End-to-end learning for detecting MYC translocations

Stephan Dooper [1] STEPHAN.DOOPER@RADBOUDUMC.NL
Konnie Hebeda [1] KONNIE.HEBEDA@RADBOUDUMC.NL
Geert Litjens [1] GEERT.LITJENS@RADBOUDUMC.NL

on behalf of
the BIGPICTURE Consortium
www.bigpicture.eu
[1] *Radboud University Medical Center, Department of Pathology, Nijmegen, the Netherlands.*

Editors: Under Review for MIDL 2022

Abstract

Recent developments have improved whole-slide image classification to the point where the entire slide can be analyzed using only weak labels, whilst retaining both local and global context. In this paper, we use an end-to-end whole-slide image classification approach using weak labels to classify MYC translocations in slides of diffuse large B-cell lymphoma. Our model is able to achieve an AUC of 0.8012, which indicates the possibility of learning relevant features for MYC translocations.

Keywords: MYC, DLBCL, End-to-end learning, whole slide image classification

1. Introduction

In recent years, digital pathology and deep learning have allowed for computational analysis of whole slide images (WSI), which has seen successes in, for example, tumor classification (Lu et al., 2021; Pinckaers et al., 2020). A further challenge explored recently is to see if deep learning can also recognize morphology representing underlying genetic changes, such as MYC translocations in diffuse large B-cell lymphomas (DLBCL) (Swiderska-Chadaj et al., 2020). The presence of these translocations generally indicates a more aggressive course of disease and can thus have significant treatment implications for patients. The presence of MYC rearrangements cannot be reliably confirmed through visual inspection on standard H&E by pathologists, and thus the ground truth is typically obtained via expensive and time-consuming fluorescence in situ hybridization (FISH) tests.

When translating this problem to the domain of machine learning, it can only be formulated as a weakly supervised classification problem, as only a single label for the whole-slide is available. Since images in digital pathology are typically gigapixels in size, computer memory becomes a bottleneck.

In this study we propose to use the streaming CNN algorithm (Pinckaers et al., 2020), which directly solves the memory bottleneck, and allows for full end-to-end training of convolutional neural networks with whole-slide images, whereas the approach of (Swiderska-Chadaj et al., 2020) used a patch-based approach with less data. We use the proposed method to classify MYC translocations in H&E stained slides of DLBCL, and show in our preliminary results that our model has the ability to learn features that are relevant to MYC translocations from morphology, achieving an AUC of 0.8012.

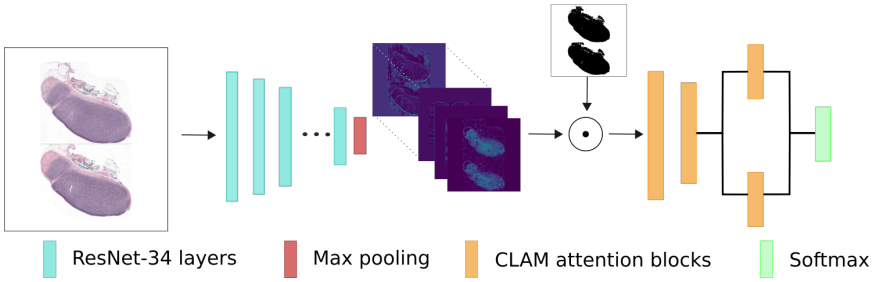

Figure 1: A whole-slide image is forwarded into the encoder. Final feature maps of the encoder are masked to reduce white space. The resulting embeddings are then forwarded to the decoder.

2. Methods

We use the streaming CNN algorithm to train the architecture described in (Lu et al., 2021) end-to-end, which consist of a ResNet encoder and an attention decoder. The advantage is that the entire model can be trained, maintaining both global and local context.

We modify the architecture in several ways to accommodate for streaming. First, we change the encoder backbone from a ResNet-50 to a ResNet-34, since the latter network takes less time to train. Secondly, we replace the global average pooling layer at the end of the network by a single 8×8 max-pool layer with stride 8 to reduce total downsampling and maintain enough spatial context for classification.

We also add a masking operation at the end of the encoder to filter out non-tissue areas before being passed to the decoder, which decreases the number of input neurons. We use a tissue background segmentation algorithm (Bándi et al., 2019) to generate the tissue masks.

The model is trained in two stages. In the first stage, the ResNet-34 encoder weights are frozen, and only the attention decoder is trained. In the second stage, we unfreeze the encoder weights and train the encoder and decoder simultaneously. The ResNet-34 model is initialized with pre-trained ImageNet weights, and the decoder weights are randomly initialized using Glorot initialization. Both stages use an Adam optimizer with a learning rate of $1e-4$ and $1e-5$, respectively.

3. Experimental results

We collected 385 H&E stained DLBCL slides, of which 156 MYC positive, from Radboud University Medical Center (Nijmegen, the Netherlands), and Rijnstate Hospital (Arnhem, the Netherlands) which were then digitized and scanned at 0.25μm. Additionally, we use two independent clinical trials from the Haemato Oncology Foundation for Adults in the Netherlands (HOVON) to act as an independent test set.

The first trial, HOVON HO84 NHL, is a randomized phase III study on the effect of early intensification of rituximab in combination with 2-weekly CHOP chemotherapy followed by rituximab maintenance in patients with diffuse large B-cell lymphoma, and contains 127 H&E slides of individual patients, of which 9 MYC positive. The second trial, HOVON HO130, is a phase II study evaluating the effect of the addition of lenalidomide to R-CHOP for patients with newly diagnosed MYC positive DLBCL and BCL-U, which contains 56 slides.

In total, the HOVON dataset consists of 183 slides, and have been gathered from and stained at various hospitals around the Netherlands. During training we oversample cases with MYC translocations to deal with the imbalance in the dataset. For regularization, we apply an augmentation strategy consisting of random vertical and horizontal flips, rotations, and color augmentations by perturbing the HSV color space. We train the model on input images of 65536×65536. To achieve this, the individual biopsies are padded/cropped to the specified size, which fit 99% of the data.

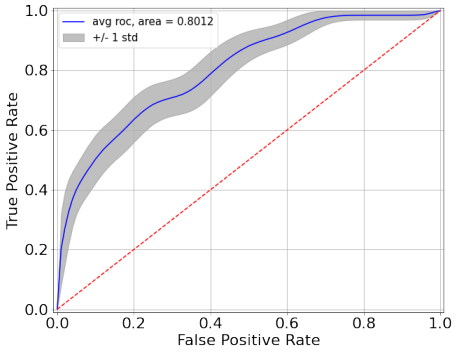

Figure 2: ROC curve on the HOVON test set.

4. Conclusion

Figure 2 shows the ROC achieved by our model on the independent test set, where the confidence intervals are generated using 10000 bootstrap samples. We achieve an average AUC of 0.8012. The output of our model indicates that it might be possible to locate features in H&E stained images of DLBCL that are relevant towards detecting MYC translocations.

Acknowledgments

This project has received funding from the Innovative Medicines Initiative 2 Joint Undertaking under grant agreement No 945358. This Joint Undertaking receives support from the European Union's Horizon 2020 research and innovation program and EFPIA (www.imi.europe.eu).

References

Péter Bándi, Maschenka Balkenhol, Bram Ginneken, Jeroen Laak, and Geert Litjens. Resolution-agnostic tissue segmentation in whole-slide histopathology images with convolutional neural networks. *PeerJ*, 7:e8242, 12 2019. doi: 10.7717/peerj.8242.

Ming Y. Lu, Drew F. K. Williamson, Tiffany Y. Chen, Richard J. Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. In *Nature Biomedical Engineering*, 2021.

Hans Pinckaers, Bram van Ginneken, and Geert Litjens. Streaming convolutional neural networks for end-to-end learning with multi-megapixel images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. doi: 10.1109/TPAMI.2020.3019563.

Zaneta Swiderska-Chadaj, Konnie Hebeda, Michiel van den Brand, and Geert Litjens. Predicting MYC translocation in HE specimens of diffuse large B-cell lymphoma through deep learning. In John E. Tomaszewski and Aaron D. Ward, editors, *Medical Imaging 2020: Digital Pathology*, volume 11320, pages 238 – 244. International Society for Optics and Photonics, SPIE, 2020. URL https://doi.org/10.1117/12.2549650.