# Self-supervised Deep Learning Methods for Medical Image Analysis

**Aiham Taleb**

## *Abstract*

Deep learning has seen widespread application in many domains, mainly for its ability to learn data representations from raw input data. Nevertheless, its success has so far been coupled with the availability of large annotated (labelled) datasets. This is a requirement that is difficult to fulfil in several domains, such as in medical imaging. Annotation costs form a barrier in extending deep learning to clinically-relevant use cases. The labels associated with medical images are scarce, since the generation of expert annotations of multimodal patient data at scale is non-trivial, expensive, and time-consuming. This substantiates the need for algorithms that learn from the increasing amounts of unlabeled data. Self-supervised representation learning algorithms offer a pertinent solution, as they allow solving real-world (downstream) deep learning tasks with fewer annotations. Self-supervised approaches leverage unlabeled samples to acquire generic features about different concepts, enabling annotation-efficient downstream task solving subsequently.

Nevertheless, medical images present multiple unique and inherent challenges for existing self-supervised learning approaches, which we seek to address in this thesis: (i) medical images are multimodal, and their multiple modalities are heterogeneous in nature and imbalanced in quantities, e.g. MRI and CT; (ii) medical scans are multi-dimensional, often in 3D instead of 2D; (iii) disease patterns in medical scans are numerous and their incidence exhibits a long-tail distribution, so it is oftentimes essential to fuse knowledge from different data modalities, e.g. genomics or clinical data, to capture disease traits more comprehensively; (iv) Medical scans usually exhibit more uniform color density distributions, e.g. in dental X-Rays, than natural images. Our proposed self-supervised methods meet these challenges, besides significantly reducing the amounts of required annotations.

We evaluate our self-supervised methods on a wide array of medical imaging applications and tasks. Our experimental results demonstrate the obtained gains in both annotation-efficiency and performance; our proposed methods outperform many approaches from related literature. Additionally, in case of fusion with genetic modalities, our methods also allow for cross-modal interpretability. In this thesis, not only we show that self-supervised learning is capable of mitigating manual annotation costs, but also our proposed solutions demonstrate how to better utilize it in the medical imaging domain. Progress in self-supervised learning has the potential to extend deep learning algorithms application to clinical scenarios.