EGU22-3691
https://doi.org/10.5194/egusphere-egu22-3691
EGU General Assembly 2022

Hydrogeological inference by adaptive sequential Monte Carlo with geostatistical resampling model proposals

Macarena Amaya[1], Niklas Linde[1], and Eric Laloy[2]
[1]Institute of Earth Sciences, University of Lausanne, Lausanne, Switzerland
[2]Institute for Environment, Health and Safety, Belgian Nuclear Research Centre, Mol, Belgium

For strongly non-linear inverse problems, Markov chain Monte Carlo (MCMC) methods may fail to properly explore the posterior probability density function (PDF). Particle methods are very well suited for parallelization and offer an alternative approach whereby the posterior PDF is approximated using the states and weights of a population of evolving particles. In addition, it provides reliable estimates of the evidence (marginal likelihood) that is needed for Bayesian model selection at essentially no cost. We consider adaptive sequential Monte Carlo (ASMC), which is an extension of annealed importance sampling (AIS). In these methods, importance sampling is performed over a sequence of intermediate distributions, known as power posteriors, linking the prior to the posterior PDF. The main advantages of ASMC with respect to AIS are that it adaptively tunes the tempering between neighboring distributions and it performs resampling of particles when the variance of the particle weights becomes too large. We consider a challenging synthetic groundwater transport inverse problem with a categorical channelized 2D hydraulic conductivity field designed such that the posterior facies distribution includes two distinct modes with equal probability. The model proposals are obtained by iteratively re-simulating a fraction of the current model using conditional multi-point statistics (MPS) simulations. We focus here on the ability of ASMC to explore the posterior PDF and compare it with previously published results obtained with parallel tempering (PT), a state-of-the-art MCMC inversion approach that runs multiple interacting chains targeting different power posteriors. For a similar computational budget involving 24 particles for ASMC and 24 chains for PT, the ASMC implementation outperforms the results obtained by PT: the models fit the data better and the reference likelihood value is contained in the ASMC sampled likelihood range, while this is not the case for PT range. Moreover, we show that ASMC recovers both reference modes, while none of them is recovered by PT. However, with 24 particles there is one of the modes that has a higher weight than the other while the approximation is improved when moving to a larger number of particles. As a future development, we suggest that including fast surrogate modeling (e.g., polynomial chaos expansion) within ASMC for the MCMC steps used to evolve the particles in-between importance sampling steps would strongly reduce the computational cost while still ensuring results of similar quality as the importance sampling steps could still be performed using the regular more costly forward solver.