

Internship Proposal: Precision-recall Divergences for Diffusion Models

Keywords: Generative AI, Denoising Diffusion Models, Precision-recall Tradeoff

Supervisors:

- Yann Chevaleyre: yann.chevaleyre@dauphine.fr
- Benjamin Negrevertne : benjamin.negrevertne@dauphine.fr
- Alexandre Vérine: alexandre.verine@ens.fr

Place : MILES, LAMSADE, Université Paris Dauphine

Context and Motivation:

Generative AI has achieved remarkable successes in the past years, both in text and image generation. These models, particularly diffusion models in image generation, have set new benchmarks for quality. However, a common limitation is their tendency towards high quality (precision) of the generated images, sacrificing diversity (recall). This phenomenon is attributed to the usual optimization objectives used during the training phase.

The two common challenges related to this topic are:

- 1- Assessing the quality (precision) and diversity (recall) of a specific generative model. This question has received a lot of attention recently [2,3,4,5]
- 2- Achieving a desired balance between image quality (precision) and diversity (recall) during the training process [6,7].

Objective:

The goal of this internship is to design new training methodologies to precisely control the precision-recall tradeoff as in [6], but for diffusion models. There are many tasks and open questions that the intern will have to tackle to achieve this objective:

- **Existing diversity promoting approaches for diffusion models.** The intern will need to conduct a review on diversity in Diffusion Models, knowing that there are very few papers on this topic. See e.g. [7].
- **Which precision-recall measure is more adapted to diffusion models?** The intern will have to conduct a thorough review of existing literature on precision-recall measures, focusing on measures that have the potential to be optimized (e.g. by gradient descent), and measures that are better suited to diffusion models.
- **What is the optimal sequential strategy for enforcing a precision-recall trade-off?** Diffusion model are temporal processes, so guiding the generation of an image to achieve a specific goal is a sequential decision-making problem. This aspect might be of high importance in designing the training methodology.
- **How do we estimate a f -divergence at any timestep of the backward process in the context of diffusion models?** Most precision-recall measures are based on f -divergences, so estimating those divergences is a crucial objective.

Bibliography

1. Jonathan Ho, Ajay Jain, Pieter Abbeel, Denoising Diffusion Probabilistic Models, Neurips 2020.
2. Cheema, Fasil, et Ruth Urner. « Precision Recall Cover: A Method for Assessing Generative Models », AISTATS, 2023
3. Kynkäänniemi, Tuomas, Tero Karras, Samuli Laine, Jaakko Lehtinen, et Timo Aila. « Improved Precision and Recall Metric for Assessing Generative Models », Neurips 2019.
4. Sajjadi, Mehdi S. M., Olivier Bachem, Mario Lucic, Olivier Bousquet, et Sylvain Gelly. « Assessing Generative Models via Precision and Recall ». Neurips 2018.
5. Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, Zaid Harchaoui. MAUVE: Measuring the Gap Between Neural Text and Human Text using Divergence Frontiers. NeurIPS 2021
6. Alexandre Verine, Benjamin Negrevergne, Muni Sreenivas Pydi, Yann Chevaleyre, Precision-Recall Divergence Optimization for Generative Modeling with GANs and Normalizing Flows, Neurips 2023.
7. Dongjun Kim, Yeongmin Kim, Se Jung Kwon, Wanmo Kang, Il-Chul Moon, Refining generative process with discriminator guidance in score-based diffusion models, ICML 2023