ABCDP: Approximate Bayesian Computation with Differential Privacy

Mijung Park.
Max Planck Institute.

25 March 2021, 11:30 UK time

Abstract
We develop a novel approximate Bayesian computation (ABC) framework, ABCDP, that produces differentially private (DP) and approximate posterior samples. Our framework takes advantage of the Sparse Vector Technique (SVT), widely studied in the differential privacy literature. SVT incurs the privacy cost only when a condition (whether a quantity of interest is above/below a threshold) is met. If the condition is met sparsely during the repeated queries, SVT can drastically reduces the cumulative privacy loss, unlike the usual case where every query incurs the privacy loss. In ABC, the quantity of interest is the distance between observed and simulated data, and only when the distance is below a threshold, we take the corresponding prior sample as a posterior sample. Hence, applying SVT to ABC is an organic way to transform an ABC algorithm to a privacy-preserving variant with minimal modification, but yields the posterior samples with a high privacy level. We theoretically analyze the interplay between the noise added for privacy and the accuracy of the posterior samples. We apply ABCDP to several data simulators and show the efficacy of the proposed framework.

References

[1] M. Park, M. Vinaroz, W. Jitkrittum. ABCDP: Approximate Bayesian Computation with Differential Privacy. arXiv:1910.05103.