

Scalable Algorithms for Bayesian Inference of Large-Scale Models from Large-Scale Data[☆]

Omar Ghattas¹, Tobin Isaac², Noémi Petra³, and Georg Stadler⁴

¹ Institute for Computational Engineering & Sciences and Departments of Geological Sciences and Mechanical Engineering, The University of Texas at Austin

² Computation Institute, University of Chicago

³ School of Natural Sciences, University of California, Merced

⁴ Courant Institute for Mathematical Sciences, New York University

One of the greatest challenges in computational science and engineering today is how to combine complex data with complex models to create better predictions. This challenge cuts across every application area within CS&E, from geosciences, materials, chemical systems, biological systems, and astrophysics to engineered systems in aerospace, transportation, structures, electronics, biomedicine, and beyond. Many of these systems are characterized by complex nonlinear behavior coupling multiple physical processes over a wide range of length and time scales. Mathematical and computational models of these systems often contain numerous uncertain parameters, making high-reliability predictive modeling a challenge. Rapidly expanding volumes of observational data—along with tremendous increases in HPC capability—present opportunities to reduce these uncertainties via solution of large-scale inverse problems.

In an inverse problem, we infer unknown model parameters (e.g., coefficients, material properties, source terms, initial or boundary conditions, geometry, model structure) from observations of model outputs. The need to quantify the uncertainty in the solution of such inverse problems has attracted widespread attention in recent years. This can be carried out in a systematic manner by casting the inverse problem within the framework of Bayesian inference. In this framework, uncertain observations and uncertain models are combined with available prior knowledge to yield a probability density in the model parameters as the solution of the inverse problem, thereby providing a rational and systematic means of quantifying uncertainties in the inference of these parameters. The resulting uncertainties in model parameters are then propagated forward through models to yield predictions with associated uncertainty. Finally, given this capability to quantify uncertainties in inverse problems, one can determine the design of the observational system (e.g., location of sensors, nature of measured quantities) that maximizes the information gain from the observations (or minimizes the uncertainty in the inferred model or subsequent prediction). This

[☆] This work was supported by AFOSR grants FA9550-12-1-0484 and FA9550-09-1-0608, DARPA/ARO contract W911NF-15-2-0121, DOE grants DE-SC0010518, DE-SC0009286, DE-11018096, DE-SC0006656, DE-SC0002710, and DE-FG02-08ER25860, and NSF grants ACI-1550593, CBET-1508713, CBET-1507009, CMMI-1028889, and ARC-0941678. Computations were performed on supercomputers at TACC, ORNL, and LLNL. We gratefully acknowledge this support.

is the optimal experimental design (OED) problem, which wraps an optimization problem around the Bayesian inverse problem.

The Markov chain Monte Carlo (MCMC) method has emerged as the method of choice for solving Bayesian inverse problems. Unfortunately, when the forward model is large and complex (e.g., when the model takes the form of an expensive-to-solve system of partial differential equations), and when the parameters are high-dimensional (as results from discretization of an infinite dimensional field such as an initial condition or heterogeneous material property), solution of Bayesian inverse problems via conventional MCMC is intractable. Moreover, addressing the meta-question of how to optimally obtain experimental data for such problems via solution of an OED problem is completely out of the question.

However, a number of advances over the past decade have brought the goal of Bayesian inference of large-scale complex models from large-scale complex data much closer. First, improvements in scalable forward solvers for many classes of large-scale models have made feasible numerous evaluations of model outputs for differing inputs. Second, sustained growth in HPC capabilities has multiplied the effects of the advances in solvers. Third, the emergence of MCMC methods that exploit problem structure (e.g., curvature of the posterior probability) has radically improved the prospects of sampling posterior distributions for inverse problems governed by expensive models. And fourth, recent exponential expansions of observational capabilities have produced massive volumes of data from which inference of large computational models can be carried out.

To overcome the prohibitive nature of Bayesian methods for high-dimensional inverse problems governed by expensive-to-solve PDEs, we exploit the fact that, despite the large size of observational data, they typically provide only sparse information on model parameters. This implicit dimension reduction is provided by low rank approximations of the Hessian of the data misfit functional, which is typically a compact operator due to ill-posedness of the inverse problem. A low rank approximation of the Hessian can be extracted efficiently in a matrix-free manner (without forming the Hessian) by a Lanczos [8, 14] or randomized SVD [4, 5, 12, 15, 21] method, requiring a number of matrix-vector products that scales only with the rank of the Hessian, and not the parameter dimension. Moreover, the rank reflects how informative the data are, i.e., how many directions in parameter space are informed by the data. Finally, each Hessian-vector product can be computed using just a pair of linearized forward/adjoint PDE solves [4, 5, 8, 9, 12, 14–17, 21, 22].

We have applied the methodology described above (for exploiting the geometric structure of the posterior) to geophysical inverse problems arising in ice sheet flow, seismic wave propagation, mantle convection, atmospheric transport, poromechanics, and subsurface flow. We are able to substantially reduce the effective parameter dimension (often by three orders of magnitude) at a cost, measured in (linearized) forward/adjoint PDE solves, that is independent of both the parameter and data dimensions [4, 5, 8, 9, 12, 14, 15, 21].

For linearized Bayesian analysis of nonlinear inverse problems, the Hessian evaluated at the point in parameter space that maximizes the posterior (i.e., the

MAP point) completely characterizes the uncertainty in inferred parameters. One can build on this idea to solve optimal experimental design problems at a cost that also does not scale with the parameter or data dimensions [1–3]. For nonlinear Bayesian inverse problems, the Hessian varies from point to point. However the low rank Hessian approximation machinery described above can still be exploited to accelerate MCMC sampling, by serving as an inverse covariance approximation for a Gaussian proposal that is tailored to the local curvature of the posterior [14, 15] (this is known as the stochastic Newton method).

The most complex inverse problem for which we have carried out Bayesian inversion involves ice sheet flow [12, 15, 16, 22]. The flow of ice from polar ice sheets such as Antarctica and Greenland is the primary contributor to projected sea level rise in the 21st century. The ice is modeled as a creeping, viscous, incompressible, non-Newtonian, shear-thinning fluid, for which we have developed custom scalable parallel solvers [13, 18, 19] on adaptively refined forest-of-octree meshes [6, 7, 10, 11], the combination of which has scaled to hundreds of billions of unknowns on up to 1.6 million cores [4, 6, 18]. One of the main difficulties faced in modeling ice sheet flow is the unknown spatially-varying Robin boundary condition that describes the resistance to sliding at the base of the ice. Satellite observations of the surface ice flow velocity can be used to infer this uncertain basal boundary condition. We have solved this ill-posed inverse problem using the (linearized) Bayesian inference machinery described above, which allows us to infer not only the unknown basal sliding parameters, but also the associated uncertainty [12]. We have demonstrated that the number of required forward solves is independent of the parameter dimension, data dimension, and number of processor cores. The largest Bayesian inverse problem solved has over one million uncertain parameters.

References

1. A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas, A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0 -sparsification, *SIAM Journal on Scientific Computing*, 36(5):A2122–A2148, 2014.
2. A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas, A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems, *SIAM Journal on Scientific Computing*, 38(1):A243–A272, 2016.
3. A. Alexanderian, P. Gloor, and O. Ghattas, On Bayesian A- and D-optimal experimental designs in infinite dimensions, *Bayesian Analysis*, 11(3):671–695, 2016.
4. T. Bui-Thanh, C. Burstedde, O. Ghattas, J. Martin, G. Stadler, and L.C. Wilcox, Extreme-scale UQ for Bayesian inverse problems governed by PDEs, *Proceedings of IEEE/ACM SC12*, 2012.
5. T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems. Part I: The linearized case, with applications to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
6. C. Burstedde, O. Ghattas, M. Gurnis, T. Isaac, G. Stadler, T. Warburton, and L.C. Wilcox, Extreme-Scale AMR, *Proceedings of ACM/IEEE SC10*, 2010.

7. C. Burstedde, L.C. Wilcox, and O. Ghattas, p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011.
8. H.P. Flath, L.C. Wilcox, V. Akcelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas, Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations, *SIAM Journal on Scientific Computing*, 33(1):407–432, 2011.
9. M. Hesse and G. Stadler, Joint inversion in coupled quasistatic poroelasticity, *Journal of Geophysical Research: Solid Earth*, 119:1425–1445, 2014.
10. T. Isaac, C. Burstedde, and O. Ghattas, Low-Cost Parallel Algorithms for 2:1 Octree Balance, *International Parallel and Distributed Processing Symposium (IPDPS 12)*, 426–437, IEEE Computer Society, 2012.
11. T. Isaac, C. Burstedde, L.C. Wilcox, and O. Ghattas, Recursive algorithms for distributed forests of octrees, *SIAM Journal on Scientific Computing*, 37(5):C497–C531, 2015.
12. T. Isaac, N. Petra, G. Stadler, and O. Ghattas, Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet, *Journal of Computational Physics*, 296(1):348–368, 2015.
13. T. Isaac, G. Stadler, and O. Ghattas, Solution of nonlinear Stokes equations discretized by high-order finite elements on nonconforming and anisotropic meshes, with application to ice sheet dynamics, *SIAM Journal on Scientific Computing*, 37(6):B804–B833, 2015.
14. J. Martin, L.C. Wilcox, C. Burstedde, and O. Ghattas, A Stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion, *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.
15. N. Petra, J. Martin, G. Stadler, and O. Ghattas, A computational framework for infinite-dimensional Bayesian inverse problems. Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems, *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555, 2014.
16. N. Petra, H. Zhu, G. Stadler, T.J.R. Hughes, and O. Ghattas, An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model, *Journal of Glaciology*, 58(211):889903, 2012.
17. V. Ratnasmamy, G. Stadler, and M. Gurnis, Adjoint-based estimation of plate coupling in a non-linear mantle flow model: theory and examples, *Geophysical Journal International*, 202(2):768–786, 2015.
18. J. Rudi, A.C.I. Malossi, T. Isaac, G. Stadler, M. Gurnis, P.W.J. Staar, Y. Ineichen, C. Bekas, A. Curioni, and O. Ghattas, An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earths Mantle, *Proceedings of IEEE/ACM SC15*, 2015.
19. J. Rudi, G. Stadler, and O. Ghattas, Weighted BFBT Preconditioner for Stokes Flow Problems with Highly Heterogeneous Viscosity, submitted, 2016.
20. J. Worthen, G. Stadler, N. Petra, M. Gurnis, and O. Ghattas, Towards adjoint-based inversion for rheological parameters in nonlinear viscous mantle flow, *Physics of the Earth and Planetary Interiors*, 234:23–34, 2014.
21. H. Zhu, S. Li, S. Fomel, G. Stadler, and O. Ghattas, A Bayesian approach to estimate uncertainty for full waveform inversion using a priori information from depth migration, *Geophysics*, 81(5):R307–R323, 2016.
22. H. Zhu, N. Petra, G. Stadler, T. Isaac, T.J.R. Hughes, and O. Ghattas, Inversion of geothermal heat flux in a thermomechanically coupled nonlinear Stokes ice sheet model, *The Cryosphere*, 10:1477–1494, 2016.