Sequential Particle Flow for Bayesian Inference

Le Song, Georgia Institute of Technology

In many real-world problems, Bayesian inference for hidden parameters needs to be conducted incrementally as observations arrive sequentially. In this setting, commonly used sequential Monte Carlos methods, such as particle filtering, are facing various challenges in high dimensions, such as the particle die-out problem. In this work, we propose a meta-learning approach which learns an ODE-based algorithm to perform incremental Bayesian inference. The key component of the learned algorithm is an operator conditioned on observations which can control the flow of particles towards the incremental Bayesian posteriors. The proposed approach is computational efficient in test time and achieve good results in high dimensional problems.