

Gradient-Based Variational Inference by Policy Search

Project Manager
Dr. Oleg Arenz

Principal Investigator
Prof. Jan Peters (PhD)

Project Term
2021 - 2022

Clusters
Lichtenberg Cluster Darmstadt

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt



Introduction

In variational inference, we want to approximate an intractable target distribution (often given as a posterior distribution) with a tractable model distribution. Applications of variational inference include learning of posterior distributions in deep neural networks as well posterior distribution over hyper-parameters of other machine learning representations. A common choice for tractable models that are still very flexible are Gaussian mixture models (GMMs). These models are particularly suited for problems with up to a few hundreds of dimensions, for example, when we want to learn a distribution over joint configurations of a robot. In this project, we extended an algorithm for learning such GMM representations that has been recently developed by our group. The original algorithm is only applicable in the blackbox setup, i.e., where no gradient of the target distribution is available. We extended our algorithm to also utilize this gradient information which again considerably increased its speed and scalability.

Methods

We incorporate gradient information by making use of Stein's Lemma, which has been recently introduced for Gaussian variational inference. We further compare our improved method with an alternative method for GMM-based variational inference, and show that these methods mainly differ in algorithmic choices that are independent of the theoretical derivations of the methods. We evaluate the impact of these design choices on the performance of the algorithm by running experiments on several test problems.

Results

In preliminary experiments, we show that our extended methods performs better than existing methods for GMM-based variational inference. In particular we show that using the gradient information improves the efficiency of our algorithm by around one order of magnitude. We also demonstrate that choosing the learning rate based on KL-constrained trust regions significantly improves the stability of the updates and the quality of the learned approximation. Furthermore, we show that dynamically adapting the number of components improves exploration and thereby helps to detect different modes of the target distribution.

Discussion

We investigated recent methods for GMM-based variational inference and showed that they are indeed closely related. Still, these methods differ in several algorithmic choices that substantially affect the performance in terms of sample efficiency and the quality of the learned approximation. We evaluated several of these choices. Based on our experiments, KL-constrained trust regions—even when applied on top of higher-order natural gradients—seem to outperform directly controlled learning rates, and dynamically adapting the number of components has clear advantages compared to optimizing GMMs with fixed number of components. Based on these insights, we proposed a novel method that combines fast first-order based estimates of the natural gradients with adaptive KL-constrained trust regions and adaptive number of components, setting a new standard for GMM based variational inference. However, we need to run additional experiments to better isolate the effects of the individual design choices.

Last Update: 2022-07-14 17:31