

Training neural networks with corrected Langevin dynamics and stochastic gradients

Adrià Garriga-Alonso @ (University of Cambridge)

R 3.07 Tue Z1 17:00-17:30

The last decade has popularised a family of Markov chain Monte Carlo (MCMC) algorithms for training Bayesian neural networks. These algorithms are numerical solutions to the Langevin stochastic differential equations of various orders, where the potential function is related to the neural network and training data. In order to computationally scale to the large training data sets common in modern machine learning (ML), the algorithms use stochastic estimates of the potential's gradient.

To ensure that MCMC converges to the target probability distribution, one needs to correct for discretisation errors using a ratio of density functions. Calculating this ratio requires a pass over the whole data set. Common wisdom in Bayesian ML holds that it is prohibitive to perform this correction at every step, so the sampler is used without correction instead. As a result, we are left with no correctness guarantees, and no way to monitor the error of our samples from the target.

In this talk I show that it is possible to exactly perform the correction step while using stochastic gradients. The exact ratio evaluation is performed only after many gradient steps, thus introducing only a small (2-5%) slowdown compared to the uncorrected algorithm, or to mainstream neural network training algorithms like stochastic gradient descent. This allows the practitioner to monitor the approximation error, or to eliminate it completely at extra computational cost. The correction is applicable to most semi-implicit discretisation schemes.