Diffusion Adapters for Disentangled Latent Space Conditioning

Lab: Computational Bioimaging and Bioinformatics, ENS

Location: ENS Paris, 45 rue d'Ulm, 75005 Paris

Duration: 6 months (starting February - April 2025)

Supervisors: Thomas Boyer, Nicolas Bourriez, Auguste Genovesio

Contacts: thomas.boyer@ens.psl.eu | nicolas.bourriez@ens.psl.eu

Context:

Recent advancements in deep learning have shown promising results in disentangling semantic information within latent spaces of encoders [1,2]. Our team has recently developed a novel architecture that can "distill" specific semantic factors from entangled latent spaces, enabling targeted semantic transfer between images. Notably, this method has demonstrated the ability to generalize to unseen conditions during training, such as transferring a previously unseen color to an object's appearance.

This breakthrough opens up exciting possibilities for more controlled and interpretable image generation and manipulation. However, several critical questions remain unexplored, particularly regarding the nature of the constructed latent space and its potential applications in state-of-the-art generative models like diffusion models [2, 3, 4].

In our context of computational biology, the long-term vision is to develop a “virtual experiment regenerator” allowing to modify selected attributes of an existing assay to generate “what if“ versions of it. Typically, cell assays datasets [7, 8, 9] consist of imagery of cell culture perturbed by thousands of different chemical or genetic perturbations that are partially replicated across different laboratories, cell lines, staining conditions, microscopes, exposure-to-perturbation time, etc. Having the ability to reliably separate (disentangle) these attributes and generate a virtual experiment with an unexisting combination of conditions would represent a major leap in biology and drug discovery.

Project Description:

This research project aims to deepen our understanding of the disentangled semantic latent spaces created by our novel architecture and explore their potential in enhancing conditional image generation. The project will focus on four main objectives:

1. Latent Space Analysis: Investigate the properties of the constructed latent space, including its smoothness and the nature of interpolations between known points. This will involve developing visualization techniques and metrics to characterize the space's structure.
2. Diffusion Model Conditioning: Explore the feasibility and effectiveness of using our disentangled latent representations to condition diffusion models. This will require adapting existing diffusion model architectures and conditioning mechanisms [5] to incorporate our semantic latent codes.
3. Adaptation for Robust Distillation on Heldout Set: Based on the findings from the previous objectives, develop techniques to enhance the performance of our approach when working with distillation on heldout data.
4. Generalization of Diffusion Models to Unseen Conditions: Assess the behavior of conditioned diffusion models when provided with latent codes representing semantic factors *not seen during training*. Analyze whether the output is scientifically meaningful, and develop methods to quantify and improve generalization.

General-domain multi-label datasets such as [10, 11] will be used as ground truth and for easier experimentation.

Your role:

As an intern, you will play a crucial role in advancing this cutting-edge research and in the team. Your responsibilities will include:

- Implementing and optimizing algorithms for latent space analysis and visualization.
- Adapting and training diffusion models to work with our disentangled semantic latent codes.
- Designing and conducting experiments to evaluate the performance and generalization capabilities of the developed models.
- Regularly collaborating with team members and participating in team meetings.
- Documenting research findings and contributing to potential publications or conference submissions.

Pre-requisites:

- Strong proficiency in Python and experience with deep learning frameworks (preferably PyTorch).
- Very solid understanding of machine learning concepts, particularly in the areas of representation learning and generative models.
- Familiarity with computer vision techniques and image processing.
- Ability to communicate effectively in English or French.

Nice to have:

- Experience with diffusion models or other advanced generative modeling techniques.
- Background in information theory or disentanglement metrics.
- Familiarity with latent space manipulation techniques.
- Interest in or experience with interpretable AI and controlled generation.

Application Process:

To apply for this internship, please submit your resume, a brief statement of interest, and any relevant project/paper examples or GitHub repositories. Don't hesitate to contact us for more details about the project or application process.

References:

1. Pan et al., *Disentangled Information Bottleneck,* Proceedings of the AAAI Conference on Artificial Intelligence 2021
2. Burgess et al., *Understanding disentangling in β-VAE,* 2018
3. Ho, J., Jain, A., & Abbeel, P. (2020). *Denoising Diffusion Probabilistic Models*. Advances in Neural Information Processing Systems, 33.
4. Sungbin Lim, EUN BI YOON, Taehyun Byun, Taewon Kang, Seungwoo Kim, Kyungjae Lee, Sungjoon Choi, Score-Based Generative Modeling through Stochastic Differential Equations, NeurIPS 2023
5. Tero Karras, Miika Aittala, Timo Aila, Samuli Laine, Elucidating the Design Space of Diffusion-Based Generative Models, NeurIPS 2022
6. Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, Björn Ommer, High-Resolution Image Synthesis with Latent Diffusion Models, CVPR 2022
7. JUMP Cell Painting datasets (Chandrasekaran et al., 2023), available from the Cell Painting Gallery on the Registry of Open Data on AWS (https://registry.opendata.aws/cellpainting-gallery/). Chandrasekaran et al., 2023: doi:10.1101/2023.03.23.534023
8. Chandrasekaran, S.N., Cimini, B.A., Goodale, A. et al. Three million images and morphological profiles of cells treated with matched chemical and genetic perturbations. Nat Methods 21, 1114–1121 (2024). https://doi.org/10.1038/s41592-024-02241-6
9. *RxRx3 dataset (Fay et al. (2023). RxRx3: Phenomics Map of Biology. bioRxiv 2023.02.07.527350), available from Recursion at rxrx.ai/rxrx3.*
10. *Burgess, Chris and Kim, Hyunjik, 3D Shapes Dataset, https://github.com/deepmind/3dshapes-dataset*
11. *DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations*
12. *Liu, Ziwei and Luo, Ping and Qiu, Shi and Wang, Xiaogang and Tang, Xiaoou, DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations, CVPR 2016*