Statistical Science
2024, Vol. 39, No. 1, 1–2
https://doi.org/10.1214/23-STS920

Editorial: Bayesian Computations in the 21st Century

Christian P. Robert and Dennis Prangle

This special section of *Statistical Science* is devoted to the future of Bayesian computational statistics, from several perspectives. It involves a large group of researchers who contributed to collective articles, bringing their own perspectives and research interests into these surveys. Somewhat paradoxically, it starts with the past—and a conference on a Gold Coast beach. Martin, Frazier, and Robert first submitted a survey on the history of Bayesian computation, written after Gael Martin delivered a plenary lecture at Bayes on the Beach, a conference held in November 2017 in Surfers Paradise, Gold Coast, Queensland, and organised by Bayesian Research and Applications Group (BRAG), the Bayesian research group headed by Kerrie Mengersen at the Queensland University of Technology (QUT). Following a first round of reviews, this paper was split into two separate articles, *Computing Bayes: From Then 'Til Now*, retracing some of the history of Bayesian computation, and *Approximating Bayes in the 21st Century*, which is both a survey and a prospective on the directions and trends of approximate Bayesian approaches (and not solely ABC). At this point, Sonia Petrone, Editor of *Statistical Science*, suggested we had a special issue on the whole issue of trends of interest and promise for Bayesian computational statistics. Joining forces, after some delays and failures to convince others to engage, or to produce multilevel papers with distinct vignettes, we eventually put together an additional four papers, where lead authors gathered further co-authors to produce this diverse picture of some incoming advances in the field. We have deliberated avoided topics which have excellent recent reviews—such as Stein's method [1], sequential Monte Carlo [2], piecewise deterministic Markov processes [3]—and topics which are still in their infancy, such as the relationship of Bayesian approaches to large language models (LLMs) and foundation models.

Within this issue, *Past, Present, and Future of Software for Bayesian Inference* from Erik Štrumbelj et al covers the state of the art in the most popular Bayesian software, reminding us of the massive impact BUGS has had on the adoption of Bayesian tools since its early introduction in the early 1990s (which one of us remember discovering at the Fourth Valencia meeting on Bayesian statistics in April 1991). This paper makes an interesting distinction between first and second generations, and a light foray of the potential third generation, maybe missing the role of LLMs in coding that are already impacting the approach to computing and the less immediate revolution brought by quantum computing. Winter et al.'s *Emerging Directions in Bayesian Computation* makes a link with machine learning techniques—albeit without looking at the scariest issue of how Bayesian inference can survive in a machine learning world! It produces an additional excursion into the blurry division between proper sampling (à la MCMC) and approximations—complementing the historical viewpoint of Martin et al. elsewhere in this issue—and articulates these aspects within a (deep) machine learning perspective, emphasizing the role of summaries produced by generative models exploiting the power of neural network computation/optimization. It emphasises the pivotal reliance on variational Bayes, which is the most active common denominator with machine learning, and covers further topics like distributed computing, opening on the important aspect of data protection and guaranteed privacy. We particularly like the clinical presentation of this paper with attention to automation and limitations. Normalizing flows actually link this paper with Heng, Bortoli and Doucet's *Diffusion Schrödinger Bridges for Bayesian Computation*, which is more a focussed review of recent advances on possibly the next generation of posterior samplers. The final paper, *Modern Bayesian experimental design* by Rainforth et al., provides a most convincing application of the methods exposed in the earlier papers in that the field of Bayesian experimental design has hugely benefited from the occurrence of such tools to become a prevalent way of designing statistical experiments in real settings.

We feel that this is an exciting time for Bayesian computing! The Monte Carlo revolution of the 1990s continues to be a huge influence on today's work, and now is complemented by a diverse range of new directions informed by modern machine learning.

Christian Robert is a Professor in Statistics, Université Paris–Dauphine, 75775 Paris Cedex 16, France (e-mail: xian@ceremade.dauphine.fr). Dennis Prangle is a Senior Lecturer in Statistics, University of Bristol, School of Mathematics, Woodland Rd, Bristol BS8 1UG, UK (e-mail: dennis.prangle@bristol.ac.uk).

ACKNOWLEDGMENTS

The authors would like to thank the Editor Sonia Petrone for her support throughout the construction of this