

# Summary

Microscopy plays a role of paramount importance in biomedical research: the microscopes have historically been exploited to capture, investigate, quantify, and analyze various living organisms and structures at scales which are very distant below the normal human visual perception. In the last forty years, the advent of modestly priced, high-performance computers, and the developments of optics and image sensors, have importantly boosted digital microscope imaging, which today is replacing the conventional film-based photomicrography as the most universally used method for microscope image acquisition and storage.

Digital image processing is not only a spontaneous extension of microscope imaging, but it is essential for a proficient downstream data analysis and interpretation of microscope images. As a matter of fact, it allows to extract quantitative information about the specimen, and, importantly, it can transform an image so that the elaborated version is much more informative. Thus, digital processing of microscope images has opened the way to new horizons of medical/biological research, attracting a large body of recent literature. In this regard, machine learning, with its intrinsic capability of acquiring knowledge from the data representation itself, has notably boosted the digital microscope image processing of the last fifteen years. Nonetheless, conventional machine learning techniques do not directly deal with raw data, but conversely rely on the data representation, which requires solid domain expertise. This is a major limitation for the robustness and generalization capabilities of a given machine learning-based system.

On the contrary, deep learning, as a method of data representational learning, is able to learn from end to end by itself, without requiring any handcrafted features or human-based data representation. Today, deep learning has evolved in many aspects of scientific fields, but nowhere is more fundamental than in the image processing branch. Here, thanks to their outstanding performance, convolutional neural networks are certainly the most widely adopted deep learning-based algorithms.

Although convolutional neural networks have been recently recognized as an innovative, useful and extremely promising paradigm also in the biomedical field, there exist at least three different reasons which have hampered their large-scale adoption in a broad range of biological/medical applications:

- i The paucity of available annotated training data. Convolutional neural networks are data-hungry supervised algorithms, and thus supposed to be fed with a large amount of high quality annotated training samples. Nonetheless, associating labels to a massive number of images may be extremely problematic in a number of real-world applications, among which the medical and computational biology domains are significant examples. In such research branches, image annotation is indeed an especially cumbersome and time-consuming task that requires solid domain expertise and, more often than not, necessitates consensus strategies to aggregate annotations from several experts to solve class variability problems. This ultimately results in a scarcity of annotated dataset, which has greatly reduced the feasibility of deep learning methodologies in the bio-domain.
- ii The absence of a principled uncertainty measure associated with convolutional nets' predictions. Providing physicians and biologists with transparent and plausible outcomes is definitively of paramount importance to avoid deep networks to appear as black boxes, that absorb data and generate output to reach a conclusion. Furthermore, a model capable of reporting its lack of confidence could identify and flag difficult diagnostic cases, reducing misclassification as well as the risk of taking misleading research directions.
- iii The needing of deep learning-based solutions able to evolve as new information emerges. This situation is of paramount importance in the bio-domain, where there is still no complete knowledge about the causes and the progression of many heterogeneous diseases. Hence, the collected data is regularly growing and also the number of classes is continuously increasing. Ultimately, this aspect claims for systems able to learn new knowledge without forgetting the existing one.

This dissertation, specifically seeks to pose a solution to the above-mentioned open challenges, and it is organized into four main themed chapters.

The first part briefly introduce the outline of the dissertation, and summarize the main objectives as well as the chief contributions.

The second chapter investigates application-independent, deep learning-based data representations for microscopic bio-images. This examination includes: (i) a comprehensive analysis of supervised and unsupervised transfer learning for convolutional nets applied to biological images; (ii) a novel methodology able to avoid negative transfer in the challenging context of renal cell carcinoma sub-typing, by combining supervised transfer learning representation capabilities with expert-based knowledge; (iii) the introduction of a self-supervised learning embedding characterized on a particularly complex biological case-study, that is COVID-19 drug discovery based on fluorescent microscopy images.

The third chapter tries indeed to solve the lack of capability of traditional deep learning models of providing usable information about the uncertainty level of the obtained predictions. The main methodological contribution lies in a novel two-module convolutional network architecture exploiting Bayesian inference to identify and discard spurious images during the training, as well as to broadcast information on the prediction confidence at inference time. Such architecture is proficiently employed in a classification scenario from the digital pathology branch of research.

Chapter number four faces the so-called “catastrophic forgetting”, that is the inherent tendency of deep learning models of forgetting previously learned information upon learning new one. This is indeed an important limitation when a model needs to be adapted to a new classification task, as for instance new classes or sub-classes of a specific tumor. In this sense, a novel deep learning architecture for incremental learning is proposed, and experimental validated on the histological classification of nine classes of colon tissue.

Lastly, the fifth chapter concludes the dissertation, providing an overall overview of the main technical contribution.