

Amortised and provably-robust simulation-based inference

Harita Dellaporta

Complex simulator-based models are now routinely used to perform inference across the sciences and engineering, but existing inference methods are often unable to account for outliers and other extreme values in data which occur due to faulty measurement instruments or human error. In this paper, we introduce a novel approach to simulation-based inference grounded in generalised Bayesian inference and a neural approximation of a weighted score-matching loss. This leads to a method that is both amortised and provably robust to outliers, a combination not achieved by existing approaches. Furthermore, through a carefully chosen conditional density model, we demonstrate that inference can be further simplified and performed without the need for Markov chain Monte Carlo sampling, thereby offering significant computational advantages, with complexity that is only a small fraction of that of current state-of-the-art approaches.