Navigating the genetic perturbation landscape: Multi-modal, causal representation learning for target discovery.

Oxford supervisor: Professor Philip Torr[1]
Novo Nordisk supervisor: Associate Professor Jesper Ferkinghoff-Borg[2], Dr Robert Kitchen[3], Dr Cesar Medina[3]

Department:
1. Department of Engineering Science
2. Novo Nordisk, Denmark
3. Novo Nordisk Research Centre Oxford

Project outline

The cardiometabolic diseases remain the major cause of mortality worldwide [1][2]. The etiology of these disorders involves the complex interaction of genetic and environmental factors that affect the physiology of multiple cell types across the human body [3]. To effectively address this complexity, new drug discovery approaches are needed to explore the vast combinatorial space of pharmacological interventions and intricate cellular phenotypes. At the Novo Nordisk Research Centre Oxford, cutting-edge technological platforms have been put in place to genetically perturb and deeply characterize the phenotype of various cellular models at scale. By combining high-throughput imaging, transcriptomics, proteomics, and cell-functional assays to characterize the outcome of genetic and pharmacological perturbations, rich target perturbation landscapes are being built. However, these landscapes do not yet cover the entire human genome, and defining the most effective way to continue exploring the potentially druggable genome remains a challenge. This scenario sets a great opportunity for investigating active learning frameworks like [4], which select data points that are most likely to enhance our understanding, thereby optimizing resource allocation and facilitating a more efficient pathway to discovery. Our goal is to exploit the vast richness of multimodal deep cell phenotyping of limited perturbation examples to guide the increasing exploration of new batches of genetic perturbations and discover new biology.

Deep generative models provide a natural framework for characterising cellular responses to genetic perturbations. For example, recent work in this area has proposed strategies to predict combinatorial perturbations outcomes [5] and how to additionally incorporate prior knowledge in the model [6]. However, existing methods are dataset-specific, lacking any mechanism to transfer knowledge across cell types, and rely on a single data modality. Leveraging insights from previous approach on text and image [7] and [8], our intention is to adapt multimodal generative models to the intricate domain of genetic sequences and expression data. This thoughtful expansion will see us harmonizing generative models, traditionally used in text and image processing, to the specific requirements of genomic data, aiming for a cohesive framework that can accurately predict cellular responses to various perturbations. The project will thus revolve around the following 3 work-packages (WP):

WP1: Representation learning. We have a track record of developing representation learning techniques, [9], [10], [11] and [12]. The focus will be on developing representations that more

accurately reflect the detailed patterns found in images, genetic sequences, and expression data. This will enhance our models' ability to reliably predict the outcomes of genetic perturbations, which is instrumental in identifying new drug targets with subtlety and precision. The prediction models will be formulated as joint probability distributions across the different data modalities to quantify prediction uncertainties and the mutual information between modalities.

WP2: Active learning. The uncertainty of the model predictions is a key ingredient to develop an active learning framework [4][13]. In an “explore”-mode, it will inform the experimentalist which perturbations and measurements to run in the next cycle that will maximize the information gain. In an “exploit”-mode, it can be used to search for novel IP opportunities by examining the neighbourhood around known IPs.

WP3: Model extensions. A promising extension of the model framework consists of using biomedical knowledge graphs (biomedKGs) [14] to contextualize and guide the models. Some of the advances in KG analytics include coupling them with large language models (LLMs) [15]. In particular, the use of in-context learning prompting, which consists of task description, prompt and demonstrations (taken from the screening results of the previously analysed batch), on a LLM that has been enhanced in its inference step with the structured information in a biomedKG.

We believe that the proposed project to develop foundational and causal deep-phenotypic cellular models, represents a significant step forward towards an AI-driven drug discovery pipeline.

Bibliography

[1] GBD 2021 Diabetes Collaborators (2023). Global, regional, and national burden of diabetes from 1990 to 2021, with projections of prevalence to 2050: a systematic analysis for the Global Burden of Disease Study 2021. Lancet, 402(10397), 203–234.

[2] Dagenais GR, Leong DP, Rangarajan S, et al. Variations in common diseases, hospital admissions, and deaths in middle-aged adults in 21 countries from five continents (PURE): a prospective cohort study. Lancet. 2020;395(10226):785-794.

[3] Priest C, Tontonoz P. Inter-organ cross-talk in metabolic syndrome. Nat Metab. 2019;1(12):1177-1188.

[4] Sun, S., Chen, L., Slabaugh, G., & **Torr, P.** (2020). Learning to sample the most useful training patches from images. *arXiv preprint arXiv:2011.12097*.

[5] Lotfollahi, M., ... & Theis, F. J. (2023). Predicting cellular responses to complex perturbations in high-throughput screens. Molecular Systems Biology, e11517.

[6] Roohani, Y., Huang, K., & Leskovec, J. (2023). Predicting transcriptional outcomes of novel multigene perturbations with gears. Nature Biotechnology, 1-9.

[7] Li, B., Qi, X., Lukasiewicz, T., & **Torr, P**. (2019). Controllable text-to-image generation. Advances in Neural Information Processing Systems, 32.

[8] Shi, Y., Paige, B., & **Torr, P.** (2019). Variational mixture-of-experts autoencoders for multi-modal deep generative models. Advances in neural information processing systems, 32.

[9] Wang, G., Wang, K., Wang, G., **Torr, P. H.,** & Lin, L. (2021). Solving inefficiency of self-supervised representation learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 9505-9515).

[10] Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., ... **Torr, P. H.** & Zhang, L. (2021). Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 6881-6890).

[11] Ma, L., Lin, C., Lim, D., Romero-Soriano, A., Dokania, P. K., Coates, M., **Torr, P. H.** & Lim, S. N. (2023). Graph Inductive Biases in Transformers without Message Passing. ICML 2023.

[12] Yue, X., Sun, S., Kuang, Z., Wei, M., **Torr, P. H.,** Zhang, W., & Lin, D. (2021). Vision transformer with progressive sampling. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 387-396).

[13] Kirsch, A., Van Amersfoort, J., & Gal, Y. (2019). Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *Advances in neural information processing systems*, *32*.

[14] Chandak P, Huang K, Zitnik M. Building a knowledge graph to enable precision medicine. Sci Data. 2023;10(1):67. Published 2023 Feb 2. doi:10.1038/s41597-023-01960-3

[15] Pan S, Luo L, Wang Y, et al. Unifying Large Language Models and Knowledge Graphs: A Roadmap. 2023 arXiv:2306.08302 .

Supervisor's recent relevant publications

[7] Li, B., Qi, X., Lukasiewicz, T., & **Torr, P**. (2019). Controllable text-to-image generation. Advances in Neural Information Processing Systems, 32.

[8] Shi, Y., Paige, B., & **Torr, P.** (2019). Variational mixture-of-experts autoencoders for multi-modal deep generative models. Advances in neural information processing systems, 32.

[9] Wang, G., Wang, K., Wang, G., **Torr, P. H.,** & Lin, L. (2021). Solving inefficiency of self-supervised representation learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 9505-9515).

[10] Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., ... **Torr, P. H.** & Zhang, L. (2021). Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 6881-6890).

[11] Ma, L., Lin, C., Lim, D., Romero-Soriano, A., Dokania, P. K., Coates, M., **Torr, P. H.** & Lim, S. N. (2023). Graph Inductive Biases in Transformers without Message Passing. ICML 2023.