The 8th European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS Congress 2022
5-9 June 2022, Oslo, Norway

Variational Bayesian Approximation of Inverse Problems using Sparse Precision Matrices

Jan Povala[1*], Ieva Kazlauskaite[2], Eky Febrianto[3], Fehmi Cirak[4] and Mark Girolami[5]

[1] Department of Mathematics, Imperial College London, London, SW7 2AZ, UK,
jan.povala@gmail.com
[2,3,4,5] Department of Engineering, University of Cambridge, Cambridge, CB2 1PZ, UK,
{ik394, evf27, fc286, mag92}@cam.ac.uk
[1,3,4,5] The Alan Turing Institute, London, NW1 2DB, UK

Keywords: *Inverse problems, Bayesian inference, variational Bayes, precision matrix, uncertainty quantification*

Inverse problems involving partial differential equations are widely used in science and engineering. Although such problems are generally ill-posed, different regularisation approaches have been developed to ameliorate this problem. Among them is the Bayesian formulation, where a prior probability measure is placed on the quantity of interest. The resulting posterior probability measure is usually analytically intractable. The Markov Chain Monte Carlo (MCMC) method has been the go-to method for sampling from those posterior measures. MCMC is computationally infeasible for large-scale problems that arise in engineering practice. Lately, Variational Bayes (VB) has been recognised as a more computationally tractable method for Bayesian inference, approximating a Bayesian posterior distribution with a simpler trial distribution by solving an optimisation problem. In this work, we argue, through an empirical assessment, that VB methods are a flexible and efficient alternative to MCMC for this class of problems. We propose a natural choice of a family of Gaussian trial distributions parametrised by precision matrices, thus taking advantage of the inherent sparsity of the inverse problem encoded in its finite element discretisation. We utilise stochastic optimisation to efficiently estimate the variational objective and assess not only the error in the solution mean but also the ability to quantify the uncertainty of the estimate. We test this on PDEs based on the Poisson equation in 1D and 2D. A Tensorflow implementation is made publicly available on GitHub.