



## Editorial: special edition on probabilistic numerics

M. Girolami<sup>1,2</sup> · I. C. F. Ipsen<sup>3</sup> · C. J. Oates<sup>2,4</sup> · A. B. Owen<sup>5</sup> · T. J. Sullivan<sup>6,7</sup>

Published online: 10 September 2019

© Springer Science+Business Media, LLC, part of Springer Nature 2019

This special edition of *Statistics and Computing* is dedicated to the emerging field of *Probabilistic Numerics*, at the interface of statistical inference and numerical analysis, and accompanies the workshop on Probabilistic Numerics held in London, 11–13 April 2018.

As traditionally understood, numerical analysis concerns itself with the design and analysis of methods for the (approximate) solution of deterministic well-posed problems such as quadrature/cubature or the solution of a differential equation. Such a numerical method can be cast, at least conceptually, as a statistical estimator for its associated quantity of interest. Often such numerical methods are accompanied by error estimates, or are designed in such a way that a user-specified tolerance is hoped to be met; these can equally be cast, at a high conceptual level, as a kind of confidence interval. Probabilistic numerics seeks to better understand and make

rigorous the use of formal statistical techniques in the numerical analysis context.

The use of statistical techniques to better characterise and understand numerical algorithms has received considerable attention in recent years. A particular success story is the widespread use of Bayesian optimisation in the training of regression and classification models in machine learning, while to a lesser extent statistical approaches to numerical cubature and the numerical solution of differential equations have been deployed in the applied context. However, there remain substantial research questions that must be addressed before the field can reach its full potential. First, these methods can involve certain modelling choices that must be user-specified (e.g. in a Bayesian context, the choice of a prior) and further research is required to assist or automate probabilistic numerical methods in this important respect. Second, until a unified statistical framework is developed that encompasses a suite of probabilistic numerical methods, it is unclear how such methods could be directly combined within the context of a complex computer simulation.

This special issue accompanies the *2018 Workshop on Probabilistic Numerics* held on 11–13 April 2018, at the Alan Turing Institute in London. This workshop brought together over thirty researchers from seven countries for three days of exciting presentations and intensive collaboration on probabilistic numerics, culminating in the panel discussion to be found online at <https://youtu.be/-NSoSCwwo-s>

The workshop was made possible by the generous financial support of the Lloyd's Register Foundation Programme on Data-Centric Engineering and the National Science Foundation, USA, under Grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute, as well as the organisational and logistical support of the Alan Turing Institute itself.

Four contributions in this special issue consider the probabilistic approach to cubature:

In *Fast Automatic Bayesian Cubature Using Lattice Sampling*, Jagadeeswaran and Hickernell consider Bayesian cubature in an adaptive setting in which scale and regularity parameters for the integrand are estimated from integrand

✉ C. J. Oates  
chris.oates@ncl.ac.uk; coates@turing.ac.uk

M. Girolami  
mgirolami@turing.ac.uk

I. C. F. Ipsen  
ipsen@ncsu.edu

A. B. Owen  
owen@stanford.edu

T. J. Sullivan  
t.j.sullivan@fu-berlin.de; sullivan@zib.de

<sup>1</sup> Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK

<sup>2</sup> The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, UK

<sup>3</sup> Department of Mathematics, North Carolina State University, Raleigh, NC 27695-8205, USA

<sup>4</sup> Newcastle University, Herschel Building, Newcastle upon Tyne NE1 7RU, UK

<sup>5</sup> Department of Statistics, Sequoia Hall, 390 Serra Mall, Stanford University, Stanford, CA 94305, USA

<sup>6</sup> Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany

<sup>7</sup> Zuse Institute Berlin, Takustraße 7, 14195 Berlin, Germany

values or are given non-informative priors. They pair low-discrepancy nodes with matching kernels that lower the computational cost to near-linear in the number of integrand evaluations.

The article *On the Positivity and Magnitudes of Bayesian Quadrature Weights*, by Karvonen, Kanagawa, and Särkkä, establishes novel theoretical results that imply stability and robustness of the Bayesian cubature method when the integration nodes are selected to minimise the posterior variance of the integral.

Ehler, Gräf, and Oates, in *Optimal Monte Carlo Integration on Closed Manifolds*, discuss Bayesian cubature based on randomly sampled integration nodes, establishing worst-case optimality of the posterior mean estimator when the integrand belongs to a Sobolev space on a closed Riemannian manifold.

In the paper *Symmetry Exploits for Bayesian Cubature Methods*, Karvonen, Särkkä, and Oates extend their previous work on Bayes–Sard cubature to take advantage of so-called “full symmetry” in the nodal set, kernel, and measure. This approach reduces the computational cost from cubic in the number of design points to cubic in the number of fully symmetric partitions of these points, which can lead to vast speedups.

Three papers in this issue consider the analysis and development of probabilistic numerical methods for ordinary differential equations (ODEs):

*Strong Convergence Rates of Probabilistic Integrators for Ordinary Differential Equations* by Lie, Stuart, and Sullivan, considers perturbation-based probabilistic solvers for ODE initial value problems. The authors extend the existing global convergence analysis for these randomised numerical methods to allow for non-Gaussian and non-independent perturbations to a deterministic numerical method, while significantly strengthening the mode of convergence so that the results can be used to substantiate accuracy estimates for any Bayesian inverse problems that use such noisy solvers.

*Adaptive Step-Size Selection for State-Space Based Probabilistic Differential Equation Solvers* by Chkrebtii and Campbell considers statistical principles for the adaptive choice of time step length in the numerical solution of an ordinary differential equation. An information-theoretic criterion (the Kullback–Leibler divergence) is used to sequentially derive an optimal sequence of discrete times at which the ODE is numerically solved; this exposes how adaptive behaviour (i.e. shorter time steps in regions of quickly changing dynamics, longer time steps where the dynamics are more sedate) can naturally arise from a purely statistical treatment.

In *Probabilistic Solutions to Ordinary Differential Equations as Nonlinear Bayesian Filtering: A New Perspective*, Tronarp, Kersting, Särkkä, and Hennig draw a connection between the numerical solution of ODEs and the extensive literature on nonlinear filtering. The authors demonstrate how

many well-established methods from the signal processing community can be brought to bear on the numerical solution of ODEs.

One paper in this issue considers the problem of signal recovery from noisy data, within a linear functional analytic context:

*Denoising by Thresholding Operator-Adapted Wavelets* by Yoo and Owhadi addresses the problem of reconstructing an unknown smooth function  $u$  from noisy data  $u + \zeta$  with prior information on the regularity of  $\mathcal{L}u$ , where  $\mathcal{L}$  is a linear operator such as a partial differential operator or a graph Laplacian. They show that the approximation of  $u$  obtained by thresholding the gamblet (operator-adapted wavelet) coefficients of  $u + \zeta$  is nearly minimax optimal (up to a multiplicative constant), and with high probability, its operator energy norm is bounded by that of  $u$  up to a constant depending on the amplitude of the noise. The proposed method is of near-linear computational complexity and can be generalised to non-homogeneous noise.

One paper in this issue provides a probabilistic perspective on (finite-dimensional) linear algebra:

The probabilistic solution of systems of linear equations is the topic of *Probabilistic Linear Solvers: A Unifying View* by Bartels, Cockayne, Ipsen, and Hennig. This paper brings together several recently proposed Bayesian probabilistic interpretations of numerical algorithms for solving linear systems. In this work, surprisingly general conditions for equivalence of these disparate methods are presented, as well as connections between probabilistic linear solvers and projection methods for linear systems, providing a probabilistic interpretation of a far more general class of iterative methods (e.g. GMRES), and probabilistic view on preconditioning.

As just described, this special issue contains a broad selection of the problems being addressed and the techniques being developed in this nascent field. The final paper in this issue, *A Modern Retrospective on Probabilistic Numerics* by Oates and Sullivan, provides a broader perspective on the field within the context of the historical development of probabilistic numerics, highlighting in particular the oft-overlooked contributions of Sul'din and Larkin in the 1960s, as well as linking this history to modern developments and prospects for the future of the field.

This special issue captures a snapshot of the state of probabilistic numerics as of 2018–2019, with the selected contributions having a particularly mathematical–statistical flavour. Considering also the broader recent literature, it is clear to us, the editors, that the field is advancing on a wide front in two senses: many numerical tasks are now receiving probabilistic treatments and contributions are coming from a broad community of statisticians, mathematicians, machine learners and computer scientists. The depth of the field is also improving, with works ranging from foundational theory to practical and even industrial applications.

That said, there is also much room for growth, especially in terms of “killer applications” of probabilistic numerics to the full spectrum of numerical tasks (beyond the current success story of Bayesian optimisation) in industrial settings.

The editors wish to express their sincere thanks to the anonymous peer reviewers, without whose expert input this special issue would not have been possible.

Mark Girolami, Cambridge  
Ilse Ipsen, Raleigh  
Chris Oates, Newcastle  
Art Owen, Stanford  
Tim Sullivan, Berlin

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