A Semi-Supervised Deep Learning Approach for Multi-Stain Foreground Segmentation in Digital Pathology

Agathe de Vulpian* ADEVULPIAN@GMAIL.COM
Valentina di Proietto* VALENTINA.DIPROIETTO@GMAIL.COM
Gauthier Roy* GAUTHIER.ROY@STUDENT-CS.FR
Saima Ben Hadj* SBENHADJ@TRIBUN.HEALTH
Rutger H.J. Fick* RFICK@TRIBUN.HEALTH

Editors: Under Review for MIDL 2022

Abstract

The analysis of whole computational pathology slides can often be accelerated by excluding background areas from the analysis. Deep learning has proven to be superior to signal processing techniques to robustly recover the foreground in HE Images. However, naively generalizing this technique to the wide variability of histological stains used in practice would require annotations in all stain domains. To avoid this, we propose a method which leverages tissue annotation from a single stain to perform foreground segmentation in slides with other non-annotated stains.

Keywords: digital pathology, foreground segmentation, semi-supervised learning, GANs.

1. Introduction

Deep learning has proven to be a valuable tool for various clinical applications in computational pathology. A preliminary step for many such applications is foreground detection. This step is crucial to avoid processing clinically irrelevant regions, thus reducing the computing time. To extract various types of tissue information pathologists use different stains, rendering different tissue structures in different colours (Alturkistani et al., 2016). A robust foreground detection model must be able to discern foreground from background no matter the contrast and color of the foreground.

To obtain the foreground in Hematoxylin & Eosin (HE) slides, deep learning-based segmentation has shown to be more effective than classic signal processing techniques (Riasatian et al., 2020). However, to create a segmentation model that generalizes to multiple stains, classic training procedures require manual annotations for each stain. Given the vast variability of stains, this would incite a considerable annotation burden. Instead, we propose a training approach that can produce models, each of which can generalize to any target stain, while only requiring annotated data for one source stain.

2. Materials and Methods

Our proposed solution is a deep learning-based semi-supervised training approach requiring annotations from a single stain domain only. Our innovation consists of (1) introducing a semi-supervised loss function imposing similar predictions for paired images of different domains and (2) using a Residual cycle GAN to generate these cross-stain paired images.

* Tribun Health, Paris, France

© 2021 A.d. Vulpian, V.d. Proietto, G. Roy, S. Ben Hadj & R.H. Fick.

2.1. Data

We use two data sets: a subset of the Cancer Genome Atlas (TCGA) data set of HE slides with tissue annotations provided by the authors of (Riasatian et al., 2020). Furthermore, a subset of the Academia and Industry Collaboration for Digital Pathology (AIDPATH) data set[1], a non-annotated dataset of unpaired HE and Immunohistochemistry (IHC) slides. We manually annotated 20 IHC slides to allow the estimation of test metrics only.

2.2. Model Training using Semi-Supervised Loss Function

We use a composite loss function to train an EfficientNet-B3 segmentation model using ADAM, consisting of a supervised and unsupervised loss:

$$\mathcal{L}(I_{\mathcal{D}_1}, I_{\mathcal{D}_2}, M_{\mathcal{D}_1}, \lambda_e | \mathcal{G}_{\mathcal{D}_2 \to \mathcal{D}_1}) = \mathcal{L}_{sup} + \lambda_e \mathcal{L}_{unsup} \tag{1}$$

$$= ||\mathcal{T}(I_{\mathcal{D}_1}) - M_{\mathcal{D}_1}||_2^2 + \lambda_e ||\mathcal{T}((I_{\mathcal{D}_2}) - \mathcal{T}(\mathcal{G}_{\mathcal{D}_2 \to \mathcal{D}_1}(I_{\mathcal{D}_2}))||_2^2. \tag{2}$$

Let $I_{\mathcal{D}_n}$ an image in domain $\mathcal{D}_n$, where $\mathcal{D}_1$ and $\mathcal{D}_2$ represent respectively HE and IHC in our case. $M_{\mathcal{D}_1}$ the annotated tissue mask in domain $\mathcal{D}_1$, $\mathcal{G}_{\mathcal{D}_2 \to \mathcal{D}_1}$ a generator that transforms IHC slides to look like HE slides. λ_e the regularization weight, which we increase linearly from 0-1 over the first 30 epochs e to avoid enforcing equality between noisy masks.

Globally, the supervised loss $\mathcal{L}_{sup}$ guides the network to minimize the difference between the predicted mask and the HE-based ground truth (Riasatian et al., 2020), while the novel unsupervised loss $\mathcal{L}_{unsup}$ guides the network to minimize the difference between two predicted masks on a target IHC-stained image and a *paired* HE-stained-looking transformation of that image. We obtain this transformation using a Residual Cycle GAN (de Bel et al., 2021) trained on unpaired HE and IHC AIDPATH slides.

3. Results

We conduct an ablation study to understand what aspect of our solution most contributed to its performance. In Figure 1 we compare the approach performance both qualitatively based on two example HE and IHC-stained slides, using the optimal F1-based threshold for each approach, and quantitatively based on the PR-AUC.

Our ablation consists of (i) our model, (ii) **Ablation 1** which only uses $\mathcal{L}_{sup}$ for training, (iii) **Ablation 2** which uses pairs of images obtained through customized color deconvolution-based stain transfer (instead of cycle GAN) for $\mathcal{G}_{\mathcal{D}_2 \to \mathcal{D}_1}$, and (iv) Otsu adaptive threshold method (Otsu and Yoshida, 1982). We apply all models at magnification 1, corresponding to a resolution of about 10 μm/pixel.

It can be seen that using classic Otsu thresholding performs badly compared to deep learning approaches, it does significantly worse on faint IHC images compared to HE images (0.883 vs 0.573). Using the HE-only Ablation 1 by (Riasatian et al., 2020) performs best on HE, but does not generalize well to IHC (0.995 vs 0.958), resulting in spotty segmentation. Ablation 2 using color deconvolution-based stain transfer improves performance on IHC

1. https://mitel.dimi.uniud.it/aidpath-db/app/login.php

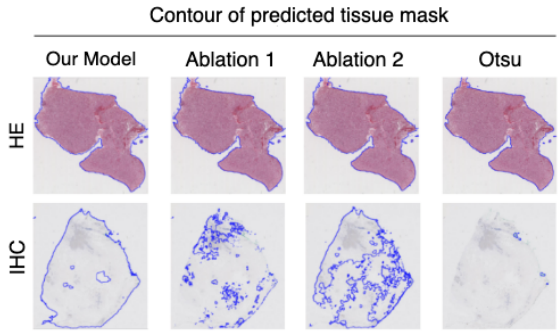

Model	PR-AUC (std)	
	HE	IHC
Our model	0.993 (0.013)	**0.984** (0.013)
Ablation 1	**0.995** (0.010)	0.958 (0.052)
Ablation 2	0.993 (0.013)	0.971 (0.017)
Otsu's method	0.883	0.573

Figure 1: (a) Contour of the predicted tissue mask in an HE and IHC slide (b) PR-AUC (and standard deviation) on HE and IHC test set

while slightly decreasing HE performance (0.993 vs 0.971). Using our approach using cycle-GAN-based stain transfer performs best on IHC while keeping HE performance stable (0.993 vs 0.984). Moreover, the qualitative results show that our method obtains the smoothest segmentation encompassing all the clinically relevant tissue, while the others - in particular Otsu - tend to miss faint parts of the image.

4. Discussion and Conclusion

In this work, we presented a training approach for foreground segmentation that is generalizable to multiple stains, while only needing annotation in one source stain domain. We demonstrated our approach transferring strong HE performance to faint IHC slides and compared to ablations of our approach. The key innovation of our approach is the composite semi-supervised training approach, using both supervised and unsupervised loss terms. Our approach has the potential to accelerate automatic computational pathology pipelines.

References

Hani A Alturkistani, Faris M Tashkandi, and Zuhair M Mohammedsaleh. Histological stains: a literature review and case study. *Global journal of health science*, 8(3):72, 2016.

Thomas de Bel, John-Melle Bokhorst, Jeroen van der Laak, and Geert Litjens. Residual cyclegan for robust domain transformation of histopathological tissue slides. *Medical Image Analysis*, 70:102004, 2021.

Takayuki Otsu and Masatoshi Yoshida. Role of initiator-transfer agent-terminator (iniferter) in radical polymerizations: Polymer design by organic disulfides as iniferters. *Die Makromolekulare Chemie, Rapid Communications*, 3(2):127–132, 1982.

Abtin Riasatian, Maral Rasoolijaberi, Morteza Babaei, and Hamid R Tizhoosh. A comparative study of u-net topologies for background removal in histopathology images. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020.