AI Foundation Models for Breast Cancer Screening: Advancing Early Detection through Artificial Intelligence

December 12, 2023

Context and Motivation

With a global estimate of 2.3 million women diagnosed with breast cancer, 685,000 deaths globally in 2020, and 7.8 million women alive who were diagnosed with breast cancer in the past 5 years, breast cancer is the world's most prevalent cancer, according to the WHO. Appropriate patient selection strategies can, therefore, greatly impact the treatment outcomes and the life expectancy of the patients. To this end, tumor phenotyping, patient stratification, and thus treatment selection is mainly performed by the assessment of histopathological slides under a variety of tissue staining protocols by experienced pathologists. To enable the automatic processing of such tissue slides, they have been digitized, resulting in extremely large images, whose automatic understanding is quite challenging, even with existing state-of-the-art methods.

Technical limitations. Deep learning (DL) models have achieved remarkable performance in a wide span of visual recognition problems in strategic areas for our society, including many areas of healthcare, such as breast cancer care. Nevertheless, DL models present important drawbacks that limit their deployment in real-life scenarios, which could affect the lives of millions of people. First, modern deep learning models do not generalize well to unseen scenarios, such as novel classes or distributions presenting a domain shift, requiring large labeled training datasets for each novel task. Nevertheless, obtaining annotations in medical images is a cumbersome process, which requires user-expertise and suffers from inter/intra-rater variability. A common practice to overcome the scarcity of labeled datasets involves fine-tuning a pre-trained model on a different problem with annotated samples of the target task. However, this technique is inefficient, as it still requires a substantial amount of labeled target samples, and fine-tunes the entire network, which increases the computational burden. Second, recent evidence exposes that high capacity neural networks suffer from poor calibration [1, 3], i.e. the confidence scores of the predictions do not reflect the real world probabilities of those predictions being true. Thus, these networks tend to produce overconfident estimates, even in situations of high uncertainty, leading to poorly calibrated and unreliable models [2]. This is further magnified if the model is trained, or adapted, under a low labeled data regime, a popular learning paradigm to alleviate the need of large labeled training datasets. This results in a major concern, which can have catastrophic consequences in critical decision-making systems, such as medical diagnosis, where downstream decisions depend on predicted probabilities. It is therefore desired to empower learning strategies that, *i)* accommodate deep models to new tasks with a few labeled samples (i.e., few-shot), and *ii)* in a computationally efficient manner, and *iii)* abstain model decisions from unreliable predictions.

Our Exciting Opportunity

The field of DL is undergoing a paradigm shift with the emerging rise of models trained on broad data, typically using self-supervision at large-scale, and which can be adapted to a wide range of downstream tasks. These models, commonly referred to as *foundation* models have been fueled by the significant improvements achieved in computer vision and natural language processing. To adapt the model to novel tasks, initial attempts updated the whole network parameters. However, how to properly and efficiently adapt these models for new tasks was not yet studied. Last, calibrating neural networks has gained increased popularity, where post-processing and training-time approaches have emerged. Despite the importance of accurately modeling the uncertainty of deep models, the study of model calibration in foundation models remains unexplored.

The proposed research program will devise innovative and robust algorithms to understand the content of histology images in the context of breast cancer. These methods will set a new state of the art in providing powerful foundation models that can be adapted efficiently to novel tasks, while yielding well-calibrated predictions. In terms of **research outcomes**, we expect the results from this research to be published in the top venues of medicine (Nature Machine Intelligence), medical image processing (MedIA, IEEE TMI, MICCAI, IPMI), computer vision (CVPR, ECCV, ICCV) and/or machine learning (NeurIPS, ICLR, ICML). Furthermore, these novel learning strategies will have a significant impact in the general areas of semantic segmentation and classification in a broad span of disciplines. In terms of **clinical outcomes**, the proposed methods will provide an invaluable tool to support clinicians in the diagnosis, treatment and follow-up of breast cancer, with a high potential to expand it to multiple diseases. The developed algorithms will have a high impact on many healthcare areas, such as personalized medicine and safe deployment of AI approaches in clinical routine.

Research Project

The main goal of this project is to create an open-sourced foundation model for the assessment of breast cancer patients in histology images that can adapt efficiently to novel tasks –both in terms of labeled data and computational complexity– while yielding robust predictions, i.e., accurate uncertainty estimates. We are expected to design and develop novel AI-based models and algorithms for training specifically tailored for analyzing the content of histology images and the efficient adaptation of trained models to novel tasks, which also improve their predictive uncertainty. Extensive experiments on real-world clinical data sets will be executed to showcase the viability of our approach, benchmark its performance, and analyze its advantages, limitations, and areas for improvement.

Research questions. Some potential research questions for our consideration are:

- How domain expert knowledge, in the form of text prompts, can be integrated in the training of large vision-language models?
- How to properly adapt large vision-language models to novel tasks, including novel classes or new domains? Does the original training objective play a crucial role in the objective selected for adaptation?
- What to update during adaptation? Is whole fine-tuning a good strategy? or do more recent techniques such as adapters offer a better alternative?
- What information (statistical) measures and techniques can be leveraged during both training and adaptation to improve the model uncertainty?

This research will provide a deeper understanding of the underlying mechanisms to properly (i.e., efficiently) adapt large vision-models to novel tasks in the context of breast cancer.

Team supervision. The École de Technologie Supérieure of Montreal (ETS) and CentraleSupélec within the Université Paris-Saclay are opening two postdoctoral fellowship of 18 months. The work will be conducted within an international and dynamic environment between ETS and CentraleSupélec locations. The potential candidates are expected to have strong background at the intersection of Machine Learning, Computer Vision and Medical Imaging, who gets inspired by sciences and the opportunities to solve complex vision problems. They should have strong programming skills and a very good understanding of data science, and Machine Learning, as well as a strong publication track in recognized venues of computer vision, machine learning and/or medical image computing.

An international and stimulating environment for research. The International Laboratory on Learning Systems (ILLS) with the supports of the CNRS and the FRQ will promote international mobility between Canada and France, respectively, to facilitate collaborations with other students and professors. Other associated institutions to promote collaboration in Canada are: McGill University and École de Technologie Supérieure (ÉTS), and the Quebec AI Institute (Mila), and in France are: Université Paris-Saclay, CNRS, and CentraleSupélec, which are all major players in AI at the international. They are involved in many research, industrial and academic projects. *Our research team is formed of* Jose Dolz (ÉTS), Associate Member to Laboratory for Image, Vision and AI (LIVIA) and ILLS, Maria Vakalopoulou (CentraleSupélec) and Stergios Christodoulidis (CentraleSupélec), Associate Members to Laboratoire Mathématiques et Informatique pour la Complexité et les Systèmes (MICS), are experts in computer vision, particularly in the medical field, and Prof. Pablo Piantanida (CNRS, CentraleSupélec), director of ILLS and Associate Member to the Quebec AI Institute, is a expert in information theory and Machine Learning.

Position Qualifications

- PhD program in Computer Science, Machine Learning, Computer Engineering, Mathematics, or related field (e.g. applied mathematics/statistics).
- Very good understanding of Machine Learning theory and techniques, as well as of computer vision.
- Strong publication track in recognized venues of computer vision (CVPR, ECCV, ICCV), machine learning (NeurIPS, ICLR, ICML) and/or medical image computing (MedIA, IEEE TMI, MICCAI).
- Good programming skills in Python (PyTorch).
- Applications/ domain-knowledge in medical image processing is a plus.
- Good communication skills in written and spoken English.
- Creativity and ability to formulate problems and solve them independently.

How to apply. Applications should be sent by email to : jose.dolz@etsmtl.ca; pablo.piantanida@mila.quebec; maria.vakalopoulou@ecp.fr and stergios.christodoulidis@centralesupelec.fr.

If you are interested, please send us the following elements as soon as possible and **not later than January 15th**:

- Detailed CV.
- Letter of motivation.
- Details of transcripts (especially M1 and M2).
- Elements of bibliography or personal achievements related to a research activity (e.g. master project, research internship subject, etc.).
- 2 references or recommendation letters.

Applications with **missing elements** will not be considered.

References

[1] Gabriel Pereyra et al. “Regularizing neural networks by penalizing confident output distributions”. In: *International Conference on Learning Representations (ICLR)*. 2017.

[2] Agustinus Kristiadi, Matthias Hein, and Philipp Hennig. “Being bayesian, even just a bit, fixes overconfidence in relu networks”. In: *International Conference on Machine Learning (ICML)*. 2020.

[3] Bingyuan Liu et al. “The Devil is in the Margin: Margin-based Label Smoothing for Network Calibration”. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022, pp. 80–88.