

Peer Review

Review of: "Active Sequential Posterior Estimation for Sample-Efficient Simulation-Based Inference"

Carlo Albert¹

1. Eawag: Das Wasserforschungs-Institut des ETH-Bereichs, Dübendorf, Switzerland

This manuscript introduces Active Learning to the APT-variant of Simulation-based Inference. The proposed method iteratively selects parameter sets with the highest expected information gain in terms of the variance of the posterior density w.r.t. the current parameter-distribution for APT's neural density estimator. The paper is generally well written, and the method is well explained.

My main criticism is that the reported performance gains (both on the benchmarks and the real-world application) do not appear to be very significant when you look at the reported standard errors. This makes the results less convincing.

Moreover, their notation is not always clear to me: Chapter 3.1: In which sense does the simulator $p(x|\theta, D)$ depend on D ? Chapter 3.2: What exactly does it mean that " $p(\theta|x, \phi)$ is the corrected posterior produced by $q_{\phi}(\theta|x)$ "? A formula might help here.

Declarations

Potential competing interests: No potential competing interests to declare.