

Gradient-Based Variational Inference by Policy Search



Project Manager
Dr. Oleg Arenz

Principal Investigator
Dr. Oleg Arenz

Project Term
2022 - 2023

Clusters
Lichtenberg II Cluster Darmstadt

Software
TensorFlow

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt

Introduction

Many problems in machine learning involve inference from intractable distributions. For example, when learning latent variable models, in maximum entropy reinforcement learning, or Bayesian inference. Important applications can be found, for example, in robotics—where the intractable distribution could be a multimodal distribution over joint-configurations that reach the desired pose or over collision-free motions that reach a given goal, or in non-amortized variational inference for latent variable models.

Variational inference (VI) aims to approximate the intractable target distribution by means of a tractable, parametric model. Variational inference is typically framed as minimizing the reverse Kullback-Leibler (KL) divergence between approximation and target distribution. In this work, we focus on a particular choice of variational distribution—a Gaussian mixture model. Gaussian mixture models are a simple yet powerful choice for a model family since they can approximate arbitrary distributions when assuming a sufficiently high number of components. Compared to more complex models, such as normalizing flows, they are more interpretable and tractable since not only sampling and evaluating GMMs is cheap, but also marginalizations and certain expectations (of linear or quadratic functions) can be computed in closed form. Furthermore, the simplicity of GMMs allows for sample-efficient learning algorithms, that is, algorithms that require relatively few evaluations of the target distribution for learning the model parameters.

Arguably the two most effective algorithms GMM-based variational inference, both apply independent natural gradient (NG) updates on each component as well as on the categorical distribution over weights. Yet, both algorithms were derived from a different perspective, have different theoretical guarantees, and even different objectives for the independent updates. Namely, iBayes-GMM uses the original GMM objective for each independent update to perform natural gradient descent also with respect to the full mixture model, whereas VIPs uses a lower bound for an expectation-maximization procedure, which yields independent objective functions for each component and the mixture weights. Their approach can be shown to converge, even when the M-Step does not consist of single natural gradient updates. However, it was not yet proven, that their proposed procedure, which does use single NG steps, also performs natural gradient descent on the full mixture.

Methods

We further explored the previous works iBayes-GMM and VIPs, and use our findings to derive a generalized method that outperforms both of them. We present a general framework for learning a GMM approximation that unifies both previous methods. Our framework uses seven modules to independently select design choices, for example, regarding how samples are selected, how natural gradients are estimated or how the learning rate or the number of components is adapted. For each design choice, we review and compare the different options that

have been used in prior works, and we discuss potential limitations. For example, VIPS uses an inefficient zero-order method for estimating natural gradients, whereas iBayesGMM updates the individual components based on samples from current GMM approximation, which can prevent component with low weight from receiving meaningful updates. We propose a novel combination of design choices and show that it significantly outperforms both prior methods. In particular, we combine KL-constrained trust regions, which have been popularized in the gradient-free reinforcement learning setting, with gradient-based estimates of the NG, use samples from each component and adapt the number of components. Test problems are used from both prior works.

Results

Although VIPS and iBayes-GMM are derived from different perspectives— where the derivations for Bayes-GMM are less general (by requiring single NG steps for the component update) but enjoy stronger guarantees (by proving natural gradient descent on the whole mixture model)—, we showed that both algorithms only differ in design choices and could have been derived from the other perspective, respectively. This unification of both perspective shows that we can derive approximate natural gradient descent algorithms also for mixtures of non-Gaussian components—where the approximation errors of the natural gradient are potentially much larger—without having to give up on convergence guarantees. Furthermore, our results are of high relevance for the practitioner, both due to our extensive study on the effects of the individual design choices—which shows that both prior works can be improved by using a combination of their design choices— and by releasing our modular framework for natural gradient GMM-based variational inference, which is well-documented and easy to use and outperforms the reference implementations for VIPS and iBayesGMM when using the respective design choices.

Discussion

The scope of this work is narrow, focusing on two specific approaches for natural-gradient GMM-based variational inference. There are of course many other models that can be applied for variational inference, and, depending on the problem setting, some of these models are highly preferable over GMMs, for example, normalizing flows should likely be preferred for high-dimensional problem settings, such as (deep) Bayesian neural networks. However, for this work we assume that we indeed want to optimize a Gaussian mixture model, for example, because we require an interpretable model with smooth gradients. Even in the field of GMM-based variational inference, alternative methods, based on boosting or the reparameterization trick are possible. By not using natural gradients, these methods can be applied more straightforwardly to sparse covariance parameterizations which can be beneficial for higher-dimensional problems. However, in the considered problem setting where we can learn GMMs with full covariance matrices, these methods are not competitive to the natural

gradient based methods described in this work.

Publications

Arenz, O; Dahlinger, P; Ye, Z; Volpp, M; Neumann, G. A Unified Perspective on Natural Gradient Variational Inference with Gaussian Mixture Models, Transactions on Machine Learning Research (2023)
<https://doi.org/10.48550/arXiv.2209.11533>

Last Update: 2023-10-10 11:57