

Supervision and Source Domain Impact on Representation Learning: A Histopathology Case Study

Milad Sikaroudi¹, Amir Safarpour¹, Benyamin Ghoghogh², Sobhan Shafiei¹,
Mark Crowley² and H.R. Tizhoosh^{1,*}, *Senior Member, IEEE*

Abstract—As many algorithms depend on a suitable representation of data, learning unique features is considered a crucial task. Although supervised techniques using deep neural networks have boosted the performance of representation learning, the need for a large sets of labeled data limits the application of such methods. As an example, high-quality delineations of regions of interest in the field of pathology is a tedious and time-consuming task due to the large image dimensions. In this work, we explored the performance of a deep neural network and triplet loss in the area of representation learning. We investigated the notion of similarity and dissimilarity in pathology whole-slide images and compared different setups from unsupervised and semi-supervised to supervised learning in our experiments. Additionally, different approaches were tested, applying few-shot learning on two publicly available pathology image datasets. We achieved high accuracy and generalization when the learned representations were applied to two different pathology datasets.

I. INTRODUCTION

With the advent of digital whole-slide image (WSI) scanners and digital pathology, a vast range of computer-vision algorithms using machine learning have been developed to process histopathology images. These technologies offer opportunities for better quantitative modeling of disease appearance and hence possibly improved prediction of disease severity and aggressiveness, and patient outcome. More specifically, machine-learning applications in digital pathology ranges from computer-aided diagnosis based on classification, detection, segmentation, and content-based image retrieval. But domain-specific limitations such as large image dimensions and insufficient amount of annotated data restrict the applicability of such approaches [1].

Current deep learning algorithms may reach or surpass human-level accuracy when a large amount of data is available. However, deep networks suffer from poor sample efficiency in stark contrast to human perception that can learn object categories after seeing just a few pictures, in some cases even one instance. Few-shot learning as the ability to learn from a few labeled samples aims to address this issue. More specifically, the knowledge is extracted from other similar problems since there is not enough data. As

a result, many methods characterize few-shot learning as a meta-learning problem.

Approaches to meta-learning can fall into three major sub-categories by whether they rely on prior knowledge about similarity, learning, or data [2]. We aimed to utilize prior knowledge about the similarity to learn robust embeddings by investigating the notion of pairwise similarity between samples. To this end, we propose a novel framework to impose abstract domain knowledge in histopathology as prior knowledge to train a triplet deep neural network.

II. RELATED WORK

Koch et al. [4] suggested the Siamese network as a choice for representation learning based on sample-wise comparisons. The Siamese networks generally consist of two similar neural networks that accept two inputs from either the same or different classes. The pair is embedded by identical neural networks first. Then the component-wise difference of the representations is passed into a comparison neural network. By doing so, the Siamese network learns to identify discrepancies between classes. Later, Hoffer and Ailon [5] proposed triplet networks for representation learning that learned the configuration among anchor, positive, and negative cases (or triplets) simultaneously. In detail, the triplet network learned to put anchor and positive samples closer while pushing the negative samples farther in the latent space. These works fall within a greater field of study called distance metric learning [6]. Medela et al. [7] employed triplet networks on the colorectal cancer slides as a source domain, and utilized their model to extract features and represent data from the healthy and tumorous colon, breast, and lung slides as the target domain. For this purpose, they utilized a VGG16 model [8] as the backbone of their triplet network while they replaced the last fully connected layer with a more compact version. They used labels provided by pathologists in the source domain to create triplets. Furthermore, due to the mismatch between the properties of source and target domain, images were adjusted through rescaling.

Gildenblat and Klaiman [9] trained a Siamese network based on a ResNet-50 architecture [10] by acknowledging adjunct patches as similar and remote tiles as non-similar cases. They both trained and tested their model on the Camelyon16 dataset [11]. Eventually, they evaluated their model in a tumor image retrieval task using the Camelyon16 test set. They reported 34% as the ratio of correctly retrieved tumorous patches in comparison with only 26% accuracy using a ResNet-50 [10] with ImageNet weights. In another

¹M. Sikaroudi, A. Safarpour, S. Shafiei and H.R. Tizhoosh are members of the Laboratory for Knowledge Inference in Medical Image Analysis (Kimia Lab), University of Waterloo, Waterloo, ON, Canada, (e-mail: {msikaroudi, asafarpo, s7shafie, tizhoosh}@uwaterloo.ca).

²B. Ghoghogh and M. Crowley are with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada, (e-mail: {bghoghogh, mcrowley}@uwaterloo.ca).

The first three authors contributed equally to this work.

*Corresponding author: H.R. Tizhoosh.

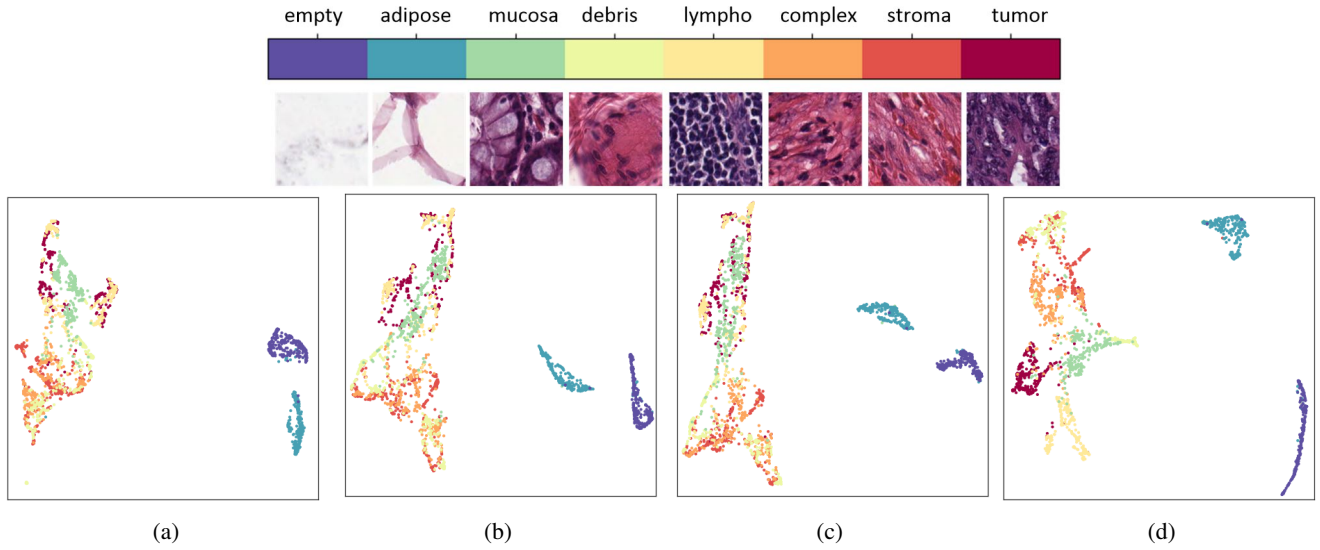


Fig. 1: 2-D representation of the D^T embedding using models trained by (a) D_1^S with the triplet loss, (b) D_2^S , (c) D_3^S , and (d) D_1^S with the cross-entropy loss. The 2-D representations are produced by UMAP [3].

work, Teh and Taylor [12] investigated the performance of a weakly-supervised framework for representation learning in digital pathology using ResNet-34 [10]. Thus, they examined the following setups for representation learning with varying target domain size: (1) Training from scratch using target domain dataset, (2) transfer learning with cross-entropy loss based on a network pre-trained on a weakly labeled source dataset, and (3) employing the metric learning approach on the same model, pre-trained with weakly labeled data. Due to the property mismatch between different datasets utilized in their study, the authors had to resize the images. Finally, they reported 92.46% accuracy on using 625 samples per class, on the CRC dataset as their best outcome [13].

In this study, we aim to address shortcomings such as the definition of the similarity, investigation of the impact of the source and target domain datasets, and the level of supervision in representation learning in digital pathology. The description of the source and target domain datasets, triplet generation algorithm, and representation learning approach are introduced in section III. Section IV contains experiments associated with the effects of the source domain, supervision, and target domain size on feature learning. Finally, the discussion and a summary of our findings are described in section V.

III. METHODOLOGY

In this work, we used two popular histopathology datasets, namely The Cancer Genome Atlas (TCGA) (available at <https://www.cancer.gov/tcga>) and colorectal cancer (CRC) [13] datasets. The TCGA dataset contains WSIs from 25 different anatomical sites, including 32 different cancer subtypes. We utilized 40% randomly selected WSIs from three organ sites, namely prostate, gastrointestinal, and lung. These organs were chosen as they are among the most commonly diagnosed cancers. These sites had a total of nine cancer subtypes, namely Prostate adenocarcinoma (PRAD), Testicular

germ cell tumors (TGCT), Oesophageal carcinoma (ESCA), Stomach adenocarcinoma (STAD), Colonic adenocarcinoma (COAD), Rectal adenocarcinoma (READ), Lung adenocarcinoma (LUAD), Lung squamous cell carcinoma (LUSC), and Mesothelioma (MESO) [14]. The CRC dataset [13] contains 5,000 histological images extracted from different tissues present in the colorectal cancer slides. The tissue types are background, adipose tissue, normal mucosal glands, debris, immune cells, complex stroma, simple stroma, and tumor epithelium. Some CRC patches of different tissue types can be found in Fig. 1.

A. Triplet Generation

Since we used a triplet network, we defined the similarity concept in a way capable of generating triplets. The triplet was comprised of anchor, neighbor, and distant tiles (patches), in which anchor and neighbor were defined as similar and anchor and distant as dissimilar pairs. Hence, the notion of similarity was abstract enough to avoid limiting the learning performance. Also, the concept of similarity should require minimal supervision to reduce the cost of triplet generation. In this study, similar to [15], we utilized spatial correlation as one of the approaches to define the similarity among patches extracted from WSIs. In other words, we assumed that similar patterns usually emerge in an adjacent neighborhood, while the dissimilar layouts often appear in the spatially remote neighborhood. More specifically, a neighbor patch was selected within a certain range of the anchor's tile center of the same WSI. On the other hand, we used several alternatives for choosing the distant patch. The distant sample was chosen from (1) the same WSI as long as it was spatially remote, (2) another WSI associated with the same cancer subtype, (3) another WSI associated with other subtypes of the same organ, or (4) another WSI associated with another organ. An example of a type 1 triplet generation is depicted in Fig. 2. For

the experiments performed on the CRC dataset [13], we selected the neighbor and distant patches from the same and different tissue types to the anchor patch, respectively. Fig. 3 shows sample triplets extracted from the TCGA and the CRC datasets [13].




Fig. 2: An example of the type 1 triplet generation from a sample WSI (from COAD subtype) from TCGA dataset.

B. Representation Learning in the Source Domain

In this paper, we aimed to analyze the effect of the source domain on a few-shot learning framework.




Fig. 3: The first two rows, sample triplets extracted from TCGA dataset: (a) anchor (STAD), (b) neighbor (STAD), (c) distant type 1 (from the same WSI), (d) distant type 2 (from different STAD WSI), (e) distant type 3 (COAD), and (f) distant type 4 (LUAD). The last row, a sample triplet generated from the CRC dataset [13]: (g) anchor, (h) neighbor, and (i) distant.

Therefore, we embedded the data in the source domain using different settings to evaluate its impact on the performance. For embedding, we used the triplet network, with a ResNet-18 backbone.

The triplet loss [16] was chosen as it evaluates similarity and dissimilarity among samples simultaneously, while the contrastive loss assesses them one by one. We implemented a triplet network [16] where three ResNet-18 networks [10] have shared weights. The triplet loss can be defined as:

$$\sum_{i=1}^b \left[\|f(x_i^a) - f(x_i^n)\|_2^2 - \|f(x_i^a) - f(x_i^d)\|_2^2 + \alpha \right]_+, \quad (1)$$

where $[z]_+ = \max(z, 0)$, b is the batch size, and x_i^a , x_i^n , and x_i^d are the i^{th} anchor, neighbor, and distant tiles in the batch, respectively. Also, $f(x)$ denotes the embedded output of the network. In the embedding space, the triplet loss pulls the neighbor toward the anchor by minimizing the first term and pushes the distant away from the anchor by maximizing the second term. The term α prevents the network from pushing the distant sample further than the margin value. As a baseline for comparison, we also used a ResNet-18 [10] trained by a standard cross-entropy loss.

C. Knowledge Transfer to the Target Domain

After training the triplet network, the target domain data was embedded using the model. Next, a classifier was trained on a portion of the target domain data and was tested against the remaining samples. We used a Support Vector Machine (SVM) as the classifier.

IV. EXPERIMENTS

For all experiments in this study, we utilized a ResNet-18 [10] backbone with Adam optimizer (learning rate of 10^{-5} and betas of 0.9 and 0.999), and a batch size of 32. Also, all images were of size 128×128 pixels, extracted from WSIs at $20\times$ magnification (CRC tiles were cropped to adjust the size). Additionally, the last layer of the model had a dimension of 128. Also, the margin value, α , was set to 0.25 for all experiments. Furthermore, during the supervised setup, our model was trained using 3,000 image tiles from the CRC dataset while for the triplet network the total number of 22,528 triplets was used (both TCGA and CRC dataset). All neural network approaches in this study were implemented using TensorFlow [17].

Triplet Loss vs. Cross-entropy Loss – First, we split the CRC dataset [13] into 60% and 40% portions, denoting D_1^S and D^T , respectively. Accordingly, we trained a ResNet-18 with two different approaches: (1) cross-entropy loss, and (2) triplet loss. To train our model with cross-entropy loss, we attached a softmax layer on top of the neural network. After training both models, the embeddings were extracted from the last layer. The embeddings of the D^T with the triplet and the cross-entropy loss are shown in Figs. 1-a and 1-d, respectively. We applied the Uniform Manifold Approximation and Projection (UMAP) [3] to visualize the 128-dimensional representations in 2D plane. The number of neighborhood parameter was set to 40 for all 2-D representations.

Source Domain Effect – The triplets extracted from the CRC training set were sampled in a supervised manner as the labels of tissues were used. However, as we described in Section III-A, the triplets of TCGA data were sampled using the spatial and tissue type information in an unsupervised manner. As a result, we trained extra two models on triplets extracted from TCGA. The first one was trained on all three anatomical sites, called D_2^S , while the second model was only trained on the gastrointestinal data from TCGA called D_3^S as the D^T was also related to the same anatomical site. Similarly, the D^T embeddings encoded by these models are shown in Figs. 1-b and 1-c.

TABLE I: Average accuracy and confidence interval of the target domain classification over the different folds of the test dataset.

Portions of Data (%)	Triplet			Cross-entropy
	D_1^S	D_2^S	D_3^S	D_1^S
5	83.00±2.18	77.00±2.52	75.00±2.98	91.00±1.65
10	88.50±0.99	88.50±0.77	88.50±0.77	87.00±1.09
25	92.80±0.41	93.20±0.24	94.20±0.29	87.00±0.35
50	94.90±0.11	93.80±0.19	94.60±0.14	89.50±0.14
100	95.90±0.03	95.75±0.07	95.95±0.09	91.90±0.09

Target Domain Size Effect – Finally, we split the D^T into chunks of 5%, 10%, 25%, and 50% using stratified sampling, besides the full set. Then we trained an SVM classifier utilizing 10-fold cross-validation on the latter subsets. For all experiments, all SVM classifiers were fine-tuned by changing the kernel (linear, RBF, sigmoid, and polynomial), and the parameters C (between 0.001, and 1000) and gamma (between 0.001 and 1000, scale and auto) to achieve the best performance. All results are reported in Table I.

V. SUMMARY AND CONCLUSIONS

As it is depicted in Fig. 1, empty and adipose tiles' representation were closely located to each other and far from other tissue textures in all models. Also, the mucosa was almost well-clustered in all embeddings because of its unique pattern. Moreover, the corresponding representations of classes with visual similarities blended in 2-D space when the triplet loss was employed. According to Table I, however, all models trained by the triplet loss scored higher in terms of mean accuracy in comparison to the model trained with the label information using cross-entropy loss, excluding 5% subset. The confidence intervals shrunk as the size of the target data increased. Furthermore, the confidence interval was relatively larger when the source domain was different from the target domain. More importantly, the highest accuracy, 95.95%, was achieved using the model which was trained on the gastrointestinal subset of the TCGA. This outcome may suggest that extensive training on a similar

dataset using weakly supervised and unsupervised methods can improve the performance of the model and generalization of the solution.

In this work, We investigated the effect of the supervision and the source domain in embedding and few-shot learning for histopathology images. Accordingly, we have shown that a network can learn a meaningful representation from histopathology images using a massive amount of the weakly labeled data currently available online. This was achieved by acknowledging the spatial correlation, anatomical information, and slide level diagnosis.

REFERENCES

- [1] D. Komura and S. Ishikawa, "Machine learning methods for histopathological image analysis," *Computational and structural biotechnology journal*, vol. 16, pp. 34–42, 2018.
- [2] J. Vanschoren, "Meta-learning: A survey," *arXiv preprint arXiv:1810.03548*, 2018.
- [3] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection for dimension reduction," *arXiv preprint arXiv:1802.03426*, 2018.
- [4] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *ICML deep learning workshop*, vol. 2, 2015.
- [5] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *International Workshop on Similarity-Based Pattern Recognition*. Springer, 2015, pp. 84–92.
- [6] J. L. Suárez, S. García, and F. Herrera, "A tutorial on distance metric learning: Mathematical foundations, algorithms and software," *arXiv preprint arXiv:1812.05944*, 2018.
- [7] A. Medela *et al.*, "Few shot learning in histopathological images: reducing the need of labeled data on biological datasets," in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*. IEEE, 2019, pp. 1860–1864.
- [8] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [9] J. Gildenblat and E. Klaiman, "Self-supervised similarity learning for digital pathology," in *MICCAI 2019 COMPAY Workshops*, 2019.
- [10] K. He *et al.*, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [11] B. E. Bejnordi *et al.*, "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer," *Jama*, vol. 318, no. 22, pp. 2199–2210, 2017.
- [12] E. W. Teh and G. W. Taylor, "Learning with less data via weakly labeled patch classification in digital pathology," *arXiv preprint arXiv:1911.12425*, 2019.
- [13] J. N. Kather *et al.*, "Multi-class texture analysis in colorectal cancer histology," *Scientific reports*, vol. 6, p. 27988, 2016.
- [14] L. A. Cooper *et al.*, "Pancancer insights from the cancer genome atlas: the pathologist's perspective," *The Journal of pathology*, vol. 244, no. 5, pp. 512–524, 2018.
- [15] N. Jean *et al.*, "Tile2vec: Unsupervised representation learning for spatially distributed data," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3967–3974.
- [16] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 815–823.
- [17] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>