A general purpose pretrained deep network for bioimages

Lab: computational bioimaging and bioinformatics, ENS
Location: ENS Paris, 45 rue d'Ulm, 75005 Paris
Contacts: Auguste Genovesio (auguste.genovesio@ens.psl.eu), Ethan Cohen (ecohen@bio.ens.psl.eu), Ihab Bendidi (bendidi@bio.ens.psl.eu) .

Context

In the past decade, broadly used deep networks such as Resnet or Inception have emerged as a standard to obtain representation of images [1][2]. In addition, transfer learning has allowed many researchers and industries to easily reuse these networks in various domains including non natural images. Transfer learning has also been used successfully in some domains such as biology or drug discovery. However, there are still several issues associated with biological images. Firstly, we have no clear intuition as to which point transfer learning is good at representing microscopy images and if it could be improved. Secondly, microscopy acquisition produces a variable number of image channels that are unrelated to those of natural images which are made of only three, highly correlated channels: RGB. For instance, in fluorescence microscopy, each channel corresponds to an emission wavelength and collects a signal corresponding to a given protein. The number of channels is also in general different from three. Those issues call for the development of a dedicated deep network that would be trained on biological images and conceived to work efficiently with any number of channels in a meaningful way. A similar initiative has been proposed in [3] for medical images. In this project we propose to develop a method that would be channel agnostic and trained on various biological dataset. We would make this general purpose pretrained model available online for the community.

This project

The project is divided in 5 steps that could be improved along the way:

Step 1) Due to the large size of available biological image datasets, a module to stream data will have to be created. It will consist of an appropriate data loader streaming data from an external website as IDR (https://idr.openmicroscopy.org/) which contains ~13M images from different publicly available experiments of cells and tissues.

Step 2) A model agnostic to the number of image channels should then be proposed and implemented. Previous work in the lab already consisted of the development of a channel-number agnostic convolutional neural network that could be used as a base. We expect to develop a so-called BioInception (or BioResnet) and possibly a

BioVisioTransformers for general purpose bioimage encoder.

Step 3) Effective self supervised and weakly supervised training strategies should also be proposed and implemented.

Step 4) We will then evaluate the efficiency of the model on different biological as well as natural image downstream tasks to show the robustness and the interest of the model. It is also expected to develop a so-called BioFID (Frechet Inception Distance [4]) or related that would be used to evaluate robustness of biological image generative models.

Step 5) If this work is shown to be of sufficient interest for the community, we will release the models to the community on a platform such as TorchHub or HuggingFace.

The candidate:

The candidate should know Python and a deep learning framework such as Pytorch.
Experience using Big Database and training giant models will be a plus.

Bibliography

[1] Kaiming et al, Deep Residual Learning for Image Recognition
[2] Szegedy et al, Rethinking the Inception Architecture for Computer Vision
[3] Alzubaidi et, MedNet: Pre-trained Convolutional Neural Network Model for the Medical Imaging Tasks
[4] Heusel et al, Gans trained by a two time-scale update rule converge to a local Nash equilibrium

Multi modal learning for drug discovery

Lab: computational bioimaging and bioinformatics, ENS
Location: ENS Paris, 45 rue d'Ulm, 75005 Paris
Contacts: Auguste Genovesio (auguste.genovesio@ens.psl.eu), Ethan Cohen (ecohen@bio.ens.psl.eu)

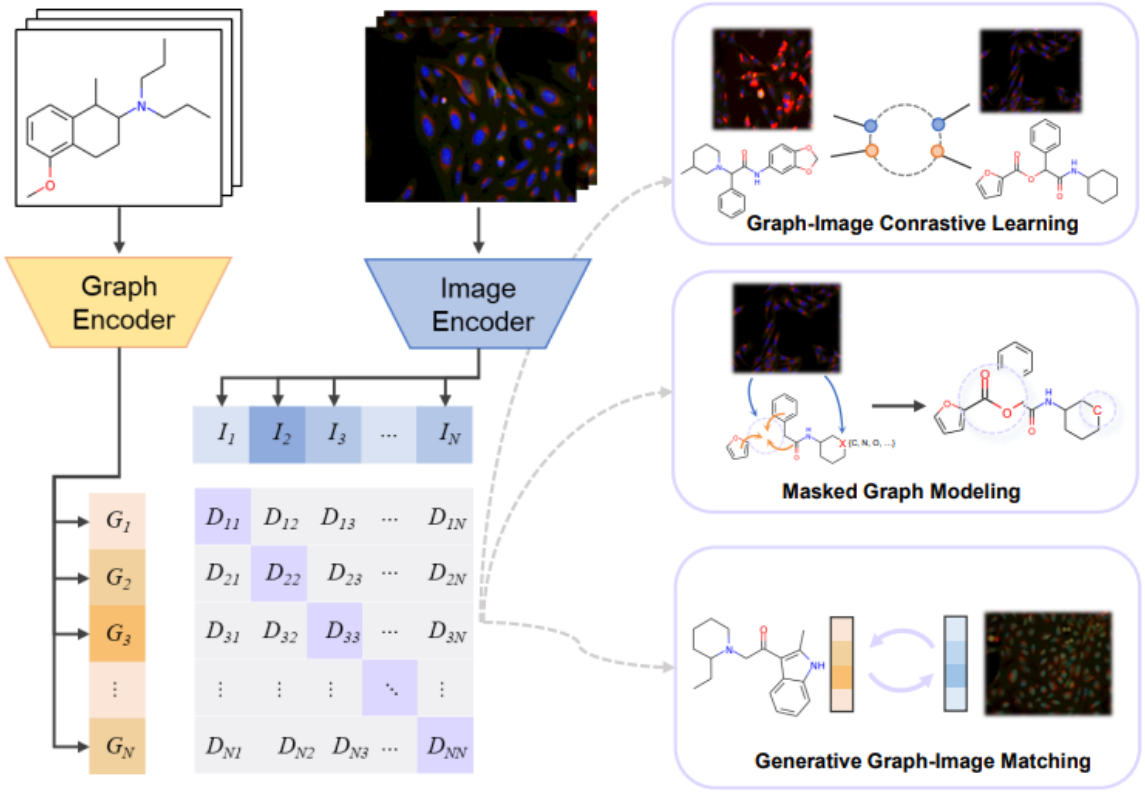

Fig a. Multi modal learning example from [2].

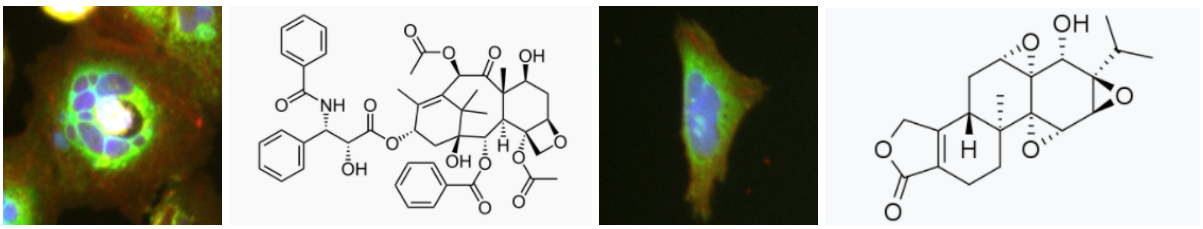

Fig b. Cellular images disturbed by different molecules produce various phenotypes.

Context

Drug discovery is a long process, it often takes more than a decade to design a new therapeutics drug that will end up on the market and cost more than 2 billion dollars. The

target-based approach that assumes the prior knowledge of a molecular target is currently the main strategy in the pharmaceutical industry. However, an alternative method, the phenotype-based approach has led to many more discoveries of first in class drugs. One of the most promising approaches with this strategy is High Content Screening. By combining robotized experiments and automated microscopy, it allows to image cells under thousands of parallel small molecule perturbations.

Background

In recent years, outstanding improvements have been made in multi-modal learning, especially for text and images data. Methods such as CLIP [1] use contrastive learning to learn mutual representations of text and images. Recent work such as [2] has adapted those methods to learn multi modal representation learning for molecules and images (see **Figure a**).

This project

This project aims at extending [2] to learn meaningful relationships between molecules and microscopy images of treated cells using the cell painting assay [4]. A non exhaustive list of improvement will be:

1. Different training scheme as well as using other representations that 2D graphs for molecules.

2. Incorporate Quality Control analysis for HCS data beforehand.

3. Incorporate a way to prevent batch effect as in [3].

The candidate:

The candidate should know Python and would ideally have some experience withPyTorch. Experience in multi-modal learning would be considered a plus.

Bibliography

[1] Radford et al, Learning transferable visual models from natural language supervision
[2] Zheng et al, Cross-modal Graph Contrastive Learning with Cellular Images
[3] Cross-Zamirski et al, Self-Supervised Learning of Phenotypic Representations from Cell Images with Weak Labels
[4] Bray et al, Cell Painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes

Conditional diffusion models for High Content Screening

Lab: computational bioimaging and bioinformatics, ENS
Location: ENS Paris, 45 rue d'Ulm, 75005 Paris
Contacts: Auguste Genovesio (auguste.genovesio@ens.psl.eu), Ethan Cohen (ecohen@bio.ens.psl.eu)

Fig a. Image generated by Stable Diffusion..

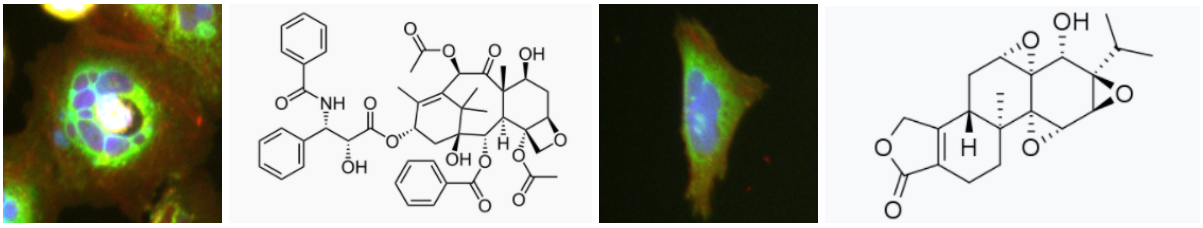

Fig b. Cellular images disturbed by different molecules produce various phenotypes.

Context

Drug discovery is a long process that often takes a decade to design and lead to new therapeutics molecules on the market. The target-based approach that assumes the prior knowledge of a molecular target is currently the main strategy in the pharmaceutical industry. However, an alternative method, the phenotype-based approach has led to many more discoveries of first in class drugs [4]. One of the major tools of this approach is High Content Screening. By combining robotized experiments and automated microscopy, it allows to image cells under thousands of parallel small molecule perturbations, generating hundreds of thousands of microscopy images of cells.

Background

In recent years, outstanding improvements have been made in multi-modal learning, especially for text and images data. Methods such as CLIP [1] uses contrastive learning methods to learn mutual representations of text and images. Other methods such as GLIDE [2] or Stable Diffusion [3] improved drastically text to image generation tasks, and allowed generations of images with high quality, fidelity, and diversity. Such methods were made possible thanks to recent improvements in Diffusion models as well as multi-modal representation learning.

This project

This project aims at extending such methods to the drug discovery context in order to predict microscopy images of cells treated by a chemical compound from the molecule structures of this compound as first investigated in [5]. The intern will therefore adapt relevant methods to the molecule to cell image generation problem. Depending on interest, the intern will have the possibility to explore different ways to process molecules (Fingerprints, Smiles, 2D graphs or 3D graphs).

The candidate:

The candidate should know Python and would ideally have some experience with PyTorch. Experience in multi-modal learning and generative modeling would be considered a plus.

Bibliography

[1] Radford et al, Learning transferable visual models from natural language supervision
[2] Nichol et al, GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models
[3] Rombach et al, High-Resolution Image Synthesis with Latent Diffusion Models
[4] Swinney et al, How were new medicines discovered?
[5] Yang et al, Mol2Image: Improved Conditional Flow Models for Molecule to Image Synthesis

Generating tissue images from single cell gene expression

Lab: computational bioimaging and bioinformatics, ENS
Location: ENS Paris, 45 rue d'Ulm, 75005 Paris
Contacts: Auguste Genovesio (auguste.genovesio@ens.psl.eu), Ihab Bendidi (bendidi@bio.ens.psl.eu) .

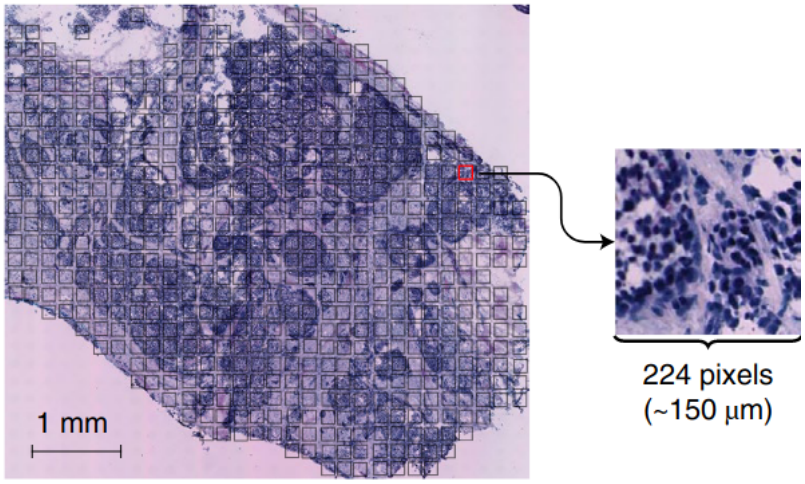

Fig a. Images of cancerogenic tissue that is coupled to its gene expression [1]

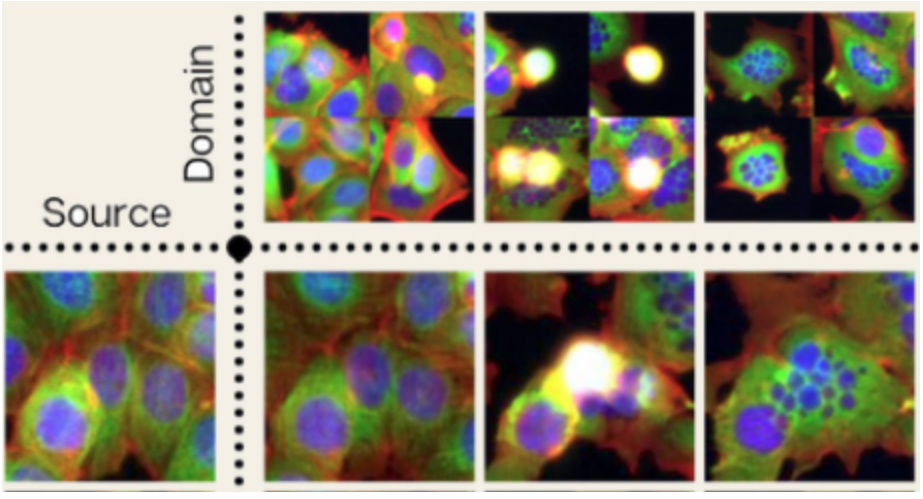

Fig b. Cell Images generated using GANs

Context

Bridging the gap between single cell transcriptomic and single cell visual phenotypes and

function is a graal in basic research in biology. Indeed, current single cell approaches rely either only on microscopy images of cells that provide insight on the spatial arrangement of their organelles and the relationships with the neighboring cells, or on single cell gene expression that provides a view on what genes are being expressed. However, deciphering how genes affect the morphology and the spatial relationships of cells can lead to the understand the gene function and the effects of their modulation by genetic tools or drug treatment. While a few companies are working toward solutions to obtain single cell genotype/phenotype coupling, none of them are yet established.

This project

While not available at the single cell level precision, public data that couples image acquisition of cancerogenic tissue to its local gene expression[1] (group of cells) are accessible. Independently, our partner Minos Biosciences made progress in generating private data that couples isolated single cell images to their gene expression. Furthermore recent research work has emerged with the aim to predict local gene expression from tissue images using deep networks [1]. However little work was yet proposed on predicting the morphology of the cell using its gene expression. Finally, generative models have shown tremendous progresses in generating realistic images and videos of all types including microscopy. This project aims to make use of the cancerogenic tissue data and the private single cell data from Minos Biosciences, to generate images of tissues from their gene expression, using generative methods proven as efficient such as StyleGAN[2] or diffusion models.

The candidate:

The candidate should know Python and would ideally have some experience with PyTorch.
Experience in GANs and image translation would be considered a plus.

Bibliography

[1] Bryan He et al, Integrating spatial gene expression and breast tumor morphology via deep learning
[2] Karras et al, A style-based generator architecture for generative adversarial networks

Diffusion Models Based Unpaired Image-to-Image Translation for Revealing Subtle Phenotypes

Lab: computational bioimaging and bioinformatics, ENS
Location: ENS Paris, 45 rue d'Ulm, 75005 Paris
Contacts: Auguste Genovesio (auguste.genovesio@ens.psl.eu), Anis Bourou (bourou@ens.psl.eu)

Context:

Spotting differences in visual cell pheneotypes from medical and biological images has many applications in fundamental research, drug discovery and medicine. In microscopy experiments, hand-crafted image analysis can be used to measure the variation of phenotypes produced by a perturbation when the last is visible. However, because the cell-to-cell variability within an image often largely overlaps the cell-to-cell variability between phenotypes, the last is often invisible. Recently, image-to-image translation methods were used to tackle this issue [3].

Image-to-image translation methods [8] consist of transferring an image from a source domain to a target domain while keeping the content of the image. This technique is used in different fields like autonomous driving cars, style transfer and biological imaging. Generative adversarial networks (GANs) [1] were extensively used to build efficient image-to-image translation methods using different strategies. However, GANs still suffer from some limitations such as the lack of diversity in the generation and an unstable training.

Diffusion models are recent generative models inspired by non-equilibrium thermodynamics [7], they are considered as the new state-of-the-art methods that generate diverse high-resolution images [2]. They are the basis of some powerful generative models like GLIDE [4], DALLE-2 [5] and Imagen [6]. In this project, we would like to build an efficient diffusion models based image-to-image translation method to better understand some neuro-devlopmental disorders.

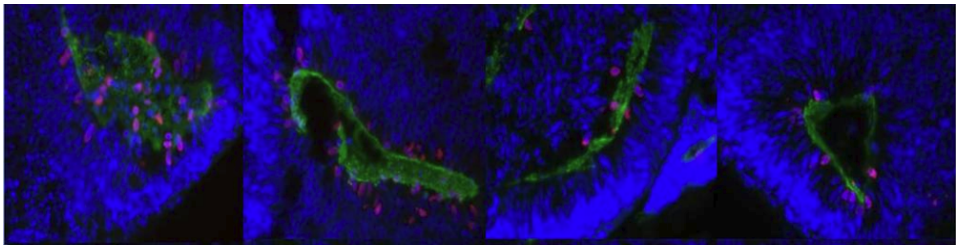

Fig1: images of brain organoids

The project:

In this project, we aim at using diffusion models for generating and translating biological images of a rare neuro-developmental disorder. The goal of the project is twofold:

- Adapting the diffusion models for the generation of new biological images. These synthetic images can be used by the biological community in other deep learning tasks (classification, segmentation...etc).
- Build an efficient image-to-image translation method based on diffusion models to reveal subtle phenotypes in our images.

The candidate:

The candidate should be familiar with Python and some deep learning frameworks like PyTorch. Having experience with some deep generative models is a plus.

References:

[1] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio.Generative adversarial networks, 2014.

[2] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020.

[3] Alexis Lamiable, Tiphaine Champetier, Francesco Leonardi, Ethan Cohen,Peter Sommer, David Hardy, Nicolas Argy, Achille Massougbodji, Elaine Del Nery, Gilles Cottrell, Yong-Jun Kwon, and Auguste Genovesio. Revealing invisible cell phenotypes with conditional generative modeling.bioRxiv,2022.

[4] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever,and Mark Chen. Glide:Towards photorealistic image generation and editing with text-guided diffusion models, 2022.

[5] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022.

[6] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho,David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022.

[7] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics,2015.

[8] Jun-Yan Zhu,Taesung Park,Phillip Isola,and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks, 2017.