More expressive amortized Bayesian inference via joint learning and self-consistency?

Stefan Radev
Rensselaer Polytechnic Institute, New York.

26 October 2023, 1.00pm UK time

Abstract
We propose "jointly amortized neural approximation" (JANA) to simultaneously learn intractable likelihood functions and posterior densities arising in Bayesian surrogate modeling and simulation-based inference. To this end, we train three complementary networks in an end-to-end fashion: 1) a summary network to compress individual data points, sets, or time series into informative embedding vectors; 2) a posterior network to learn an amortized approximate posterior; and 3) a likelihood network to learn an amortized approximate likelihood. The interaction between these networks opens a new route to amortized marginal likelihood and posterior predictive estimation. Furthermore, we observe that we can increase the efficiency of amortized inference by leveraging the marginal likelihood as a proxy for the self-consistency of Bayes' rule. Upon perfect approximation, the marginal likelihood is constant across all parameter values by definition. However, approximation error leads to undesirable variance in the marginal likelihood estimates across different parameter values. We formulate violations of this symmetry as a loss function applicable in both likelihood-based and simulation-based scenarios. We demonstrate the utility of these approaches on illustrative Bayesian models, including both simple examples and representative cases.