

Unsupervised Latent Stain Adaptation for Computational Pathology

Daniel Reisenbüchler^{1,*,@}, Lucas Luttner^{1,*}, Nadine S. Schaad²,
Friedrich Feuerhake², and Dorit Merhof^{1,3}

¹ Institute of Image Analysis and Computer Vision, University of Regensburg,
Regensburg, Germany

² Institute of Pathology, Hannover Medical School, Hannover, Germany

³ Fraunhofer Institute for Digital Medicine MEVIS, Bremen, Germany

*These authors contributed equally to this work.

@Correspondence: daniel.reisenbuechler@informatik.uni-regensburg.de

Abstract. In computational pathology, deep learning (DL) models for tasks such as segmentation or tissue classification are known to suffer from domain shifts due to different staining techniques. Stain adaptation aims to reduce the generalization error between different stains by training a model on source stains that generalizes to target stains. Despite the abundance of target stain data, a key challenge is the lack of annotations. To address this, we propose a joint training between artificially labeled and unlabeled data including all available stained images called Unsupervised Latent Stain Adaptation (ULSA). Our method uses stain translation to enrich labeled source images with synthetic target images in order to increase the supervised signals. Moreover, we leverage unlabeled target stain images using stain-invariant feature consistency learning. With ULSA we present a semi-supervised strategy for efficient stain adaptation without access to annotated target stain data. Remarkably, ULSA is task agnostic in patch-level analysis for whole slide images (WSIs). Through extensive evaluation on external datasets, we demonstrate that ULSA achieves state-of-the-art (SOTA) performance in kidney tissue segmentation and breast cancer classification across a spectrum of staining variations. Our findings suggest that ULSA is an important framework for stain adaptation in computational pathology.

Keywords: Semi-supervised Learning · Stain Adaptation · Whole Slide Image · Transfer Learning · Segmentation · Classification

1 Introduction

Recent advances in DL for computational pathology have shown promising results for a wide range of applications, from cancer and biomarker detection [13] to tissue structure segmentation [3]. However, large-scale studies have shown that the effectiveness of DL techniques in histology is highly dependent on the availability of labeled data [16]. Despite its theoretical promise, acquiring a sufficient number of expert annotations remains challenging. In the realm of digital

pathology, image datasets often consist of sequential slides stained with various techniques, each providing different insights into the same region of interest. Despite variations in staining protocols, these slides often share a significant amount of consistent information. However, expert annotations may be available for one type of staining but may be lacking for others, which are often accessible in large quantities without labels. Generating expert annotations for multiple staining techniques for the same analysis tasks would be extremely time-consuming. In the era of foundation models [11], we also prefer generalized DL models that are robust to data shifts instead of domain experts. In this paper, we question how to tailor a DL model trained for a specific task to handle variations in staining within the distribution of target stains, for which no annotations are available. This can be accomplished by incorporating unlabeled data during the training phase. The aspect of stain adaptation across different inter-staining techniques has not been sufficiently explored so far. Despite efforts to develop stain-to-stain translation techniques, their effectiveness is typically evaluated either visually by experts or through translation metrics [20]. Previous research has not focused on directly incorporating unlabeled target stain images into the training process, but only used them for translation [2,7]. Here, we present ULSA, a semi-supervised strategy designed for joint training of all staining data for the first time. We introduce a framework that integrates unlabeled target stain images into supervised training by maintaining the supervised learning signal for synthetic target stainings generated through Cycle GAN (cGAN) inference [19]. Feature-wise stain-adaptation enables using unlabeled target data and enforces feature consistency across stains. By combining these key ingredients, we propose a new method for efficient stain adaptation that outperforms current SOTA approaches. Our novelties can be summarized as

- (1) **Unsupervised Stain Adaptation.** ULSA leverages target stain data in a supervised and unsupervised manner using only annotated source stains. We propose a framework for training stain-invariant models for digital pathology.
- (2) **Feature Consistency Learning.** We maximize cosine-similarity between hierarchical features across stains to achieve stain-invariance on feature level.
- (3) **Task Agnostic Framework.** ULSA is applicable for classification and segmentation training of stain-invariant models.
- (4) **Outperforming SOTA.** Our approach outperforms methods from stain-translation, DL based augmentation and semi-supervised learning slightly for source stains and by a large margin for target stains. Our approach needs only 10% of labels to achieve the same performance as SOTA trained on all data.

2 Method

Stain adaptation aims to minimize the generalization error in task performance between a source stain $s \in S$ and a target stain $t \in T$ (Fig. 1a). In particular, a parameterized model m_θ trained on labeled source staining data x_S^L should ideally maintain task performance on other unlabeled target stains x_T^U where no labels are available. We propose to address this challenge by incorporating




Fig. 1. (a). Problem statement of unsupervised stain adaptation. (b). Stain-invariant feature consistency learning. (c). Artificial images generated by cGAN.

unlabeled target stainings through (I) a cGAN model to augment labeled images into target stains that inherit the same annotation, (II) unsupervised stain adaptation to jointly train on all stains with supervised and unsupervised objectives including all stains, followed by (III) stain-invariant feature consistency learning (FCL) by unsupervised matching of latent representations between stains. The overall method is outlined in Fig. 2.

(I) cGAN stain augmentation. We pretrained and used cGANs, which we define as stain translation function $\mathcal{G}(x_S^L) = x_{S \cup T}^L$ to synthetically augment source training data into target stainings. This process is structure preserving [19], thus each target stain image inherits the label corresponding to the associated source image x_S^L used for translation. Fig. 1c shows exemplary translation results. This strategy aims to increase the labeled training dataset and thus the supervising signal to achieve stain-invariance on prediction level.

(II) Unsupervised stain adaptation. Since labeled data are given only in source staining $x^L \sim p_S(x^L)$, we would ideally desire a mapping from source to target samples. We approximate this mapping by using $|S| \cdot |T|$ distinct cGAN augments to obtain additional labeled target samples $x^L \sim \hat{p}_t(x^L | x^U)$. With $|\cdot|$ we denote the cardinality of a set of stains. Note that all inferred samples x^L inherit the label y associated with x^L . Thus, we approximate $p_{S \cup T}$ by a mixture

$$\hat{p}_{S \cup T}(x^L) = \frac{1}{|S| + |T|} \left(\sum_{s \in S} p_s(x^L) + \sum_{t \in T} \hat{p}_t(x^L) \right).$$




Fig. 2. ULSA model. Labeled source stains are translated into synthetic target stain data to obtain supervision for image-wise stain-invariance. We extract features for a real and stain translated noised image of the target stain data, where we maximize cosine similarity to achieve unsupervised feature-wise stain-invariance in latent space.

In addition, we leverage unlabeled data by unsupervised learning. Given an unlabeled image $x_{t \in T}^{U,1}$, we translate it to $\tilde{x}_{t \in T}^{U,1} = \mathcal{I} \left(\mathcal{R} \left(x_{t \in T}^{U,1} \mid x_{t \in T}^{U,2} \right) \right)$ by Reinhard translation $\mathcal{R}$ and noise injection $\mathcal{I}$. Note that $\mathcal{R}$ receives another unlabeled random sampled image $x_{t \in T}^{U,2}$ from target stains $t \in T$ as reference for subsequent translations. Fig. 1b include example inputs. We used Reinhard normalization as Macenko's method has much higher runtime leading to computational overhead, see Supplementary Material (SM). Finally, we enforce the model to embed images invariant to stain translations by maximizing cosine similarity in unsupervised loss $\mathcal{L}_U$, see (III). For this reason, we used light Gaussian blurring for noise injection, see SM for other choices. In summary, we define our objective as

$$\min_{\theta} \mathcal{L}(\theta \mid x^L, y, x^U) = \mathbb{E}_{x^L \sim p_{S \cup T}(x^L)} [\mathcal{L}_S(\theta \mid x^L, y)] + \lambda \mathbb{E}_{x^U \sim p_{S \cup T}(x^U)} [\mathcal{L}_U(\theta \mid x^U)]$$

where $\mathcal{L}_S$ is a supervised loss. With minimizing $\mathcal{L}_S$ using labeled images of synthetic target stains, we aim for stain-adaptation on prediction level. We use multi-class and binary cross-entropy loss for segmentation and classification, respectively. We set equal weight $\lambda = 1$, obtained by performing a grid search as described in SM. In each iteration we compute $\mathcal{L}_S$ on a batch of labeled data and compute feature consistency with $\mathcal{L}_U$ on a batch of unlabeled data. Next, we detail the unsupervised loss term associated with feature consistency learning.

(III) Stain-invariant feature consistency learning. For unsupervised learning, we aim to calculate similar feature representations for an input $x = x_{t \in T}^{U,1}$ and a stain translated noised version $\tilde{x} = \tilde{x}_{t \in T}^{U,1}$. To enforce this for every downsampling block in a model, we apply non-parametric 2D adaptive average pooling. In Einstein notation, this tensor operation leads to $B, C_i, H_i, W_i \rightarrow B, C_i$ where B, C, H, W and i refer to batch-size, channels, height, width and the model block, respectively (Fig. 1b). We aim to maximize the cosine similarity between

all hierarchical features of a model m_θ . In this way, we enforce feature similarity by updating the model such that extracted image features become similar in latent space. We define features as model outputs up to block i after 2D adaptive pooling, $f_{\bar{\theta},i} = m_{\bar{\theta},i}(x)$, $\tilde{f}_{\theta,i} = m_{\theta,i}(\tilde{x})$. For brevity, index θ, i also refers to parameter optimization from input layer until block i . Following the literature [10,18], we do not backpropagate gradients for non-augmented input, denoted with $\bar{\theta}$. Hence, we define our unsupervised objective $\mathcal{L}_u$ as

$$\mathcal{L}_U(\theta) = -\frac{1}{b} \sum_{i \in b} \frac{\mathbf{f}_{\bar{\theta},i} \cdot \tilde{\mathbf{f}}_{\theta,i}}{\|\mathbf{f}_{\bar{\theta},i}\|_2 \cdot \|\tilde{\mathbf{f}}_{\theta,i}\|_2},$$

where $\|\cdot\|_2$ denotes the Euclidean norm and b the number of downsampling blocks. With minimizing $\mathcal{L}_U$, we enforce stain-invariance at the feature level.

3 Experiments

We applied our proposed method to kidney tissue segmentation and cancer classification. We measured the performance using the dice score and the area under the receiver operating characteristic curve (AUROC), respectively. Next, we describe our datasets and more information about statistics can be found in SM.

3.1 Datasets and Comparable Methods

For slide tiling, we used a modified version of the CLAM preprocessing pipeline [8] and manually selected color thresholds as different stains require adjustments in tissue detection. All images were processed at a size of 512×512 px.

Kidney segmentation datasets. Our internal annotated train $n_{train}^{PAS} = 2,100$ and validation $n_{val}^{PAS} = 160$ datasets consist of annotated patches extracted from periodic acid-Schiff (PAS) stained WSIs with $20\times$ magnification. The annotation masks contain the classes tubule, glomerulus, glomerular tuft, artery, arterial lumen, and vein. For external testing we included annotated glomerulus images from PAS stained WSIs provided by the HuBMAP consortium [4]. After processing the raw data we obtained $n_{ext,hubmap}^{PAS} = 2,670$ samples. Moreover, we included test data from the NEPTUNE [1] study, containing $n_{ext,neptune}^{HE} = 402$, $n_{ext,neptune}^{PAS} = 1,176$, $n_{ext,neptune}^{SIL} = 688$ Silver (SIL) and $n_{ext,neptune}^{TRI} = 817$ Trichrome (TRI) stain samples. The classes glomerulus, glomerular tuft, artery and tubule were annotated and images were taken from WSIs with $40\times$ magnification and rescaled to $20\times$ magnification. We further processed unlabeled WSIs at $20\times$ magnification from the KPMP database [6]. We obtained $n_{unlabeled}^{HE} = 385,670$, $n_{unlabeled}^{PAS} = 409,554$, $n_{unlabeled}^{SIL} = 538,412$, $n_{unlabeled}^{TRI} = 415,822$ tiles.

Breast cancer classification datasets. Our cancer classification datasets contain $n_{train}^{HE} = 1,950$, $n_{val}^{HE} = 287$ images obtained from HE stained WSIs and tiled at $40\times$ magnification level. Our test set contains Cytokeratin (CK5) and Cluster of differentiation (CD8) stains with $n_{test}^{CK5} = 1,000$ and $n_{test}^{CD8} = 1,000$ samples. Sample sizes and binary label ratios were equalized for test sets after preprocessing and excess data were held out of experiments to avoid data leakage at patient level. Additionally we have unlabeled datasets containing $n_{unlabeled}^{HE} = 117,369$, $n_{unlabeled}^{CK5} = 41,735$, $n_{unlabeled}^{CD8} = 52,492$ samples. We split the data at slide-level so that each patient appears exclusively in one dataset.

Comparable methods. We selected comparable methods from the domains of stain translation, unsupervised augmentation, and semi-supervised consistency training. **(1) Baseline.** For comparison, we include a naive approach where we train a model on source stains without access to target stains, thus obtaining a lower bound for stain adaptation. **(2) Reinhard.** Reinhard’s method is a stain normalization technique that adjusts the color appearance of histopathology images by aligning them with a reference color space [12]. **(3) Macenko.** Macenko’s method [9] is a stain normalization technique that standardizes the color appearance of images by mapping them to a reference color space. It has been shown that this method provides the best performance across various colorization techniques for downstream tasks [7]. **(4) cGAN Augmentation.** This method generates synthetic images with diverse staining variations in order to translate between stainings and train stain-invariant models [2]. **(5) FixMatch.** A semi-supervised learning approach that combines labeled and unlabeled data by enforcing consistency between predictions made on unlabeled samples based on confidence scores [14]. **(6) Unsupervised Data Augmentation (UDA).** UDA is a semi-supervised learning technique designed for classification tasks that leverages augmentations on unlabeled data to improve model performance by minimizing probability distributions between two versions of an image. Originally proposed for natural images [18] it has been adapted to histology images [5].

3.2 Implementation

cGAN augmentation. We initially started with hyperparameters following the literature [2] and performed a grid search with details noted in SM. By visually evaluating the stain-translation results, we selected a model trained with 300 epochs and a learning rate of $1.5e - 4$. For each stain translation we used 10,000 unlabeled images from KPMP. For source stains, we used unlabeled HE data for cancer classification, and our labeled PAS stained training set for segmentation. Each stain translation training and inference task took about 8 hours.

Segmentation and classification. For all experiments we used a ResNet-50 as classification model and encoder model for U-Net in segmentation. We tuned hyperparameters on validation sets and set a learning rate decay of $1e - 10$ and early stopping for 5 and 10 consecutive epochs with no decrease in validation loss.

We employed AdamW as optimizer and set the initial learning rate to $1e-04$ and the weight decay to $1e-05$. Images were resized to a scale of 224×224 px. We used a total batch-size of 128 for all experiments. For semi-supervised learning we batched data to 32 and 96 for labeled and unlabeled data. All models were initialized with ImageNet weights and experiments were run on a single Nvidia A100 GPU. We measured a maximum training time of 10 and 7 hours across methods (except Macenko) for segmentation and classification, respectively.

4 Results and Discussion

In this section, we report the results and analyze the number of labels required for stain adaptation in kidney tissue segmentation and breast cancer classification.

Table 1. Dice and AUROC scores for segmentation and classification tasks, respectively. We report the mean and standard deviation across three consecutive runs. All methods (except Baseline) access unlabeled target stains in the trainings phase.

Method	Segmentation					Classification		
	Intra-stain	Inter-stain				Inter-stain		
	PAS	TRI	HE	SIL	Overall	CK5	CD8	Overall
Baseline	87.4 (0.66)	59.4 (4.19)	43.2 (4.01)	76.4 (1.67)	62.1 (3.24)	86.6 (9.61)	90.5 (7.10)	88.6 (8.35)
Reinhard [12]	87.3 (0.50)	64.8 (6.50)	40.2 (3.23)	77.9 (2.57)	64.3 (4.39)	89.8 (3.65)	94.1 (1.48)	91.9 (2.32)
Macenko [7]	85.0 (0.80)	71.6 (2.60)	48.3 (4.70)	81.6 (0.50)	70.3 (2.29)	89.5 (2.90)	93.4 (2.06)	91.4 (2.48)
cGAN [2]	84.2 (0.92)	69.9 (1.44)	46.0 (4.71)	79.0 (1.63)	68.1 (2.18)	87.1 (6.24)	88.0 (3.32)	87.5 (3.67)
FixMatch [14]	87.2 (0.51)	64.8 (2.67)	40.3 (7.28)	78.2 (1.02)	64.5 (3.05)	88.3 (2.25)	94.3 (1.37)	91.3 (1.78)
UDA [5]	87.2 (0.54)	64.9 (1.84)	45.8 (2.49)	77.5 (0.60)	65.4 (1.53)	89.7 (2.19)	92.8 (1.97)	91.3 (1.47)
Ours	87.9 (0.33)	74.1 (0.94)	53.6 (1.46)	81.8 (0.61)	72.6 (0.93)	91.6 (4.63)	94.6 (1.40)	93.1 (3.01)

Measuring stain adaptation. In the area of stain adaptation, we aim to at least maintain performance on source stains (Intra-stain, Tab. 1) and maximize performance on target stains (Inter-stain, Tab. 1). For segmentation on source data our method is on par with other methods, and also shows slightly improved performance. Note that all other methods decrease their performance on source stains compared to the baseline. More importantly, we increase the Dice score for target stains (Inter-stain, Overall) by more than 10 and 2.3 compared to naive

training and Macenko colorization, the best comparable method. Note, however, that although Macenko provides the second best stain adaptation result, source performance is reduced by 2.4 and 2.9 compared to baseline and ULSA. In the case of classification, we did not have further annotated intra-stain data for testing, but report results for target data. Our method increases the AUROC by 4.5 and 1.2 compared to the naive and best comparable methods. Overall, these results demonstrate efficient stain adaptation by increasing the target while not only maintaining but increasing source performance (Fig. 1a). Additionally, we trained all methods on different fractions of labeled data and tested their performance on targets in segmentation (Fig. 3a). Interestingly, ULSA shows great stain adaptation even on very little annotated data and is on par with the best comparable method (on full data) using only 10% of labeled data.

Ablation study. We measured the influence of different components of our method in segmentation (Fig. 3b). By dropping either cGAN or FCL components, overall target performance decreases. Using hierarchical features (ULSA) instead of features from the last block (LB FCL) increases the scores. We also initialized our ULSA method with pretrained weights from a foundation model for histology (FM ULSA), obtained by large-scale learning of HE images [17]. With this setup, we demonstrate that ImageNet pretrained weights yield better performance. This suggests that when building foundation models, stain variations should be considered to avoid catastrophic forgetting of learned morphologies.




Fig. 3. (a). Performance for different fractions of labeled data. (b). Ablation study.

5 Conclusion

We proposed ULSA, a novel SOTA strategy to reduce stain generalization errors in computational pathology tasks. Our semi-supervised learning strategy uses annotated data for both source and artificial target stains. In addition, we incorporate unlabeled data for stain-invariant feature consistency learning. Finally,

joint optimization of supervised and unsupervised objectives enables efficient stain adaptation. We empirically demonstrated that ULSA training increases performance on unlabeled target stains in patch level segmentation and classification. This suggests that ULSA is a task agnostic framework. We further showed ULSA achieves efficient stain adaptation even in settings with scarce labels. A potential limitation of our approach is that even if the performance for unseen inter-stains is improved, augmentation may not translate the correct marker information from other stains. A possible example is immunohistochemical (IHC) staining, where specific immune cells are highlighted. This could potentially affect downstream applications in certain scenarios not seen in this study. Future work could compare other augmentation strategies such as HistAuGAN [15].

Acknowledgements. This work was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) under project number 445703531 and the German Federal Ministry of Education and Research (Bundesministerium für Bildung und Forschung, BMBF) under project numbers 01IS21067A and 01IS21067B. The authors gratefully acknowledge the computational and data resources provided by the Leibniz Supercomputing Centre (www.lrz.de).

Disclosure of Interests. The authors declare that they have no conflicts of interest related to this work.

References

1. Barisoni, L., Nast, C.C., Jennette, J.C., Hodgin, J.B., Herzenberg, A.M., Lemley, K.V., Conway, C.M., Kopp, J.B., Kretzler, M., Lienczewski, C., Avila-Casado, C., Bagnasco, S., Sethi, S., Tomaszewski, J., Gasim, A.H., Hewitt, S.M.: Digital pathology evaluation in the multicenter nephrotic syndrome study network (neptune). *Clinical Journal of the American Society of Nephrology* **8**(8), 1449–1459 (Aug 2013). <https://doi.org/10.2215/cjn.08370812>, <http://dx.doi.org/10.2215/CJN.08370812>
2. Bouteldja, N., Hölscher, D.L., Klinkhammer, B.M., Buelow, R.D., Lotz, J., Weiss, N., Daniel, C., Amann, K., Boor, P.: Stain-independent deep learning-based analysis of digital kidney histopathology. *The American Journal of Pathology* **193**(1), 73–83 (Jan 2023). <https://doi.org/10.1016/j.ajpath.2022.09.011>, <http://dx.doi.org/10.1016/j.ajpath.2022.09.011>
3. Bouteldja, N., Klinkhammer, B.M., Bülow, R.D., Droste, P., Otten, S.W., Freifrau von Stillfried, S., Moellmann, J., Sheehan, S.M., Korstanje, R., Menzel, S., Bankhead, P., Mietsch, M., Drummer, C., Lehrke, M., Kramann, R., Floege, J., Boor, P., Merhof, D.: Deep learning-based segmentation and quantification in experimental kidney histopathology. *Journal of the American Society of Nephrology* **32**(1), 52–68 (Nov 2020). <https://doi.org/10.1681/asn.2020050597>, <http://dx.doi.org/10.1681/ASN.2020050597>
4. Howard, A., Lawrence, A., Sims, B., Tinsley, E., Kazmierczak, J., Borner, K., Godwin, L., Novaes, M., Culliton, P., Holland, R., Watson, R., Ju, Y.: Hubmap - hacking the kidney (2020), <https://kaggle.com/competitions/hubmap-kidney-segmentation>
5. Jiang, Y., Sui, X., Ding, Y., Xiao, W., Zheng, Y., Zhang, Y.: A semi-supervised learning approach with consistency regularization for tumor histopathological images analysis. *Frontiers in Oncology* **12** (Jan 2023). <https://doi.org/10.3389/fonc.2022.1044026>, <http://dx.doi.org/10.3389/fonc.2022.1044026>
6. Kidney Precision Medicine Project: Kidney Precision Medicine Project Data. Accessed September 01, 2023. <https://www.kpmp.org>, the results here are in whole or part based upon data generated by the Kidney Precision Medicine Project. Funded by the National Institute of Diabetes and Digestive and Kidney Diseases
7. Lampert, T., Merveille, O., Schmitz, J., Forestier, G., Feuerhake, F., Wemmer, C.: Strategies for training stain invariant cnns. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). IEEE (Apr 2019). <https://doi.org/10.1109/isbi.2019.8759266>, <http://dx.doi.org/10.1109/ISBI.2019.8759266>
8. Lu, M.Y., Williamson, D.F., Chen, T.Y., Chen, R.J., Barbieri, M., Mahmood, F.: Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering* **5**(6), 555–570 (2021)
9. Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Guan, X., Schmitt, C., Thomas, N.E.: A method for normalizing histology slides for quantitative analysis. In: 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. IEEE (Jun 2009). <https://doi.org/10.1109/isbi.2009.5193250>, <http://dx.doi.org/10.1109/ISBI.2009.5193250>
10. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **41**(8), 1979–1993 (2018)

11. Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J., Rajpurkar, P.: Foundation models for generalist medical artificial intelligence. *Nature* **616**(7956), 259–265 (Apr 2023). <https://doi.org/10.1038/s41586-023-05881-4>, <http://dx.doi.org/10.1038/s41586-023-05881-4>
12. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Computer graphics and applications* **21**(5), 34–41 (2001)
13. Reisenbüchler, D., Wagner, S.J., Boxberg, M., Peng, T.: Local attention graph-based transformer for multi-target genetic alteration prediction. In: *Lecture Notes in Computer Science*, pp. 377–386. Springer Nature Switzerland (2022)
14. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence (2020)
15. Wagner, S.J., Khalili, N., Sharma, R., Boxberg, M., Marr, C., de Back, W., Peng, T.: Structure-preserving multi-domain stain color augmentation using style-transfer with disentangled representations. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021* (2021)
16. Wagner, S.J., Reisenbüchler, D., West, N.P., Niehues, J.M., Zhu, J., Foersch, S., Veldhuizen, G.P., Quirke, P., Grabsch, H.I., van den Brandt, P.A., Hutchins, G.G., Richman, S.D., Yuan, T., Langer, R., Jenniskens, J.C., Offermans, K., Mueller, W., Gray, R., Gruber, S.B., Greenson, J.K., Rennert, G., Bonner, J.D., Schmolze, D., Jonnagaddala, J., Hawkins, N.J., Ward, R.L., Morton, D., Seymour, M., Magill, L., Nowak, M., Hay, J., Koelzer, V.H., Church, D.N., Matek, C., Gelpert, C., Peng, C., Zhi, C., Ouyang, X., James, J.A., Loughrey, M.B., Salto-Tellez, M., Brenner, H., Hoffmeister, M., Truhn, D., Schnabel, J.A., Boxberg, M., Peng, T., Kather, J.N., Church, D., Domingo, E., Edwards, J., Glimelius, B., Gogenur, I., Harkin, A., Hay, J., Iveson, T., Jaeger, E., Kelly, C., Kerr, R., Maka, N., Morgan, H., Oien, K., Orange, C., Palles, C., Roxburgh, C., Sansom, O., Saunders, M., Tomlinson, I.: Transformer-based biomarker prediction from colorectal cancer histology: A large-scale multicentric study. *Cancer Cell* **41**(9), 1650–1661.e4 (Sep 2023). <https://doi.org/10.1016/j.ccell.2023.08.002>, <http://dx.doi.org/10.1016/j.ccell.2023.08.002>
17. Wang, X., Du, Y., Yang, S., Zhang, J., Wang, M., Zhang, J., Yang, W., Huang, J., Han, X.: Retccl: Clustering-guided contrastive learning for whole-slide image retrieval. *Medical Image Analysis* **83**, 102645 (Jan 2023). <https://doi.org/10.1016/j.media.2022.102645>, <http://dx.doi.org/10.1016/j.media.2022.102645>
18. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation for consistency training (2020)
19. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. pp. 2242–2251 (2017). <https://doi.org/10.1109/ICCV.2017.244>
20. Zingman, I., Frayle, S., Tankoyeu, I., Sukhanov, S., Heinemann, F.: A comparative evaluation of image-to-image translation methods for stain transfer in histopathology (2023)