MultiPathGAN: Structure Preserving Stain Normalization using Unsupervised Multi-domain Adversarial Network with Perception Loss

Haseeb Nazki*
hn33@st-andrews.ac.uk
School of Computer Science, University of St Andrews
St Andrews, UK

Ognjen Arandjelović
oa7@st-andrews.ac.uk
School of Computer Science, University of St Andrews
St Andrews, UK

InHwa Um
ihu@st-andrews.ac.uk
School of Medicine, University of St Andrews
St Andrews, UK

David Harrison
david.harrison@st-andrews.ac.uk
School of Medicine, University of St Andrews
St-Andrews, UK

arXiv:2204.09782v2 [eess.IV] 2 Aug 2022

ABSTRACT

Histopathology relies on the analysis of microscopic tissue images to diagnose disease. A crucial part of tissue preparation is staining whereby a dye is used to make the salient tissue components more distinguishable. However, differences in laboratory protocols and scanning devices result in significant confounding appearance variation in the corresponding images. This variation increases both human error and the inter-rater variability, as well as hinders the performance of automatic or semi-automatic methods. In the present paper we introduce an unsupervised adversarial network to translate (and hence normalize) whole slide images across multiple data acquisition domains. Our key contributions are: (i) an adversarial architecture which learns across multiple domains with a *single* generator-discriminator network using an information flow branch which optimizes for perceptual loss, and (ii) the inclusion of an additional feature extraction network during training which guides the transformation network to keep all the structural features in the tissue image intact. We: (i) demonstrate the effectiveness of the proposed method firstly on H&E slides of 120 cases of kidney cancer, as well as (ii) show the benefits of the approach on more general problems, such as flexible illumination based natural image enhancement and light source adaptation.

CCS CONCEPTS

• **Applied computing** → **Imaging**.

KEYWORDS

Digital pathology, multi-domain image translation, generative adversarial networks, semantic structure

1 INTRODUCTION

Digitization of pathological and histopathological slides in the medical imaging community has driven a significant increase in the development of computer aided diagnostic (CAD) systems [10, 32, 45]. One of the important steps in the process of tissue preparation for whole slide image (WSI) acquisition is staining, whereby histochemical stains such as hematoxylin and eosin (H&E) are used to highlight the intensity of tissue slice components and thus make tissue structures distinct from each other [5]. Staining facilitates the use of WSIs in manual or automated analysis by pathologists or CAD algorithms as a means of further the insight into the symptoms or mechanisms associated with a particular disease. However, batches of these slides collected from different groups or laboratories exhibit diverse stain styles due to high variability in staining protocol, dye manufacturers and scanners used to acquire these WSIs [3]. This variability may not greatly hinder the analysis of the tissue structure by an experienced pathologist; however, it significantly reduces the accuracy of existing CAD systems owing to their limited ability for generalization from data used to train them [38]. Meanwhile, given the massive amount of gigapixel-sized WSI data [12], there is a growing demand to build fast, automated, and scalable pipelines for large-scale image analysis. In particular, normalization methods increase the accuracy of machine learning (ML) algorithms that use stain-normalized images as input to a pre-trained deep network [1, 16, 31].

Considering the importance of the problem, it is not surprising that a number of attempts at solving it have been described in the literature. Simple attempts which focused on color matching [34, 46], that is the color-channel alignment of a novel WSI with those of a reference template, can lead to unrealistic looking results as a consequence of the lack of learning over a representative corpus and the reliance on a single target image. On the other hand, while stain-separation methods [10, 29] do consider each staining channel independently for normalization, they fail to take into account the spatial features and the structure of the components in the tissue. Moreover, they rely on a skillfully chosen reference template image which can significantly affect the outcome [37]. A related machine learning based approach employs a sparse autoencoder [22], dividing an input image into multiple tissue regions and then independently normalizing each region using a single template. It can be readily seen that this approach suffers from the same limitation as that noted before, namely that a single reference template fails to capture the significant variation in the appearance of WSIs within a single lab.

More recently, deep learning-based approaches for stain normalization [4, 11, 24, 25, 36, 37, 42] using style transfer and image to image translation with adversarial learning methods [7, 14, 21, 47]

*Corresponding Author

have made strides in first establishing and then repeatedly improving the state-of-the-art. However, all of the existing adversarial network based methods are capable of translating between only two specific domains which makes them unsuitable for the inherently multi-domain challenge posed by the real-world clinical practice. Some of these approaches even rely on supervision, learning a mapping between the two domains given a labelled set of input-output pairs which are not always readily available, particularly for WSI datasets. Moreover, while these methods transfer the style content between the domains efficiently, they too struggle at preserving the structure of the tissue components. Previous efforts to develop algorithms capable of preserving the anatomical structure of tissue in a WSI while also translating the style between domains in an unsupervised setting have shown decent results [11, 13, 26, 31]; nevertheless these methods too remain constrained by being able to deal with two specific domains only. Wagner et al. [43] use a GAN architecture based on disentangled representations [28], which produces realistic looking high level structures but which poorly semantically correspond to the input.

In this work, we introduce MultiPathGAN, a unified deep learning adversarial network for translating histopathological WSIs obtained from different pathology labs/scanners across multiple domains for stain normalization while preserving the structure of the internal tissue components. The data adaptive nature of the MultiPathGAN means that it can also be readily employed in other applications wherein it is critical to constrain the effects of adversarial learning in a manner that retains the salient input image structure, while only translating the style content across multiple domains. This alleviates the problem of variation in the feature representation of training data by having a unified representation that can map many-to-many stain style domains instead of a simple one-to-one representation. Importantly, our network takes in training data from multiple domains, and learns the mappings between *all of them* using only a single generator as proposed by Choi et al. [6]. However, unlike StarGAN [6] which translates images between multiple domains without effectively preserving the content of the WSI while changing the style-related part of the inputs, we also achieve the preservation of the fine salient anatomical structure using an auxiliary feature extraction network and perception loss to minimize the perceptual distance [27] between a real and the corresponding generated (synthetic), fake image. To summarize, we make the following key contributions:

- We propose MultiPathGAN: an unsupervised multi-domain image translation network based on adversarial learning to normalize the stain variations between multiple WSI domains. Training a single generator on multiple domains increases scalability and robustness as it allows the translation process to be steered towards any desired domain given its target label from any large set of available domain classes.
- To preserve the semantic information and capture the geometric and structural patterns of the source image, we introduce an additional feature extraction network in order to regulate multi-domain adversarial learning in a multi-domain setting. Consisting of a pre-trained convolutional classifier, the feature extractor aids the generator to optimize on the perceptual loss by providing perceptual distance between the extracted feature maps of the source and target WSI domains.

2 METHOD

In this section we first describe in broad strokes the proposed framework (MultiPathGAN) for multi-domain image-to-image translation aimed at preserving the structure of the salient input image content while effectively integrating it with the style content of the target. Following this coarse overview of the method, each of its constituent components is elaborated on in detail.

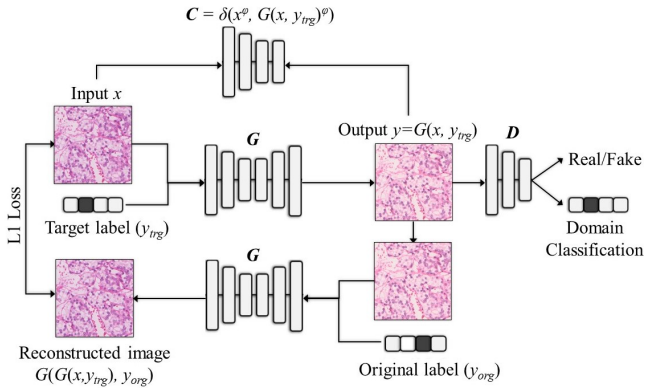

Figure 1: Information flow in MultiPathGAN. Generator G generates output image y, given input image x and the randomly generated target domain label (y_{trg}). Discriminator D distinguishes real from fake images, and simultaneously assures that y is classifiable as being in the target domain. The input x is also reconstructed from y to enforce self-consistency. Concurrently, x and y are also passed through a common feature extraction network C which provides feature maps for perceptual distance minimization.

2.1 MultiPathGAN

One of the key ideas in MultiPathGAN lies in the use of an auxiliary feature extraction network. This network is used to integrate perceptual loss in a multi-domain adversarial network, thereby introducing a semantic relationship between the input and the output, and drives the network to learn how to retain the salient structural elements of the input.

As illustrated in Fig. 1, MultiPathGAN trains a single generator G to learn mappings between multiple domains using target labels. To achieve this, we train G to translate an input image x into an output image y conditioned on the randomly generated target domain label (y_{trg}), such that $G(x, y_{trg}) \rightarrow y$. The target domain label y_{trg} is generated randomly so that G learns to translate the input image flexibly [6]. Further, an auxiliary classifier allows a single discriminator to control multiple domains by producing probability distributions over both sources and domain labels [6]. Note that the single generator is used twice, first to translate the original image into the target domain and then to reconstruct the original image from the translated image. Simultaneously, the input image

x and output image y are fed to a pre-trained deep convolutional classifier C to obtain feature maps $x^{\phi}, G(x, y_{trg})^{\phi}$ and obtain the mean squared distance between these activations [27]. This distance is used in an additional perception loss function to help preserve the perceptually subtle salient content in the input image as it is translated to the output domain [27, 33]. In the subsections which follow we provide detailed descriptions of all loss functions used in training the network that make this preservation possible.

2.1.1 Adversarial Loss Function. As in previous work, G tries to generate an image $G(x, y_{trg})$ conditioned on both the input image x and the target domain label y_{trg}. At the same time, a discriminator D tries to distinguish between real and fake images, x and $G(x, y_{trg})$ respectively. To stabilize the training process and generate higher quality output of each mapping that matches the empirical distribution of the target domain, while focusing on the source domain, we adopt the Wasserstein adversarial loss [2, 30] with gradient penalty [15]:

$$\mathcal{L}_{Adv} = \mathbb{E}_x[D_S(x)] - \mathbb{E}_{x, y_{trg}}[D_S(G(x, y_{trg}))] - \lambda_{gp}\mathbb{E}_{\hat{x}}[(\|\nabla_{\hat{x}} D_S(\hat{x}\|_2 - 1)^2] \quad (1)$$

where λ_{gp} is the coefficient for gradient penalty and $\hat{x}$ is uniformly sampled along a straight line between a real and a generated image pair [15]. D_S refers to the probability distribution over the sources given by discriminator D.

2.1.2 Domain Classification Loss Function. We also adopt the domain classification loss [6] using an auxiliary classifier on top of D when optimizing G and D to ensure that the input x translated into an output image y is properly classified to the target domain y_{trg}. This objective can be broken down into two terms: a domain classification loss calculated using real images to optimize D, and a domain classification loss calculated using fake images to optimize G:

$$\mathcal{L}_C^r = \mathbb{E}_{x, y_{org}}[-logD_C(y_{org}|x)] \quad (2)$$

$$\mathcal{L}_C^f = \mathbb{E}_{x, y_{trg}}[-logD_C(y_{trg}|G(x, y_{trg})] \quad (3)$$

where D_C is the probability distribution calculated by D over domain labels. In (2), D tries to minimize the objective by learning to classify a real image x to its corresponding original domain y_{org}. At the same time, in (3), G tries to minimize the objective to generate images that could be classified as the target domain y_{trg}.

2.1.3 Reconstruction Loss Function. To ensure that the translated images preserve the high level features of the input images while changing their domain specific characteristics, we apply a cycle consistency loss [6, 47] to the generator G:

$$\mathcal{L}_{Cyc} = \mathbb{E}_{x, y_{trg}, y_{org}}\|x - G(G(x, y_{trg}), y_{org})\|_1 \quad (4)$$

where G takes in the translated image $G(x, y_{trg}) \rightarrow y$ and the original domain label y_{org} as input and tries to reconstruct the original image x using the $L1$ norm in the loss function.

2.1.4 Perceptual Loss Function. In addition to the reconstruction loss used to retain the high level features, we also include a perceptual loss which preserves the fine structure of the input image during domain transfer as the perceptual distance is minimized [23, 27] between the generated $G(x, y_{trg})$ and the original image x. This distance is calculated by feeding both the original and the transformed image to a pre-trained classification network C and extracting their feature maps $x^{\phi}, G(x, y_{trg})^{\phi}$ respectively, at spatial resolution ϕ:

$$\mathcal{L}_P = \delta(x^{\phi}, G(x, y_{trg})^{\phi}) \quad (5)$$

The mean squared distance is calculated for each feature map and the total perceptual distance is then a weighted sum of the $L2$ distances. By minimizing this distance, we implicitly compel the transformed image to contain the same structural meaning as the original image as perceived by a pre-trained classification network.

2.1.5 Final Objective. The final objective used to train the discriminator D and the generator G can be written as follows:

$$\mathcal{L}_D = -\mathcal{L}_{Adv} + \lambda_C \mathcal{L}_C^r \quad (6)$$

$$\mathcal{L}_G = \mathcal{L}_{Adv} + \lambda_C \mathcal{L}_C^f + \lambda_{Cyc}\mathcal{L}_{Cyc} + \lambda_P \mathcal{L}_P \quad (7)$$

where λ_C, λ_{Cyc} and λ_P are the hyper-parameters to regulate the individual loss functions $\mathcal{L}_C$, $\mathcal{L}_{Cyc}$ and $\mathcal{L}_P$ respectively.

2.2 Implementation

2.2.1 Network Architecture and Training. In our network, the generator G is based on the six block ResNet [17] backbone and a pair of convolutional and transposed convolutional layers with the stride size of two on either side, respectively. We use instance normalization for the generator, which has been shown to improve performance [39]. For the discriminator D, we use PatchGANs [21] with no normalization and Leaky ReLU with a negative slope of 0.01 to classify the overlapping real and fake patches. Finally, for our feature extraction network C, we utilize a pre-trained 34-layer residual network [17]. We use the activations from the last convolutional layer of the ResNet feature extraction network to concentrate on the fine changes in these higher levels, propagated by its skip connections, which thus allows us to determine a better perceptual distance between the generated $G(x, y_{trg})$ and original image x.

For our experiments, we set the values of $\lambda_{gp} = 10$, $\lambda_C = 1$, $\lambda_{Cyc} = 10$ and $\lambda_P = 0.75$. We use a mini-batch size of 16 and a base learning rate of 0.0001 for both G and D which decays after every 10 epochs. Further, we use the Adam optimizer [9] with $\beta_1 = 0.5$ and $\beta_2 = 0.999$ and following Gulrajani et al. [15], we perform one G update for every five D updates over 80 epochs which takes about 21 hours to train our network on a single Nvidia DGX-1 instance.

3 EVALUATION

3.1 Dataset.

H&E slides were prepared from a cohort of 120 cases of kidney cancer obtained from the Pathology Archives, Lothian NHS. Sections were cut at 3 μm and routinely stained for H&E before being coverslipped. Slides containing both cancerous and non-cancerous "normal" tissue were distributed to different laboratories and scanned locally on their scanners of choice before all images were collated at <removed for review>. Ethics approval was provided by Lothian NHS Biorepository (ES/15/0094). More specifically, we used WSI acquired from Hamamatsu NanoZoomer S60 at $20\times$ $(0.42\mu m/pix)$ magnification, Hamamatsu NanoZoomer 2.0-HT at $20\times$ $(0.45\mu m/pix)$ and Leica Aperio AT2 at $20\times$ $(0.50\mu m/pix)$

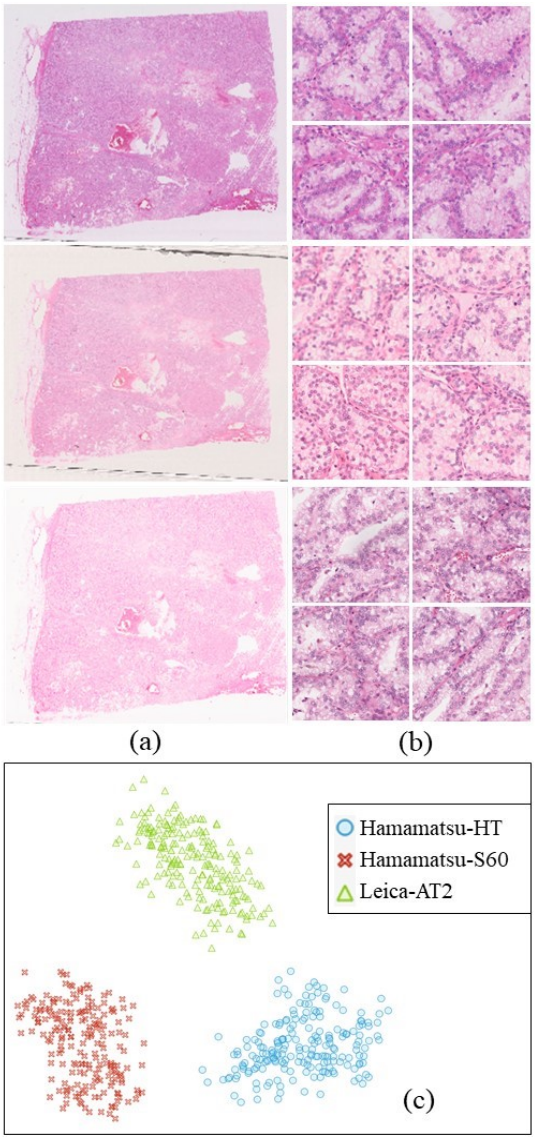

Figure 2: (a) WSI samples from the three domains, (b) the corresponding patches, and (c) TSNE visualization of our dataset samples taken from the three WSI domains.

magnifications, respectively, to create three dataset domains. The WSIs were then split into 256×256 non-overlapping patches. To train our MultiPathGAN, we used 10 WSIs for each domain and randomly extracted 1250 training and 200 testing sample patches for each class. We also used an additional domain of WSIs from Hitachi HV-F202SCL at $20\times(0.22\mu m/pix)$ magnification, unseen to the network at the training time. This domain was used to evaluate the generalization of our network and its ability to adapt to the high variance data space of WSIs in histopathological studies.

Fig. 2 shows examples of WSIs and their respective tiles from our dataset. Even though the variation in the acquired data is readily apparent to the naked eye, we corroborate this in Fig. 2(c) using t-Distributed Stochastic Neighbor Embedding (t-SNE) visualizations, which help visualize high dimensional image data by assigning each data point a location in a two or three dimensional map [40]. In our case, to concentrate on the specific features resulting in these variations in the WSI domains, the image embeddings are extracted from the penultimate layer of a pre-trained deep convolutional neural network trained on images of paintings by various artists. We chose to use the painters dataset specifically, to aid the network in distinguishing between the color spaces of different WSI domains, as well as to help us analyse the performance of our method by facilitating a meaningful visualization of the WSI clusters and the distances between them [41].

Although the primary motivation behind our work was biomedical in nature, in order to demonstrate the effectiveness of the proposed method on natural images, we also evaluate it on Virtual Image Dataset for Illumination Transfer (VIDIT) [18]. We derive five domains from the illumination setting at five different (2500K, 3500K, 4500K, 5500K and 6500K) color temperatures. All the domains have 2400 samples and we use the train:test split of 9:1 at both 256 and 1024 input-output resolutions.

3.2 Evaluation Metrics

To quantify how well different methods preserve the salient structures present in the original patches, we use Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM) [20], Multi-Scale Structural Similarity Index (MS-SSIM) [44] and Haar Wavelet-Based Perceptual Similarity Index (HaarPSI) [35] measures. While PSNR depends purely on the pixel-wise differences in pixel value between the two compared images, SSIM, MS-SSIM and HaarPSI are more correlated with the perceptual differences between them. These metrics have high sensitivity to detect distortions between images as well as the differences arising due to luminance and contrast changes.

3.3 Experimental Results

3.3.1 Qualitative Analysis. We train and evaluate MultiPathGAN on our dataset with three WSI domains. Fig. 3 shows typical examples of translation results. As confirmed independently by a pair of experienced pathologists, observe that our method provides high quality of translation, appropriately adapting the style of images (that is the image acquisition characteristics of specific domains), while correctly preserving the salient anatomical structures (cell types and morphology, etc.).

In addition, we also translate patches from the fourth unseen WSI domain to the three domains from our dataset used while training MultiPathGAN. As can be seen from Fig. 4, MultiPathGAN can effectively translate, and hence normalize, data from an alien subset of the WSI space into a known domain. This shows that normalizing unseen subsets of WSI color space using MultiPathGAN could be achieved without retraining the whole network. However, considering that this is an unseen domain, the inverse translation cannot be achieved without domain specific training, i.e. without the parameters which correspond to this domain, along with the corresponding target labels.

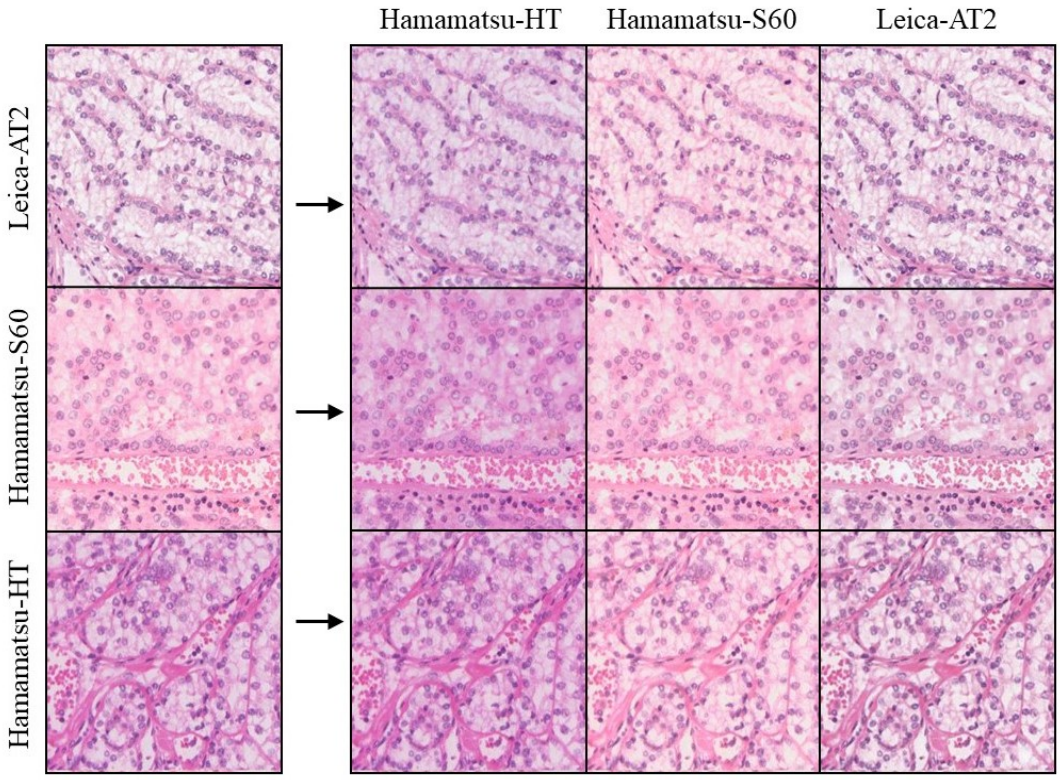

Figure 3: Multi-domain translation results using MultiPathGAN on our WSI dataset for inter-domain normalization. The column on the left shows the input WSI patch while the remaining columns on the right show the translated output for the three domains as per the target label. This is achieved using a single generator.

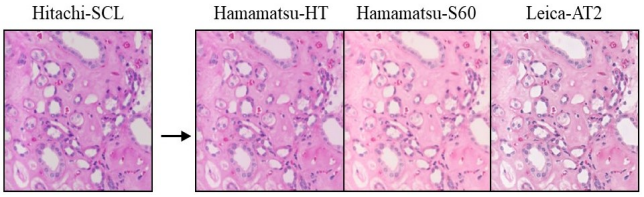

Figure 4: Translations from the unseen Hitachi-SCL (left-most column) domain to our three dataset domains used for training MultiPathGAN generator.

We also use TSNE to visualize the effect of our translation upon the normalization of WSI patches in the deep neural network feature space. From Fig. 5 it can be seen that the patches normalized using MultiPathGAN match well with the distinctive distribution of the target class. This can provide an intuition for the improved performance in terms of evaluation metrics as discussed in the next subsection. In addition to Fig. 4, we can also use Fig. 5(d) to extrapolate the effectiveness of MultiPathGAN in translating an unseen WSI domain into a known WSI domain without the need to retrain the entire network.

Finally, in Fig. 6 we compare our cross-domain MultiPathGAN to the existing methods in the literature and observe that our method not only provides high quality and perceptually meaningful WSI normalization results, but also that it can normalize multiple WSI domains simultaneously, using a single generator in a single training session. While the method of Runz et al. [36] effectively translates the style between two domains, it fails at preserving the quality of the final image and resulting in blur. Mahapatra et al. [31] have shown that using semantic guidance in addition to adversarial and cycle reconstruction loss in a CycleGAN, we can preserve the detailed structural information, however, their method is restricted to only two WSI domains. In their attempt, Vasilejic et al. [42] have used StarGAN to translate between multiple staining modalities, but when we consider stain variations in a single modality, it fails at preserving the fine structure of the tissue components in the translated output. In addition to the StarGAN, we also show results from its newer variant by Choi et al. [7] and found that the generator produces images that are perceptually valid but slightly blurry and completely out of context with regards to the input WSI tissue structure. Translated samples from Wagner et al. [43] look realistic when compared to StarGAN but it struggles at preserving the detailed structural information and introduces objects that are out of context from the original image.

3.3.2 Quantitative Analysis. For the quantitative evaluation, we first report the performance metrics for translations between two

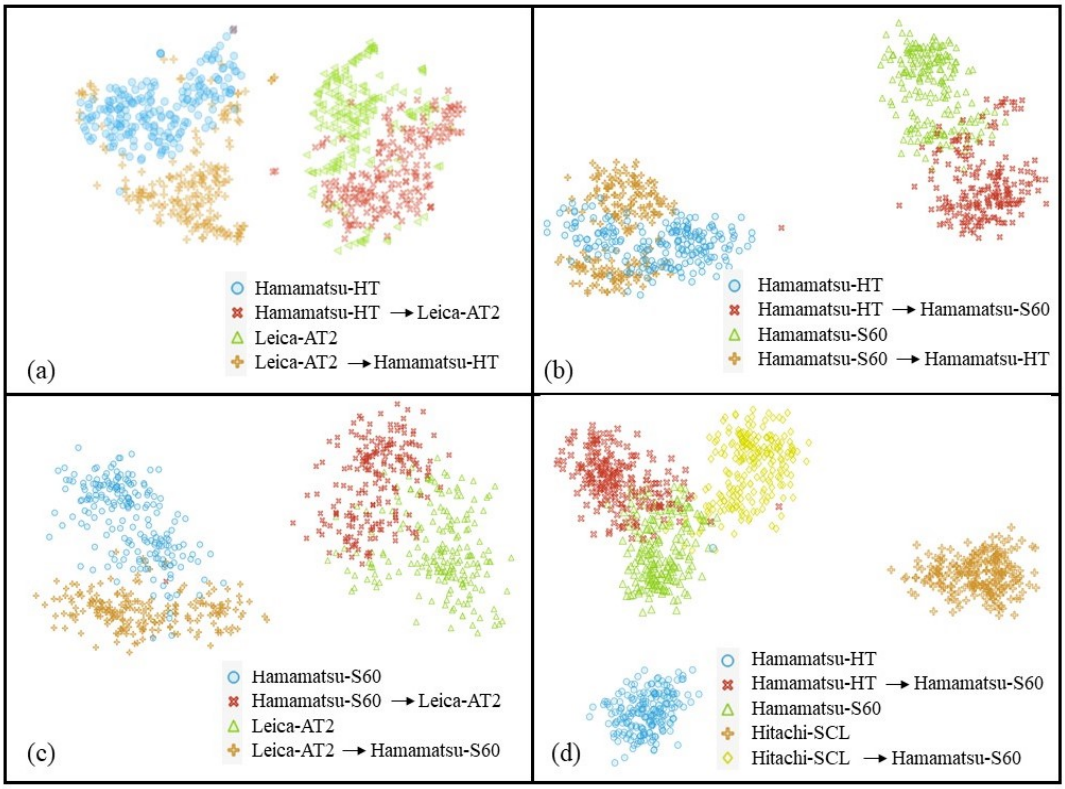

Figure 5: For comprehensibility, we only show TSNE representation for real and translated samples between two domains at a time (a)(b)(c). We also show the effectiveness of MultiPathGAN at translating unseen domain (Hitachi-SCL) to one of the domains used during the training phase (d).

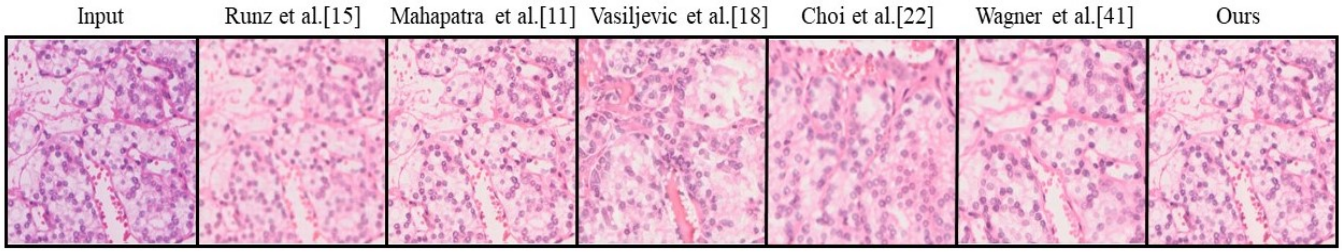

Figure 6: Comparison with previous methods when translating from Hamamatsu-HT → Hamamatsu-S60.

domains, namely those of the generated image $G(x, y_{trg})$ and of the original image x, for different methods in Table 1. We compute the evaluation metrics between a real image and the same image translated to other domains in all comparative assessments. As can be seen, our model performs well with all the metrics, indicating that our model produces the most realistic WSI patches by preserving the fabrication of the tissue as well as normalizing the stains. These outputs match the style of the target domain and preserve the structure of the input patch, thereby confirming and corroborating the findings of our previous qualitative analysis shown in Fig. 6.

To evaluate our model on the unseen data, we translate between the three test domains including the additional unseen class and

Table 1: Comparison with previous methods. We calculate PSNR, SSIM, MS-SSIM and HaarPSI between the test set of Hamamatsu-HT patches and Hamamatsu-HT → Hamamatsu-S60 translations.

Method	PSNR	SSIM	MS-SSIM	HaarPSI
Runz et al. [36]	21.16	0.80	0.92	0.77
Mahapatra et al. [31]	21.96	0.93	0.96	0.88
Vasiljevic et al. [42]	14.73	−0.20	0.01	0.53
Choi et al. [7]	16.19	0.20	0.15	0.38
Wagner et al. [43]	18.92	0.58	0.64	0.55
MuliPathGAN (Ours)	**22.96**	**0.96**	**0.98**	**0.93**

Table 2: Comparison of translation performance of MultiPathGAN between seen and unseen domain pairs.

Domain Transfer	PSNR	SSIM	MS-SSIM	HaarPSI
HT-to-S60	22.96	0.96	0.98	0.94
AT2-to-S60	27.21	0.95	0.95	0.96
S60-to-HT	26.98	0.94	0.95	0.95
AT2-to-HT	35.69	0.97	0.99	0.98
S60-to-AT2	22.63	0.96	0.97	0.93
HT-to-AT2	36.47	0.98	0.99	0.98
SCL-to-S60	18.50	0.90	0.94	0.85
SCL-to-HT	23.77	0.90	0.93	0.87
SCL-to-AT2	19.32	0.88	0.94	0.87

Abbreviations: HT: Hamamatsu-HT, S60: Hamamatsu-S60, AT2: Leica-AT2 and SCL: Hitachi-SCL

compare them with each other using the similarity metrics. From Table 2 we can notice that our model performs well with the domain transfer between the test data for all the cases. However, when we translate an unseen domain (i.e. HitachiSCI in this case) to one of the domains used at the time of training MultiPathGAN, there is a slight decrease in the performance in terms of the quantitative evaluation metrics. This difference is due to the lack of knowledge about the distance of the parameter distribution of our trained network from the distribution of the unseen domain and would remain constant for all such domains, provided they come from a similar locus. Nevertheless, these translations for unseen domains still deliver a good result and increasing the number of data domains to train MultiPathGAN would make the generator invariant to these unseen cohorts. This in return would further improve the performance of MultiPathGAN at normalizing unseen WSI domains.

3.3.3 Analysis of Individual Loss Components. To evaluate the effect of individual loss components in our network, we add them successively to our general baseline architecture while training on our dataset. We illustrate their effects in Fig. 7, where we normalize the input domain of our test data by translating it to the other domain (Hamamatsu-HT → Hamamatsu-S60). Moreover, Table 3 provides the results in terms of our chosen evaluation metrics for these configurations. We start with our baseline architecture consisting of our Generator G with a ResNet backbone, a PatchGAN discriminator D and an adversarial loss $\mathcal{L}_{Adv}$.

As can be seen from Fig. 7(a) and Table 3, and as expected from previous work, using adversarial loss undeniably helps translate the style of the input to the target output. However, it is also clear that without any semantic supervision it cannot produce a result which retains the correct semantic content, that is the relevant anatomical, structural information in the present case. Adding semantic guidance using reconstruction loss produces images whereby the generated output image correlates well with the input data, as can be seen in Fig. 7(b) and Table 3; nevertheless, it fails to preserve the fine-grained semantic information. This effect can be seen in Fig. 8, where we show the generator producing images with artifacts which do not exist in their corresponding input images when we only use reconstruction loss. Although adding classification loss forces the network to classify the input image into the corresponding target image which aids at rendering it into the style of the target domain, it again fails at preserving the structure of the input image; see Fig. 7(c) and Table 3. However, the effectiveness of classification loss can be observed in its absence in Fig. 7(d) where we can see the unified effect of reconstruction and perception loss. The resultant images highly correlate with the input in terms of both content and style, resulting in the least perceptual distance (see Table 3), therefore demonstrating that the network fails to translate it to the target domain without classification loss.

Table 3: Comparison of different MultiPathGAN losses.

Loss	PSNR	SSIM	MS-SSIM	HaarPSI
$\mathcal{L}_{Adv}$	18.20	0.39	0.49	0.45
$\mathcal{L}_{Adv} + \mathcal{L}_{Cyc}$	22.04	0.85	0.93	0.90
$\mathcal{L}_{Adv} + \mathcal{L}_{Cyc} + \mathcal{L}_{C}$	14.05	−0.22	0.03	0.56
$\mathcal{L}_{Adv} + \mathcal{L}_{Cyc} + \mathcal{L}_{P}$	37.76	0.97	0.99	0.99
$\mathcal{L}_{Adv} + \mathcal{L}_{C} + \mathcal{L}_{P}$	22.79	0.93	0.97	0.91
$\mathcal{L}_{Adv} + \mathcal{L}_{Cyc} + \mathcal{L}_{C} + \mathcal{L}_{P}$	22.96	0.96	0.98	0.93

We also train MultiPathGAN in the absence of reconstruction loss and observe a good performance with the generator trained on perception loss producing images that display tissue structures identical to the input in terms of fine-grained features, as can be seen in Fig. 7(d). However, as we can observe from Table 3, adding reconstruction loss preserves global semantic composition of the output image unnoticeable to the naked eye, resulting in better performance in terms of the evaluation metrics. Lastly, we observe that using the final MultiPathGAN loss produces images which are semantically indistinguishable from the input image (see Table 3) while at the same time also closer to the target domain in terms of its style composition, as illustrated by the example in Fig. 7(f).

3.3.4 Qualitative results on VIDIT. Finally, we demonstrate empirically that our method can maintain structural consistency between the original images and transferred images irrespective of the discrepancy in domain composition and style. As can be seen from Fig. 9, MultiPathGAN can synthetically render a scene as if re-illuminated using an unseen lighting setup, which is useful in a variety of tasks, for example in reference-based image relighting. More results demonstrating the effectiveness of MultiPathGAN and its robustness to lighting conditions on VIDIT can be found in the supplementary material.

4 DISCUSSION

Considering the widespread use of convolutional neural networks in CAD systems, stain normalization plays a crucial part in the interoperability and improved accuracy of their recommendations and results. MultiPathGAN provides a reliable approach to remove variations in the acquired images by providing clinically meaningful output images, normalized at the pixel resolution comparable to real biomedical images, for further assessment. There are several reasons for the improved robustness, flexibility, and accuracy of MultiPathGAN in WSI normalization when compared to all previously proposed methods. First, unlike these, MultiPathGAN learns the mapping from the entire dataset instead of relying on a single reference image, thus avoiding the complexity of choosing reference images. Second, most existing methods that use adversarial learning methods focus on using GANs to translate only between

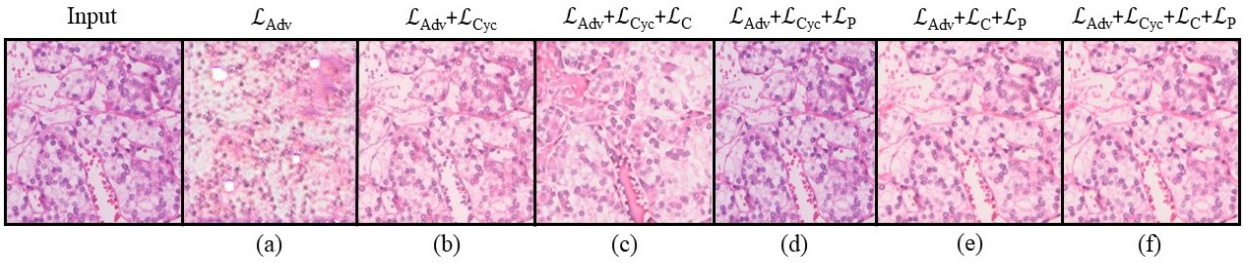

Figure 7: Visual comparison of different loss functions when translating from Hamamatsu-HT → Hamamatsu-S60.

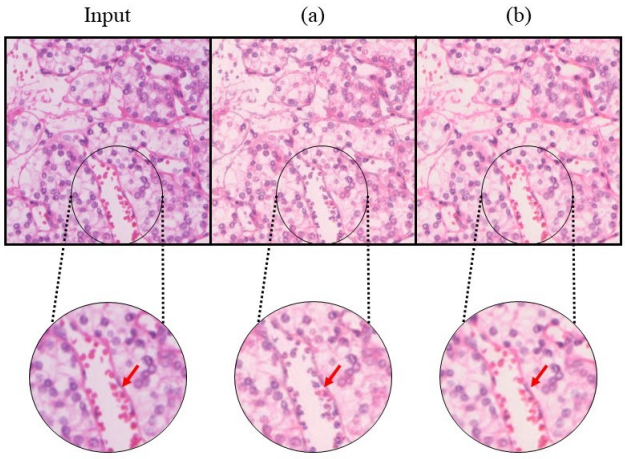

Figure 8: Effect of adding Perception loss $\mathcal{L}_P$ (b) when compared to Reconstruction loss $\mathcal{L}_{Cyc}$ (a).

two domains, thus limiting the potential of GANs to translate between multiple domains simultaneously using a single generator and style code [6]. Moreover, all existing methods are ineffective, in that even if multiple generators are used to translate between different domain pairs, the learning of the global features available to each generator is done separately. Hence, the entirety of information available within the training data corpus is not used to its full extent. This inability to fully utilize the training data by jointly training the domains inevitably limits the quality of generated samples and the ability of the generator to successfully normalize unseen WSI domains. Therefore, in MultiPathGAN we can say that the shared data from each domain helps to learn domain-invariant features which produces a regularization effect, thus facilitating better generalization to unseen WSI samples.

Lastly, we analysed different loss functions and introduced a pre-trained auxiliary ResNet feature extraction network to calculate a perceptual distance. This distance is employed in the additional perceptual loss function used while training the generator to preserve fine details of the input WSI patches in the normalized output. The smaller features are better propagated to higher convolutional levels in residual networks owing to skip connections, such that, there is an implicit weighing between fine and large features. We use the activations from the final layers of the ResNet feature extraction network to enable us to retain the finer details of the original image during the training process. This helps us to produce images which preserve the fine structure of the input image while effectively translating to the style of the target image domain.

The applications of MultiPathGAN are not limited to WSI stain normalization. Indeed, any application wherein there is high value in the exact semantic content of images can benefit from the method's ability to preserve such content while translating between multiple, even unseen domains.

5 CONCLUSIONS

To summarise, we presented the state-of-the-art adversarial model, MultiPathGAN, an unsupervised multi-domain image-to-image translation network that outperforms all other stain normalization methods devised for histopathological whole slide imagery. Our main contributions are: (1) domain invariance, whereby we can easily expand the set of domains (including to unseen ones) for normalization using only a single generator-discriminator, and (2) compelling the output image to preserve the salient anatomical tissue structure of the input image using an auxiliary feature extraction network and perception loss. We also demonstrated that our model can be readily adopted in other image-to-image translation applications in which it is important to preserve the structural content of the input image while only changing its style. In closing, MultiPathGAN has the potential to enhance the performance of medical image based diagnostic systems wherein annotations are not readily available, and its key strength stems from a meaningful constrainment of adversarial networks and their tendency to induce unwanted diversity which is manifest in undesirable image artifacts.

6 SUPPLEMENTARY MATERIAL

Additional Qualitative Results

Fig. 10 and Fig. 11 show additional images generated using MultiPathGAN at 256×256 resolution for stain normalization on our kidney tissue WSI dataset and scene re-illumination on VIDIT dataset, respectively. Fig. 12, shows translation results for VIDIT dataset generated at 256×256, 512×512 and 1024×1024 resolutions.

Additional Quantitative Results

We compare the performance metrics for MultiPathGAN trained on images at 256×256, 512×512 and 1024×1024 resolutions from the VIDIT dataset. We compute these metrics between the translated output and the real image at the target color temperature as depicted

Figure 9: Translation results using MultiPathGAN trained on five different color temperature domains from VIDIT.

in Fig. 12. As can be seen from Table 4, our model performs better according to all metrics at 256×256 resolution. We also compare the performance of MultiPathGAN trained at different λ_P values on our WSI dataset in Table 5.

Table 4: Comparison of translation performance of MultiPathGAN trained at 256×256, 512×512 and 1024×1024 resolutions using VIDIT dataset.

I/O Dimension	PSNR	SSIM	MS-SSIM	HaarPSI
256×256	**30.67**	**0.94**	**0.98**	**0.88**
512×512	24.67	0.86	0.94	0.78
1024×1024	20.90	0.73	0.92	0.63

Table 5: Comparison of translation performance of MultiPathGAN trained at various values of λ_P.

λ_P	PSNR	SSIM	MS-SSIM	HaarPSI
0.25	22.60	0.95	0.98	0.92
0.50	22.54	0.95	0.98	0.92
0.75	**22.96**	**0.96**	**0.98**	**0.93**
1.00	22.67	0.95	0.98	0.93

Mean Color Difference

To measure the color difference between the input image, output image and the target image, we calculate the difference of euclidean distances between their corresponding pixels in the CIELAB color space. [8, 19]. The resultant mean difference with higher values indicates a better match in terms of color value between the generated fake image and the target image. This value for MultiPathGAN comes out to be -0.0036 followed closely by Vasiljevic et al. [42][-0.0039]. For Runz et al. [36], Mahapatra et al. [31] and Wagner et al. [43] this metric value was -0.0054, -.0055 and -0.0056 respectively.

REFERENCES

[1] Andreea Anghel, Milos Stanisavljevic, Sonali Andani, Nikolaos Papandreou, Jan Hendrick Rüschoff, Peter Wild, Maria Gabrani, and Haralampos Pozidis. 2019. A high-performance system for robust stain normalization of whole-slide images in histopathology. *Frontiers in medicine* (2019), 193.

[2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*. PMLR, 214–223.

[3] Babak Ehteshami Bejnordi, Nadya Timofeeva, Irene Otte-Höller, Nico Karssemeijer, and Jeroen AWM van der Laak. 2014. Quantitative analysis of stain variability in histology slides and an algorithm for standardization. In *Medical Imaging 2014: Digital Pathology*, Vol. 9041. International Society for Optics and Photonics, 904108.

[4] Aïcha BenTaieb and Ghassan Hamarneh. 2017. Adversarial stain transfer for histopathology image analysis. *IEEE transactions on medical imaging* 37, 3 (2017), 792–802.

[5] Peter D Caie, Neofytos Dimitriou, and Ognjen Arandjelović. 2021. Precision medicine in digital pathology via image analysis and machine learning. In *Artificial intelligence and deep learning in pathology*. Elsevier, 149–173.

[6] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8789–8797.

[7] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. 2020. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8188–8197.

[8] Christine Connolly and Thomas Fleiss. 1997. A study of efficiency and accuracy in the transformation from RGB to CIELAB color space. *IEEE transactions on image processing* 6, 7 (1997), 1046–1048.

[9] Kingma Da. 2014. A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[10] Christel Daniel, Marcial García Rojo, Jacques Klossa, Vincenzo Della Mea, David Booker, Bruce A Beckwith, and Thomas Schrader. 2011. Standardizing the use of whole slide images in digital pathology. *Computerized Medical Imaging and Graphics* 35, 7-8 (2011), 496–505.

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

[12] Neofytos Dimitriou and Ognjen Arandjelovic. 2021. Magnifying networks for images with billions of pixels. *arXiv preprint arXiv:2112.06121* (2021).

[13] Michael Gadermayr, Vitus Appel, Barbara M Klinkhammer, Peter Boor, and Dorit Merhof. 2018. Which way round? A study on the performance of stain-translation for segmenting arbitrarily dyed histological images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 165–173.

[14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).

[15] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of wasserstein gans. *Advances in neural information processing systems* 30 (2017).

[16] Laxmi Gupta, Barbara M Klinkhammer, Peter Boor, Dorit Merhof, and Michael Gadermayr. 2019. GAN-based image enrichment in digital pathology boosts segmentation accuracy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 631–639.

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[18] Majed El Helou, Ruofan Zhou, Johan Barthas, and Sabine Süsstrunk. 2020. VIDIT: Virtual image dataset for illumination transfer. *arXiv preprint arXiv:2005.05460* (2020).

[19] Bernhard Hill, Th Roger, and Friedrich Wilhelm Vorhagen. 1997. Comparative analysis of the quantization of color spaces on the basis of the CIELAB color-difference formula. *ACM Transactions on Graphics (TOG)* 16, 2 (1997), 109–154.

[20] Alain Hore and Djemel Ziou. 2010. Image quality metrics: PSNR vs. SSIM. In *2010 20th international conference on pattern recognition*. IEEE, 2366–2369.

[21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.

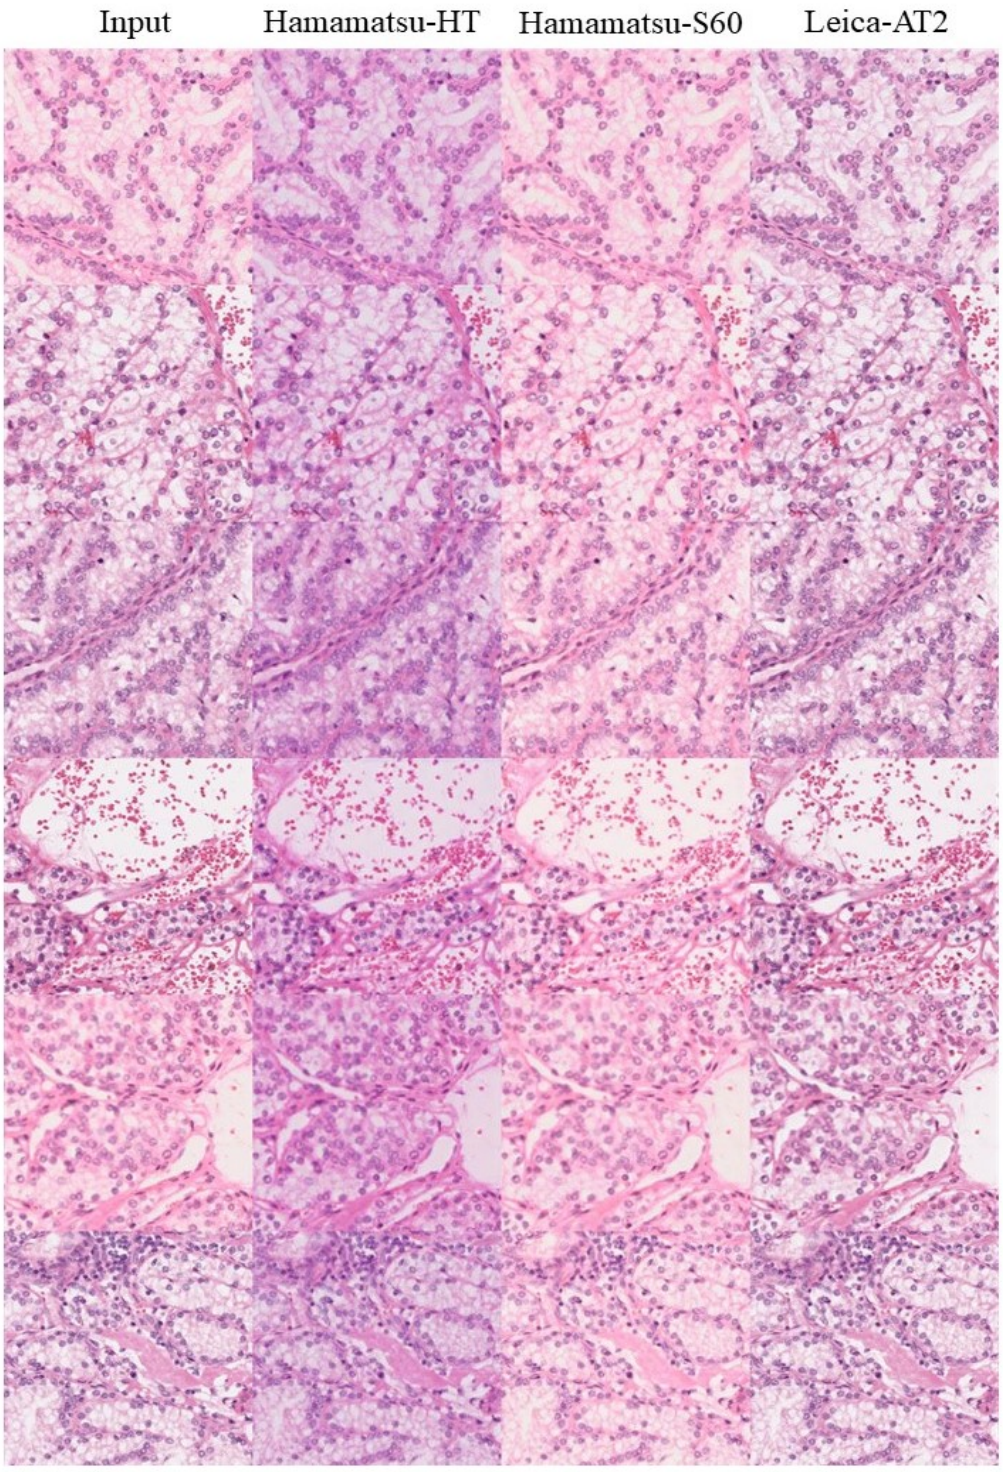

Figure 10: Additional translation results using MultiPathGAN on our WSI dataset for inter-domain normalization. The column on the left shows the random input WSI patch while the remaining columns on the right show the translated output for the three domains as per the target label.

Figure 11: Additional translation results using MultiPathGAN trained on five different color temperature domains from VIDIT at 256×256 **resolution. The column on the left shows the input scene at random color temperature while the remaining columns on the right show the translated output for different illumination temperatures.**

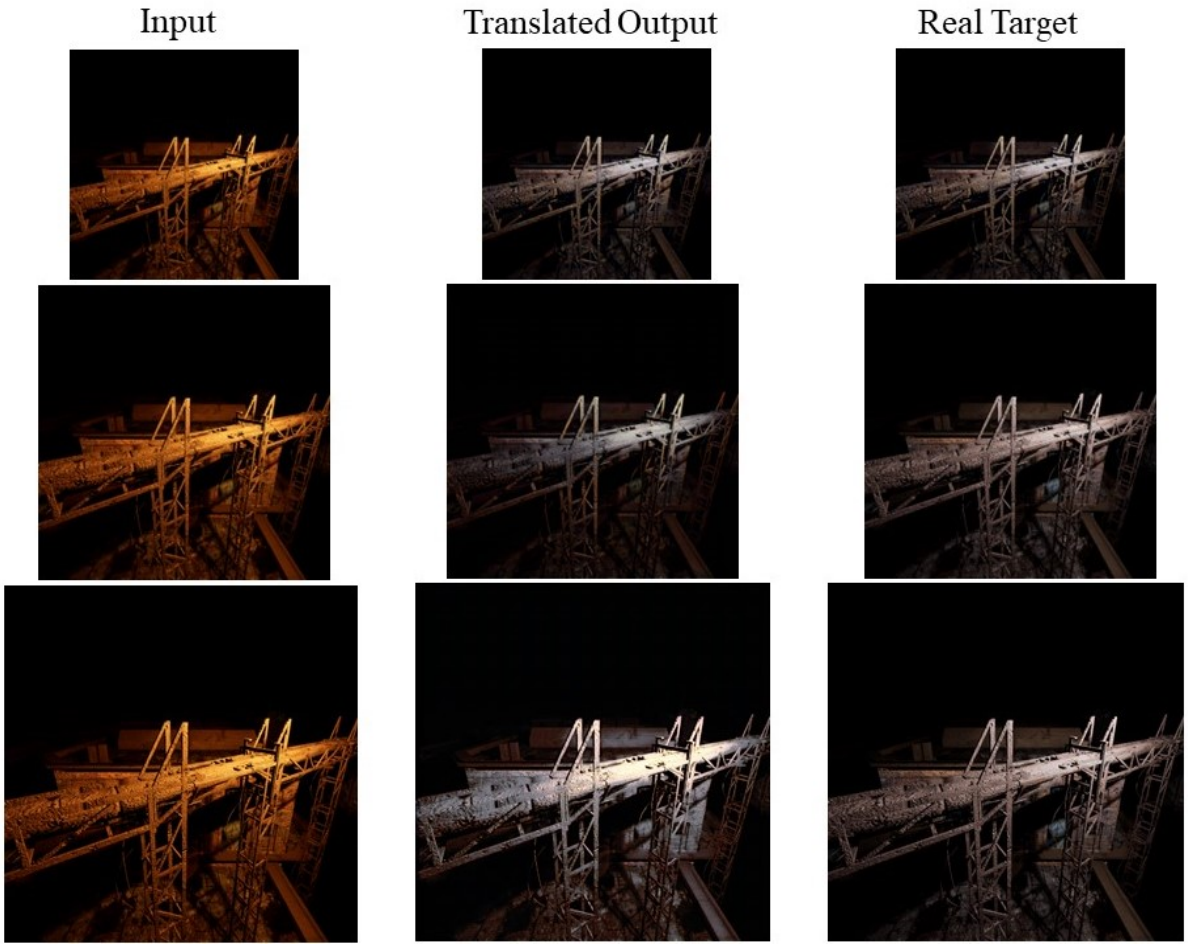

Figure 12: Translation results between 2500K and 6500K color temperatures for VIDIT dataset generated at 256×256, 512×512 and 1024×1024 resolutions.

[22] Andrew Janowczyk, Ajay Basavanhally, and Anant Madabhushi. 2017. Stain normalization using sparse autoencoders (StaNoSA): application to digital pathology. *Computerized Medical Imaging and Graphics* 57 (2017), 50–61.

[23] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*. Springer, 694–711.

[24] Hongtao Kang, Die Luo, Weihua Feng, Shaoqun Zeng, Tingwei Quan, Junbo Hu, and Xiuli Liu. 2021. Stainnet: a fast and robust stain normalization network. *Frontiers in Medicine* 8 (2021).

[25] Jing Ke, Yiqing Shen, Xiaoyao Liang, and Dinggang Shen. 2021. Contrastive Learning Based Stain Normalization Across Multiple Tumor in Histopathology. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 571–580.

[26] Amal Lahiani, Nassir Navab, Shadi Albarqouni, and Eldad Klaiman. 2019. Perceptual embedding consistency for seamless reconstruction of tilewise style transfer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 568–576.

[27] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4681–4690.

[28] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. 2020. Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision* 128, 10 (2020), 2402–2417.

[29] Marc Macenko, Marc Niethammer, James S Marron, David Borland, John T Woosley, Xiaojun Guan, Charles Schmitt, and Nancy E Thomas. 2009. A method for normalizing histology slides for quantitative analysis. In *2009 IEEE international symposium on biomedical imaging: from nano to macro*. IEEE, 1107–1110.

[30] Lucie Charlotte Magister and Ognjen Arandjelović. 2021. Generative Image Inpainting for Retinal Images using Generative Adversarial Networks. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2835–2838.

[31] Dwarikanath Mahapatra, Behzad Bozorgtabar, Jean-Philippe Thiran, and Ling Shao. 2020. Structure preserving stain normalization of histopathology images using self supervised semantic guidance. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 309–319.

[32] Rossana CN Melo, Maximilian WD Raas, Cinthia Palazzi, Vitor H Neves, Kássia K Malta, and Thiago P Silva. 2020. Whole slide imaging and its applications to histopathological studies of liver disorders. *Frontiers in medicine* (2020), 310.

[33] Haseeb Nazki, Sook Yoon, Alvaro Fuentes, and Dong Sun Park. 2020. Unsupervised image translation using adversarial networks for improved plant disease recognition. *Computers and Electronics in Agriculture* 168 (2020), 105117.

[34] Erik Reinhard, Michael Adhikhmin, Bruce Gooch, and Peter Shirley. 2001. Color transfer between images. *IEEE Computer graphics and applications* 21, 5 (2001), 34–41.

[35] Rafael Reisenhofer, Sebastian Bosse, Gitta Kutyniok, and Thomas Wiegand. 2018. A Haar wavelet-based perceptual similarity index for image quality assessment. *Signal Processing: Image Communication* 61 (2018), 33–43.

[36] Marlen Runz, Daniel Rusche, Stefan Schmidt, Martin R Weihrauch, Jürgen Hesser, and Cleo-Aron Weis. 2021. Normalization of HE-stained histological images using cycle consistent generative adversarial networks. *Diagnostic Pathology* 16, 1 (2021), 1–10.

[37] M Tarek Shaban, Christoph Baur, Nassir Navab, and Shadi Albarqouni. 2019. Staingan: Stain style transfer for digital histological images. In *2019 Ieee 16th international symposium on biomedical imaging (Isbi 2019)*. IEEE, 953–956.

[38] David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen Van Der Laak. 2019. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical image analysis* 58 (2019), 101544.

[39] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022* (2016).

[40] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).

[41] Laurens JP Van der Maaten and Eric O Postma. 2010. Texton-based analysis of paintings. In *Applications of Digital Image Processing XXXIII*, Vol. 7798. SPIE, 107–120.

[42] Jelica Vasiljević, Friedrich Feuerhake, Cédric Wemmert, and Thomas Lampert. 2021. Towards histopathological stain invariance by unsupervised domain augmentation using generative adversarial networks. *Neurocomputing* 460 (2021), 277–291.

[43] Sophia J Wagner, Nadieh Khalili, Raghav Sharma, Melanie Boxberg, Carsten Marr, Walter de Back, and Tingying Peng. 2021. Structure-preserving multi-domain stain color augmentation using style-transfer with disentangled representations. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 257–266.

[44] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. 2003. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, Vol. 2. Ieee, 1398–1402.

[45] JD Webster and RW Dunstan. 2014. Whole-slide imaging and automated image analysis: considerations and opportunities in the practice of pathology. *Veterinary pathology* 51, 1 (2014), 211–223.

[46] Xingzhi Yue, Neofytos Dimitriou, and Ognjen Arandjelovic. 2019. Colorectal cancer outcome prediction from H&E whole slide images using machine learning and automatically inferred phenotype profiles. *In Proc. International Conference on Bioinformatics and Computational Biology* (2019), 139–149.

[47] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.