



HAL
open science

A new color augmentation method for deep learning segmentation of histological images

Yang Xiao, Etienne Decencière, Santiago Velasco-Forero, Hélène Burdin, Thomas Bornschlögl, Françoise Bernerd, Emilie Warrick, Thérèse Baldeweck

► **To cite this version:**

Yang Xiao, Etienne Decencière, Santiago Velasco-Forero, Hélène Burdin, Thomas Bornschlögl, et al.. A new color augmentation method for deep learning segmentation of histological images. 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI), Apr 2019, Venise, France. hal-02167903

HAL Id: hal-02167903

<https://minesparis-psl.hal.science/hal-02167903>

Submitted on 28 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A NEW COLOR AUGMENTATION METHOD FOR DEEP LEARNING SEGMENTATION OF HISTOLOGICAL IMAGES

Yang Xiao¹, Etienne Decencière¹, Santiago Velasco-Forero¹, H el ene Burdin²
Thomas Bornschl ogel³, Fran oise Bernerd³, Emilie Warrick³ and Th er ese Baldeweck³

¹ MINES ParisTech, PSL Research University, Centre for Mathematical Morphology, France

²ADCIS SA, Saint-Contest, France

³L'Or al Research and Innovation, Aulnay-sous-Bois, France

ABSTRACT

This paper addresses the problem of labeled data insufficiency in neural network training for semantic segmentation of color-stained histological images acquired via Whole Slide Imaging. It proposes an efficient image augmentation method to alleviate the demand for a large amount of labeled data and improve the network's generalization capacity. Typical image augmentation in bioimaging involves geometric transformation. Here, we propose a new image augmentation technique by combining the structure of one image with the color appearance of another image to construct augmented images on-the-fly for each training iteration. We show that it improves performance in the segmentation of histological images of human skin, and also offers better results when combined with geometric transformation.

Index Terms— color-stained slide, deep learning, segmentation, color transfer, histopathology, Fontana Masson, skin.

1. INTRODUCTION

Histological images of plant and animal cell tissues allow us to explore their structures and functions. Image segmentation is a crucial first step in many image analysis tasks, especially in histopathology, and aims at identifying accurately the presence, number, distribution, size, or morphology of certain tissue features (specific cells, nuclei, ...). In the dermatological field, it is used in a range of applications including melanoma detection and the assessment of histopathological damage of the skin [1]. With the recent advent of digital and whole slide imaging, the number and the size of acquired images are growing up and there is a need of finding ways to also adapt the throughput of image quantification.

A review of segmentation methods of color-stained histological images of pathological skin (lymphoma) has been presented in [2], which covers various methods based on regions, thresholding, K-means, graph-cut, and watershed transform. Recent advances in deep learning have enabled automatic methods for image segmentation using convolutional neural networks (CNN) [3–5]. By extending CNN to fully convolutional networks (FCN), we can train a network that segments arbitrary-sized images without redundant computation [6–8]. Nevertheless, these deep neural networks usually require large training sets to achieve an acceptable performance, while the generation of the segmentation ground-truth necessary for supervised learning is very time-consuming. Another challenge in the segmentation of histological images using deep learning is that the network generalization could be influenced by the complex tissue structures and inconsistencies in sample preparation [9].

The aim of this paper is to show that a deep neural network can learn a satisfactory segmentation model with relatively few data, thanks to a convenient image augmentation method. This is obtained by using an image augmentation technique that exploits the color transformation between different images, with a specific attention on the stained components within each sample. This allows the network to learn invariance to such variation, without the need to see these transformations in the labeled data [10], which is particularly important in the segmentation of histological images of human skin since the color variation is one of the most common variations [11, 12]. It is shown below that such transformations can be efficiently implemented.

To augment the available labeled data for training, some people make use of simple geometric transformation such as image rotation and translation for achieving the invariance to irrelevant spatial factors [5, 6], while others explore the combination of geometric and photometric augmentation techniques to increase the robustness to differing illumination color and intensity [13]. In this paper, we propose a novel image augmentation method working in the color space of the images, and combine it with existing geometric augmentation techniques for increased variation generation. We demonstrate fast and accurate results on histological images of human skin and we provide a direct comparison with other methods.

The presented work contributions are twofold: (1) It proposes a new image augmentation method adapted to the histological images with various color appearances; (2) Experimental results illustrate the good performance of the proposed method, which outperforms traditional ones in deep learning frameworks.

2. MATERIAL

2.1. Histological image of human skin

Skin is an epithelial tissue which possesses a specific layered structure: a layer of *stratum corneum* (SC) located on the top of a layer of Living Epidermis (LE), and the dermis (see in Fig 1). Besides, the three interfaces between these layers are respectively named Surface, Internal Epidermis Boundary (IEB), and Dermal-Epidermal Junction (DEJ).

In this paper we deal with histological images of normal and lesional human skin (Fontana Masson staining). Our aim is to segment *stratum corneum* (labeled SC) and living epidermis (labeled LE): for that purpose, other components out of SC and LE were labeled as Background (BG).

Table 1: Description of databases

Database	Number	Image size in pixels	Memory
Database1	76	$0.5 \times 10^6 - 1.6 \times 10^7$	835 Mb
Database2	52	$0.7 \times 10^6 - 1.1 \times 10^7$	420 Mb




Fig. 1: Histological image of lesional human skin (with high DEJ structural deformations) showing its main compartments and corresponding boundaries.

2.2. Database description

In order to test the generalization capabilities of different models, and in particular the model using the proposed color augmentation method, two databases have been collected. The first one, **Database1** (76 images), is used for network training. It contains images from two clinical studies including paired lesional and non-lesional samples. Among the 76 images in Database1, we randomly selected 35 images for training (26 images) and validation (9 images), while the remaining images (41) are used for testing. Moreover, in order to assure the independence of images from different subsets, images from the same histological samples were distributed to the same subset since they have very similar appearances.

The second one, **Database2** (52 images), is used to evaluate the network generalization capacity. It contains images coming from a third clinical study, which incorporates different color appearances compared to Database1. Table 1 summarizes the characteristics of these two databases.

3. METHOD

The field of data augmentation is not new, and various data augmentation techniques have been applied to specific problems. In image classification, data augmentation methods artificially create training images by altering available images [14]. Previous works [15–17] have shown its effectiveness to reduce overfitting, thus increasing the quality of generalization on new data. As data augmentation should be adapted to the intrinsic nature of training samples, the proposed image augmentation method focus on the color transformation of the stained components contained in different histological images.

The main idea of color augmentation, inspired by Reinhard [18], is to impose one image’s color characteristics on another using a statistical analysis. In our method, instead of transferring the color of the whole image, the transformation is limited to the stained components in the histological images. Such transformation should augment the color variations of training images used in the training process. Based on these augmented images, the neural network’s performance has been largely improved on the histological images of human skin from outside the training set, where various color appearances are present.

3.1. Lab color space

The CIE-Lab color space endows the color space with a perceptually meaningful measure of Euclidean distance as color similarity, and it is related to the RGB color space through a complex transformation equation [19]. In RGB color space, the color information is separated into three channels but the same three channels also encode brightness information. On the contrary, in Lab color space, the lightness channel L is independent of color information and only encodes brightness, while the other two channels are chromatic yellow-blue and red-green opponent channels.

Another advantage of Lab color space is that the selection of stained components can be achieved by a simple thresholding operation in the lightness channel. According to the Beer-Lambert law mentioned in [12], the transmission of light through a material can be modeled as

$$I = I_0 e^{-\alpha c x} \quad (1)$$

where I_0 is the intensity of the incident light, I is the intensity of light after passing through the medium, α is the absorption coefficient, c the concentration of absorbing substance, and x the distance traveled through the medium. α and x are assumed to be constant for a specimen and a given stain, while c can vary between different images and within the same image. Thus, in the histological images, the principal stained components with a large concentration of absorbing substance should have a lower intensity compared to other parts where fewer or no absorbing substance is present. In the Lab color space, as the brightness information is encoded in the L channel, these components can be extracted by selecting the pixels such that their L channel values are lower than a threshold value (set as $0.86 \times L_h$ in our work, with L_h being the highest possible value of the L channel).

3.2. Color transfer

In the area of automatic image analysis, several works have been proposed to address the problem of stain inconsistency by pre-processing images using stain normalization techniques, in which all images from a dataset are mapped to a user-selected reference image [11, 20, 21]. As this method is very sensitive to the choice of reference image, others proposed to normalize the stains in an adversarial frameworks [22], which eliminates the need for an expert to pick a representative reference image [23, 24].

In this work, we propose a stain-focused image augmentation technique to augment training images using color -matching. The color variation between different histological images mainly comes from the stains, while the background remains bright (see for images (a), (b) and (d) in Fig 2). Thus, the color transformation should be applied to the stains rather than the whole image. With this method, we aim to transfer the color appearance of a histological image towards another one without modifying the background. More specifically, for each training image, a target image is randomly selected from the training set at each iteration, and an augmented image is generated through a color transfer defined as below:

$$C_{\text{transferred}} = C_{\text{original}} - \bar{C}_{\text{original}} + \bar{C}_{\text{target}} \quad (2)$$

where \bar{C} is the mean of channel C calculated over the principal stained components within an image. This translation is applied to each channel of the image. Using this image augmentation technique, the augmented image takes the color appearance from the target image while its structure is the same as in the original image (see Fig 2).

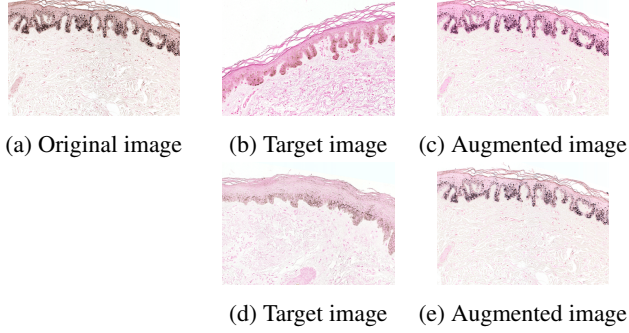


Fig. 2: Examples of two augmented images (c and e) obtained from an original image (a) using two different target images (b and d). This augmentation is applied to crops of size 512×512 in training.

3.3. Data preparation and preprocessing

To use the data augmentation method proposed in Section 3, images in RGB color space are transformed into Lab color space for augmentation using the color transfer. Besides, we use a geodesic reconstruction to cope with non-local information within fully convolutional networks [25]. Then, as the histological images in the datasets are of huge yet variable sizes, crops of size 512×512 are extracted from them and used for training (158 crops) and validation (53). When necessary in an experiment, image augmentation methods are applied to these crops on-the-fly during training. Moreover, crops containing only the background are removed since no interface appears in them.

3.4. Network training

In this paper, we consider a U-Net architecture [8], a typical FCN used in biomedical imaging, which consists of an "encoding-decoding" architecture for extracting high-level information without losing the object details. Here, the window size is 3×3 for convolution and 2×2 for max-pooling and upsampling. In our network, four layers of downsampling/upsampling are contained in the encoding/decoding path, with the output number of filters being 16 after the first convolution layer. At the end of the decoding path, a 1×1 convolution with the sigmoid activation is applied to have a channel dimension equaling the number of classes in the segmentation task. Thus, three probability maps of the same spatial dimensions as the input image is obtained at the output of the network.

Besides the network's architecture, another essential element in deep learning is the loss function. Inspired by [26], who proposed a differentiable version of the Jaccard distance to measure the dissimilarity between two sets, we encode the ground truth into one-hot vectors and define the loss function as below:

$$L_J = 1 - \frac{\sum_{i,j,k} (t_{ijk} p_{ijk})}{\sum_{i,j,k} t_{ijk}^2 + \sum_{i,j,k} p_{ijk}^2 - \sum_{i,j,k} (t_{ijk} p_{ijk})} \quad (3)$$

where $t_{ijk} = 1$ if the true class of pixel I_{ij} in the input image is k ($t_{ijk} = 0$ otherwise), and p_{ijk} represents an estimated probability that this pixel belongs to class k .

3.5. Post-processing

After training the network on crops of size 512×512 , we can apply it to images of arbitrary size as long as they can be passed on GPU

Table 2: Training results on Database1

Method	Time	Best train loss	Best val loss
No aug	14 mins	0.0060	0.0307
Geo	44 mins	0.0133	0.0243
Color	22 mins	0.0114	0.0272
Mix	58 mins	0.0178	0.0211

for network prediction. Another constraint on the image size is that it has to be the multiples of 16 (2^4) since four 2×2 max-pooling layers are included in the network. Then, as proposed in [25], the following post-processing is applied. For the *stratum corneum* and the living epidermis, only the largest connected component is kept; for the background, the connected components touching the top or the bottom of the image are kept. Based on these components, the final segmentation result is obtained through a watershed transform.

4. EXPERIMENTAL RESULTS

To illustrate the performance of the proposed augmentation method, we trained four networks with different methods on Database1. Several experiments were conducted to define the parameters for geometric transformation: 5 degrees of rotation range, 0.1 total width for horizontal shift range, 0.1 total height for vertical shift range, random horizontal flip, and interpolated by nearest value. 'No aug' refers to no augmentation; 'Geo' refers to augmentation with geometric transformation; 'Color' refers to augmentation with the proposed color transfer and 'Mix' combines our color transfer with geometric transformation to augment the variability of the labeled data.

The networks were implemented using Keras with TensorFlow backend and were trained on an NVIDIA Titan-X GPU of 11 Gb memory. All networks were trained using the same architecture and loss function. Besides, we used Adadelta optimizer with the default parameters proposed by [27], during 200 epochs and patience of 50. For the learning process, online learning (one training sample for each iteration) was applied while the augmented images were constructed on-the-fly.

4.1. Results on the test set from Database1 (41 images)

In Table 2, we present the best training and validation losses obtained during the training process. Convergence seems satisfactory in most cases. However, a problem of overfitting appeared in the un-augmented training process. Data augmentation has mitigated overfitting, while the mixed augmentation achieved the best validation loss with only a tiny increase compared to the training loss.

For evaluating the network's performance on the test set, we calculated a Jaccard index [28] for each class, also known as the Intersection over Union (IoU), which is commonly used in the evaluation of medical image segmentation. Besides, as histological images of human skin possess a specific layered structure, correct segmentation results should contain three interfaces as shown in Fig 1. Therefore, for each interface within a slide, we can calculate a mean spatial distance \bar{D} between the interface in the ground truth I_{gt} and the interface predicted by the network I_{pre} , defined as:

$$\bar{D} = \frac{1}{2} \left[\frac{1}{|I_{gt}|} \sum_{p \in I_{pre}} d(p, I_{gt}) + \frac{1}{|I_{pre}|} \sum_{p \in I_{gt}} d(p, I_{pre}) \right] \quad (4)$$

where $d(\cdot)$ is the Euclidean distance calculated in pixels. Finally, for




Fig. 3: Three examples of segmentation results on Database2 obtained with networks trained with various image augmentation techniques (BG in black, SC in gray, and LE in white). From left to right: original image, ground-truth, no augmentation, geometric transformation, color transfer, combined augmentation.

Table 3: Test results on Database1

Method	Jaccard index (per class)			\bar{D} (per interface)		
	BG	SC	LE	Surface	IEB	DEJ
No aug	0.97	0.77	0.87	54.6	311.5	7.9
Geo	0.99	0.82	0.89	16.3	30.9	5.1
Color	0.99	0.89	0.91	2.1	9.7	4.3
Mix	0.99	0.89	0.92	2.0	9.7	3.6

Table 4: Test results on Database2

Method	Jaccard index (per class)			\bar{D} (per interface)		
	BG	SC	LE	Surface	IEB	DEJ
No aug	0.92	0.47	0.70	156.6	992.5	58.5
Geo	0.95	0.74	0.70	19.2	321.5	72.8
Color	0.99	0.91	0.91	1.8	9.0	5.1
Mix	0.99	0.92	0.92	1.4	8.0	4.2

each interface, a mean distance averaged over all the test set is calculated to determine the segmentation’s quality. In the case where a certain interface is not detected in the segmentation result, the corresponding mean distance would take the value of the image diagonal as a penalization.

Table 3 presents the test results on the test set from Database1. While the Jaccard index of the BG class is equally improved using different data augmentation techniques, the networks trained with color transfer give better results for the Jaccard indexes of SC and LE compared to the un-augmented training and the augmented training using geometric transformation.

In terms of mean distances, networks trained with data augmentation introduce a great improvement compared to the un-augmented training. Moreover, networks trained with the proposed method of color augmentation give better results than geometric transformation, while the best result is provided by the combination of these two augmentation techniques.

4.2. Results on Database2 used as test set (52 images)

In the real application, the trained networks would be applied to images from different studies. In order to verify the networks generalization, we applied them to Database2, which consists of images coming from a different study than Database1. Fig 3 shows examples of segmentation results on this database.

Quantitative evaluation results are given in Table 4. Firstly, the mean distance of IEB detected by the un-augmented training network is much larger than the other networks, which is a consequence of the segmentation results where no corresponding IEB is present (see the third image of the first line in Fig 3). Secondly, an improvement by a large margin has been achieved by networks using the

color transfer, especially in the aspect of mean distances. The performance difference between the network using only geometric augmentation and the networks using color augmentation is much larger in the case of Database2. This shows that the proposed method brings a welcome generalization capacity to the models dealing with this kind of histological data. Last but not least, the best result on Database2 is given by the network trained with the combined augmentation, which is also true for Database1.

5. CONCLUSIONS

We present an image augmentation method to automatically increase the color appearance variety of color-stained histological images used for neural network training, which yields improved performance in comparison to a typical augmentation technique. We evaluated the networks trained with various methods on histological images of human skin, where the color transformation is one of the most important variations between different images. By applying the networks to a generalization database containing histological images coming from a different study, a satisfactory result was obtained by using our method.

The main constraint with our algorithm is that its performance is very dependent on the color variety existing in the training set, which could be insufficient to cover all the variations observed in the real application. We consider extending our method of color augmentation with elastic deformation [8] to increase the appearance variety existing in the training images. Additional future work may include combining image augmentation with stain normalization techniques, in which a generative adversarial network, such as [24, 29, 30], could be used to automatize the whole process.

6. REFERENCES

- [1] J. Haggerty, X. Wang, A. Dickinson, C. O'Malley, and E. Martin, "Segmentation of epidermal tissue with histopathological damage in images of haematoxylin and eosin stained human skin," *BMC medical imaging*, vol. 14, pp. 7, 2014.
- [2] T. Azevedo Tosta, L. Neves, and M. do Nascimento, "Segmentation methods of H&E-stained histological images of lymphoma: A review," *Informatics in Medicine Unlocked*, vol. 9, pp. 35–43, 2017.
- [3] K. Fukushima, "Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position," *Biological Cybernetics*, vol. 36, no. 4, pp. 193–202, Apr 1980.
- [4] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural Comput.*, vol. 1, no. 4, pp. 541–551, Dec. 1989.
- [5] D. Cireřan, A. Giusti, L. Gambardella, and J. Schmidhuber, "Deep neural networks segment neuronal membranes in electron microscopy images," in *Proceedings NIPS'12, USA*, 2012, Curran Associates Inc.
- [6] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," *CoRR*, vol. abs/1411.4038, 2014.
- [7] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *CoRR*, vol. abs/1511.00561, 2015.
- [8] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI (3)*. 2015, vol. 9351 of *Lecture Notes in Computer Science*, pp. 234–241, Springer.
- [9] D. Komura and S. Ishikawa, "Machine learning methods for histopathological image analysis," *CoRR*, vol. abs/1709.00786, 2017.
- [10] H. Zhang, M. Cisse, Y. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *ICLR*, 2018.
- [11] D. Magee, D. Treanor, D. Crellin, M. Shires, K. Smith, K. Mohee, and P. Quirke, "Colour normalisation in digital histopathology images," in *MICCAI Workshop*, 2009.
- [12] J. Vicory, H. D. Couture, N. E. Thomas, D. Borland, J. S. Marron, J. T. Woosley, and M. Niethammer, "Appearance normalization of histology slides," *Comput Med Imaging Graph*, vol. 43, pp. 89–98, 2015.
- [13] S. Hauberg, O. Freifeld, A. Boesen Lindbo Larsen, J. Fisher, and L. Kai Hansen, "Dreaming more data: Class-dependent distributions over diffeomorphisms for learned data augmentation," in *AISTATS*, 2016.
- [14] Luis Perez and Jason Wang, "The effectiveness of data augmentation in image classification using deep learning," *CoRR*, vol. abs/1712.04621, 2017.
- [15] C. Nader Vasconcelos and B. Nader Vasconcelos, "Increasing deep learning melanoma classification by classical and expert knowledge based image transforms," *CoRR*, vol. abs/1702.07025, 2017.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS'12, USA*, 2012, pp. 1097–1105, Curran Associates Inc.
- [17] Yan Xu, Ran Jia, Lili Mou, Ge Li, Yunchuan Chen, Yangyang Lu, and Zhi Jin, "Improved relation classification by deep recurrent neural networks with data augmentation," in *Proc. COLING 2016*, 2016, pp. 1461–1470.
- [18] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley, "Color transfer between images," *IEEE Comput. Graph. Appl.*, vol. 21, no. 5, pp. 34–41, Sept. 2001.
- [19] F. W. Billmeyer, "Color science: Concepts and methods, quantitative data and formulae, 2nd ed., by gunter wyszecki and w. s. stiles, john wiley and sons, new york, 1982," *Color Research & Application*, vol. 8, no. 4, pp. 262–263.
- [20] M. Macenko, M. Niethammer, J. S. Marron, D. Borland, J. T. Woosley, Xiaojun Guan, C. Schmitt, and N. E. Thomas, "A method for normalizing histology slides for quantitative analysis," in *ISBI 2009*, June 2009, pp. 1107–1110.
- [21] A. M. Khan, N. Rajpoot, D. Treanor, and D. Magee, "A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 6, pp. 1729–1738, June 2014.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NIPS 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds., pp. 2672–2680. Curran Associates, Inc., 2014.
- [23] A. Bentaieb and G. Hamarneh, "Adversarial stain transfer for histopathology image analysis," *IEEE Transactions on Medical Imaging*, vol. 37, no. 3, pp. 792–802, March 2018.
- [24] M. Tarek Shaban, Christoph Baur, Nassir Navab, and Shadi Albarqouni, "Staingan: Stain style transfer for digital histological images," *CoRR*, vol. abs/1804.01601, 2018.
- [25] E. Decenciere, S. Velasco-Forero, F. Min, J. Chen, G. Gauthier H. Burdin, B. Lay, T. Bornschloegl, and T. Baldeweck, "Dealing with topological information within a fully convolutional neural network," in *ACIVS*, 2018.
- [26] Yading Yuan, Ming Chao, and Yeh-Chi Lo, "Automatic skin lesion segmentation using deep fully convolutional networks with jaccard distance," *IEEE Transactions on Medical Imaging*, vol. 36, pp. 1876–1886, 2017.
- [27] M. D. Zeiler, "Adadelta: An adaptive learning rate method," *CoRR*, vol. abs/1212.5701, 2012.
- [28] Maxim Berman and Matthew B. Blaschko, "Optimization of the jaccard index for image segmentation with the lovász hinge," *CoRR*, vol. abs/1705.08790, 2017.
- [29] F. Mahmood, D. Borders, R. Chen, G. N. McKay, K. J. Salimian, A. S. Baras, and N. J. Durr, "Deep adversarial training for multi-organ nuclei segmentation in histopathology images," *CoRR*, vol. abs/1810.00236, 2018.
- [30] L. Hou, A. Agarwal, D. Samaras, T. M. Kurç, R. R. Gupta, and J. H. Saltz, "Unsupervised histopathology image synthesis," *CoRR*, vol. abs/1712.05021, 2017.