Machine learning (ML) is one of the most important areas in the field of artificial intelligence (AI), which is increasingly present in our daily lives. From the beginning, ML algorithms have played an important role in the development of computer-aided diagnostic systems aimed at improving the efficiency and accuracy of experts. In this context, medical imaging has been of particular interest, as computer vision (CV) techniques can automatically perform pattern recognition tasks to associate certain biomedical structures with a specific disease.

Over time, different imaging modalities have been used to address a wide range of diseases under the umbrella of CV. In this thesis, we focus on two important research areas at the forefront of medical imaging: digital pathology and ophthalmology. Specifically, we use histological images to assist pathologists in the diagnosis of prostate and bladder cancer, and optical coherence tomography (OCT) data to help ophthalmologists in glaucoma decision making. We propose different cutting-edge solutions based on traditional and deep learning methods, as well as hybrid approaches to leverage the strengthens of each.

In order to exploit the potential of AI in diagnostic imaging, we address several learning paradigms to cover different supervision scenarios. For histological imaging, we propose fully supervised methods for the segmentation and classification of prostate-specific structures, as well as fully unsupervised techniques for bladder-specific histological pattern recognition. Moreover, for glaucoma assessment, we rely on recurrent learning to detect glaucoma from SD-OCT volumes, and on few-shot learning to determine the severity level of glaucoma from circumpapillary B-scans.

In the prostate-based studies, we provide a comparison between hand-driven and deep learning methods for identifying the earliest stage of prostate cancer. The conventional ML approach shows better performance than deep learning in distinguishing between artefacts (false glands), benign and pathological glands, as hand-crafted features allow the computation of spatial hierarchies and orientations that are essential in this multi-class scenario. The proposed end-to-end system contributes to the accurate localisation and classification of prostate histological structures, achieving an accuracy of 88.30% in discriminating normal and Gleason grade 3 glands. In contrast, the proposed deep unsupervised algorithm far outperforms other conventional clustering algorithms in the classification of muscle-invasive bladder cancer (MIBC). Here, we resort to high-resolution histological samples stained with immunohistochemistry techniques to self-recognise non-tumour, mild and infiltrative MIBC patterns. The proposed model achieves a multi-class accuracy of 90.31% without incurring prior annotation steps, which bridges the gap with respect to training the model on labelled data.

Regarding glaucoma detection from SD-OCT volumes, we propose the combination of convolutional neural networks (CNNs) with long short-term memory (LSTM) units to find glaucoma-specific spatial dependencies between adjacent 2D slides of the SD-OCT cube. Important contributions to glaucoma detection are included in the architectures of both the slide-level feature extractor and the volume-based predictive model. The proposed recurrent learning system improves on other state-of-the-art approaches based on 3D architectures, reaching an accuracy of 81.25% in discerning between healthy and glaucomatous SD-OCT volumes. Delving deeper into glaucoma assessment, we address a novel learning strategy to discriminate, for the first time, between different levels of glaucoma severity from circumpapillary OCT B-scans. We propose a new hybrid backbone to optimise the feature extraction process and embed it in a novel few-shot learning scenario based on dynamic prototypical neural networks (PNN). The convolutional coefficients of the backbone are refined during model training according to the prototypical latent feature assignment, leading to higher performance compared to dense layers activated by softmax. At test time, the proposed model achieves accuracies of 96.97% and 87.88% in detecting and grading glaucoma, respectively.

In short, the AI methods proposed in this thesis contributes to the diagnosis of prostate and bladder cancer from histological images, as well as to glaucoma detection from OCT samples, making use of ML algorithms under different supervision scenarios.