Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)

Prof. Gerhard Neumann

Project Type

- [x] Master Thesis
- [ ] Bachelor Thesis
- [ ] Research Project

Supervisors

Denis Blessing

denis.blessing@kit.edu

Philipp Dahlinger

philipp.dahlinger@kit.edu

Difficulty

Algorithmic

Math

Application

Amortized Variational Inference for Bayesian Meta-Learning using Langevin Diffusions

Description

Multi-task learning aims to leverage inductive biases learned on a meta-dataset of similar tasks for improved data efficiency on unseen target tasks of similar structure. Most approaches to multi-task meta-learning rely on Variational Inference (VI) with an amortized, factorized Gaussian variational distribution.

A recent study by Volpp et al. (2022) [4] demonstrated that employing more expressive variational distributions produces tighter evidence lower bounds, enhances the efficiency of optimizing marginal likelihood, and leads to improved uncertainty estimation. Nevertheless, they employed non-amortized VI, necessitating the training of separate models for each task.

This thesis aims to harness recent advancements in diffusion-based VI [3, 1] and apply it in an amortized manner to Bayesian multi-task meta-learning. Consequently, we can leverage a single model that serves as the variational distribution for all tasks.

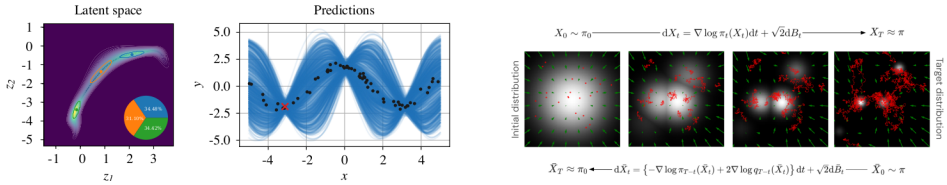

Tasks

The tasks in this project will involve:

- Literature Review. Getting familiar with multi-task meta-learning [4], neural processes [2] and diffusion for VI [3, 1].
- Implementation. Extend the code-base from [4] such that we can use amortized VI schemes.
- Benchmarking. Compare our method against competing approaches.

References

[1] Arnaud Doucet, Will Grathwohl, Alexander G Matthews, and Heiko Strathmann. Score-based diffusion meets annealed importance sampling. *Advances in Neural Information Processing Systems*, 35:21482–21494, 2022.

[2] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.

[3] Tomas Geffner and Justin Domke. Langevin diffusion variational inference. In *International Conference on Artificial Intelligence and Statistics*, pages 576–593. PMLR, 2023.

[4] Michael Volpp, Philipp Dahlinger, Philipp Becker, Christian Daniel, and Gerhard Neumann. Accurate bayesian meta-learning by accurate task posterior inference. In *The Eleventh International Conference on Learning Representations*, 2022.