

Adversarial Bayesian Simulation

Yuexi Wang

University of Chicago.

24 November 2022, 1.30pm UK time

Abstract

In the absence of explicit or tractable likelihoods, Bayesians often resort to approximate Bayesian computation (ABC) for inference. In this talk, we will cover two summary-free ABC approaches, both inspired by adversarial learning. The first one adopts a classification-based KL estimator to quantify the discrepancy between real and simulated datasets. We consider the traditional accept/reject kernel as well as an exponential weighting scheme which does not require the ABC acceptance threshold. In the second paper, we develop a Bayesian GAN (B-GAN) sampler that directly targets the posterior by solving an adversarial optimization problem. B-GAN is driven by a deterministic mapping learned on the ABC reference by conditional GANs. Once the mapping has been trained, iid posterior samples are obtained by filtering noise at a negligible additional cost. We propose two post-processing local refinements using (1) data-driven proposals with importance reweighting, and (2) variational Bayes. For both methods, we support our findings with frequentist-Bayesian theoretical results and highly competitive performance in empirical analysis. (Joint work with Veronika Rockova)

References

- [1] Wang, Y., Ročková, V. (2022). Adversarial Bayesian Simulation. Preprint at arXiv [arXiv:2208.12113](https://arxiv.org/abs/2208.12113).