arXiv:2206.09769v3 [cs.CV] 30 Jun 2022

Test-time image-to-image translation ensembling improves out-of-distribution generalization in histopathology

Marin Scalbert[1,2], Maria Vakalopoulou[1], and Florent Couzinié-Devy[2]

[1] MICS, CentraleSupélec - Université Paris-Saclay, Gif-sur-Yvette, France
{marin.scalbert, maria.vakalopoulou}@centralesupelec.fr
[2] VitaDX International, Paris, France
f.couzinie-devy@vitadx.com

Abstract. Histopathology whole slide images (WSIs) can reveal significant inter-hospital variability such as illumination, color or optical artifacts. These variations, caused by the use of different protocols across medical centers (staining, scanner), can strongly harm algorithms generalization on unseen protocols. This motivates the development of new methods to limit such loss of generalization. In this paper, to enhance robustness on unseen target protocols, we propose a new test-time data augmentation based on multi domain image-to-image translation. It allows to project images from unseen protocol into each source domain before classifying them and ensembling the predictions. This test-time augmentation method results in a significant boost of performances for domain generalization. To demonstrate its effectiveness, our method has been evaluated on two different histopathology tasks where it outperforms conventional domain generalization, standard/H&E specific color augmentation/normalization and standard test-time augmentation techniques. Our code is publicly available at https://gitlab.com/vitadx/articles/test-time-i2i-translation-ensembling.

Keywords: Test-time data augmentation · image-to-image translation · domain generalization · Generative Adversarial Networks

1 Introduction

Histopathology is the gold standard to diagnose most types of cancer. Sampled tissues are processed following specific protocols (fixation, staining, scanning) to obtain whole slide images (WSIs). It enables tissue types/cells structures differentiation and highlights abnormalities or cancer indicators. However, these protocols substantially vary across different medical institutions resulting in significant inter-hospital variability at the WSIs level. Such variability originating from illumination and/or color variations or optical artifacts can even occur within the same hospital [11] as well as over time [18].

Intra and inter-hospital variability prevent models from generalizing well on unseen hospitals [21]. This makes domain generalization one of the primordial

and studied task of computational pathology. To circumvent the drop of generalization performances in histopathology applications, recent methods have relied on standard color augmentations [21], H&E specific data augmentations [20,5,23], stain normalization [14,13,21,16] or domain generalization (DG) techniques [19,2,17,15]. In histopathology, data augmentation has proven to be one of the most efficient and simple technique to close the gap due to domain shift [21]. Some data augmentation/unsupervised domain adaptation (UDA) methods are even exploiting image-to-image translation methods to learn stain-invariant models [23,22]. However, UDA methods need access to unlabeled data from the target domain which is not the case in the DG setting and must be trained every time we want to predict on a new target domain. Currently, data augmentation methods only exploit image-to-image translation models at training which as shown by our study is not optimal.

In this paper, we propose for the first time, the use of image-to-image translation for test-time augmentation (TTA) designing a tailored ensemble strategy. In particular, the method is based upon the multi-domain image-to-image translation model StarGanV2 [3]. At test-time, it projects images from unseen domains to the source domains, classify the projected images and ensemble their predictions. Additionally, several ensembling strategies have been explored. The proposed TTA does not rely on any prior on the domain shifts, learning them directly from the training data. Our method, operating at test-time, can be easily combined with other DG and/or data augmentation techniques. To demonstrate its effectiveness, it has been evaluated on two different histopathology tasks and has shown better generalization over previous conventional DG, color augmentation/normalization, and standard TTA techniques [19,2,15,17,4,20].

2 Method

In the problem of multi-source DG, we are given S source domains $\{\mathcal{D}_1, \ldots, \mathcal{D}_S\}$. Each source domain $\mathcal{D}_i$ with domain label d_i contains $n_{\mathcal{D}_i}$ labeled examples $\{(x_j, y_j) \mid 1 \leq j \leq n_{\mathcal{D}_i}\}$. The goal is to learn a robust model from the S labeled source domains so that it generalizes well on an unseen target domain $\mathcal{D}_T$. In the following sections the different components of our method are detailed.

2.1 StarGanV2 Architecture

StarGanV2 is a multi-domain image-to-image translation model capable of translating an input image x into any source domain without requiring the input image domain label: this is what allows us to use it on unseen target domain. It is considered as a style based image-to-image translation model meaning that the generator takes as inputs the image that needs to be translated but also a domain-dependant style vector helping the generator to translate the input image into the corresponding domain. Style vectors can be extracted from a reference style image or from a random latent code (Figure 1).

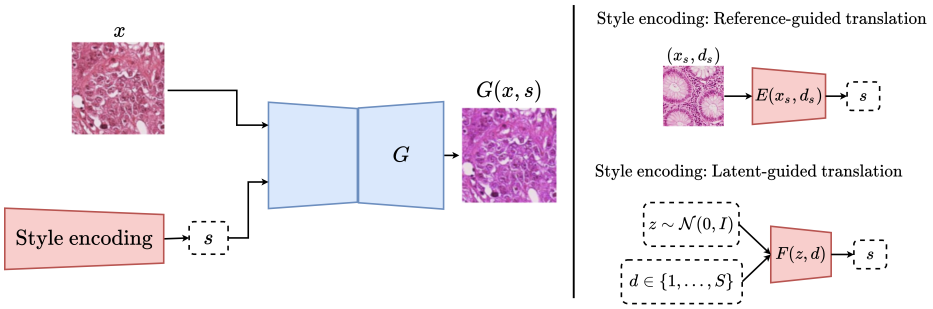

Fig. 1: The two different style encodings: on right top the reference-guided and right bottom the latent-guided image-to-image translation are presented.

More specifically, given an image x and a style vector s from a source domain d, the generator G translates the image x into an image $G(x, s)$ reflecting visual attributes of the domains d. s can be computed either from a style encoder E given a reference style image with its domain (x_s, d_s) (i.e $s = E(x_s, d_s)$) or from a mapping network F given a random latent code z along with a domain label d (i.e $s = F(z, d)$). Our proposed TTA method relies on latent-guided image-to-image translation due to the lighter size and lower computational overhead of F compared to E. Additionally, some of the explored ensembling strategies exploit the different domain discriminators used during the StarGanV2 training. To better preserve the structure of the input image in the translated images, the perceptual domain invariant loss [6] is added to the overall loss of StarGanV2. Examples of translated images are provided in Fig. 1 of supplementary material.

2.2 Test-time image-to-image translation ensembling

Once a StarGanV2 and a classifier C have been trained separately on the source data, our TTA can be used at inference. An overview of our TTA method is presented in Figure 2. Given a test image x, we sample S random latent code vectors $\{z_1, \ldots, z_S\} \sim \mathcal{N}(0, I)$ with respective domain label $\{d_1, \ldots, d_S\}$. Each latent code vector along with its domain label are fed to the mapping network F to produce the different style vectors $\{s_1, \ldots, s_S\} = \{F(z_1, d_1), \ldots, F(z_S, d_S)\}$. Then, the image x, along with each of the style vectors $\{s_1, \ldots, s_S\}$, are fed to the StarGanV2 generator G translating x into the images $\{x_1, \ldots, x_S\} = \{G(x, s_1), \ldots, G(x, s_S)\}$. From each x_i that should have now similar characteristics as real images from domain d_i, we compute the predictions of the classifier $\hat{y}_i = C(x_i)$ and the different StarGanV2 domain discriminators score $\hat{d}_i = D_i(x_i)$ reflecting how well the image x has been projected into domain d_i. Finally, from the classifier predictions $\{\hat{y}_1, \ldots, \hat{y}_S\}$ and the domain discriminators score $\{\hat{d}_1, \ldots, \hat{d}_S\}$, we propose three different strategies to ensemble the predictions into a final prediction $\hat{y}$.

Naive ensembling. In this method, domain discriminators scores $\{\hat{d}_1, ..., \hat{d}_S\}$ are discarded and the predictions $\{\hat{y}_1, \ldots, \hat{y}_S\}$ are simply averaged.

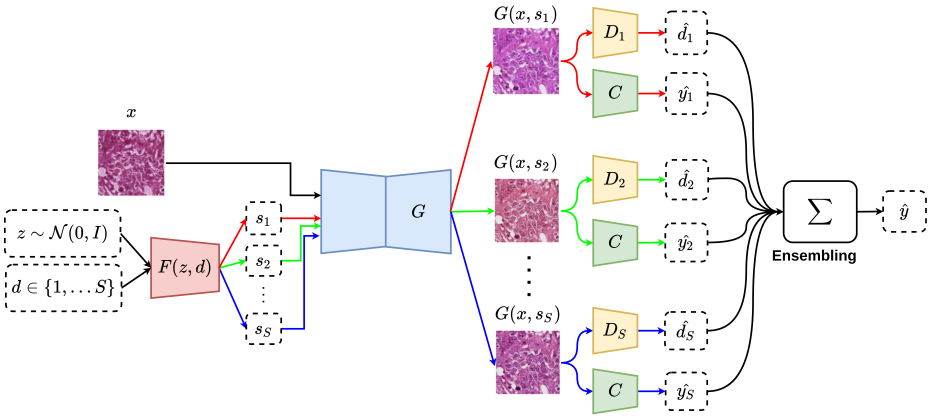

Fig. 2: Test-time data augmentation with latent-guided StarGanV2 image-to-image translation

Discriminator based top-k ensembling. A domain discriminator score should reflect how well a test image has been projected into the corresponding source domain. Following this intuition, we propose to ensemble only k predictions associated to the images with the top-k domain discriminators scores:

$$\hat{y} = \frac{1}{S'} \sum_{s=1}^{S'} \hat{y}_s \quad ; \quad S' = \underset{\substack{I \subseteq \{1,\dots,S\} \\ |I|=k}}{\arg\max} \sum_{i \in I} \hat{d}_i \tag{1}$$

Discriminator based weighted ensembling. In the discriminator based top-k ensembling, although the predictions associated to the lowest domain discriminator scores have been discarded from the ensembling, the k kept predictions contribute evenly even if there might be significant score disparities. Preferably, we would like to consider all the predictions to unlock the full power of ensembling while reducing the importance of predictions associated to badly projected images. That is why, we propose here to ensemble all the predictions $\{\hat{y}_1, \dots, \hat{y}_S\}$ with respective weights $\{\alpha_1, \dots, \alpha_S\}$ based on domain discriminators scores $\{\hat{d}_1, \dots, \hat{d}_S\}$:

$$\hat{y} = \sum_{s=1}^{S} \alpha_s \hat{y}_s \quad ; \quad \alpha_s = \frac{e^{\frac{\hat{d}_s}{T}}}{\sum_{j=1}^{S} e^{\frac{\hat{d}_j}{T}}} \quad \forall s \in \{1, \dots, S\} \tag{2}$$

The temperature T in Eq. 2 is a hyperparameter controlling how peaky the predictions weights distribution should be. When $T \to \infty$, predictions weights tend to be uniform ($\alpha_s = \frac{1}{S}$) which is equivalent to the naive ensembling. While,

when $T \to 0$, predictions weights tend to be a one-hot encoded vector which is equivalent to the discriminator based top-k ensembling with $k = 1$.

3 Experiments

3.1 Datasets

Patch classification of lymph node section WSIs. For this DG task, we use the standard Camelyon17 WILDS dataset [10] containing patches of WSIs of lymph node sections from patients with potentially metastatic breast cancer. WSIs have been collected from 5 medical centers and splits have been made with respect to the medical center. The goal of the task is to predict whether or not patches contain tumorous tissue and to generalize well on patches from unseen centers. The distribution (#hospitals, #WSIs, #patches) for the `train`, `id_val`, `val` and `test` splits are respectively $(3, 30, 302436)$, $(3, 30, 33560)$, $(1, 10, 34904)$ and $(1, 10, 85054)$. `train` and `id_val` hospitals are identical while `val` and `test` hospitals are both distinct and different from `train` and `id_val` hospitals. To be able to exploit the full potential of ensembling, we decide to choose the different WSIs as the different source domains ($S = 30$) rather than the number of distinct hospitals. Moreover, considering WSIs as domain can be convenient when you have no more information than the WSI itself.

Tissue type classification in colorectal histological images. For this DG task, we use 3 different datasets of colorectal cancer histological images namely: Kather16 [9], Kather19 [8] and CRC-TP [7]. Kather16 contains 5000 patches from 10 H&E WSIs spread into 8 classes. Kather19 includes 100000 patches from several H&E WSIs spread into 9 classes. CRC-TP comprises a total of 196000 patches extracted from 20 H&E WSIs and spread over 7 classes. Given that the hospital or slide information are not provided for Kather19 and CRC-TP, we only use Kather16 as the train set and Kather19 with CRC-TP as test sets. Because class definitions can vary from one dataset to another, we follow the class grouping suggested by [1]. Additionally, we remove the class *complex stroma* from Kather16 and CRC-TP which do not share the same definition. Finally, our problem consists in classifying patches from colorectal cancer histological images into 7 classes (*tumor*, *normal*, *stroma*, *lymphocytes*, *debris*, *adipose* and *background*) and to generalize well on unseen domains (Kather19, CRC-TP). Since our training set (Kather16) does not provide hospital information for each patch, we considered WSIs as source domains ($S = 10$).

3.2 Method evaluation

In this section, the proposed TTA with the different ensembling methods are evaluated on the two different tasks. We report a baseline (base) corresponding to cross-entropy minimization on the source domains. Additionally, we also evaluate

	Lymph node patch classification			Colorectal tissue type classification	
	`id_val`	`val`	`test`	Kather19	CRC-TP
Base	**98.5(0.1)**	80.3(4.4)	62.7(7.4)	70.8(1.8)	43.6(3.7)
SOTA domain generalization methods					
CORAL [19]		86.2(1.4)	59.5(7.7)		
IRM [2]		86.2(1.4)	64.2(8.1)		
Group DRO [15]		85.5(2.2)	68.4(7.3)		
FISH [17]		83.9(1.2)	74.7(7.1)		
Color normalization					
Macenko [13]		81.6(1.2)	92.5(1.8)	64.9(3.3)	51.6(1.2)
Train-time data augmentation					
Base + RandAugment [4]		90.6(1.2)	82.0(7.4)	72.6(3.2)	51.0(5.1)
Base + H&E staining color jitter [20]		88.0(4.2)	91.6(1.9)	72.4(3.3)	47.6(4.8)
Base + StarGanV2 data aug	**98.4(0.0)**	89.6(0.7)	76.4(4.5)	66.3(3.3)	37.7(6.1)
Test-time data augmentation					
Base + StarGanV2 data aug + geometric TTA	**98.5(0.1)**	90.0(0.6)	76.5(4.2)	66.2(3.5)	37.6(5.9)
Base + color jittering + color jittering TTA	97.4(0.1)	91.2(0.4)	77.2(1.2)	**74.5(2.5)**	**51.7(6.1)**
Our method					
Base + StarGanV2 data aug + naive ens	96.9(0.3)	**92.8(0.2)**	**94.0(1.2)**	**72.9(2.0)**	**54.4(3.2)**
Base + StarGanV2 data aug + disc top-k ens	97.3(0.2)	92.1(0.2)	**94.0(1.2)**	**72.9(2.0)**	**54.4(3.2)**
Base + StarGanV2 data aug + disc weighted ens	97.4(0.2)	**92.7(0.2)**	**94.0(1.2)**	**72.9(2.0)**	**51.6(3.3)**

Table 1: Accuracy on Camelyon17 WILDS, weighted F1-score on Kather19 and CRC-TP. *Naive ens*, *disc top-k ens* and *disc weighted ens* referred respectively to naive ensembling, discriminator based top-k ensembling and discriminator based weighted ensembling. Best performances are highlighted in **bold** and standard deviations are specified in parenthesis.

a method named StarGanV2 data aug where StarGanV2 is used only for train-time data augmentation. The performances on Camelyon17, Kather19, CRC-TP are respectively averaged over 10, 5 and 5 independent runs and reported on Table 1. Implementation details about the StarGanV2 and classifier are provided in Table 1 of supplementary material.

State-of-the-art results on Camelyon17 WILDS Our method obtains state-of-the-art results on Camelyon17 WILDS. Results of competing methods are reported in the "SOTA domain generalization methods", "Color normalization" and the two first lines of the "Train-time data augmentation" row. On the `val` and `test` splits, our TTA method with naive ensembling outperforms the previous best method by +2.2% and +1.5% respectively. It is worth mentioning that our method learns the domains variations from the data while the two previous best methods on `test` use expert priors (H&E staining color jittering and Macenko color normalization). Therefore, in theory, our method could generalize on other medical imaging modalities. Compared to the best DG method FISH, that do not use any priors on the domain change, our method leads to much better generalization with +8.9% and +19.3% accuracy improvements. The TTA method performs also well on source domain data (`id_val` split) even if performances are slightly worse than the baseline.

Better robustness on Kather19 and CRC-TP The second task is more challenging as only 4400 images are available in Kather16 to train the StarGanV2

and the classifier. The experiments on these datasets have two objectives: show that our method generalizes on another problem/dataset and that it can work even with relatively small size datasets. Previous works usually use Kather16 for evaluation making performances comparisons impossible. However, the results on Kather19 and CRC-TP shows that our method increases the baseline results significantly: +2.1% and +11.8% respectively. These improvements are not as large as in Camelyon17 which is probably due to a less robust StarGanV2 model. This hypothesis is reinforced by the bad performances of train-time data augmentation using the StarGanV2.

3.3 Ablation studies

Analyzing data augmentation effects We have studied independently the effects of the train-time and test-time data augmentations based on the same StarGanV2. The train-time data augmentation with StarGanV2 alone helps to make the model more robust on Camelyon17 (+9.3% and +13.7% for `val` and `test` splits). However, on Kather16, where the dataset is small, it leads to loss of performances (−4.5% and −5.9%). On both tasks, when train-time data augmentation based on StarGanV2 is combined with our TTA method, generalization performances improve considerably.

To evaluate the impact of our TTA, geometric TTA and color jittering TTA have been evaluated. Both perform quite well: in particular, the color jittering prior seems to be correct for 3 out of the 4 target datasets making it as good as our method on the colorectal datasets and nearly as good on the Camelyon17 `val` split. However, color jittering TTA fails on the Camelyon17 `test` split (−16.8% compared to our TTA method). This suggests that the augmented target domain distribution is not overlapping enough the source domains distributions and that projecting images onto source domains is a better approach to achieve generalization. A t-SNE [12] plot provided in Fig. 2 of supplementary material confirms that projecting target examples onto the different source domains is one of the key to achieve better generalization.

Comparison of the ensembling strategies The naive ensembling has slightly better results than the two other proposed ensembling strategies. This seems counter-intuitive but two reasons might explain this result. Firstly, domain discriminators learned during the StarGanV2 training may not be adapted for the classification tasks. Secondly, the small number of WSIs in each training dataset (10 and 30) might be problematic: selecting a subset of projected images predictions when their number is already small, even if images are not well projected, might remove some of the advantages of standard ensembling. To explore this question, performances of discriminator based top-k ensembling and weighted ensembling with respect to their hyperparameters k and T are reported on Figure 3. A small temperature T corresponds to top-1 ensembling while a large T is equivalent to naive ensembling or top-k with the biggest possible k.

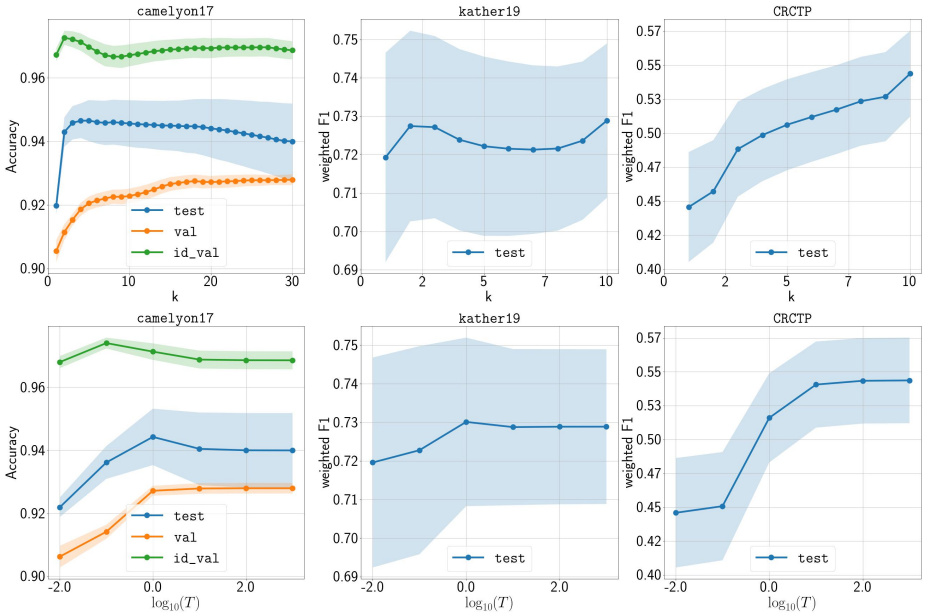

Fig. 3: Performances of *discriminator based top-k ensembling* with respect to k (1^{st} row) and *discriminator based weighted ensembling* with respect to T (2^{nd} row) for Camelyon17 WILDS, Kather19 and CRC-TP datasets.

On the four target datasets (`val`, `test` splits of Camelyon17, Kather19 and CRC-TP), the general trend is that the more predictions we consider in the top-k ensembling or the higher the temperature T for the weighted ensembling the better the performances are. Increasing k or T in the first steps enables significant gains while for higher values the performances continue to improve or decrease only slightly. It should be noted that the decrease in performances occurs on the Camelyon17 `test` split where the number of domains is the largest in our experiments (30). In this case, the naive ensembling does not lead to the best performance and it seems that predictions of badly projected images decrease the performance of the ensembling method. This observation suggests that when a bigger number of available domains (WSIs) is available, the top-k and weighted ensembling methods might generalize better.

Currently, when dealing with many domains our TTA can face a computational overhead. In Table 2 of supplementary material, time complexity, scalability of the approach are investigated and suggestions are proposed to limit such overhead.

Acknowledgments This work was partially supported by the ANR Hagnodice ANR-21-CE45-0007. Experiments have been conducted using HPC resources

from the "Mésocentre" computing center of CentraleSupélec and École Normale Supérieure Paris-Saclay supported by CNRS and Région Île-de-France.

4 Conclusion

To tackle problems of domain generalization in histopathology, we have proposed a new TTA technique based on the multi-domain image-to-image translation model StarGanV2 and explored three different ensembling strategies. Experiments have been conducted on several histopathology datasets where the method has proven to be more efficient than previous color augmentation/normalization, TTA and DG methods. Even in the low data regime and when WSIs are coming from a single hospital, our method still performs well. Finally, since the method is performed at test-time, it can be combined with any other DG or data augmentation techniques applied at training-time. In the future, we plan to explore StarGanV2 training on unrelated histopathology data (multi organs/cancers) and exploit it only for train-time data augmentation to check if even better protocol invariance could be achieved.

References

1. Abbet, C., Studer, L., Fischer, A., Dawson, H., Zlobec, I., Bozorgtabar, B., Thiran, J.P.: Self-rule to adapt: Generalized multi-source feature learning using unsupervised domain adaptation for colorectal cancer tissue detection. arXiv preprint arXiv:2108.09178 (2021)
2. Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant risk minimization. arXiv preprint arXiv:1907.02893 (2019)
3. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: Stargan v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8188–8197 (2020)
4. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 702–703 (2020)
5. Faryna, K., van der Laak, J., Litjens, G.: Tailoring automated data augmentation to h&e-stained histopathology. In: Medical Imaging with Deep Learning. PMLR (2021)
6. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proceedings of the European conference on computer vision (ECCV). pp. 172–189 (2018)
7. Javed, S., Mahmood, A., Fraz, M.M., Koohbanani, N.A., Benes, K., Tsang, Y.W., Hewitt, K., Epstein, D., Snead, D., Rajpoot, N.: Cellular community detection for tissue phenotyping in colorectal cancer histology images. Medical image analysis **63**, 101696 (2020)
8. Kather, J.N., Krisam, J., Charoentong, P., Luedde, T., Herpel, E., Weis, C.A., Gaiser, T., Marx, A., Valous, N.A., Ferber, D., et al.: Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study. PLoS medicine **16**(1), e1002730 (2019)

9. Kather, J.N., Weis, C.A., Bianconi, F., Melchers, S.M., Schad, L.R., Gaiser, T., Marx, A., Zöllner, F.G.: Multi-class texture analysis in colorectal cancer histology. Scientific reports **6**(1), 1–11 (2016)
10. Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., et al.: Wilds: A benchmark of in-the-wild distribution shifts. In: International Conference on Machine Learning. pp. 5637–5664. PMLR (2021)
11. Lafarge, M.W., Pluim, J.P., Eppenhof, K.A., Moeskops, P., Veta, M.: Domain-adversarial neural networks to address the appearance variability of histopathology images. In: Deep learning in medical image analysis and multimodal learning for clinical decision support, pp. 83–91. Springer (2017)
12. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research **9**(11) (2008)
13. 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. pp. 1107–1110. IEEE (2009)
14. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. IEEE Computer graphics and applications **21**(5), 34–41 (2001)
15. Sagawa, S., Koh, P.W., Hashimoto, T.B., Liang, P.: Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. arXiv preprint arXiv:1911.08731 (2019)
16. Shaban, M.T., Baur, C., Navab, N., Albarqouni, S.: Staingan: Stain style transfer for digital histological images. In: 2019 Ieee 16th international symposium on biomedical imaging (Isbi 2019). pp. 953–956. IEEE (2019)
17. Shi, Y., Seely, J., Torr, P.H., Siddharth, N., Hannun, A., Usunier, N., Synnaeve, G.: Gradient matching for domain generalization. arXiv preprint arXiv:2104.09937 (2021)
18. Stacke, K., Eilertsen, G., Unger, J., Lundström, C.: Measuring domain shift for deep learning in histopathology. IEEE journal of biomedical and health informatics **25**(2), 325–336 (2020)
19. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: European conference on computer vision. pp. 443–450. Springer (2016)
20. Tellez, D., Balkenhol, M., Otte-Höller, I., van de Loo, R., Vogels, R., Bult, P., Wauters, C., Vreuls, W., Mol, S., Karssemeijer, N., et al.: Whole-slide mitosis detection in h&e breast histology using phh3 as a reference to train distilled stain-invariant convolutional networks. IEEE transactions on medical imaging **37**(9), 2126–2136 (2018)
21. Tellez, D., Litjens, G., Bándi, P., Bulten, W., Bokhorst, J.M., Ciompi, F., Van Der Laak, J.: Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. Medical image analysis **58**, 101544 (2019)
22. Vasiljević, J., Feuerhake, F., Wemmert, C., Lampert, T.: Towards histopathological stain invariance by unsupervised domain augmentation using generative adversarial networks. Neurocomputing **460**, 277–291 (2021)
23. Wagner, S.J., Khalili, N., Sharma, R., Boxberg, M., Marr, C., Back, W.d., Peng, T.: Structure-preserving multi-domain stain color augmentation using style-transfer with disentangled representations. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 257–266. Springer (2021)

Test-time image-to-image translation ensembling improves out-of-distribution generalization in histopathology (Supplementary material)

Marin Scalbert[1,2], Maria Vakalopoulou[1], and Florent Couzinié-Devy[2]

[1] MICS, CentraleSupélec - Université Paris-Saclay, Gif-sur-Yvette, France
{marin.scalbert, maria.vakalopoulou}@centralesupelec.fr
[2] VitaDX International, Paris, France
f.couzinie-devy@vitadx.com

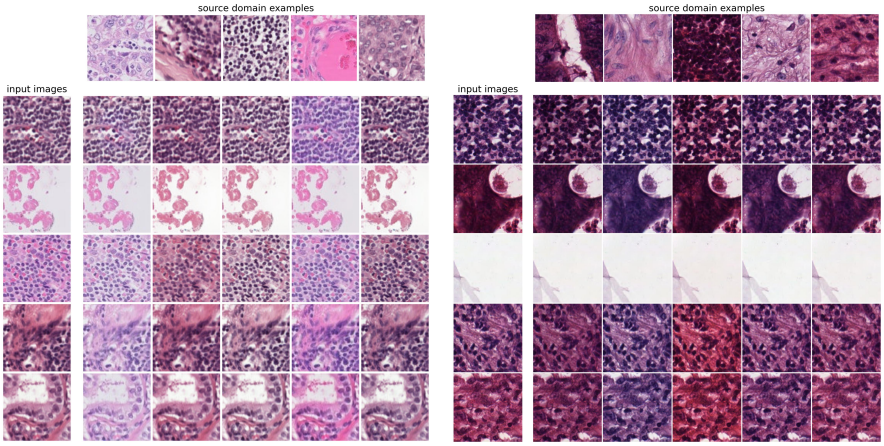

Fig. 1: Translated images examples from Camelyon17 (left grid) and Kather16 (right grid). In each grid, the first column shows input images fed to StarGanV2 and the first row represents random examples from 5 target domains.

	Camelyon17	Kather16
image resolution	(128, 128)	(256, 256)
S	30	10
style vector dimension	64	64
$\lambda_{adv}, \lambda_{cyc}, \lambda_{ds}, \lambda_{percep}, \lambda_{sty}, \lambda_{gp}$	1	1
batch size	96	64
training steps	50K	40K
optimizer	Adam	Adam
learning rate	$1e^{-4}$	$1e^{-4}$
GPUs	4 Nvidia V100	4 Nvidia V100

	Camelyon17	Kather16
image resolution	(128, 128)	(256, 256)
model	DenseNet121	ResNet18
pretraining	None	ImageNet
batch size	32	32
training steps	40K	4K
optimizer	Adam	Adam
learning rate	$1e^{-4}$	$1e^{-4}$
GPUs	1 Nvidia V100	1 Nvidia V100

Table 1: StarGanV2 (left table) and classifier (right table) hyperparameters for Camelyon17 and Kather16 datasets.

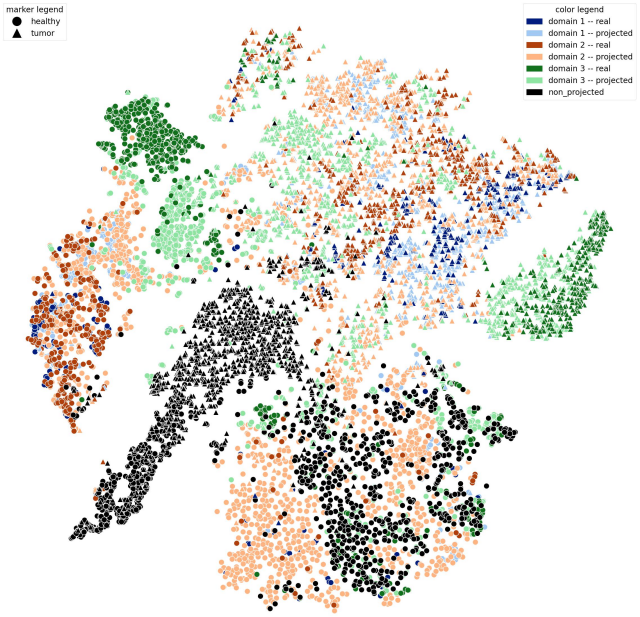

Fig. 2: t-SNE of classifier last features representation for Camelyon17 `test` split and 3 random source domains. Non projected target examples (black) tend to be isolated. However, when projected into a specific source domain (light color), they come closer to real examples from the domain (darker color). Additionally, this projection seems to preserve the class label: examples are still clustered by class labels (healthy/tumor).

TTA method	Inference time per batch(ms)
No TTA	15.7
Geometric TTA	69.7
H&E Color Jitter TTA	144.0
StarGanV2 TTA	667.0
Lighter StarGanV2 TTA	105.6

Table 2: Inference times per batch for different TTA methods on Camelyon17 WILDS images with a 32 batch size and 10 domains. Geometric and H&E Color Jitter TTA are indeed faster: our method is approximatively and respectively 10 and 4 times slower. However, our TTA leads much larger performances gains. Computational overhead of the method could be mitigated by using a more modern generator architecture without modifying the method (Lighter StarGanV2 TTA is 6 times faster) or by grouping the WSIs based on their protocol.