag{129}$$

$$\leq \frac{R}{\sigma^2 h + R} \tag{130}$$

$$\stackrel{R \sim h^p}{=} \frac{Kh^p}{Kh + Kh^p} \tag{131}$$

$$\leq Kh^{(p-1)\vee 0}, \tag{132}$$

which is sharp because $\inf_n P_{11}^-(nh) \geq Kh$ is sharp (due to Eqs. (10) and (6)). And since, for $\beta^{(0)}$, maximizing over both $P_{01}^-(nh)$ and $P_{11}^-(nh)$ at the same time does not yield a sharp bound (while above in Eqs. (129) and (119) the maximization over just one quantity does), we prove Eq. (48) by inductively showing that

$$|\beta^{(0)}(nh)| \leq \hat{\beta}h, \quad \forall n \in \mathbb{N}, \tag{133}$$

with $\hat{\beta} := \left(\frac{2K_0}{\sigma^2} + \frac{1}{2}\right) \vee 1 > 0,$ (134)

where $K_0 > 0$ is the constant from Assumption 2. The constant $\hat{\beta}$ is independent of n and a possible choice for K in Eq. (48).

The base case ($n = 1$) follows from

$$|\beta^{(0)}(h)| = \frac{|P_{01}^-(h)|}{P_{11}^-(h) + R} \tag{135}$$

$$\stackrel{\text{eq. (10)}}{\leq} \frac{|P_{01}(0)| + hP_{11}(0) + \frac{\sigma^2}{2}h^2}{\sigma^2 h} \tag{136}$$

$$\stackrel{\text{Ass. 2}}{\leq} \left(\frac{2K_0}{\sigma^2} + \frac{1}{2}\right)h \tag{137}$$

$$\leq \hat{\beta}h. \tag{138}$$

In the following inductive step ($n - 1 \rightarrow n$) we, to avoid notational clutter, simply denote $P^-(n - 1)h_{ij}$ by P_{ij}^- which leaves us—by Eq. (11), (10) and (15)—with the following term to bound:

$$|\beta^{(0)}(nh)| = \frac{|P_{01}^-(nh)|}{P_{11}^-(nh) + R} \tag{139}$$

$$\leq \frac{|P_{01}^-| \alpha(nh) + hP_{11}^- \alpha(nh) + \frac{\sigma^2}{2}h^2}{P_{11}^- \alpha(nh) + \sigma^2 h + R}, \tag{140}$$

with $\alpha(nh) = \left(1 - \frac{P_{11}^-}{P_{11}^- + R}\right) = \frac{R}{P_{11}^- + R}$. Application of the inductive hypothesis (i.e., $P_{01}^- \leq \hat{\beta}(P_{11}^- + R)$) yields, after some rearrangements, that

$$\begin{aligned} |\beta^{(0)}(nh)| &\leq \frac{\hat{\beta}(P_{11}^- + R)h\alpha(nh) + hP_{11}^- \alpha(nh) + \frac{\sigma^2}{2}h^2}{P_{11}^- \alpha(nh) + \sigma^2 h + R} \\ &= \frac{2\hat{\beta}P_{11}^- R + \sigma^2 h(P_{11}^- + R) + 2P_{11}^- R + 2\hat{\beta}R^2}{2(P_{11}^- R + \sigma^2 h(P_{11}^- + R) + P_{11}^- R + R^2)}h \\ &= \frac{2(\hat{\beta} + 1)\Lambda_1 + \Lambda_2 + 2\hat{\beta}\Lambda_3}{4\Lambda_1 + 2\Lambda_2 + 2\Lambda_3}h, \end{aligned} \tag{141}$$

with $\Lambda_1 := 2P_{11}^- R$, $\Lambda_2 := \sigma^2 h(P_{11}^- + R)$, and $\Lambda_3 := R^2$.

Now, application of $\hat{\beta} \geq 1$ yields $|\beta^{(0)}(nh)| \leq \hat{\beta}h$, which completes the inductive proof of Eq. (133). This implies Eq. (48), which is sharp because it is the order of $\beta^{(0)}$ in the steady state Eq. (43), for all $p \in [0, \infty]$. Now, insertion of Eq. (48) into Eq. (127) immediately yields Eq. (47), which—by Eq. (127)—inherits the sharpness of Eq. (48). \square

I Proof of Lemma 11

Proof For all $n \in [T/h + 1]$, we can estimate

$$\delta^{(1)}(nh) = \|m^{(1)}(nh) - f(m^{(0)}(nh))\| \tag{142}$$

$$= \|\Psi_h^{(1)}(m((n - 1)h) - f(m^{(0)}(nh)))\| \tag{143}$$

$$\begin{aligned} &\leq \underbrace{\|\Psi_h^{(1)}(m((n - 1)h) - f(m^{-,(0)}(nh)))\|}_{=: J_1(h)} \\ &\quad + \underbrace{\|f(m^{-,(0)}(nh)) - f(m^{(0)}(nh))\|}_{=: J_2(h)}, \end{aligned} \tag{144}$$

bound J_1 , using the definition Eq. (14) of $\Psi_h^{(1)}(m((n - 1)h))$ as well as the definition Eq. (13) of $r(nh)$, by

$$J_1(h) = \|m^{-,(1)}(nh) - f(m^{-,(0)}(nh))\| \tag{145}$$

$$\begin{aligned} &+ \beta^{(1)}(nh) \left[f(m^{-,(0)}(nh)) - m^{-,(1)}(nh) \right] \\ &\leq \|1 - \beta^{(1)}(nh)\| \|r(nh)\| \end{aligned} \tag{146}$$

$$\stackrel{\text{eq. (49)}}{\leq} Kh^{(p-1)\vee 0} \|r(nh)\| \tag{147}$$

and bound J_2 , by exploiting L -Lipschitz continuity of f , inserting the definition Eq. (14) of $\Psi_h^{(0)}(m((n - 1)h))$ and applying Eq. (48) to $\|\beta^{(0)}(nh)\|$,

$$J_2(h) \leq L \|m^{(0)}(nh) - m^{-,(0)}(nh)\| \tag{148}$$

$$\leq L \|\beta^{(0)}(nh)\| \|r(nh)\| \tag{149}$$

$$\stackrel{\text{eq. (48)}}{\leq} Kh \|r(nh)\|. \tag{150}$$

Altogether, after inserting these bounds into Eq. (144),

$$\delta^{(1)}(nh) \leq (Kh^{(p-1)\vee 0} + Kh) \|r(nh)\| \tag{151}$$

$$\leq Kh^{((p-1)\vee 0)\wedge 1} \|r(nh)\| \tag{152}$$

$$\stackrel{\text{eq. (31)}}{\leq} Kh^{(p\vee 1)\wedge 2} \tag{153}$$

$$+ (Kh^{((p-1)\vee 0)\wedge 1} + Kh^{(p\vee 1)\wedge 2}) \delta^{(1)}((n-1)h) \\ =: \bar{T}(\delta^{(1)}((n-1)h)). \tag{154}$$

As $p \geq 1$ (by Assumption 4), BFT is applicable for all sufficiently small $h > 0$ such that $Kh^{((p-1)\vee 0)\wedge 1} + Kh^{(p\vee 1)\wedge 2} < 1$ and so \bar{T} is a contraction with a unique fixed point δ^∞ of order

$$\delta^\infty \leq \frac{Kh^{(p\vee 1)\wedge 2}}{1 - (Kh^{((p-1)\vee 0)\wedge 1} + Kh^{(p\vee 1)\wedge 2})} \tag{155}$$

$$\leq Kh^{(p\vee 1)\wedge 2}. \tag{156}$$

We proceed with showing by induction that, for all $n \in [T/h]$,

$$\delta^{(1)}(nh) \leq \delta^{(1)}(0) \vee 2\delta^\infty. \tag{157}$$

The base case $n = 0$ is trivial. For the inductive step, we distinguish two cases. If $\delta^{(1)}((n-1)h) \leq \delta^\infty$, then $\bar{T}(\delta^{(1)}((n-1)h)) < 2\delta^\infty$, since

$$\bar{T}(\delta^{(1)}((n-1)h)) - \delta^\infty \leq \left| \delta^\infty - \bar{T}(\delta^{(1)}((n-1)h)) \right| \tag{158}$$

$$< \delta^\infty - \underbrace{\delta^{(1)}((n-1)h)}_{\geq 0} \tag{159}$$

$$\leq \delta^\infty. \tag{160}$$

In this case,

$$\delta^{(1)}(nh) \stackrel{\text{eq. (154)}}{\leq} \bar{T}(\delta^{(1)}((n-1)h)) \tag{161}$$

$$< 2\delta^\infty \tag{162}$$

$$\leq \delta^{(1)}(0) \vee 2\delta^\infty, \tag{163}$$

where the last inequality follows from the inductive hypothesis. In the other case, namely $\delta^{(1)}((n-1)h) > \delta^\infty$, it follows that

$$\delta^{(1)}(nh) - \delta^\infty \stackrel{\text{eq. (154)}}{\leq} \bar{T}(\delta^{(1)}((n-1)h)) - \delta^\infty \tag{164}$$

$$\leq \left| \bar{T}(\delta^{(1)}((n-1)h)) - \delta^\infty \right| \tag{165}$$

$$\leq \left| \delta^{(1)}((n-1)h) - \delta^\infty \right| \tag{166}$$

$$= \delta^{(1)}((n-1)h) - \delta^\infty, \tag{167}$$

which, after adding δ^∞ and applying the inductive hypothesis, completes the inductive step. Hence, Eq. (157) holds. Since this bound is uniform in n , inserting the orders of $\delta^{(1)}(0)$ from Lemma 7 and of δ^∞ from Eq. (155) yields Eq. (50). \square

J Proof of Theorem 15

Proof Again, w.l.o.g. $d = 1$. We first show that the bounds Eqs. (67) and (68) hold and then argue that they are sharp. The recursion for $P_{00}^-(nh)$ is given by

$$P_{00}^-((n+1)h) \stackrel{\text{eqs. (10),(6)}}{=} P_{00}(nh) + 2hP_{01}(nh) \\ + h^2P_{11}(nh) + \frac{\sigma^2}{3}h^3 \tag{168}$$

$$= P_{00}^-(nh) - \beta^{(0)}(nh)P_{01}^-(nh) + \frac{\sigma^2}{3}h^3, \\ + 2hR\beta^{(0)}(nh) + h^2R\beta^{(1)}(nh) \tag{169}$$

where we used $P_{00}(nh) = P_{00}^-(nh) - \beta^{(0)}P_{01}^-(nh)$ and $P_{11}(nh) = R\beta^{(1)}(nh)$ (both due to Eq. (15) and Eq. (11)), as well as $P_{01}(nh) = R\beta^{(0)}(nh)$ (see Eq. (127)), for the last equality in Eq. (169). By $P_{01}^-(nh) \leq P_{01}(nh)$ and $|\beta^{(1)}| \leq 1$ (due to Eq. (11)), application of the triangle inequality to Eq. (169) yields

$$P_{00}^-((n+1)h) \leq P_{00}^-(nh) + \left| \beta^{(0)}(nh) \right| |P_{01}(nh)| \\ + 2hR \left| \beta^{(0)}(nh) \right| + h^2R + \frac{\sigma^2}{3}h^3, \tag{170}$$

which, by Eqs. (47) and (48), implies

$$P_{00}^-((n+1)h) \leq P_{00}^-(nh) + Kh^{(p+2)\wedge 3}. \tag{171}$$

This, by $N = T/h$, implies Eq. (67). Since $P_{00}(nh) \leq P_{00}^-(nh)$, this bound is also valid for P_{00} , i.e., Eq. (68) holds. The bound Eq. (67) is sharp, since, e.g., when the covariance matrices are in the steady state, the covariance matrix keeps growing by a rate of $Kh^{(p+2)\wedge 3}$ for all sufficiently small $h > 0$, since the only negative summand in Eq. (169) is given by

$$\beta^{\infty,(0)} P_{01}^{\infty} = S_1(h) \times S_2(h) \times S_3(h) \in \Theta(h^{5 \wedge \frac{3p+7}{2}}), \tag{172}$$

where the factors have, due to $R \equiv Kh^p$, the following orders:

$$S_1(h) = \frac{1}{2}h^2 \in \Theta(h^2), \tag{173}$$

$$S_2(h) = \sqrt{(\sigma^2 h)^2 + 4(\sigma^2 h)R}, \in \Theta(h^{1 \wedge \frac{p+1}{2}}), \tag{174}$$

$$S_3(h) = ((\sigma^2 h) + 2R)\sqrt{(\sigma^2 h)^2 + 4(\sigma^2 h)R} + (\sigma^2 h)^2 + 4(\sigma^2 h)R \in \Theta(h^{2 \wedge (p+1)}). \tag{175}$$

The orders in Eqs. (173) to (175) imply the order in Eq. (172). Hence, the sole negative summand $-\beta^{\infty,(0)} P_{01}^{\infty}$ of Eq. (169) is in $\Theta(h^{5 \wedge \frac{3p+7}{2}})$ and thereby of higher order than the remaining positive summands of Eq. (169):

$$\underbrace{2hR}_{\in \Theta(h^{p+1})} \underbrace{\beta^{\infty,(0)}(nh)}_{\in \Theta(h)} \in \Theta(h^{p+2}), \tag{176}$$

$$\underbrace{h^2 R}_{\in \Theta(h^{p+2})} \underbrace{\beta^{\infty,(1)}(nh)}_{\in \Theta(1), \text{ by eq. (44)}} \in \Theta(h^{p+2}), \tag{177}$$

$$\frac{\sigma^2}{3} h^3 \in \Theta(h^3). \tag{178}$$

Hence, for all sufficiently small $h > 0$, it still holds in the steady state that $P_{00}^-(n+1)h - P_{00}^-(nh) \geq Kh^{(p+2) \wedge 3}$, and therefore Eq. (67) is sharp. The sharpness of Eq. (67) is inherited by Eq. (68) since, in the steady state, by Eqs. (15) and (11), $P_{00}(nh) = P_{00}^-(nh) - \beta^{(0),\infty} P_{01}^{-,\infty}$ and the subtracted quantity $\beta^{(0),\infty} P_{01}^{-,\infty}$ is—as shown above—only of order $\Theta(h^{5 \wedge \frac{3p+7}{2}})$. \square

References

Abdulle, A., Garegnani, G.: Random time step probabilistic methods for uncertainty quantification in chaotic and geometric numerical integration. *Stat. Comput.* **30**(4), 907–932 (2020)

Anderson, B., Moore, J.: *Optimal Filtering*. Prentice-Hall, Englewood Cliffs (1979)

Callier, F.M., Desoer, C.A.: *Linear System Theory*. Springer, Berlin (1991)

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

Clark, D.S.: Short proof of a discrete Gronwall inequality. *Discret. Appl. Math.* **16**(3), 279–281 (1987)

Cockayne, J., Oates, C.J., Sullivan, T.J., Girolami, M.: Bayesian probabilistic numerical methods. *SIAM Rev.* **61**(4), 756–789 (2019)

Conrad, P.R., Girolami, M., Särkkä, S., Stuart, A., Zygalakis, K.: Statistical analysis of differential equations: introducing probability measures on numerical solutions. *Stat. Comput.* **27**(4), 1065–1082 (2017)

Davis, H.T.: *Introduction to Nonlinear Differential and Integral Equations*. Dover Publications, New York (1962)

Diaconis, P.: Bayesian numerical analysis. *Stat. Decis. Theory Rel. Top.* **IV**(1), 163–175 (1988)

Gessner, A., Gonzalez, J., Mahsereci, M.: Active multi-information source Bayesian quadrature. In: *Uncertainty in Artificial Intelligence (UAI)* (2019)

Hairer, E., Nørsett, S., Wanner, G.: *Solving Ordinary Differential Equations I—Nonstiff Problems*. Springer, Berlin (1987)

Hennig, P., Osborne, M.A., Girolami, M.: Probabilistic numerics and uncertainty in computations. *Proc. R. Soc. Lond. A* **471**(2179), 20150142 (2015)

Jazwinski, A.: *Stochastic Processes and Filtering Theory*. Academic Press, Cambridge (1970)

Karatzas, I., Shreve, S.: *Brownian Motion and Stochastic Calculus*. Springer, Berlin (1991)

Kersting, H., Hennig, P.: Active uncertainty calibration in Bayesian ODE solvers. In: *Uncertainty in Artificial Intelligence (UAI)* (2016)

Kersting, H., Mahsereci, M.: A Fourier state space model for Bayesian ODE filters. In: *Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models, ICML* (2020)

Kersting, H., Krämer, N., Schiegg, M., Daniel, C., Tiemann, M., Hennig, P.: Differentiable likelihoods for fast inversion of ‘likelihood-free’ dynamical systems. In: *International Conference on Machine Learning (ICML)* (2020)

Lancaster, P., Rodman, L.: *Algebraic Riccati Equations*. Oxford Science Publications, Oxford (1995)

Law, K., Stuart, A., Zygalakis, K.: *Data Assimilation: A Mathematical Introduction, Texts in Applied Mathematics*, vol. 62. Springer, Cham (2015)

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

Loscalzo, F.R., Talbot, T.D.: Spline function approximations for solutions of ordinary differential equations. *SIAM J. Numer. Anal.* **4**, 433–445 (1967)

Magnani, E., Kersting, H., Schober, M., Hennig, P.: Bayesian Filtering for ODEs with Bounded Derivatives. [arXiv:1709.08471](https://arxiv.org/abs/1709.08471) [csNA] (2017)

Maybeck, P.S.: *Stochastic Models, Estimation, and Control*. Academic Press, Cambridge (1979)

Nordsieck, A.: On numerical integration of ordinary differential equations. *Math. Comput.* **16**, 22–49 (1962)

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

O’Hagan, A.: Bayes–Hermite quadrature. *J. Stat. Plann Inference* **29**(3), 245–260 (1991)

O’Hagan, A.: Some Bayesian Numerical Analysis. *Bayesian statistics*, 4 (Peñíscola, 1991), pp. 345–363. Oxford Univ. Press, New York (1992)

Øksendal, B.: *Stochastic Differential Equations: An Introduction with Applications*, 5th edn. Springer, Berlin (2003)

Poincaré, H.: *Calcul des probabilités*. Gauthier-Villars, Paris (1896)

Rasmussen, C., Williams, C.: *Gaussian Processes for Machine Learning*. MIT Press, London (2006)

Reich, S., Cotter, C.: *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press, New York (2015)

Ritter, K.: *Average-Case Analysis of Numerical Problems. Lecture Notes in Mathematics*, vol. 1733. Springer-Verlag, Berlin (2000)

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, Cambridge (2013)

Särkkä, S., Solin, A.: *Applied Stochastic Differential Equations*. Cambridge University Press, Cambridge (2019)

- Schober, M., Duvenaud, D., Hennig, P.: Probabilistic ODE solvers with Runge–Kutta means. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2014)
- Schober, M., Särkkä, S., Hennig, P.: A probabilistic model for the numerical solution of initial value problems. *Stat. Comput.* **29**(1), 99–122 (2019)
- Skilling, J.: *Bayesian solutions of ordinary differential equations. Maximum Entropy and Bayesian Methods*, Seattle (1991)
- Solak, E., Murray-Smith, R., Leithead, W.E., Leith, D.J., Rasmussen, C.E.: Derivative observations in Gaussian process models of dynamic systems. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2003)
- Teschl, G.: *Ordinary Differential Equations and Dynamical Systems*. American Mathematical Society, Providence (2012)
- Teymur, O., Zygalakis, K., Calderhead, B.: Probabilistic linear multistep methods. In: Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R., Lee, D.D. (eds.) *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 4314–4321. Curran Associates Inc (2016)
- Teymur, O., Lie, H.C., Sullivan, T.J., Calderhead, B.: Implicit probabilistic integrators for ODEs. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2018)
- Tronarp, F., Kersting, H., Särkkä, S., Hennig, P.: Probabilistic solutions to ordinary differential equations as nonlinear Bayesian filtering: a new perspective. *Stat. Comput.* **29**(6), 1297–1315 (2019)
- Tronarp, F., Kersting, H., Särkkä, S., Hennig, P.: Bayesian ode solvers: The maximum a posteriori estimate. [arXiv:2004.00623](https://arxiv.org/abs/2004.00623) [mathNA] (2020)
- Wang J, Cockayne, J., Oates, C.: On the Bayesian solution of differential equations. In: *Proceedings of the 38th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering* (2018)
- Xiaoyue, X., Briol, F.X., Girolami, M.: Bayesian quadrature for multiple related integrals. In: *International Conference on Machine Learning (ICML)* (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.