

Amortized Variational Inference by Policy Search



Project Manager
Dr. Oleg Arenz

Researchers
Daniel Musekamp

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

Variational inference with Gaussian mixture models (GMMs) can be used to learn highly tractable approximations of complex, multi-modal distributions. These approximations are applied in robotics, for example, to obtain fast predictions of diverse solutions for inverse kinematics (IK) or path planning problems. The most efficient approaches for GMM-based variational inference use decompositions of the objective function to update the individual components independently using natural gradients. While these approaches are much more effective than naive gradient descent, they cannot be used to learn conditional GMMs, where the predicted solutions can depend on contexts, such as goal positions, which limits their usefulness in many robotic applications.

Methods

Instead, we consider conditional GMMs for robotics, which can be trained offline and efficiently predict multi-modal distributions based on contexts. In this work we focus on Inverse Kinematics applications, where the context is given by desired end effector positions, but our work is directly applicable to other applications such as motion planning or grasping where the context could correspond, for example, to goal locations, images or point clouds. To enable efficient training of conditional GMMs, we extend the prior work VIPS to train neural networks for predicting the GMM parameters based on the context, instead of directly optimizing the weights, means and covariance matrices. VIPS is a general method for variational inference (VI) and, thus, our contributions could be also useful in different fields, for

example for training multimodal encoders for variational autoencoders (VAE) or policies for reinforcement learning, e.g. within soft actor-critic. We introduce amortized VIPS (aVIPS) for training conditional GMMs for Variational Inference. aVIPS makes use of several sound and non-trivial techniques to make the training more effective, such as a differentiable projection layer to enforce well-conditioned covariance matrices, adaptive information-geometric trust regions for improving the stability of optimization and to gain more diverse mixtures, as well as sampling importance resampling (SIR) to make the training more efficient. We evaluate aVIPS for inverse kinematics by predicting multimodal distributions over joint configurations for high-DOF robots, such as Pal Robotics TALOS, and compare it to competing approaches based on GANs and normalizing flows.

Results

We evaluated aVIPS for inverse kinematics, where it achieves competitive performance compared to recent methods based on less tractable models. aVIPS shows lower errors on two robot experiments, Panda and Valkyrie, however does not reach the accuracy of a normalizing flow baseline on the Atlas robot. Compared to a vanilla GMM, aVIPS reaches higher entropy and generates more diverse mixtures. Furthermore, the ELBO was significantly increased on Atlas. By removing obsolete components, aVIPS learns much smaller models. On the Valkyrie task, a single Gaussian model is a strong baseline that surpasses all other models in terms of accuracy and ELBO, as the 4-DOF chain lacks redundancy for reaching the target pose. On the other tasks, GMMs clearly outperform a single Gaussian, not only in terms of entropy, but also with respect to the accuracy. The improved accuracy shows that GMMs can explore the output space better: if one component is stuck in a local optima, another component may still be able to reach a good mode due to different initialization.

Discussion

We could not achieve the same variability compared to normalizing flows, which is a limitation that we want to address in future work. Although our trained GMM made use of all components, for most contexts only 2-3 components were used. Yet, as shown in the Valkyrie experiment, it is possible to train a single Gaussian to cover the whole context space when using sufficiently deep networks. To make training of more diverse GMMs feasible, we plan to introduce a shared network. Furthermore, we want to investigate additional applications of our method in the fields of motion planning and grasping.

Last Update: 2023-10-10 11:53