

ON THE COMPUTATIONAL COMPLEXITY OF HIGH-DIMENSIONAL MCMC IN NON-LINEAR
INVERSE PROBLEMS

Richard Nickl

University of Cambridge, UK
Nickl@maths.cam.ac.uk

We discuss recent results, both positive and negative, about the run-time of Markov chain Monte Carlo (MCMC) algorithms that target posterior distributions arising from high-dimensional non-linear statistical regression models with Gaussian process priors. Prototypical applications include inverse problems with partial differential equations (PDEs). We show that cold-start local MCMC may not work (have ‘exponential in dimension’ runtime) even for target measures that are radially strictly decreasing away from their unique mode, but that warm-start Langevin type MCMC can achieve polynomial runtime for posterior computation under certain ‘gradient stability’ conditions that can be verified in a large class of relevant PDE models.

Joint work with Sven Wang (MIT), and for the lower bounds further with Afonso Bandeira (ETH Zürich), Antoine Maillard (ETH Zürich)..