

Inverse Problems and Uncertainty Quantification

By Marco Iglesias
and Andrew M. Stuart

Quantifying uncertainty in the solution of inverse problems is an exciting area of research in the mathematical sciences, one that raises significant challenges at the interfaces between analysis, computation, probability, and statistics. The reach in terms of applicability is enormous, with diverse problems arising in the physical, biological, and social sciences, such as weather prediction, epidemiology, and traffic flow.

Loosely speaking, inverse problems confront mathematical models with data so that we can deduce the inputs needed to run the models; knowledge of these inputs can then be used to make predictions, and even to devise control strategies based on the predictions. Both the models and the data are typically uncertain, as are the resulting deductions and predictions; as a consequence, any decisions or control strategies based on the predictions will be greatly improved if the uncertainty is made quantitative.

Bayesian Approach to Inverse Problems

A mathematical model of an experiment is a set of equations relating inputs u to outputs y . Inputs represent physical variables that can be adjusted before the experiment takes place; outputs represent quantities that can be measured as a result of the experiment. For the *forward problem* the mathematical model G is used to predict the outputs y of an experiment from given inputs u . For the *inverse problem* the mathematical model is used to make inferences about inputs u that would result in given measured outputs y [6,10]; in practice, many inverse problems are *ill-posed* in the sense of Hadamard: They fail to satisfy at least one of the criteria for well-posedness—existence, uniqueness, and continuous dependence of the solution on data. When the measured data is subject to noise, and/or when the mathematical model is imperfect, it is important to quantify the uncertainty inherent in any inferences and predictions made as part of the solution to the inverse problem. The Bayesian approach to inverse problems allows us to undertake this task in a principled fashion. When properly applied, this formulation of inverse problems can simultaneously regularize any ill-posedness present.

In the Bayesian approach the pair (u, y) is considered a random variable, and the solution of the inverse problem is the conditional probability distribution of the random variable u given y , denoted $u|y$. A formula for this conditional probability distribution is given by *Bayes' rule*, which states that

$$\mathbb{P}(u|y) \propto \mathbb{P}(y|u) \mathbb{P}(u).$$

In words: The *posterior* $\mathbb{P}(u|y)$, which is what we know about the unknown u given the data y , is proportional to the *likelihood* $\mathbb{P}(y|u)$, which measures how likely the observed data is for given inputs u , multiplied by the *prior* $\mathbb{P}(u)$, which describes our knowledge of the unknown prior to the acquisition of data. We describe a concrete application below, and the reader may wish to revisit our overarching mathematical presentation with that example in mind.

Bayes' rule, and its application, are illustrated schematically in Figure 1. The *prior* is denoted by the black dotted curve in the INPUT space, with MODEL G defining its likelihood; data y is then used to obtain the *posterior*, denoted by the red curve in the INPUT space.

What does this update in probability distributions, from the prior to the posterior, do for us? Our answer is found in the predictions of uncertainty in the QUANTITY OF INTEREST q : Without data (under the *prior*), we make the prediction denoted

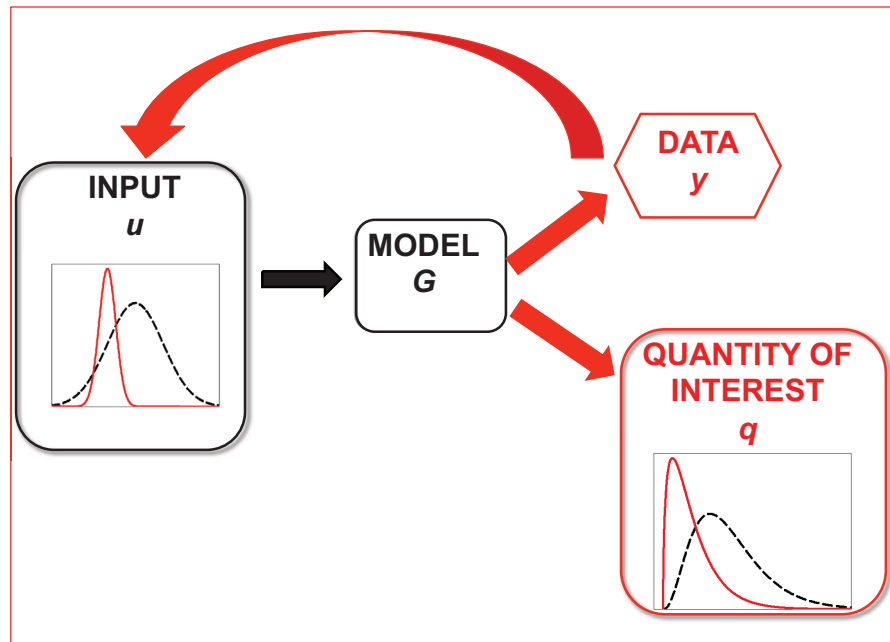


Figure 1. Uncertainty quantification in Bayesian inversion.

by the black dotted curve; with data (under the *posterior*), we can make a more refined prediction, denoted by the red curve.

Notice that all predictions are equipped with uncertainty. Furthermore, in this illustration, the Bayesian approach to the inverse problem has reduced the uncertainty, as manifest in the spread of the posterior probability distribution for the quantity of interest, reflecting the extra information obtained from the data. The reader will immediately see the benefits of this reduced uncertainty in any number of applications, including the examples of weather prediction, epidemiology, and traffic flow mentioned earlier, as well as in the increasingly vast number of application domains in which quantitative models and noisy data are available.

Practical Problem: CO₂ Capture and Storage

What, then, are the challenges for applied mathematics? To get some insight into where the challenges emerge, we consider the use of carbon capture and storage to facilitate global mitigation of the greenhouse effect [3]. Suppose, for example, that we are interested in assessing the economic viability and environmental impact of injecting CO₂ into the subsurface.

A typical CO₂ storage site could be a depleted oil/gas field or a deep saline aquifer (see Figure 2). The mathematical model in this case consists of the partial differential equations that describe the plume of injected CO₂ in the subsurface. Important inputs to the model are the permeability of the storage site, along with other geologic features, such as existing faults and fractures. Natural outputs comprise the measurements of bottom-hole pressure from the injection well and, possibly, of surface deformation from satellite data and GPS devices. Because the subsurface is not directly observable, the problem of inferring its properties (inputs) from measurements (outputs) is particularly important: With accurate inference, decisions can be made on the basis of a variety of quantities of interest, concerning both the safety and the financial feasibility of the storage site; for example, one might wish to assess the




Figure 2. Geologic storage of CO₂.

potential groundwater contamination due to leakage of CO₂.

The challenges inherent in this application become apparent when we consider what is hidden in the deceptively simple Bayes' law stated above. In this example, the likelihood itself is defined through solution of the forward model, which is a coupled set of conservation laws (PDEs) describing the physics of multi-phase flow in a porous medium. The probability distribution on the input space of permeabilities lives on a space of functions; in practice, this means that we will be representing probabilities on very high-dimensional spaces. To probe the posterior probability distribution, then, we need to take on the challenge of solving complex PDEs over an enormous space of input permeabilities. Similar daunting challenges emerge in a vast range of applications. Nonetheless, the last decade has seen considerable advances in the Bayesian approach to inverse problems, and the book [11] has played a major role in establishing the viability of the approach on today's computers.

Because our subject is in its infancy, with growing numbers of applications and a range of methodologies, considerable long-term challenges remain. A key modelling question, which will be very application-specific, concerns the choice of prior for the unknown. A key computational question concerns ways to probe the posterior distribution with sufficient accuracy that we can




Figure 3. Draws from various permeability priors: piecewise-constant layer model with fault (left), piecewise-continuous layer model (centre), and piecewise-continuous channel model (right).

compute the posterior probability distribution on quantities of interest. Furthermore, these modelling and computational questions interact.

The subject is in need of sustained input from applied mathematicians who can help to guide the development of algorithms, through analysis of their complexity and through computational innovation. This work needs to be done in the context of classes of application-specific prior models. Success in this area requires an appreciation for analysis (e.g., of PDE-based forward models), computation (e.g., high-dimensional integration), probability (e.g., in specifying random field pri-

ors), and statistics (e.g., in exploiting data in the design of algorithms to explore the posterior). In each case the work needs to be guided by application-specific modelling considerations.

Returning to the example of CO₂ storage, we consider priors for the subsurface permeability that deliver different "typical" functions, as displayed in Figure 3, left and centre. On the left, the permeability is piecewise-constant, and the unknown parameters define the position of the layers of different materials that constitute the subsurface, together with the position of the fault and the permeability values within each layer. In the centre image, the permeability is defined through the location of a single interface; the considerable variability above and below the interface is represented by a function, and not just a single value. Prior models for which these two different permeabilities are "typical" will be quite different and will lead to different computational considerations; moreover, if prior knowledge is scarce, the prior might need to incorporate both permeabilities as "typical," possibly along with those of other types, such as that displayed in Figure 3 (right). Details of the mathematical formulation of such problems, and references to the engineering literature in which these models were originally developed can be found in [9].

Once the prior and forward model are specified, the posterior is defined via Bayes' rule, and we then move on to the computational task of exploring the posterior and computing expectations with respect to it. Monte Carlo–Markov Chain, or MCMC, is a natural methodology for studying these problems, and the last decade has seen considerable progress in the theory and practice of these methods in high dimensions [5]. Nevertheless, vanilla Monte Carlo-based methods, whilst enormously flexible, are hampered by their $N^{-1/2}$ convergence rate [2], meaning that computational complexity (cost per unit error) can be rather excessive. As a consequence, we are likely to see the development of multi-level Monte Carlo [8] and quasi-Monte Carlo [2] methods in the context of inverse problems. In addition, the use of generalized polynomial chaos methodologies and their relatives [1,4,7,12] is likely to be transferred into the context of inverse problems, as exemplified in [13]; again, these methods hold the possibility of improving on the complexity of Monte

Carlo methods. Interested readers can find further details in the articles [14,15] and the references therein.

In this burgeoning field, the numerous opportunities for research in applied mathematics are driven by both the enormous numbers and types of applications, together with the wide range of areas in the mathematical sciences from which contributions are needed. The subject is at a tipping point, where computational power is starting to allow the exploration of quite complex Bayesian models, and the opportunity for impact is high. In summary, this is an excellent area for applied mathematicians who are looking for new research challenges.

Acknowledgments

The authors are grateful to EPSRC, ERC, and ONR for financial support that led to the research underpinning this article. AMS is PI on the EPSRC-funded Programme Grant

EQUIP: <http://www2.warwick.ac.uk/fac/sci/math/research/grants/equip/>.

References

- [1] I. Babuska, R. Tempone, and G. Zouraris, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM J. Numer. Anal., 42 (2004), 800–825.
- [2] R. Caflisch, *Monte Carlo and quasi-Monte Carlo methods*, Acta Numer., 7 (1998), 1–49.
- [3] A. Chadwick, *CCS: Between a rock and a hard place?*, Greenhouse Gases: Science and Technology, 2 (2011), 99–101.
- [4] A. Cohen, R. DeVore, and C. Schwab, *Convergence rates of best N -term Galerkin approximations for a class of elliptic SPDEs*, Found. Comput. Math., 10 (2010), 615–646.
- [5] S. Cotter, G. Roberts, A.M. Stuart, and D. White, *MCMC methods for functions: Modifying old algorithms to make them faster*, Stat. Sci., 28 (2013), 424–446.
- [6] H. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Kluwer, Philadelphia, 1996.
- [7] R.G. Ghanem and P.D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer, New York, 1991.
- [8] M. Giles, *Multilevel Monte Carlo path simulation*, Oper. Res., 56 (2008), 607–617.
- [9] M. Iglesias, K. Lin, and A.M. Stuart, *Well-posed Bayesian geometric inverse problems arising in subsurface flow*, arxiv.org/abs/1401.5571, to appear in Inverse Problems, special issue on Bayesian inversion, 2014.
- [10] J.B. Keller, *Inverse problems*, Amer. Math. Monthly, 83 (1976), 107–118.
- [11] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, 160, Applied Mathematical Sciences, Springer-Verlag, New York, 2005.
- [12] C. Schwab and C.J. Gittelsohn, *Sparse tensor discretizations of high-dimensional parametric and stochastic PDEs*, Acta Numer., 20 (2011), 291–467.
- [13] C. Schwab and A.M. Stuart, *Sparse deterministic approximation of Bayesian inverse problems*, Inverse Problems, 28 (2012); arxiv.org/abs/1103.4522.
- [14] A.M. Stuart, *Inverse problems: A Bayesian perspective*, Acta Numer., 19 (2010), 451–559.
- [15] A.M. Stuart, *Inverse Problems and Uncertainty Quantification*, invited lecture, ICM 2014, Seoul; http://homepages.warwick.ac.uk/~masdr/TALKS/stuart_SIAMUQ.pdf.

Marco Iglesias is an assistant professor in the School of Mathematical Sciences, Nottingham University, UK. Andrew Stuart is a professor of mathematics at the Mathematics Institute, Warwick University, UK. This article is based, in part, on a lecture delivered by Stuart at the 2014 SIAM Conference on Uncertainty Quantification, held in Savannah, Georgia.