Digital Health - Machine Learning, Chair of Prof. Dr. Lippert
Master's Thesis Proposal

Contact: benjamin.bergner@hpi.de

Digital Engineering · Universität Potsdam

Master Thesis: Gigapixel Medical Image Classification

Image classification is one of the most common problems in deep learning and has great practical value. For instance, in medical imaging, it can help medical professionals detect abnormal findings, such as cancerous tissue. While images used in deep learning are usually small (224x224 pixels or ~0.05 MP is common for the ImageNet dataset), medical images are usually much larger. Images used in radiology, such as X-rays or MR scans, often have a resolution of several megapixels, while whole-slide images used in digital pathology can have a resolution in the gigapixel range. From an engineering perspective, this is problematic because the GPUs on which the models run are limited by memory. Although there is little research addressing this important problem (e.g. [1,2,3]), a simple and efficient solution may enable novel digital health use cases that would otherwise require expensive pixel-wise expert annotations.

In this thesis, you will develop and evaluate a method capable of performing classification on large-scale images. Along the way, you will gain experience in processing medical images and applications. You can use existing ideas from the chair to get started. To be successful, you should have a solid understanding of deep learning and hands-on experience with a major framework, such as PyTorch (both of which you can, for example, acquire in our deep learning course). If you are interested in this topic, contact benjamin.bergner@hpi.de

References

[1] Katharopoulos, Angelos, and François Fleuret. "Processing megapixel images with deep attention-sampling models." International Conference on Machine Learning. PMLR, 2019.
[2] Pinckaers, Hans, Bram van Ginneken, and Geert Litjens. "Streaming convolutional neural networks for end-to-end learning with multi-megapixel images." arXiv preprint arXiv:1911.04432
[3] Papadopoulos, Athanasios, Paweł Korus, and Nasir Memon. "Hard-Attention for Scalable Image Classification." arXiv preprint arXiv:2102.10212 (2021).