Distributed Dual Averaging Variational Inference

Shiyuan He

Renmin University of China
E-mail: heshiyuan@ruc.edu.cn

Abstract: Many modern machine learning algorithms rely on Bayesian probabilistic models whose posterior is difficult to compute. Variational Inference~(VI) is a popular rescue for the intractable posterior. Despite its popularity, VI lacks the scalability to huge dataset, especially when the dataset is distributed stored on distinct servers. In this work, we presents a simple-to-implement yet efficient algorithm to tackle the problem. The algorithm depends on distributed dual averaging, and only the cumulant natural gradient is exchanged among servers to ensure global convergence. The algorithm is applicable to a wide range of Bayesian models, and only assumes the variational distribution belongs to the exponential family. Moreover, it exploits the Riemannian geometry of the exponential family for efficient update. We test the algorithm over several real world datasets, and demonstrate its increased speed of convergence.