Reinforcement Learning for Optimizing Diffusion Models

Internship Proposal for academic year 2024-2025

- Intended for: Students in second year of master in machine learning or related topic.
- Duration: 5 or 6 months, preferably from April 2025.
- Supervisors: Olivier Cappé (CSD, DI-ENS, Ecole Normale Supérieure—PSL, CNRS), Alexandre Verine (CSD, DI-ENS, Ecole Normale Supérieure—PSL, CNRS)[*]
- Location: Institut PR[AI]RIE - Paris Santé Campus, 75015 Paris.
- Follow up: Priority will be given to candidates interested in pursuing a PhD thesis on a related topic.

Diffusion models have recently emerged as a leading class of generative models, particularly excelling in image generation [1, 2]. These models operate by transforming a noise signal into a complex data distribution through an iterative denoising process. The denoising step is performed by a neural network trained to predict the next stage in the process. However, training such models is computationally intensive and time-consuming. Consequently, once the denoising models are trained, various methods have been developed to optimize performance by focusing solely on enhancing the iterative generative process [3, 4, 5].

The goal of this internship is to reframe the generative process as a Markov Decision Process (MDP) and leverage reinforcement learning techniques to optimize the generation process using ideas borrowed from the field of inverse reinforcement learning [6, 7].

References

[1] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. arXiv, December 2020. NeurIPS 2020.

[2] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. October 2022. ICLR 2022.

[3] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space of Diffusion-Based Generative Models, October 2022. NeurIPS 2022.

[4] Dongjun Kim, Yeongmin Kim, Se Jung Kwon, Wanmo Kang, and Il-Chul Moon. Refining Generative Process with Discriminator Guidance in Score-based Diffusion Models. April 2023. ICML 2023.

[5] Yilun Xu, Mingyang Deng, Xiang Cheng, Yonglong Tian, Ziming Liu, and Tommi Jaakkola. Restart Sampling for Improving Generative Processes, November 2023. NeurIPS 2023.

[6] Tianwei Ni, Harshit Sikchi, Yufei Wang, Tejus Gupta, Lisa Lee, and Benjamin Eysenbach. f-IRL: Inverse Reinforcement Learning via State Marginal Matching, December 2020. CoRL 2020.

[7] Siddhant Agarwal, Peter Stone, Ishan Durugkar, and Amy Zhang. f -Policy Gradients: A General Framework for Goal Conditioned RL using f -Divergences. 2023. NeurIPS 2023.

[*]Contact: olivier.cappe@cnrs.fr, alexandre.verine@ens.fr