

A Role for Symmetry in the Bayesian Solution of Differential Equations

Junyang Wang^{*}, Jon Cockayne[†], and Chris. J. Oates^{*,‡}

Abstract. The interpretation of numerical methods, such as finite difference methods for differential equations, as point estimators suggests that formal uncertainty quantification can also be performed in this context. Competing statistical paradigms can be considered and Bayesian probabilistic numerical methods (PNMs) are obtained when Bayesian statistical principles are deployed. Bayesian PNM have the appealing property of being closed under composition, such that uncertainty due to different sources of discretisation in a numerical method can be jointly modelled and rigorously propagated. Despite recent attention, no exact Bayesian PNM for the numerical solution of ordinary differential equations (ODEs) has been proposed. This raises the fundamental question of whether exact Bayesian methods for (in general nonlinear) ODEs even exist. The purpose of this paper is to provide a positive answer for a limited class of ODE. To this end, we work at a foundational level, where a novel Bayesian PNM is proposed as a proof-of-concept. Our proposal is a synthesis of classical Lie group methods, to exploit underlying symmetries in the gradient field, and non-parametric regression in a transformed solution space for the ODE. The procedure is presented in detail for first and second order ODEs and relies on a certain strong technical condition – existence of a solvable Lie algebra – being satisfied. Numerical illustrations are provided.

Keywords: probabilistic numerics, ordinary differential equations, Lie groups.

1 Introduction

Numerical methods underpin almost all of scientific, engineering and industrial output. In the abstract, a numerical task can be formulated as the approximation of a quantity of interest

$$Q : \mathcal{Y} \rightarrow \mathcal{Q},$$

subject to a finite computational budget. The true underlying state $y^\dagger \in \mathcal{Y}$ is typically high- or infinite-dimensional, so that only limited information

$$A : \mathcal{Y} \rightarrow \mathcal{A} \tag{1}$$

is provided and exact computation of $Q(y^\dagger)$ is prohibited. For example, numerical integration aims to approximate an integral $Q(y^\dagger) = \int y^\dagger(t) dt$ given the values $A(y^\dagger) = \{(x_i, y^\dagger(x_i))\}_{i=1}^n$ of the integrand y^\dagger on a finite number of abscissa $\{x_i\}_{i=1}^n$. Simi-

^{*}School of Mathematics, Statistics and Physics, Newcastle University, UK, j.wang68@newcastle.ac.uk

[†]Department of Statistics, University of Warwick, UK, j.cockayne@warwick.ac.uk

[‡]The Alan Turing Institute, UK, chris.oates@newcastle.ac.uk

larly, a numerical approximation to the solution $Q(y^\dagger) = y^\dagger$ of a differential equation $dy/dx = f(x, y(x))$, $y(x_0) = y_0$, must be based on at most a finite number of evaluations of f , the gradient field. In this viewpoint a numerical method corresponds to a map $b : \mathcal{A} \rightarrow \mathcal{Q}$, as depicted in Figure 1a, where $b(a)$ represents an approximation to the solution of the differential equation based on the information $a \in \mathcal{A}$.

The increasing ambition and complexity of contemporary applications is such that the computational budget can be *extremely* small compared to the precision that is required at the level of the quantity of interest. As such, in many important problems it is not possible to reduce the numerical error to a negligible level. Fields acutely associated with this challenge include climate forecasting (Wedi, 2014), computational cardiology (Chabiniok et al., 2016) and molecular dynamics (Perilla et al., 2015). In the presence of non-negligible numerical error, it is unclear how scientific interpretation of the output of computation can proceed. For example, a *posteriori* analysis of traditional numerical methods can be used to establish hard upper bounds on the numerical error, but these bounds typically depend on an unknown global constant. In the case of ODEs, this may be the maximum value of a norm $\|f\|$ of the gradient field (see e.g. Estep, 1995). If $\|f\|$ were known, it would be possible to provide a hard bound on numerical error. However, in the typical numerical context all that is known is that $\|f\|$ is finite. One could attempt to approximate $\|f\|$ with cubature, but that itself requires a numerical cubature method whose error is required to obey a known bound. In general, therefore, there are no hard error bounds without global information being *a priori* provided (Larkin, 1974). Our aim in this paper is to consider, as an alternative to traditional numerical analysis, an exact Bayesian framework for solution uncertainty quantification in the ordinary differential equation context.

1.1 Probabilistic Numerical Methods

The field of *probabilistic numerics* dates back to Larkin (1972) and a modern perspective is provided in Hennig et al. (2015); Oates and Sullivan (2019). Under the abstract framework just described, numerical methods can be interpreted as point estimators in a statistical context, where the state y^\dagger can be thought of as a latent variable in a statistical model, and the ‘data’ consist of information $A(y^\dagger)$ that does not fully determine the quantity of interest $Q(y^\dagger)$ but is indirectly related to it. Hennig et al. (2015) provide an accessible introduction and survey of the field. In particular, they illustrated how PNM can be used to quantify uncertainty due to discretisation in important scientific problems, such as astronomical imaging.

Let the notation $\Sigma_{\mathcal{Y}}$ denote a σ -algebra on the space \mathcal{Y} and let $\mathcal{P}_{\mathcal{Y}}$ denote the set of probability measures on $(\mathcal{Y}, \Sigma_{\mathcal{Y}})$. A probabilistic numerical method (PNM) is a procedure which takes as input a ‘belief’ distribution $\mu \in \mathcal{P}_{\mathcal{Y}}$, representing epistemic uncertainty with respect to the true (but unknown) value y^\dagger , along with a finite amount of information, $A(y^\dagger) \in \mathcal{A}$. The output is a distribution $B(\mu, A(y^\dagger)) \in \mathcal{P}_{\mathcal{Q}}$ on $(\mathcal{Q}, \Sigma_{\mathcal{Q}})$, representing epistemic uncertainty with respect to the quantity of interest $Q(y^\dagger)$ after the information $A(y^\dagger)$ have been processed. For example, a PNM for an ordinary differential equation (ODE) takes an initial belief distribution defined on the solution space of the differential equation, together with information arising from a finite number of




Figure 1: Diagrams for a numerical method. (a) The traditional viewpoint of a numerical method is equivalent to a map b from a finite-dimensional information space \mathcal{A} to the space of the quantity of interest \mathcal{Q} . (b) The probabilistic viewpoint treats approximation of $Q(y^\dagger)$ in a statistical context, described by a map $B(\mu, \cdot)$ from \mathcal{A} to the space of probability distributions on \mathcal{Q} . The probabilistic numerical method (A, B) is Bayesian if and only if (b) is a commutative diagram.

evaluations of the gradient field, plus the initial condition of the ODE, to produce a distribution over either the solution space of the ODE, or perhaps some derived quantity of interest. In this paper, the measurability of A and Q will be assumed.

Despite computational advances in this emergent field, until recently there had not been an attempt to establish rigorous statistical foundations for PNM. In Cockayne et al. (2019) the authors argued that Bayesian principles can be adopted. In brief, this framework requires that the input belief distribution μ carries the semantics of a Bayesian agent’s prior belief and the output of a PNM agrees with the inferences drawn when the agent is rational. To be more precise recall that, in this paper, information is provided in a deterministic¹ manner through (1) and thus Bayesian inference corresponds to conditioning of μ on the level sets of A . Let $Q_\# : \mathcal{P}_Y \rightarrow \mathcal{P}_Q$ denote the push-forward map associated to Q . i.e. $Q_\#(\mu)(S) = \mu(Q^{-1}(S))$ for all $S \in \Sigma_Q$. Let $\{\mu^a\}_{a \in \mathcal{A}} \subset \mathcal{P}_Y$ denote the disintegration, assumed to exist,² of $\mu \in \mathcal{P}_Y$ along the map A .

Definition 1. A probabilistic numerical method (A, B) with $A : \mathcal{Y} \rightarrow \mathcal{A}$ and $B : \mathcal{P}_Y \times \mathcal{A} \rightarrow \mathcal{P}_Q$ for a quantity of interest $Q : \mathcal{Y} \rightarrow \mathcal{Q}$ is Bayesian if and only if $B(\mu, a) = Q_\#(\mu^a)$ for all $\mu \in \mathcal{P}_Y$ and all $a \in \mathcal{A}$.

This definition is intuitive; the output of the PNM should coincide with the marginal distribution for $Q(y^\dagger)$ according to the disintegration element $\mu^a \in \mathcal{P}_Y$, based on the information $a \in \mathcal{A}$ that was provided. The definition is equivalent to the statement that Figure 1b is a commutative diagram. In Cockayne et al. (2019) the map A was termed an *information operator* and the map B was termed a *belief update operator*; we adhere to these definitions in our work. The Bayesian approach to PNM confers several important benefits:

¹It is of course possible to perform Bayesian inference in the noisy-data context, but for the ODEs considered in this paper we assume that one can obtain noiseless evaluations of the gradient field.

²The reader unfamiliar with the concept of a disintegration can treat μ^a as a technical notion of the ‘conditional distribution of y given $A(y) = a$ ’ when reading this work. The *disintegration theorem*, Thm. 1 of Chang and Pollard (1997), guarantees existence and uniqueness up to a $A_\# \mu$ -null set under the weak requirement that \mathcal{Y} is a metric space, Σ_Y is the Borel σ -algebra, μ is Radon, Σ_A is countable generated and Σ_A contains all singletons $\{a\}$ for $a \in \mathcal{A}$.

- The input μ and output $B(\mu, a)$ belief distributions can be interpreted, respectively, as a *prior* and (marginal) *posterior*.³ As such, they automatically inherit the stronger formal semantics and philosophical foundations that underpin the Bayesian framework and, in this sense, are well-understood (see e.g. Gelman and Shalizi, 2013).
- The definition of Bayesian PNM is operational. Thus, if we are presented with a prior μ and information a then there is a unique Bayesian PNM and it is constructively defined.
- The class of Bayesian PNM is closed under composition, such that uncertainty due to different sources of discretisation can be jointly modelled and rigorously propagated. This point will not be discussed further in this work, but we refer the interested reader to Section 5 of Cockayne et al. (2019).

Nevertheless, the strict definition of Bayesian PNM limits scope to design convenient computational algorithms and indeed several proposed PNM are not Bayesian (see Table 1 in Cockayne et al., 2019). The challenge is two-fold; for a Bayesian PNM, the elicitation of an appropriate prior distribution μ and the exact computation of its disintegration $\{\mu^a\}_{a \in \mathcal{A}}$ must both be addressed. In the next section we argue that – perhaps as a consequence of these constraints – a strictly Bayesian PNM for the numerical solution of an ODE does not yet exist.

1.2 Existing Work for ODEs

The first PNM (of any flavour) for the numerical solution of ODEs, or which we are aware, was due to Skilling (1992). Two decades later, this problem is receiving renewed critical attention as part of the active development of PNM. The aim of this section is to provide a high-level overview of existing work and to argue that existing methods do not adhere to the definition of Bayesian PNM.

Notation The notational convention used in this paper is that the non-italicised y denotes a generic function, whereas the italicised y denotes a scalar value taken by the function y . The notation y^\dagger is reserved for the true solution to an ODE. Throughout, the underlying state space \mathcal{Y} is taken to be a space occupied by the true solution of the ODE, i.e. $y^\dagger \in \mathcal{Y}$.

Skilling (1992)

The first paper on this topic, of which we are aware, was Skilling (1992). This will serve as a prototypical PNM for the numerical solution of an ODE. Originally described as

³Indeed, if the set $\mathcal{Y}^a = \{y \in \mathcal{Y} : A(y) = a\}$ is not measure zero under μ , then μ^a is the conditional distribution defined by restricting μ to the subset \mathcal{Y}^a ; $\mu^a(y) = 1[y \in \mathcal{Y}^a]\mu(y)/\mu(\mathcal{Y}^a)$. The theory of disintegrations generalises the conditional distribution μ^a to cases where \mathcal{Y}^a is a null set.

‘Bayesian’ by the author, we will argue that, at least in the strict sense of Definition 1, it is not a Bayesian PNM. Consider a generic univariate first-order initial value problem

$$\frac{dy}{dx} = f(x, y(x)), \quad x \in [x_0, x_T], \quad y(x_0) = y_0. \quad (2)$$

Throughout this paper all ODEs that we consider will be assumed to be well-defined and admit a unique solution $y^\dagger \in \mathcal{Y}$ where \mathcal{Y} is some pre-specified set. In this paper the quantity of interest $Q(y^\dagger)$ will either be the solution curve y^\dagger itself or the value $y^\dagger(x_T)$ of the solution at a specific input (in this section it will be the former). The approach outlined in Skilling (1992) allows for a general prior $\mu \in \mathcal{P}_{\mathcal{Y}}$. The gradient field f is treated as a ‘black box’ oracle that can be queried at a fixed computational cost. Thus we are provided with evaluations of the gradient field $[f(x_0, y_0), \dots, f(x_n, y_n)]^\top \in \mathbb{R}^{n+1}$ for certain input pairs $\{(x_i, y_i)\}_{i=0}^n$.

This approach of treating evaluations of the gradient field as ‘data’ will be seen to be a common theme in existing PNM for ODEs and theoretical support for this framework is rooted in the field of information-based complexity (Traub and Woźniakowski, 1992). Let $a_i = f(x_i, y_i)$ and $a^i = [a_0, \dots, a_i]$. The selection of the input pairs (x_i, y_i) on which f is evaluated is not constrained and several possibilities, of increasing complexity, were discussed in Skilling (1992). To fix ideas, the simplest such approach is to proceed iteratively as follows:

- (0.1) The first pair (x_0, y_0) is fully determined by the initial condition of the ODE.
 - (0.2) The oracle then provides one piece of information, $a_0 = f(x_0, y_0)$.
 - (0.3) The prior μ is updated according to a_0 , leading to a belief distribution μ_0 which is just the disintegration element μ^{a^0} .
- (1) A discrete time step $x_1 = x_0 + h$, where $h = \frac{x_T - x_0}{n} > 0$, is performed and a particular point estimate $y_1 = \int y(x_1) d\mu_0(y)$ for the unknown true value $y^\dagger(x_1)$ is obtained. This specifies the second pair (x_1, y_1) .

The process continues similarly, such that at time step $i - 1$ we have a belief distribution $\mu_{i-1} = B(\mu, a^{i-1}) \in \mathcal{P}_{\mathcal{Y}}$, where the general belief update operator B is yet to be defined, and the following step is performed:

- (i) Let $x_i = x_{i-1} + h$ and set $y_i = \int y(x_i) d\mu_{i-1}(y)$.

The final output is a probability distribution $\mu_n = B(\mu, a^n) \in \mathcal{P}_{\mathcal{Y}}$. Now, strictly speaking, the method just described is *not* a PNM in the concrete sense that we have defined. Indeed, the final output μ_n is a deterministic function of the values a^n of the gradient field that were obtained. However, in the absence of additional assumptions on the global smoothness of the gradient field, the values of $f(x, y)$ outside any open neighbourhood of the true solution curve $\mathcal{C} = \{(x, y) : y = y^\dagger(x), x \in [x_0, x_T]\}$ do not determine the solution of the ODE and, conversely, the solution of the ODE provides no information

about the values of the gradient field outside any open neighbourhood of the true solution curve \mathcal{C} . Thus it is not possible, in general, to write down an information operator $A : \mathcal{Y} \rightarrow \mathcal{A}$ that reproduces the information a^n when applied to the solution curve $y^\dagger(\cdot)$ of the ODE.

The approach taken in Skilling (1992) was therefore to posit an *approximate* information operator \hat{A} and a particular belief update operator B , which are now described. The approximate information operator is motivated by the intuition that if $y^\dagger(x_i)$ is well-approximated by y_i at the abscissa x_i then $\frac{dy^\dagger}{dx}(x_i)$ should be well-approximated by $f(x_i, y_i)$. That is, the following approximate information operator \hat{A} was constructed:

$$\hat{A}(y) = \left[\frac{dy}{dx}(x_0), \dots, \frac{dy}{dx}(x_n) \right]^\top. \quad (3)$$

Of course, $\hat{A}(y^\dagger) \neq a^n$ in general. To acknowledge the approximation error, Skilling (1992) proposed to model the information with an approximate likelihood:

$$\frac{d\mu_n}{d\mu_0}(y) = \prod_{i=1}^n \frac{d\mu_i}{d\mu_{i-1}}(y) \quad (4)$$

$$\frac{d\mu_i}{d\mu_{i-1}}(y) \propto \exp\left(-\frac{1}{2\sigma^2} \left(\frac{dy}{dx}(x_i) - f(x_i, y_i)\right)^2\right). \quad (5)$$

This was referred to in Skilling (1992) as simply a “likelihood” and, together with $\mu_0 = \mu^{a^0}$, the output μ_n is completely specified. Here σ is a fixed positive constant, however in principle a non-diagonal covariance matrix can also be considered. The negative consequences of basing inferences on an approximate information operator \hat{A} are potentially twofold. First, recall that values of the gradient field that are not contained on the true solution curve of the ODE do not, in principle, determine the true solution curve y^\dagger . It is therefore unclear if these values should be taken into account at all. Second, in the special case where the gradient field f does not depend the second argument then the quantities $\frac{dy}{dx}(x_i)$ and $f(x_i, y_i)$ are identical. From this perspective, μ_n represents inference under a mis-specified likelihood, since information is treated as erroneous when it is in fact exact. The use of a mis-specified likelihood violates the *likelihood principle* and implies violation of the Bayesian framework. This confirms, through a different argument, that the approach of Skilling (1992) cannot be Bayesian in the strict sense of Definition 1.

Schober et al. (2014, 2019); Teymur et al. (2016, 2018)

After Skilling (1992), several authors have proposed improvements to the above method. The approach of Schober et al. (2014) considered (5) in the $\sigma \downarrow 0$ limit. In order that exact conditioning can be performed in this limit, the input belief distribution μ was restricted to be a k -times integrated Wiener measure on the solution space of the ODE. The tractability of the integrated Wiener measure leads to a closed-form characterisation of the posterior and enables computation to be cast as a Kalman filter.

This direction of research can be motivated by the following fact: For $k \in \{1, 2\}$ the authors prove that if the input pair (x_1, y_1) is taken as $y_1 = \int y(x_1) d\mu_0(y)$, as indicated in Section 1.2, then the smoothing estimate $\hat{y}_1 = \int y(x_1) d\mu_1(y)$, i.e. the posterior mean for $y(x_1)$ based on information a^1 , coincides with the deterministic approximation to $y^\dagger(x_1)$ that would be provided by a k -th order Runge-Kutta method. As such, theoretical guarantees such as local convergence order are inherited. For $k = 3$ it was shown that the same conclusion can be made to hold, *provided* that the input pair (x_1, y_1) is selected in a manner that is no longer obviously related to μ_0 . In all cases the identification with a classical Runge-Kutta method does not extend beyond iteration $n = 1$. Similar connections to multistep methods of Nordsieck and Adams form were identified, respectively, in Schober et al. (2019) and Teymur et al. (2016, 2018). The approach of Schober et al. (2014) is not Bayesian in the sense of Definition 1, which can again be deduced from dependence on values of the gradient field away from the true solution curve, so that the likelihood principle is violated.

Kersting and Hennig (2016)

The subsequent work of Kersting and Hennig (2016) attempted to elicit an appropriate non-zero covariance matrix for use in (5), in order to encourage uncertainty estimates to be better calibrated. Their proposal consisted of the approximate likelihood

$$\frac{d\mu_i}{d\mu_{i-1}}(y) \propto \exp\left(-\frac{1}{2} \left(\frac{\frac{dy}{dx}(x_i) - m_i}{\sigma_i}\right)^2\right) \tag{6}$$

$$m_i = \int f(x_i, y(x_i)) d\mu_{i-1}(y) \tag{7}$$

$$\sigma_i^2 = \int (f(x_i, y(x_i)) - m_i)^2 d\mu_{i-1}(y). \tag{8}$$

This can be viewed as the predictive marginal likelihood for the value $f(x_i, y(x_i))$ based on μ_{i-1} . From a practical perspective, the approach is somewhat circular as the integrals in (7) and (8) involve the black-box gradient field f and are therefore cannot be computed. The authors suggested a number of ways that these quantities could be numerically approximated,⁴ which involve evaluating $f(x_i, y_i)$ at one or more values y_i that must be specified. The overall approach again violates the likelihood principle and is therefore not Bayesian in the sense of Definition 1.

Chkrebtii et al. (2016)

The original work of Chkrebtii et al. (2016) is somewhat related to Kersting and Hennig (2016), however instead of using the mean of the current posterior as input to the gradient field, the input pair (x_i, y_i) was selected by sampling y_i from the marginal

⁴One such method is *Bayesian quadrature*, another PNM wherein the integrand f is modelled as uncertain until it is evaluated. This raises separate philosophical challenges, as one must then ensure that the statistical models used for $y(\cdot)$ and $f(x_i, \cdot)$ are logically consistent. In Kersting and Hennig (2016) these functions were simply modelled as independent.

distribution for $y(x_i)$ implied by μ_{i-1} . The approximate likelihood in this approach was taken as follows:

$$\begin{aligned}\frac{d\mu_i}{d\mu_{i-1}}(y) &\propto \exp\left(-\frac{1}{2}\left(\frac{\frac{dy}{dx}(x_i) - f(x_i, y_i)}{\sigma_i}\right)^2\right) \\ m_i &= \int \frac{dy}{dx}(x_i) d\mu_{i-1}(y) \\ \sigma_i^2 &= \int \left(\frac{dy}{dx}(x_i) - m_i\right)^2 d\mu_{i-1}(y).\end{aligned}$$

Compared to (6), (7) and (8), this approach does not rely on integrals over the unknown gradient field. However, the approach also relies on the approximate information operator in (3) and is thus not Bayesian according to Definition 1.

Conrad et al. (2017); Abdulle and Garegnani (2018)

The approaches proposed in Conrad et al. (2017); Abdulle and Garegnani (2018) are not motivated in the Bayesian framework, but instead seek to introduce a stochastic perturbation into a classical numerical method. Both methods focus on the quantity of interest $Q(y^\dagger) = y^\dagger(x_T)$. In the simple context of (2), the method of Conrad et al. (2017) augments the explicit Euler method with a stochastic perturbation:

$$y_i = y_{i-1} + hf(x_{i-1}, y_{i-1}) + h^2\epsilon_i, \quad x_i = x_{i-1} + h, \quad i = 1, \dots, n.$$

The distribution of the sequence $(\epsilon_i)_{i=1}^n$ must be specified. In the simplest case where the ϵ_i are modelled as independent, say with $\epsilon_i \sim \rho$, the canonical flow map $\Phi_i : \mathbb{R} \rightarrow \mathbb{R}$ of the explicit Euler method, defined as $\Phi_i(z) = z + hf(x_i, z)$, is replaced by the probabilistic counterpart $\Psi_i : \mathcal{P}_{\mathbb{R}} \rightarrow \mathcal{P}_{\mathbb{R}}$ given by

$$\Psi_i(\nu)(dz) = \int \rho\left(\frac{dz - \Phi_i(\tilde{z})}{h^2}\right) \nu(d\tilde{z})$$

through which stochasticity can be propagated. The output of the method of Conrad et al. (2017) is then $B = \Psi_n \circ \dots \circ \Psi_1 \delta(y_0)$, where $\delta(z)$ denotes an atomic distribution on $z \in \mathbb{R}$. For the case where each ρ_i has zero mean, it can be shown that the mean of B equals $\Phi_n \circ \dots \circ \Phi_1(y_0)$, which is exactly the deterministic approximation produced with the explicit Euler method.

This framework can be practically problematic, since ϵ_i is charged with modelling the extrapolation error and such errors are not easily modelled as independent random variables – Section 2.8 of Higham (2002) is devoted to this point. Thus, if for example $f(x, y) = y$, the true linearisation error at step i is $e^{x_i} - e^{x_{i-1}}$ so that the ‘true’ sequence $(\epsilon_i)_{i=1}^n$ in this case is monotonic and exponentially unbounded. The challenge of designing a stochastic model for the sequence $(\epsilon_i)_{i=1}^n$ that reflects the highly structured nature of the error remains unresolved. On the other hand, the mathematical properties of this method are now well-understood (Lie et al., 2018, 2019). The proposal of Abdulle and Garegnani (2018) was to instead consider randomisation of the inputs $\{x_i\}_{i=0}^T$ in the context of a classical numerical method, also outside of the Bayesian framework.

Cockayne et al. (2019); Tronarp et al. (2019)

The survey just presented begs the question of whether a Bayesian PNM for ODEs can exist at all. A first step toward this goal was taken in Cockayne et al. (2019), where it was argued that an information operator can be constructed if the vector field f is brought to the left-hand-side in (2). Specifically, they proposed the information operator

$$\tilde{A}(y) = \left[\frac{dy}{dx}(x_0) - f(x_0, y(x_0)), \dots, \frac{dy}{dx}(x_n) - f(x_n, y(x_n)) \right]^\top$$

for which the ‘data’ are trivial; $\tilde{a}^n = 0$. It was rigorously established that the approximate likelihood

$$\frac{d\mu_{i,\sigma}}{d\mu_{i-1,\sigma}}(y) = \exp\left(-\frac{1}{2\sigma^2} \left(\frac{dy}{dx}(x_i) - f(x_i, y(x_i))\right)^2\right)$$

leads to an exact Bayesian PNM in the limit: $\mu_{n,\sigma} \xrightarrow{\mathcal{F}} \mu^{\tilde{a}^n}$ as $\sigma \downarrow 0$ for $\tilde{A}_\# \mu$ -almost all $\tilde{a}^n \in \mathbb{R}^{n+1}$. Here $\xrightarrow{\mathcal{F}}$ denotes convergence in an integral probability metric defined by a suitable set \mathcal{F} of test functions (see Section 4 of Cockayne et al., 2019). However, the dependence of the information operator \tilde{A} on f means that this cannot be used as the basis for a practical method. Indeed, unless f depends linearly on its second argument and conjugacy properties of the prior can be exploited, the posterior cannot easily be characterised. Approximate techniques from nonlinear filtering were proposed to address this challenge in Tronarp et al. (2019).

1.3 Our Contribution

The comprehensive literature review in the previous section reveals not only that no Bayesian PNM has yet been proposed, but also that such an endeavour may be fundamentally difficult. Indeed, a theme that has emerged with existing PNM for ODEs, which can be traced back to Skilling (1992), is the use of approximate and subjective forms for the likelihood. The complex, implicit relationship between the latent ODE solution y^\dagger and the data $f(x_i, y_i)$ arising from the gradient field appears to preclude use of an exact likelihood. Of course, violation of the likelihood principle is not traditionally a concern in the design of a numerical method, yet if the strictly richer output that comes with a Bayesian PNM is desired, then clearly adherence to the likelihood principle is important. It is therefore natural to ask the question, “under what conditions can exact Bayesian inference for ODEs be made?”.

This paper presents a proof-of-concept PNM for the numerical solution of a (limited) class of ODEs that is both (a) Bayesian in the sense of Definition 1 and (b) can in principle be implemented. The method being proposed is indicated in Figure 2 and its main properties are as follows:

- The classical theory of Lie groups is exploited, for the first time in the context of PNM, to understand when an ODE of the form in (2) can be transformed into an ODE whose gradient field is a function of the independent state variable only, reducing the ODE to an integral.




Figure 2: Schematic of our proposed approach. An n th order ODE that admits a solvable Lie algebra can be transformed into n integrals, to which exact Bayesian probabilistic numerical methods can be applied. The posterior measure on the transformed space is then pushed back through the inverse transformation onto the original domain of interest.

- For ODEs that admit a solvable Lie algebra, our proposal can be shown to simultaneously perform exact Bayesian inference on both the original and the Lie-transformed ODE. Crucially, as we explain later, to identify a Lie algebra only high-level *a priori* information about the ODE is required. The case of first- and second-order ODEs is presented in detail, but the method itself is general.
- In general the specification of prior belief can be difficult. The prior distributions that we construct are guaranteed to respect aspects of the structure of the ODE. As such, our priors are, to some extent, automatically adapted to the ODE at hand as opposed to being arbitrarily posited.
- In addition to the benefits conferred in the Bayesian framework, detailed in Section 1.1 and in Cockayne et al. (2019), the method being proposed can be computationally realised. On the other hand, there is a cost in terms of the run-time of the method that is substantially larger than existing, non-Bayesian approaches (especially classical numerical methods). As such, we consider this work to be a proof-of-concept rather than an applicable Bayesian PNM.

The remainder of the paper is structured as follows: Section 2 is dedicated to a succinct review of Lie group methods for ODEs. In Section 3, Lie group methods are exploited to construct priors over the solution space of the ODE whenever a solvable Lie algebra is admitted and exact Bayesian inference is performed on a transformed version of the original ODE which takes the form of an integral. Numerical experiments are reported in Section 4. Finally, some conclusions and recommendations for future research are drawn in Section 5.

2 Background

This section provides a succinct overview of classical Lie group methods, introduced in the 19th century by Sophus Lie in the differential equation context (Hawkins, 2012). Lie

developed the fundamental notion of a *Lie group of transformations*, which roughly correspond to maps that take one solution of the ODE to another. This provided a formal generalisation of certain algebraic techniques, such as dimensional analysis and transformations based on spatial symmetries, that can sometimes be employed to algebraically reduce the order of an ODE.

This section proceeds as follows: First, in Section 2.1 we introduce a one-parameter Lie group of transformations and then, in Section 2.2, we explain what it means for a curve or a surface to be transformation-invariant. In Section 2.3 we recall consequences of Lie's theory in the ODE context. Last, in Section 2.4 the generalisation to multi-parameter Lie groups is indicated. Our development is heavily influenced by Bluman and Anco (2002) and we refer the reader to their book when required.

2.1 One-Parameter Lie Groups of Transformations

The purpose of this section is to recall essential definitions, together with the *first fundamental theorem of Lie*, which relates a Lie group of transformations to its infinitesimal generator. In what follows we consider a fixed domain $D \subset \mathbb{R}^d$ and denote a generic state variable as $x = (x_1, \dots, x_d) \in D$.

Definition 2 (One-Parameter Group of Transformations). *A one-parameter group of transformations on D is a map $X : D \times S \rightarrow D$, defined on $D \times S$ for some set $S \subset \mathbb{R}$, together with a bivariate map $\phi : S \times S \rightarrow S$, such that the following hold:*

- (1) *For each $\epsilon \in S$, the transformation $X(\cdot, \epsilon)$ is a bijection on D .*
- (2) *(S, ϕ) forms a group with law of composition ϕ .*
- (3) *If ϵ_0 is the identity element in (S, ϕ) , then $X(\cdot, \epsilon_0)$ is the identity map on D .*
- (4) *For all $x \in D$, $\epsilon, \delta \in S$, if $x^* = X(x, \epsilon)$, $x^{**} = X(x^*, \delta)$, then $x^{**} = X(x^*, \phi(\epsilon, \delta))$.*

In what follows we continue to use the shorthand notation $x^* = X(x, \epsilon)$. The notion of a *Lie group* additionally includes smoothness assumptions on the maps that constitute a group of transformations. Recall that a real-valued function is *analytic* if it can be locally expressed as a convergent power series.

Definition 3 (One-Parameter Lie Group of Transformations). *Let X , together with ϕ , form a one-parameter group of transformations on D . Then we say that X , together with ϕ , form a one-parameter Lie group of transformations on D if, in addition, the following hold:*

- (5) *S is a (possibly unbounded) interval in \mathbb{R} .*
- (6) *For each $\epsilon \in S$, $X(\cdot, \epsilon)$ is infinitely differentiable in D .*
- (7) *For each $x \in D$, $X(x, \cdot)$ is an analytic function on S .*
- (8) *ϕ is analytic in $S \times S$.*

Without the loss of generality it will be assumed, through re-parametrisation if required, that S contains the origin and $\epsilon = 0$ is the identity element in (S, ϕ) . The definition is illustrated through three examples:

Example 1 (Translation in the x-Axis). *The one-parameter transformation $x_1^* = x_1 + \epsilon$, $x_2^* = x_2$ for $\epsilon \in \mathbb{R}$ forms a Lie group of transformations on $D = \mathbb{R}^2$ with group composition law $\phi(\epsilon, \delta) = \epsilon + \delta$.*

Example 2 (Rotation Group). *The one-parameter transformation $x_1^* = x_1 \cos(\epsilon) - x_2 \sin(\epsilon)$, $x_2^* = x_1 \sin(\epsilon) + x_2 \cos(\epsilon)$ for $\epsilon \in \mathbb{R}$ again forms a Lie group of transformations on $D = \mathbb{R}^2$ with group composition law $\phi(\epsilon, \delta) = \epsilon + \delta$.*

Example 3 (Cyclic group C_p). *Let $D = \{1, 2, 3, \dots, p\}$. Let $S = \mathbb{Z}$. For $n \in D$ and $m \in S$, let $X(n, m) = n + m \pmod{p}$. Then X , together with $\phi(a, b) = a + b$, defines a one parameter group of transformations on D , but is not a Lie group of transformations since (5) is violated.*

The first fundamental theorem of Lie establishes that a Lie group of transformations can be characterised by its infinitesimal generator, defined next:

Definition 4 (Infinitesimal Transformation). *Let X be a one-parameter Lie group of transformations. Then the transformation $x^* = x + \epsilon \xi(x)$,*

$$\xi(x) = \left. \frac{\partial X(x, \epsilon)}{\partial \epsilon} \right|_{\epsilon=0},$$

is called the infinitesimal transformation associated to X and the map ξ is called an infinitesimal.

Definition 5 (Infinitesimal Generator). *The infinitesimal generator of a one-parameter Lie group of transformations X is defined to be the operator $X = \xi \cdot \nabla$ where ξ is the infinitesimal associated to X and $\nabla = (\frac{\partial}{\partial x_1}, \frac{\partial}{\partial x_2}, \dots, \frac{\partial}{\partial x_n})$ is the gradient.*

Example 4 (Ex. 1, continued). *For Ex. 1, we have*

$$\xi(x) = \left(\left. \frac{dx_1^*}{d\epsilon} \right|_{\epsilon=0}, \left. \frac{dx_2^*}{d\epsilon} \right|_{\epsilon=0} \right) = (1, 0)$$

so the infinitesimal generator for translation in the x-axis is $X = \frac{\partial}{\partial x_1}$.

Example 5 (Ex. 2, continued). *Similarly, the infinitesimal generator for the rotation group is $X = -x_2 \frac{\partial}{\partial x_1} + x_1 \frac{\partial}{\partial x_2}$.*

The first fundamental theorem of Lie provides a constructive route to obtain the infinitesimal generator from the transformation itself:

Theorem 1 (First Fundamental Theorem of Lie; see pages 39–40 of Bluman and Anco (2002)). *A one parameter Lie group of transformations X is characterised by the initial value problem:*

$$\frac{dx^*}{d\tau} = \xi(x^*), \quad x^* = x \text{ when } \tau = 0, \quad (9)$$

where $\tau(\epsilon)$ is a parametrisation of ϵ which satisfies $\tau(0) = 0$ and, for $\epsilon \neq 0$,

$$\tau(\epsilon) = \int_0^\epsilon \frac{\partial \phi(a, b)}{\partial b} \Big|_{(a,b)=(\delta^{-1}, \delta)} d\delta.$$

Here δ^{-1} denotes the group inverse element for δ .

Since (9) is translation-invariant in τ , it follows that without loss of generality we can assume a parametrisation $\tau(\epsilon)$ such that the group action becomes $\phi(\tau_1, \tau_2) = \tau_1 + \tau_2$ and, in particular, $\tau^{-1} = -\tau$. In the remainder of the paper, for convenience we assume that all Lie groups are parametrised such that the group action is $\phi(\epsilon_1, \epsilon_2) = \epsilon_1 + \epsilon_2$.

The next result can be viewed as a converse to Theorem 1, as it shows how to obtain the transformation from the infinitesimal generator. All proofs are reserved for Supplemental Section A.1 (Wang et al., 2019).

Theorem 2. *A one parameter Lie group of transformations with infinitesimal generator X is equivalent to $x^* = e^{\epsilon X}x$, where $e^{\epsilon X} = \sum_{k=0}^\infty \frac{1}{k!} \epsilon^k X^k x$.*

The following is immediate from the proof of Theorem 2:

Corollary 1. *If F is infinitely differentiable, then $F(x^*) = e^{\epsilon X}F(x)$.*

2.2 Invariance Under Transformation

In this section we explain what it means for a curve or a surface to be invariant under a Lie group of transformations and how this notion relates to the infinitesimal generator.

Definition 6 (Invariant Function). *A function $F : D \rightarrow \mathbb{R}$ is said to be invariant under a one parameter Lie group of transformations $x^* = X(x, \epsilon)$ if $F(x^*) = F(x)$ for all $x \in D$ and $\epsilon \in S$.*

Based on the results in Section 2.1, one might expect that invariance to a transformation can be expressed in terms of the infinitesimal generator of the transformation. This is indeed the case:

Theorem 3. *A differentiable function $F : D \mapsto \mathbb{R}$ is invariant under a one parameter Lie group of transformations with infinitesimal generator X if and only if $XF(x) = 0$ for all $x \in D$.*

Theorem 4. *For a function $F : D \mapsto \mathbb{R}$ and a one parameter Lie group of transformations $x^* = X(x, \epsilon)$, the relation $F(x^*) = F(x) + \epsilon$ holds for all $x \in D$ and $\epsilon \in S$ if and only if $XF(x) = 1$ for all $x \in D$.*

The following definition is fundamental to the method proposed in Section 3:

Definition 7 (Canonical Coordinates). *Consider a coordinate system $r = (r_1(x), \dots, r_n(x))$ on D . Then any one parameter Lie group of transformations $x^* = X(x, \epsilon)$ induces a transformation of the coordinates $r_i^* = r_i(x^*)$. The coordinate system r is called canonical for the transformation if $r_1^* = r_1, \dots, r_{n-1}^* = r_{n-1}$ and $r_n^* = r_n + \epsilon$.*

Example 6 (Ex. 2, continued). For the rotation group in Ex. 2, we have canonical coordinates $r_1(x_1, x_2) = \sqrt{x_1^2 + x_2^2}$, $r_2(x_1, x_2) = \arctan(x_2/x_1)$.

In canonical coordinates, a one parameter Lie group of transformations can be viewed as a straight-forward translation in the r_n -axis. The existence of canonical coordinates is established in Thm. 2.3.5-2 of Bluman and Anco (2002). Note that Thms. 3 and 4 imply that $Xr_i^* = 0$ for $i = 1, 2, \dots, n-1$, $Xr_n^* = 1$.

Definition 8 (Invariant Surface). For a function $F : D \rightarrow \mathbb{R}$, a surface defined by $F(x) = 0$ is said to be invariant under a one parameter Lie group of transformation $x^* = X(x, \epsilon)$ if and only if $F(x^*) = 0$ whenever $F(x) = 0$ for all $x \in D$ and $\epsilon \in S$.

The invariance of a surface, as for a function, can be cast in terms of an infinitesimal generator:

Corollary 2. A surface $F(x) = 0$ is invariant under a one parameter Lie group of transformations with infinitesimal generator X if and only if $XF(x) = 0$ whenever $F(x) = 0$.

2.3 Symmetry Methods for ODEs

The aim of this section is to relate Lie transformations to ODEs for which these transformations are admitted. These techniques form the basis for our proposed method in Section 3.

For an ODE of the form in (2), one can consider the action of a transformation on the coordinates (x, y) ; i.e. a special case of the above framework where the generic coordinates x_1 and x_2 are respectively the independent (x) and dependent (y) variables of the ODE. It is clear that such a transformation also implies some kind of transformation of the derivatives $y_m := \frac{d^m y}{dx^m}$. Indeed, consider a one-parameter Lie group of transformations $(x^*, y^*) = (X(x, y; \epsilon), Y(x, y; \epsilon))$. Then we have from the chain rule that $y_m^* := \frac{d^m y^*}{d(x^*)^m}$ is a function of x, y, y_1, \dots, y_m and we denote $y_m^* = Y_m(x, y, y_1, \dots, y_m; \epsilon)$. As an explicit example:

$$y_1^* = \frac{dy^*}{dx^*} = \frac{\frac{\partial Y(x, y; \epsilon)}{\partial x} + y_1 \frac{\partial Y(x, y; \epsilon)}{\partial y}}{\frac{\partial X(x, y; \epsilon)}{\partial x} + y_1 \frac{\partial X(x, y; \epsilon)}{\partial y}} =: Y_1(x, y, y_1; \epsilon).$$

In general:

$$y_m^* = \frac{\frac{\partial y_{m-1}^*}{\partial x} + y_1 \frac{\partial y_{m-1}^*}{\partial y} + y_2 \frac{\partial y_{m-1}^*}{\partial y_1} + \dots + y_m \frac{\partial y_{m-1}^*}{\partial y_{m-1}}}{\frac{\partial X(x, y; \epsilon)}{\partial x} + y_1 \frac{\partial X(x, y; \epsilon)}{\partial y}} =: Y_m(x, y, y_1, \dots, y_m; \epsilon).$$

In this sense a transformation defined on (x, y) can be naturally extended to a transformation on (x, y, y_1, y_2, \dots) as required.

Definition 9 (Admitted Transformation). An m th order ODE $F(x, y, y_1, \dots, y_m) = 0$ is said to admit a one parameter Lie group of transformations $(x^*, y^*) = (X(x, y; \epsilon),$

$Y(x, y; \epsilon)$ if the surface F defined by the ODE is invariant under the Lie group of transformations, i.e. if $F(x^*, y^*, y_1^*, \dots, y_m^*) = 0$ whenever $F(x, y, y_1, \dots, y_m) = 0$.

Example 7. Clearly any ODE of the form $\frac{dy}{dx} = F(x)$ admits the transformation $(x^*, y^*) = (x, y + \epsilon)$.

Our next task is to understand how the infinitesimal generator of a transformation can be extended to act on derivatives y_m .

Definition 10 (Extended Infinitesimal Transformation). *The m th extended infinitesimals of a one parameter Lie group of transformations $(x^*, y^*) = (X(x, y; \epsilon), Y(x, y; \epsilon))$ are defined as the functions $\xi, \eta, \eta^{(1)}, \dots, \eta^{(m)}$ for which the following equations hold:*

$$\begin{aligned} x^* &= X(x, y; \epsilon) &&= x + \epsilon\xi(x, y) + O(\epsilon^2) \\ y^* &= Y(x, y; \epsilon) &&= y + \epsilon\eta(x, y) + O(\epsilon^2) \\ y_1^* &= Y_1(x, y, y_1; \epsilon) &&= y_1 + \epsilon\eta^{(1)}(x, y, y_1) + O(\epsilon^2) \\ &\vdots && \\ y_m^* &= Y_m(x, y, y_1, \dots, y_m; \epsilon) &&= y_m + \epsilon\eta^{(m)}(x, y, y_1, y_2, \dots, y_m) + O(\epsilon^2). \end{aligned}$$

It can be shown straightforwardly via induction that

$$\eta^{(m)}(x, y, y_1, y_2, \dots, y_m) = \frac{d^m \eta}{dx^m} - \sum_{k=0}^m \frac{m!}{(m-k)!k!} y_{m-k-1} \frac{d^k \xi}{dx^k}, \tag{10}$$

where $\frac{d}{dx}$ denotes the full derivative with respect to x , i.e. $\frac{d}{dx} = \frac{\partial}{\partial x} + y_1 \frac{\partial}{\partial y} + \sum_{k=2}^{m+1} y_k \frac{\partial}{\partial y_{k-1}}$. It follows that $\eta^{(m)}$ is a polynomial in y_1, y_2, \dots, y_m with coefficients linear combinations of ξ, η and their partial derivatives up to the m th order.

Definition 11 (Extended Infinitesimal Generator). *The m th extended infinitesimal generator is defined as*

$$\begin{aligned} X^{(m)} &= \xi_m(x, y, y_1, \dots, y_m) \cdot \nabla \\ &= \xi(x, y) \frac{\partial}{\partial x} + \eta(x, y) \frac{\partial}{\partial y} + \eta^{(1)}(x, y) \frac{\partial}{\partial y_1} + \dots + \eta^{(m)}(x, y, y_1, \dots, y_m) \frac{\partial}{\partial y_m}, \end{aligned}$$

where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_m})$ is the extended gradient.

The following corollaries are central to the actual computation of the admitted Lie groups of an ODE.

Corollary 3. *A differentiable function $F : D_m \rightarrow \mathbb{R}$ where D_m is the phase space containing elements of the form (x, y, y_1, \dots, y_m) , is invariant under a one parameter Lie group of transformations with an extended infinitesimal generator $X^{(m)}$ if and only if $X^{(m)}F(x, y, y_1, \dots, y_m) = 0$ for all $(x, y, y_1, \dots, y_m) \in D_m$.*

Corollary 4 (Infinitesimal Criterion for Symmetries Admitted by an ODE). *A one parameter Lie group of transformations is admitted by the m th order ODE $F(x, y, y_1, \dots, y_m) = 0$ if and only if its extended infinitesimal generator $X^{(m)}$ satisfies $X^{(m)}F(x, y, y_1, \dots, y_m) = 0$ whenever $F(x, y, y_1, \dots, y_m) = 0$.*

2.4 Multi-Parameter Lie Groups and Lie Algebras

To leverage the full power of Lie symmetry methods for ODEs of order $m \geq 2$, we need to consider multiple Lie symmetries which are collectively described by a *Lie algebra*. Fortunately, the notion of a multi-parameter Lie group of transformations is a natural generalisation from the one parameter case. Thus, this last section of background material concerns the generalisation of the definitions in Section 2.1 to the case of a multi-parameter Lie group. The associated Lie algebra will also be defined.

Definition 12 (Multi-Parameter Lie Group of Transformations). *The set of transformations $x^* = X(x, \epsilon)$ where $x_i^* = X_i(x, \epsilon)$ and $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_r) \in S \subset \mathbb{R}^r$ is called a r -parameter Lie group of transformations if it satisfies the same axioms as in the one parameter case, but with law of composition $\phi(\epsilon, \delta) = (\phi_1(\epsilon, \delta), \dots, \phi_r(\epsilon, \delta))$, and (without loss of generality) $\epsilon = (0, 0, \dots, 0)$ as the group identity element.*

Definition 13 (Infinitesimal Matrix). *The appropriate generalisation for the infinitesimal transformation is the infinitesimal matrix $\Xi = [\xi_{ij}]$, whose entries are defined as $\xi_{ij}(x) = \left. \frac{\partial X_j(x, \epsilon)}{\partial \epsilon_i} \right|_{\epsilon=0}$.*

Definition 14 (Infinitesimal Generator). *An r -parameter Lie group of transformations is associated with r infinitesimal generators, X_i , defined as $X_i = X_i(x) = \sum_{j=1}^d \xi_{ij}(x) \frac{\partial}{\partial x_j}$.*

The first fundamental theorem of Lie can be generalised to the multi-parameter case. In particular, it can be shown that an r -parameter Lie group of transformations is characterised by the set of its r infinitesimal generators. The generalisation is straightforward and so, for brevity, we refer the reader to pages 39–40 of Bluman and Anco (2002).

Next we explain how the collection of infinitesimal generators forms a Lie algebra. This relies on the basic facts that the set \mathcal{D} of differential operators on D is a vector space over \mathbb{R} (i.e. $\lambda X + \mu Y \in \mathcal{D}$ for all $X, Y \in \mathcal{D}$ and $\lambda, \mu \in \mathbb{R}$) and that differential operators can be composed (i.e. $XY \in \mathcal{D}$ for all $X, Y \in \mathcal{D}$).

Definition 15 (Commutator). *The commutator of two infinitesimal generators X_i and X_j is defined as $[X_i, X_j] = X_i X_j - X_j X_i$.*

Theorem 5 (Second Fundamental Theorem of Lie; see page 78 of Bluman and Anco (2002)). *Consider an r -parameter Lie group of transformations and let \mathcal{L} denote the linear span of the infinitesimal generators X_1, \dots, X_r in \mathcal{D} . Let $[\cdot, \cdot] : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{D}$ denote the unique bilinear operator that agrees with Def. (15) on the set of infinitesimal generators. i.e.*

$$\left[\sum_{i=1}^r \lambda_i X_i, \sum_{j=1}^r \mu_j X_j \right] = \sum_{i=1}^r \sum_{j=1}^r \lambda_i \mu_j (X_i X_j - X_j X_i). \quad (11)$$

Then $[\cdot, \cdot]$ maps into \mathcal{L} . i.e. the right hand side of (11) belongs to \mathcal{L} for all $\lambda, \mu \in \mathbb{R}^r$.

Example 8. *Consider the two parameter group of transformations on $D = \mathbb{R}^2$ given by $(x^*, y^*) = (x + x\epsilon + x^2\delta, y + y\epsilon + y^2\delta)$. The infinitesimal generators corresponding to*

δ and ϵ , respectively, are $X_1 = x^2 \frac{\partial}{\partial x} + y^2 \frac{\partial}{\partial y}$, $X_2 = x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y}$. It can be directly verified that $[X_1, X_2] = -X_1$.

The space \mathcal{L} , defined in Thm. 5, satisfies the properties of an r -dimensional Lie algebra \mathcal{L} , defined next:

Definition 16 (Lie Algebra). *An r -dimensional vector space \mathcal{L} over \mathbb{R} together with a bilinear operator $[\cdot, \cdot] : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ is called an r -dimensional Lie algebra if the following hold:*

- (1) *Alternativity: $[X, X] = 0$ for all $X \in \mathcal{L}$*
- (2) *Jacobi Identity: $[X, [Y, Z]] + [Y, [Z, X]] + [Z, [X, Y]] = 0$ for all $X, Y, Z \in \mathcal{L}$*

In general, for the methods presented in Section 3 to be applied, existence of an n -parameter Lie group of transformations is not in itself sufficient; we require the existence of an n -dimensional solvable Lie sub-algebra, defined next:

Definition 17 (Normal Lie Sub-algebra). *Consider a Lie sub-algebra \mathcal{J} of a Lie algebra \mathcal{L} with bilinear operator $[\cdot, \cdot]$, i.e. a subset $\mathcal{J} \subset \mathcal{L}$ such that, when equipped with the restriction of $[\cdot, \cdot]$ to $\mathcal{J} \times \mathcal{J}$, is itself a Lie algebra and, in particular, $[X, Y] \in \mathcal{J}$ for all $X, Y \in \mathcal{J}$. Then \mathcal{J} is said to be normal if, in addition, $[X, Y] \in \mathcal{J}$ for all $X \in \mathcal{J}, Y \in \mathcal{L}$.*

Definition 18 (Solvable Lie Algebra). *An r -dimensional Lie algebra \mathcal{L} is called solvable if there exists a chain of sub-algebras $\mathcal{L}^1 \subset \mathcal{L}^2 \subset \dots \subset \mathcal{L}^{q-1} \subset \mathcal{L}^r =: \mathcal{L}$ such that \mathcal{L}^{i-1} is a normal sub-algebra of \mathcal{L}^i for $i = 2, 3, \dots, r$.*

For low-order ODEs, the existence requirement for an admitted Lie group of transformations is more restrictive than the requirement that the associated Lie algebra is solvable. Indeed, we have the following result:

Theorem 6. *All two-dimensional Lie Algebras are solvable.*

This completes our review of background material. The exact Bayesian PNM developed in Section 3 for an n th order ODE require the existence of an admitted n -parameter Lie group of transformations with a solvable Lie algebra. In practice we therefore require some high-level information on the gradient field f , in order to establish which transformations of the ODE may be admitted. In addition, the requirement of a solvable Lie algebra also limits the class of ODEs for which our exact Bayesian methods can be employed. Nevertheless, this class of ODEs is sufficiently broad to have merited extensive theoretical research (Bluman and Anco, 2002) and the development of software (Baumann, 2013).

3 Methods

In this section our novel Bayesian PNM is presented. The method relies on high-level information about the gradient field f and, in Section 3.1, we discuss how such information can be exploited to identify any Lie transformations that are admitted by the

ODE. In the case of a first order ODE, any non-trivial transformation is sufficient for our method and an explicit information operator is provided for this case, together with recommendations for prior construction, in Section 3.2. Together, the prior and the information operator uniquely determine a Bayesian PNM, as explained in Section 1.1. In the general case of an m th order ODE, we require that an m -dimensional solvable Lie algebra is admitted by the ODE. The special case $m = 2$ is treated in detail, with an explicit information operator and guidance for prior construction provided in Section 3.3. In the Supplemental Section A.2 the selection of input pairs (x_i, y_i) to the gradient field is discussed.

3.1 From an ODE to its Admitted Transformations

For the methods proposed in this paper, transformations admitted by the ODE, together with their infinitesimal generators, must first be obtained. The algorithm for obtaining infinitesimal generators follows as a consequence of Cor. 4. Indeed, suppose we have a m th order ODE of the form $y_m - f(x, y, y_1, \dots, y_{m-1}) = 0$. Then, by Cor. 4, a transformation with infinitesimal generator X is admitted by the ODE if and only if:

$$X^{(m)}(y_m - f(x, y, y_1, \dots, y_{m-1})) = 0 \quad \text{whenever} \quad y_m = f(x, y, y_1, \dots, y_{m-1}). \quad (12)$$

In infinitesimal notation, (12) is equivalent to

$$\eta^{(m)}(x, y, y_1, \dots, y_{m-1}, y_m) = \xi \frac{\partial f}{\partial x} + \eta \frac{\partial f}{\partial y} + \sum_{k=1}^{m-1} \eta^{(k)} \frac{\partial f}{\partial y_k}. \quad (13)$$

The direct solution of (13) recovers any transformations that are admitted.

In the common scenario where $f(x, y, y_1, \dots, y_{m-1})$ is a polynomial in y_1, y_2, \dots, y_{m-1} , the algorithm just described, for identification of admitted transformations, can be fully automated (c.f. Baumann, 2013). Indeed, from Def. 10 it follows that the extended infinitesimals $\eta^{(k)}$ for $k \in 1, 2, 3, \dots, m$ are polynomial in y_1, y_2, \dots, y_k . Thus, by substituting $y_m = f(x, y, y_1, \dots, y_{m-1})$, (12) too must be a polynomial in y_1, y_2, \dots, y_{m-1} . Moreover, the coefficients of this polynomial must vanish because (12) holds for arbitrary values of $x, y, y_1, \dots, y_{m-1}$. This argument, of setting coefficients to zero, leads to a system of linear partial differential equations (overdetermined when $m \geq 2$) for $\xi(x, y)$ and $\eta(x, y)$, which can be exactly solved to retrieve all the infinitesimal generators of the ODE. The same strategy can often be applied beyond the polynomial case and explicit worked examples of this procedure are now provided:

Example 9 (First Order ODE). *Consider the class of all first order ODEs of the form $\frac{dy}{dx} = f(x, y(x))$, $f(x, y) = F\left(\frac{y}{x}\right)$. From (10), we have $\eta^{(1)} = \eta_x + (\eta_y - \xi_x)y_1 - \xi_y(y_1)^2$ so (13) becomes $\eta_x + (\eta_y - \xi_x)f - \xi_y(f)^2 = \xi \frac{\partial f}{\partial x} + \eta \frac{\partial f}{\partial y}$ and thus $\eta_x + (\eta_y - \xi_x)F\left(\frac{y}{x}\right) - \xi_y F\left(\frac{y}{x}\right)^2 = -\xi F'\left(\frac{y}{x}\right) \frac{y}{x^2} + \eta F'\left(\frac{y}{x}\right) \frac{1}{x}$. For this equation to hold for general F , the coefficients of F , F^2 and F' must vanish: $\eta_x = 0$, $\eta_y - \xi_x = 0$, $\xi_y = 0$, $-\xi \frac{y}{x^2} + \eta \frac{1}{x} = 0$. This is now a linear system of partial differential equations in (ξ, η) which is easily solved; namely $\xi = x, \eta = y$. The associated infinitesimal generator is $X = x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y}$.*

Example 10 (Second Order ODE). *The infinitesimal generators for the second order ODE*

$$(x - y(x)) \frac{d^2y}{dx^2} + 2 \frac{dy}{dx} \left(\frac{dy}{dx} + 1 \right) + \left(\frac{dy}{dx} \right)^{3/2} = 0 \tag{14}$$

are derived in Supplementary Section A.1.

3.2 The Case of a First Order ODE

In this section we present our approach for a first order ODE. This allows some of the more technical details associated to the general case to be omitted, due to the fact that any one-dimensional Lie algebra is trivial. The main result that will allow us to construct an exact Bayesian PNM is as follows:

Theorem 7 (Reduction of a First Order ODE to an Integral). *If a first order ODE*

$$\frac{dy}{dx} = f(x, y(x)) \tag{15}$$

admits a one parameter Lie group of transformations, then there exists coordinates $r(x, y)$, $s(x, y)$ such that

$$\frac{ds}{dr} = G(r) \tag{16}$$

for some explicit function $G(r)$.

Note that the transformed ODE in (16) is nothing more than an integral, for which exact Bayesian PNM have already been developed (e.g. Briol et al., 2019; Karvonen et al., 2018). At a high level, as indicated in Figure 2, our proposed Bayesian PNM performs inference for the solution $s(r)$ of (16) and then transforms the resultant posterior back into the original (x, y) -coordinate system. Our PNM is therefore based on the information operator

$$A(y) = [G(r_0), \dots, G(r_n)]^\top \in \mathcal{A} = \mathbb{R}^{n+1} \tag{17}$$

which corresponds indirectly to $n + 1$ evaluations of the original gradient field f at certain input pairs (x_i, y_i) . The selection of the inputs r_i is discussed in Section A.2.

The transformation of a first order ODE is clearly illustrated in the following:

Example 11 (Ex. 9, continued). *Consider the first order ODE $\frac{dy}{dx} = f(x, y(x))$, $f(x, y) = F\left(\frac{y}{x}\right)$. Recall from Ex. 9 that this ODE admits the one parameter Lie group of transformations $x^* = \alpha x$, $y^* = \alpha y$ for $\alpha \in \mathbb{R}$ and the associated infinitesimal generator is $X = x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y}$. Solving the pair of partial differential equations $Xr = 0$, $Xs = 1$ yields the canonical coordinates $s = \log y$, $r = \frac{y}{x}$. The transformed ODE is then $\frac{ds}{dr} = \frac{F(r)}{-r^2 + rF(r)} =: G(r)$. Thus an evaluation $G(r)$ corresponds to an evaluation of $f(x, y)$ at an input (x, y) such that $r = \frac{y}{x}$.*

Two important points must now be addressed: First, the approach just described cannot be Bayesian unless it corresponds to a well-defined prior distribution $\mu \in \mathcal{P}_{\mathcal{Y}}$ in the original coordinate system \mathcal{Y} . This precludes standard (e.g. Gaussian process)




Figure 3: Illustration of the implicit prior principle: A prior elicited for the function $s(r)$ in the transformed coordinate system (r, s) must be supported on functions $s(r)$ that correspond to well-defined functions $y(x)$ in the original coordinate system (x, y) . Thus the situation depicted would not be allowed.

priors in general, as such priors assign mass to functions in (r, s) -space that do not correspond to well-defined functions in (x, y) -space (see Figure 3). Second, any prior that is used ought to be consistent with the Lie group of transformations that the ODE is known to admit. To address each of these important points, we propose two general principles for prior construction in this work. The first principle is the *implicit prior* principle. This ensures that a prior specified in the transformed coordinates (r, s) can be safely transformed into a well-defined distribution on \mathcal{Y} . For such an implicit prior to be well-defined we need to understand when a function in (r, s) space maps to a well-defined function in the original (x, y) domain of interest. Let \mathcal{S} denote the image of \mathcal{Y} under the canonical coordinate transformation.

Principle 1 (Implicit Prior). *A distribution $\nu \in \mathcal{P}_{\mathcal{S}}$ on the transformed solution space \mathcal{S} corresponds to a well-defined implicit prior $\mu \in \mathcal{P}_{\mathcal{Y}}$ provided that $x(r, s(r))$ is strictly monotone as a function of r .*

Example 12 (Ex. 11, continued). *For the ODE in Ex. 11, with canonical coordinates $s = \log y$, $r = \frac{y}{x}$, if $x \in [x_0, x_T] = [1, x_T]$ and $y \in (0, \infty)$, then the region in the (r, s) plane corresponding to $[1, x_T] \times (0, \infty)$ in the (x, y) plane is $(0, \infty) \times \mathbb{R}$. Now,*

$$\frac{dx(r, s(r))}{dr} = \frac{\partial x}{\partial r} + \frac{\partial x}{\partial s} \frac{ds(r)}{dr} = \frac{rs'(r)\exp(s(r)) - \exp(s(r))}{r^2}.$$

Thus $\frac{dx}{dr} > 0$ if and only if $s'(r) > \frac{1}{r}$ and the invariant prior principle requires that we respect the constraint $\log(r) \leq s(r) \leq \log(r) + \log(x_T)$ for all $r > 0$. The set \mathcal{S} must therefore consist of differentiable functions s defined on $r \in (0, \infty)$ and satisfying $\log(r) \leq s(r) \leq \log(r) + \log(x_T)$.

Now we turn to the second important point, namely that the prior ought to encode knowledge about any Lie transformations that are known to be admitted by the ODE. In working on the transformed space \mathcal{S} , it become clear how to construct a prior measure in which this knowledge is encoded. Our second principle for prior specification states that equal prior weight should be afforded to all curves that are identical up to a Lie transformation:

Principle 2 (Invariant Prior). *A distribution $\nu \in \mathcal{P}_{\mathcal{S}}$ on the transformed solution space \mathcal{S} is said to be invariant provided that $\nu(S) = \nu(S + \epsilon)$ where the elements of $S + \epsilon$ are the elements of S after a vertical translation; i.e. $s(\cdot) \mapsto s(\cdot) + \epsilon$ and both $S, S + \epsilon \in \Sigma_{\mathcal{S}}$.*

Our recommendation is that, when possible, both the implicit prior principle and the invariant prior principle should be enforced. However, in practice it seems difficult to satisfy both principles and our empirical results in Section 4 are based on implicit priors that are not invariant.

3.3 The Case of a Second Order ODE

In this section we present our approach for a second order ODE. The study of second order ODEs is particularly important, since Newtonian mechanics is based on ODEs of second order. The presentation is again simplified relative to the general case of an m th order ODE, this time by virtue of the fact that any two dimensional Lie algebra is guaranteed to be solvable (Thm. 6). The main result that will allow us to construct an exact Bayesian PNM is as follows:

Theorem 8 (Reduction of a Second Order ODE to Two Integrals). *If a second order ODE*

$$\frac{d^2y}{dx^2} = f\left(x, y(x), \frac{dy}{dx}\right) \tag{18}$$

admits a two parameter Lie group of transformations, then there exists coordinates $r(x, y), s(x, y)$ such that

$$\frac{ds}{dr} = G(r) \tag{19}$$

for some implicitly defined function G . The function G is explicitly related to the solution of a second equation of the form

$$\frac{d\tilde{s}}{d\tilde{r}} = H(\tilde{r}) \tag{20}$$

for some explicit function $H(\tilde{r})$.

Note that the ODE in (18) is reduced to two integrals, namely (19) and (20). At a high level, our proposed Bayesian PNM performs inference for the solution $s(r)$ of (19) and then transforms the resultant posterior back into the original (x, y) -coordinate system. However, because G in (19) depends on the solution $\tilde{s}(\tilde{r})$ of (20), we must also estimate $\tilde{s}(\tilde{r})$ and for this we need to evaluate H . Our PNM is therefore based on the information operator

$$A(y) = [G(r_0), \dots, G(r_n), H(\tilde{r}_0), \dots, H(\tilde{r}_n)]^T \in \mathcal{A} = \mathbb{R}^{2(n+1)}$$

which corresponds indirectly to $2(n + 1)$ evaluations of f , the original gradient field. The extension of our approach to a general m th order ODE proceeds analogously, with $\mathcal{A} = \mathbb{R}^{m(n+1)}$. The use of Thm. 8 is illustrated in Example 14 in the Supplement.

The two principles of prior construction that we advocated in the case of a first order ODE apply equally to the case of a second- and higher-order ODE. It therefore remains only to discuss the selection of the specific inputs r_i (and \tilde{r}_i in the case of a second order ODE) that are used to define the information operator A . This discussion is again reserved for Supplemental Section A.2.

4 Numerical Illustration

In this section the proposed Bayesian PNM is empirically illustrated. Recall that we are not advocating these methods for practical use, rather they are to serve as a proof-of-concept for demonstrating that exact Bayesian inference can *in principle* be performed for ODEs, albeit at considerable effort; a non-trivial finding that helps to shape ongoing research and discussion in this nascent field.

The case of a first order ODE is considered in Section 4.1 and the second order case is contained in Section 4.2. In both cases, scope is limited to verifying the correctness of the procedures, as well as indicating how conjugate prior distributions can be constructed.

4.1 A First Order ODE

This section illustrates the empirical performance of the proposed method for a first order ODE.

ODE To limit scope we consider first order ODEs of the form

$$\frac{dy}{dx} = F\left(\frac{y(x)}{x}\right), \quad x \in [1, x_T], \quad y(1) = y_0. \quad (21)$$

Note that admitted transformation and associated canonical coordinates for this class of ODE have already been derived in Ex. 9, Ex. 11 and Ex. 12.

Prior In constructing a prior $\mu \in \mathcal{P}_Y$ we refer to the implicit prior principle in Section 3.2. Indeed, recall from Ex. 11 that the ODE in (21) can be transformed into an ODE of the form

$$\frac{ds}{dr} = G(r), \quad r \in (0, \infty), \quad s(y_0) = \log(y_0).$$

Then our approach constructs a distribution $\nu \in \mathcal{P}_S$ where, from Ex. 12, \mathcal{S} is the set of differentiable functions s defined on $r \in (0, \infty)$ and satisfying

$$\log(r) \leq s(r) \leq \log(r) + \log(x_T). \quad (22)$$

To ensure monotonicity in the implicit prior principle, we take $\frac{ds}{dr} > 0$, which translates into the requirement that

$$\frac{ds}{dr} > \frac{1}{r}. \quad (23)$$

If (23) holds, then ν induces a well-defined distribution $\mu \in \mathcal{P}_Y$. Note that the constraints in (22) and (23) preclude the direct use of standard prior models, such as Gaussian




Figure 4: Experimental results, first order ODE: The black curves represent samples from the posterior, whilst the exact solution is indicated in red. The blue curves represent a constraint on the (r, s) domain that arises when the implicit prior principle is applied. The number n of gradient evaluations is indicated. Top: results in the (r, s) domain. Bottom: results in the (x, y) domain.

processes. However, it is nevertheless possible to design priors that are convenient for a given set of canonical coordinates. Indeed, for the canonical coordinates r, s in our example, we can consider a prior of the form

$$s(r) = \log(r) + \log(x_T)\zeta(r),$$

where the function $\zeta : (0, \infty) \rightarrow \mathbb{R}$ satisfies

$$\zeta(y_0) = 0, \quad \zeta(r) \leq 1, \quad \frac{d\zeta}{dr} \geq 0. \tag{24}$$

For this experiment, the approach of López-Lopera et al. (2018) was used as a prior model for the monotone, bounded function ζ ; this requires that a number, N , of basis functions is specified – for brevity we defer the detail to Appendix A.3.

The prior just described incorporates the symmetric structure of the ODE, in the sense that the independent variable $r = \frac{y}{x}$ is the first canonical coordinate of the infinitesimal generator of the Lie group of transformations of the original ODE in 11. In other words, r is a variable fixed by the Lie group of transformations of the ODE (in this case $x^* = \alpha x$, $y^* = \alpha y$, so $r^* = r$). Because the prior is defined on functions $s(r)$ of r , this means the prior itself is also unchanged by the Lie group of transformations of the ODE, so that the prior effectively incorporates the symmetric structure of the ODE.

Results To obtain empirical results we consider the ODE with $F(r) = r^{-1} + r$ and $y_0 = 1$, $x_T = 5$. The posterior distributions that were obtained as the number n of data points was increased were sampled and plotted in the (r, s) and (x, y) planes in Figure 4. In each case a basis of size $N = 2n$ was used. Observe that the implicit prior principle ensures that all curves in the (x, y) plane are well-defined functions (i.e. there is at most one y value for each x value). Observe also that the posterior mass appears to contract to the true solution y^\dagger of the ODE as the number of evaluations n of the gradient field is increased.

4.2 A Second Order ODE

This section illustrates the empirical performance of the proposed method for a second order ODE.

ODE Consider again the second order nonlinear ODE in (14) together with the initial condition $y(x_0) = y_0$, $\frac{dy}{dx}(x_0) = y'_0$.

Prior It is shown in Ex. 14 in the Supplement that 14 can be reduced to a first order ODE in (s, r) with $-\frac{1}{x_0} - r \leq s \leq -\frac{1}{x_T} - r$. The implicit prior principle in this case requires that $\frac{ds}{dr} > -1$. Thus we are led to consider a parametrisation of the form

$$s(r) = -\frac{1}{x_0} - r + \left(\frac{1}{x_0} - \frac{1}{x_T} \right) \zeta(r),$$

where the function ζ again satisfies the conditions in (24). The approach of López-Lopera et al. (2018) was therefore again used as a prior model.

For this example an additional level of analytic tractability is possible, as described in detail in Ex. 14 in the Supplement. Thus we need only consider an information operator of the form $A(\mathbf{y}) = [G(r_0), \dots, G(r_n)]$.

Results The posterior distributions that were obtained are plotted in the (r, s) plane and the (x, y) plane in Figure 5. A basis of size $N = 2n$ was used, with $[y_0, y'_0] = [-10, 1]$, $[x_0, x_T] = [5, 10]$. Observe that the implicit prior principle ensures that all curves in the (x, y) plane are well-defined functions (i.e. there is at most one y value for each x value). The true solution appears to be smoother than the samples, even for 50




Figure 5: Experimental results, second order ODE: The black curves represent samples from the posterior in the (r, s) plane (left) and (x, y) plane (right), whilst the exact solution is indicated in red. The blue curves represent a constraint on the domain that arises when the implicit prior principle is applied. The number of gradient evaluations was $n = 50$.

gradient evaluations, which suggests that the prior was somewhat conservative in this context.

5 Conclusion

This paper presented a foundational perspective on PNM. It was first argued that there did not exist a Bayesian PNM for the numerical solution of ODEs. Then, to address this gap, a prototypical Bayesian PNM was developed. The Bayesian perspective that we have put forward sheds light on foundational issues which will need to be addressed going forward:

Foundation of PNM As explained in Section 1.2, existing PNM for ODEs each take the underlying state space \mathcal{Y} to be the solution space of the ODE. This appears to be problematic, in the sense that a generic evaluation $f(x_i, y_i)$ of the gradient field cannot be cast as information $A(y^\dagger)$ about the solution y^\dagger of the ODE unless the point (x_i, y_i) lies exactly on the solution curve $\{(x, y^\dagger(x)) : x \in [x_0, x_T]\}$. As a consequence, all existing PNM of which we are aware violate the likelihood principle and are therefore not strictly Bayesian. The assumption of a solvable Lie algebra, used in this work, can be seen as a mechanism to ensure the existence of an exact information operator A , so that the likelihood principle can be obeyed. However, for a general ODE it might be more natural to take the underlying state space to be a set \mathcal{F} of permitted gradient fields and the quantity of interest $Q(f)$ to map a gradient field f to the solution of the associated ODE. This would make the information operator A trivial but evaluation of the push-forward $Q_{\#}\mu^a$ would require the exact solution operator of the ODE. However, the reliance on access to an oracle solver Q makes this philosophically somewhat distinct from PNM.

Limitations of Bayesian PNM The proposed method was intended as a proof-of-concept and it is therefore useful to highlight the aspects in which it is limited. First, when an m th order ODE admits an r -parameter Lie group of transformations with $r > m$, there is an arbitrariness to the particular m -dimensional sub-group of transformations that are selected. Second, the route to obtain transformations admitted by the ODE demands that some aspects of the gradient field f are known, in contrast to other work in which f is treated as a black-box. For instance, in Ex. 11 we used the fact that f can be expressed as $f(x, y) = F(\frac{y}{x})$, although knowledge of the form of F was not required. Third, the class of ODEs for which a solvable Lie algebra is admitted is relatively small. On the other hand, references such as Bluman and Anco (2002) document important cases where our method could be applied. Fourth, the principles for prior construction that we identified do not entail a unique prior and, as such, the question of prior elicitation must still be addressed.

Outlook The goal of providing rigorous and exact statistical uncertainty quantification for the solution of an ODE is, we believe, important and will continue to be addressed. Traditional numerical methods have benefitted from a century of research effort and, in comparison, Bayesian PNM is an under-developed field. For example, the limited existing work on PNM for ODEs, such as Skilling (1992); Schober et al. (2014); Chkrebtii et al. (2016); Kersting and Hennig (2016); Schober et al. (2019); Kersting et al. (2018); Tronarp et al. (2019), does not attempt to provide adaptive error control (though we note promising ongoing research in that direction by Chkrebtii and Campbell, 2019). Nevertheless, the case for developing Bayesian numerical methods – which shares some parallels with the case for Bayesian statistics as opposed to other inferential paradigms – is clear, as argued in Diaconis (1988) and Hennig et al. (2015). The insights we have provided in this paper serve to highlight the foundational issues pertinent to Bayesian PNM for ODEs. Indeed, our proof-of-concept highlights that performing exact Bayesian inference for ODEs may be extremely difficult. This in turn provides motivation for the continued development of ‘approximately Bayesian’ approaches to PNM, which in Section 1.2 we surveyed in detail.

Supplementary Material

Supplementary material (DOI: [10.1214/19-BA1183SUPP](https://doi.org/10.1214/19-BA1183SUPP); .pdf).

References

- Abdulle, A. and Garegnani, G. (2018). “Random time step probabilistic methods for uncertainty quantification in chaotic and geometric numerical integration.” arXiv:1801.01340. 1064
- Baumann, G. (2013). *Symmetry Analysis of Differential Equations with Mathematica®*. Springer Science & Business Media. MR1753536. doi: <https://doi.org/10.1007/978-1-4612-2110-4>. 1073, 1074

- Bluman, G. and Anco, S. (2002). *Symmetry and Integration Methods for Differential Equations*. Springer. MR1914342. 1067, 1068, 1070, 1072, 1073, 1082
- Briol, F.-X., Oates, C. J., Girolami, M., Osborne, M. A., and Sejdinovic, D. (2019). “Probabilistic integration: A role in statistical computation? (with discussion and rejoinder).” *Statistical Science*, 34(1): 1–42. MR3938962. doi: <https://doi.org/10.1214/18-STS683>. 1075
- Chabiniok, R., Wang, V. Y., Hadjicharalambous, M., Asner, L., Lee, J., Sermesant, M., Kuhl, E., Young, A. A., Moireau, P., Nash, M. P., et al. (2016). “Multiphysics and multiscale modelling, data–model fusion and integration of organ physiology in the clinic: ventricular cardiac mechanics.” *Interface Focus*, 6(2): 20150083. 1058
- Chang, J. T. and Pollard, D. (1997). “Conditioning as disintegration.” *Statistica Neerlandica*, 51(3): 287–317. MR1484954. doi: <https://doi.org/10.1111/1467-9574.00056>. 1059
- Chkrebtii, O. and Campbell, D. (2019). “Adaptive step-size selection for state-space based probabilistic differential equation solvers.” *Statistics and Computing*. To appear. doi: <https://doi.org/10.1007/s11222-019-09899-5>. 1082
- Chkrebtii, O., Campbell, D. A., Girolami, M. A., and Calderhead, B. (2016). “Bayesian uncertainty quantification for differential equations (with discussion).” *Bayesian Analysis*, 11(4): 1239–1267. MR3577378. doi: <https://doi.org/10.1214/16-BA1017>. 1063, 1082
- Cockayne, J., Oates, C., Sullivan, T., and Girolami, M. (2019). “Bayesian probabilistic numerical methods.” *SIAM Review*. To appear. 1059, 1060, 1065, 1066
- Conrad, P. R., Girolami, M., Särkkä, S., Stuart, A., and Zygalakis, K. (2017). “Statistical analysis of differential equations: introducing probability measures on numerical solutions.” *Statistics and Computing*, 27(4): 1065–1082. MR3627563. doi: <https://doi.org/10.1007/s11222-016-9671-0>. 1064
- Diaconis, P. (1988). “Bayesian numerical analysis.” *Statistical Decision Theory and Related Topics IV(1)*, 1988. MR0927099. 1082
- Estep, D. (1995). “A posteriori error bounds and global error control for approximation of ordinary differential equations.” *SIAM Journal on Numerical Analysis*, 32(1): 1–48. MR1313704. doi: <https://doi.org/10.1137/0732001>. 1058
- Gelman, A. and Shalizi, C. R. (2013). “Philosophy and the practice of Bayesian statistics.” *British Journal of Mathematical and Statistical Psychology*, 66(1): 8–38. MR3044854. doi: <https://doi.org/10.1111/j.2044-8317.2011.02037.x>. 1060
- Hawkins, T. (2012). *Emergence of the theory of Lie groups: An essay in the history of mathematics 1869–1926*. Springer Science & Business Media. MR1771134. doi: <https://doi.org/10.1007/978-1-4612-1202-7>. 1066
- Hennig, P., Osborne, M. A., and Girolami, M. (2015). “Probabilistic numerics and uncertainty in computations.” *Proceedings of the Royal Society A*, 471(2179): 20150142. MR3378744. doi: <https://doi.org/10.1098/rspa.2015.0142>. 1058, 1082

- Higham, N. J. (2002). *Accuracy and Stability of Numerical Algorithms*. SIAM. MR1927606. doi: <https://doi.org/10.1137/1.9780898718027>. 1064
- Karvonen, T., Oates, C. J., and Särkkä, S. (2018). “A Bayes-Sard cubature method.” In *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS2018)*, 5882–5893. 1075
- Kersting, H. and Hennig, P. (2016). “Active uncertainty calibration in Bayesian ODE solvers.” In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI 2016)*, 309–318. 1063, 1082
- Kersting, H., Sullivan, T., and Hennig, P. (2018). “Convergence rates of Gaussian ODE filters.” arXiv:1807.09737. 1082
- Larkin, F. (1972). “Gaussian measure in Hilbert space and applications in numerical analysis.” *The Rocky Mountain Journal of Mathematics*, 2: 379–421. MR0303193. doi: <https://doi.org/10.1216/RMJ-1972-2-3-379>. 1058
- Larkin, F. M. (1974). “Probabilistic error estimates in spline interpolation and quadrature.” In *Information Processing 74 (Proc. IFIP Congress, Stockholm, 1974)*, 605–609. MR0408208. 1058
- Lie, H., Stuart, A., and Sullivan, T. (2019). “Strong convergence rates of probabilistic integrators for ordinary differential equations.” *Statistics and Computing*. To appear. doi: <https://doi.org/10.1007/s11222-019-09898-6>. 1064
- Lie, H. C., Sullivan, T., and Teckentrup, A. L. (2018). “Random forward models and log-likelihoods in Bayesian inverse problems.” *SIAM/ASA Journal on Uncertainty Quantification*, 6(4): 1600–1629. MR3876869. doi: <https://doi.org/10.1137/18M1166523>. 1064
- López-Lopera, A. F., Bachoc, F., Durrande, N., and Roustant, O. (2018). “Finite-dimensional Gaussian approximation with linear inequality constraints.” *SIAM/ASA Journal of Uncertainty Quantification*, 6(3): 1224–1255. MR3857898. doi: <https://doi.org/10.1137/17M1153157>. 1079, 1080
- Oates, C. J. and Sullivan, T. J. (2019). “A modern retrospective on probabilistic numerics.” *Statistics and Computing*. To appear. doi: <https://doi.org/10.1007/s11222-019-09902-z>. 1058
- Perilla, J. R., Goh, B. C., Cassidy, C. K., Liu, B., Bernardi, R. C., Rudack, T., Yu, H., Wu, Z., and Schulten, K. (2015). “Molecular dynamics simulations of large macromolecular complexes.” *Current Opinion in Structural Biology*, 31: 64–74. 1058
- Schober, M., Duvenaud, D. K., and Hennig, P. (2014). “Probabilistic ODE solvers with Runge-Kutta means.” In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS 2014)*, 739–747. 1062, 1063, 1082
- Schober, M., Särkkä, S., and Hennig, P. (2019). “A probabilistic model for the numerical solution of initial value problems.” *Statistics and Computing*, 29: 99–122. MR3905543. doi: <https://doi.org/10.1007/s11222-017-9798-7>. 1062, 1063, 1082

- Skilling, J. (1992). “Bayesian solution of ordinary differential equations.” In *Maximum Entropy and Bayesian Methods*, 23–37. Springer. 1060, 1061, 1062, 1065, 1082
- Teymur, O., Lie, H. C., Sullivan, T., and Calderhead, B. (2018). “Implicit Probabilistic Integrators for ODEs.” In *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS 2018)*. 1062, 1063
- Teymur, O., Zygalakis, K., and Calderhead, B. (2016). “Probabilistic linear multistep methods.” In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)*, 4321–4328. 1062, 1063
- Traub, J. and Woźniakowski (1992). “Perspectives on Information-Based Complexity.” *Bulletin of the American Mathematical Society*, 26(1): 29–52. MR1102756. doi: <https://doi.org/10.1090/S0273-0979-1992-00240-9>. 1061
- Tronarp, F., Kersting, H., Särkkä, S., and Hennig, P. (2019). “Probabilistic solutions to ordinary differential equations as non-linear Bayesian filtering: A new perspective.” *Statistics and Computing*. To appear. doi: <https://doi.org/10.1007/s11222-019-09900-1>. 1065, 1082
- Wang, J., Cockayne, J., and Oates, C. J. (2019). “Supplementary material.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/19-BA1183SUPP>. 1069
- Wedi, N. P. (2014). “Increasing horizontal resolution in numerical weather prediction and climate simulations: illusion or panacea?” *Philosophical Transactions of the Royal Society A*, 372(2018): 20130289. MR3209207. doi: <https://doi.org/10.1098/rsta.2013.0289>. 1058

Acknowledgments

The authors are grateful to Mark Craddock, François-Xavier Briol and Tim Sullivan for discussion of this work, as well as to an Associate Editor and two Reviewers for their challenging but constructive feedback. JW was supported by the EPSRC Centre for Doctoral Training in Cloud Computing for Big Data at Newcastle University, UK. CJO was supported by the Lloyd’s Register Foundation programme on data-centric engineering at the Alan Turing Institute, UK. This material was based upon work partially supported by the National Science Foundation under Grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.