

Distilling importance sampling

Dennis Prangle.

University of Newcastle, UK.

09 April 2020, 11.30 UK time

Abstract

To be efficient, importance sampling requires the selection of an accurate proposal distribution. This talk describes learning proposal distributions using optimisation over a flexible family of densities developed in machine learning: normalising flows. In a likelihood-free context, training data is generated by running ABC importance sampling with a large bandwidth parameter, and this is “distilled” by using it to train the normalising flow. Over many iterations of importance sampling and optimisation, the bandwidth is slowly reduced until an importance sampling proposal for a good ABC approximation to the posterior is generated. The method will be demonstrated for likelihood-free inference on a queueing model. In this example we infer the parameters, and also the random variables used to simulate data. Thus we effectively learning to control the simulator to produce simulations closely matching the observations.

References

- [1] D. Prangle (2019). Distilling important sampling, [arXiv:1910.03632](https://arxiv.org/abs/1910.03632).