

# The Spatially-Correlative Loss for Various Image Translation Tasks

Chuanxia Zheng    Tat-Jen Cham  
School of Computer Science and Engineering  
Nanyang Technological University, Singapore  
{chuanxia001, astjcham}@ntu.edu.sg

Jianfei Cai  
Department of Data Science & AI  
Monash University, Australia  
Jianfei.Cai@monash.edu

## Abstract

We propose a novel spatially-correlative loss that is simple, efficient and yet effective for preserving scene structure consistency while supporting large appearance changes during unpaired image-to-image (I2I) translation. Previous methods attempt this by using pixel-level cycle-consistency or feature-level matching losses, but the domain-specific nature of these losses hinder translation across large domain gaps. To address this, we exploit the spatial patterns of self-similarity as a means of defining scene structure. Our spatially-correlative loss is geared towards only capturing spatial relationships within an image rather than domain appearance. We also introduce a new self-supervised learning method to explicitly learn spatially-correlative maps for each specific translation task. We show distinct improvement over baseline models in all three modes of unpaired I2I translation: single-modal, multi-modal, and even single-image translation. This new loss can easily be integrated into existing network architectures and thus allows wide applicability. The code is available at <https://github.com/lyndonzheng/F-LSeSim>.

## 1. Introduction

I2I translation refers to the task of modifying an input image to fit the *style / appearance* of the target domain, while preserving the original *content / structure* (as shown in Fig. 1: *horse*  $\rightarrow$  *zebra*); learning to assess the *content* and *style* correctly is thus of central importance. While GANs [15] have the ability to generate images that adhere to the overall dataset distribution, it is still difficult to preserve scene structure during translation when image-conditional GANs are optimized with purely adversarial loss.

To mitigate the issue of scene structure discrepancies, a few loss functions for comparing the content between input and output images have been proposed, including (a) *pixel-level* image reconstruction loss [23, 46, 6] and cycle-consistency loss [29, 56, 52]; (b) *feature-level* perceptual loss [12, 25] and PatchNCE loss [39]. However,

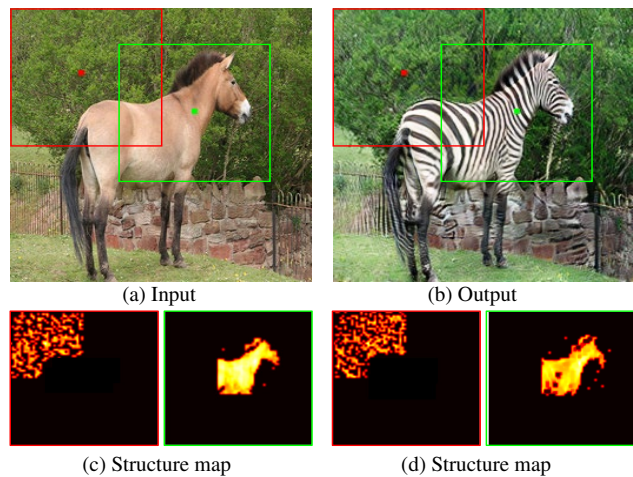


Figure 1. **Our learned spatially-correlative representation** encodes local scene structure based on self-similarities. Despite vast appearance differences between the *horse* and *zebra*, when the scene structures are identical (*i.e.* same poses), the spatial patterns of self-similarities are as well.

these losses still have several limitations. First, pixel-level losses do *not* explicitly decouple structure and appearance. Second, feature-level losses help but continue to conflate domain-specific structure and appearance attributes. Finally, most feature-level losses are calculated using a fixed ImageNet [11] pre-trained network (*e.g.* VGG16 [47]), which will not correctly adapt to arbitrary domains.

In this work, we aim to design a *domain-invariant* representation to precisely express scene structure, rather than using original pixels or features that couple both appearance and structure. To achieve this, we propose to revisit the idea of *self-similarity*. Classically, low-level self-similarity has been used for matching [43] and image segmentation [44], while feature-level self-similarity in deep learning manifests as self-attention maps [51]. We propose to go further, to advance an assumption that *all* regions within same categories exhibit some form of self-similarity. For instance, while the horse and zebra in Fig. 1 appear very different, there is obvious visual self-similarity in their own re-




Figure 2. **Comparison of unpaired I2I translation methods with various content losses.** (a) The cycle-consistency loss [29, 56, 52] in a two-sided framework. (b) Pixel-level image reconstruction loss [46] and feature-level matching loss [25]. (c) Various indirect relationships [5, 13] between the input and output. (d) Our spatially-correlative loss based on a learned spatially-correlative map.

gions. We believe a network can learn deeper representations of self-similarities (beyond just visual ones) that can encode intact object shapes, even when there are variations in appearances within an object. Then through estimating such co-occurrence signals in self-similarity, we can explicitly represent the structure as multiple *spatially-correlative* maps, visualized as heat maps in Fig.1 (c) and (d). Based on this within-shape self-similarity, we propose then that *a structure-preserving image translation will retain the patterns of self-similarity in both the source and translated images, even if appearances themselves change dramatically.*

Our basic spatially-correlative map, called *FSeSim*, is obtained by computing the **Fixed Self-Similarity** of features extracted from a pre-trained network. While this basic version achieved comparative or even better results than state-of-the-art methods [56, 13, 39] on some tasks, the generality is limited because features extracted from an ImageNet pre-trained network are biased towards photorealistic imagery.

To obtain a more general spatially-correlative map, the **Learned Self-Similarity**, called *LSeSim*, is presented by using a form of contrastive loss, in which we explicitly encourage homologous structures to be closer, regardless of their appearances, and reciprocally dissociate dissimilar structures even they have similar appearances. To do this, the model learns a domain-invariant spatially-correlative map, where having the same scene structure leads to similar maps, even if the images are from different domains.

There are several advantages of using the proposed F/LSeSim loss: (a) In contrast to the existing losses that directly compare the loss on pixels [56] or features [25], F/LSeSim captures the domain-invariant structure representation, regardless of the absolute pixel values; (b) Through contrastive learning, the LSeSim learns a spatially-correlative map for a specific image translation task, rather than features extracted from a fixed pre-trained network, as in *e.g.* perceptual loss [25], contextual loss [34]; (c) The translation model is more efficient and faster than the widely used cycle-consistency architectures, because our F/LSeSim explicitly encodes the structure, bypassing the expensive multi-cycle looping; (d) As we show in Fig. 5, our F/LSeSim correctly measures the structural distance

even when the two images are in completely different domains; (e) Finally, our F/LSeSim can easily be integrated into various frameworks. In our experiments, we directly used the generator and discriminator architectures of CycleGAN [56], MUNIT [22] and StyleGAN [26, 27] for extensive I2I translation tasks. The experimental results show that our model outperformed the existing both one-sided translation methods [5, 2, 13, 39] and two-sided translation methods [56, 52, 22].

## 2. Related Work

Existing unpaired I2I translation either use cycle-consistency loss in a two-sided framework [29, 56, 52], or other forms of pixel-level and feature-level losses in a one-sided framework [5, 2, 13] for preserving content (Fig. 2).

**Two-Sided Unsupervised Image Translation.** *Cycle-consistency* has become a de facto loss in most works, whether the cycles occur in the image domain [56, 29, 52, 9, 20, 31], or in latent space [57, 22, 32]. However, without explicit constraints, the content in a translated image can be easily distorted [31]. Furthermore, the cycle-based methods require auxiliary generators and discriminators for the reverse mapping.

**One-Sided Unsupervised Image Translation.** To avoid cycle-consistency artifacts, DistanceGAN [5] and GcGAN [13] pre-define an implicit distance in a one-sided framework. In contrast, the feature-level losses [25, 34] evaluate the content distance in a deep feature space, which have been applied in both style transfer [14, 25, 34, 55] and image translation [6, 49, 40, 24]. However, the underlying assumption that high-level semantic information is solely determined in feature space does not always hold. Furthermore, these features are extracted from a fixed pre-trained network (*e.g.* VGG16 [47]). While the latest CUT [39] learns a PatchNCE loss for a specific task, the distance used is directly computed from extracted features, and will still be affected by domain-specific peculiarities.

**Contrastive Representation Learning.** Driven by the potential of discriminative thought, a series of self-

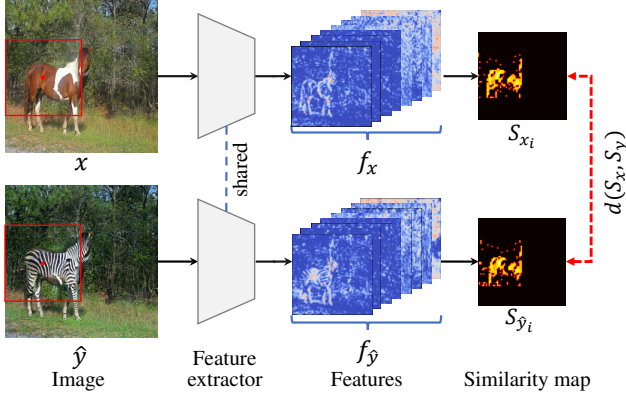


Figure 3. **An example of computing spatially-correlative loss from self-similarity maps.** Image  $x$  and corresponding translated image  $\hat{y}$  are first fed into the feature extractor. We then compute the local self-similarity for each query point. Here, we show one example for the red query point.

supervised methods [19, 50, 38, 3, 17, 7, 16] have emerged in recent years. These self-supervised methods learn robust features by associating “positive” pairs and dissociating “negative” pairs. CUT [39] first introduced contrastive learning for unsupervised I2I translation. While we utilize a patch-wise contrastive loss within an image in a similar manner to CUT, we propose a systematic way to learn a structure map that excludes appearance attributes.

### 3. Methods

As shown in Fig. 2, given a collection of images  $\mathcal{X} \subset \mathbb{R}^{H \times W \times C}$  from a particular domain (e.g. horse), our main goal is to learn a model  $\Phi$  that receives the image  $x \in \mathcal{X}$  as input and transfers it into the target domain  $\mathcal{Y} \subset \mathbb{R}^{H \times W \times C}$  (e.g. zebra), in a manner that retains the original scene structure but converts the appearance appropriately.

Here, we focus on designing a loss function that measures the structural similarity between the input image  $x$  and the translated image  $\hat{y} = \Phi(x)$ . However, unlike most existing approaches that directly attempt to evaluate the structural similarity between input and translated images at some deep feature level, we will instead compute the *self-similarity* of deep features *within each image*, and then *compare the self-similarity patterns* between the images.

In subsequent sections, we investigate two losses, *fixed self-similarity (FSeSim)* and *learned self-similarity (LSeSim)*. In the first instance, we directly compare the self-similarity patterns of features extracted from a fixed pre-trained network (e.g. VGG16 [47]). In the second instance, we additionally introduce a structure representation model that learns to correctly compare the self-similarity patterns, in which we use the contrastive infoNCE loss [38] to learn such a network without label supervision.

#### 3.1. Fixed Self-Similarity (FSeSim)

We first describe our fixed spatially-correlative loss. Given an image  $x$  in one domain and its corresponding translated image  $\hat{y}$  in another, we extract the features  $f_x$  and  $f_{\hat{y}}$  using a simple network (e.g. VGG16 [47]). Instead of directly computing the feature distance  $\|f_x - f_{\hat{y}}\|_p$ , we compute the self-similarity in the form of a map. We call this a *spatially-correlative map*, formally:

$$S_{x_i} = (f_{x_i})^T (f_{x_*}) \quad (1)$$

where  $f_{x_i}^T \in \mathbb{R}^{1 \times C}$  is the feature of a query point  $x_i$  with  $C$  channels,  $f_{x_*} \in \mathbb{R}^{C \times N_p}$  contains corresponding features in a patch of  $N_p$  points, and  $S_{x_i} \in \mathbb{R}^{1 \times N_p}$  captures the feature spatial correlation between the query point and other points in the patch. We show one query example in Fig. 3, where the spatially-correlative map for the query patch is visualized as a heat map. Note that unlike the original features that would still encode domain-specific attributes such as color, lighting and texture, the self-similarity map only captures the spatially-correlative relationships.

Next we represent the structure of the whole image as a collection of multiple spatially-correlative maps  $S_x = [S_{x_1}; S_{x_2}; \dots; S_{x_s}] \in \mathbb{R}^{N_s \times N_p}$ , where  $N_s$  is the numbers of sampled patches. This is a semi-sparse representation, but is more computationally efficient. We then compare the multiple structure similarity maps between the input  $x$  and the translated image  $\hat{y}$ , as follows:

$$\mathcal{L}_s = d(S_x, S_{\hat{y}}) \quad (2)$$

where  $S_{\hat{y}}$  are corresponding spatially-correlative maps in the target domain. Here, we consider two forms for  $d(\cdot)$ , the  $L_1$  distance  $\|S_x - S_{\hat{y}}\|_1$  and the cosine distance  $\|1 - \cos(S_x, S_{\hat{y}})\|$ . The former term strongly encourages the spatial similarity to be consistent at all points in a patch, while the latter term supports pattern correlation without concern for differences in magnitude.

#### 3.2. Learned Self-Similarity (LSeSim)

Although our FSeSim provides strong supervision for structure consistency, it does *not* explicitly learn a structure representation for a specific translation task. As opposed to existing feature-level losses [25, 34] that only utilize the features from a fixed pre-trained network, we propose to additionally learn a structure representation network for each task that expresses the learned self-similarity, or LSeSim.

In order to learn such a model *without supervision*, we consider the self-supervised contrastive learning that associates similar features, while simultaneously dissociates different features. Following PatchNCE [39], we build our contrastive loss at patch level, except here the pairs for comparison are our spatially-correlative maps, rather than the original features in existing works [19, 7, 16, 39].

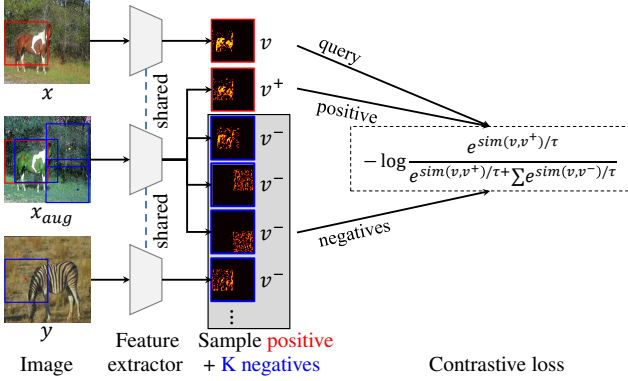


Figure 4. **Patchwise contrastive learning for the learned self-similarity.** Three images are fed into the feature extractor, in which two images,  $x$  and  $x_{aug}$ , are homologous with the same structure but varied appearances, and  $y$  is another randomly sampled image. For each query patch in  $x$ , the “positive” sample is the corresponding patch in  $x_{aug}$ , and all other patches are considered as “negative” samples.

To help generate pairs of similar patch features for self-supervised learning, we create augmented images by applying structure-preserving transformations.

Formally, let  $\mathbf{v} = S_{x_i} \in \mathbb{R}^{1 \times N_p}$  denotes the spatially-correlative map of the “query” patch. Let  $\mathbf{v}^+ = S_{\hat{x}_i} \in \mathbb{R}^{1 \times N_p}$  and  $\mathbf{v}^- \in \mathbb{R}^{K \times N_p}$  be “positive” and “negative” patch samples, respectively. The query patch is positively paired with a patch in the same position  $i$  within an augmented image  $x_{aug}$ , and negatively paired to patches sampled from other positions in  $x_{aug}$ , or patches from other images  $y$ . The number of negative patches used is  $K = 255$ .

Our LSeSim design is illustrated in Fig. 4. The contrastive loss is given by:

$$\mathcal{L}_c = -\log \frac{e^{sim(\mathbf{v}, \mathbf{v}^+) / \tau}}{e^{sim(\mathbf{v}, \mathbf{v}^+) / \tau} + \sum_{k=1}^K e^{sim(\mathbf{v}, \mathbf{v}_k^-) / \tau}} \quad (3)$$

where  $sim(\mathbf{v}, \mathbf{v}^+) = \mathbf{v}^T \mathbf{v}^+ / \|\mathbf{v}\| \|\mathbf{v}^+\|$  is the cosine similarity between two spatially-correlative maps, and  $\tau$  is a temperature parameter. To minimize this loss, our network encourages the corresponding patches with the same structure to be close even they have very different visual appearances, which fits in with the goal of image translation. Note that, this contrastive loss is only used for optimizing the structure representation network. The spatially-correlative loss for the generator is always the loss in eq. (2).

### 3.3. Full Objective

Overall, we train the networks by jointly minimizing the following losses:

$$\begin{aligned} \mathcal{L}_D &= -\mathbb{E}_{y \sim p_d} [\log D(y)] - \mathbb{E}_{\hat{y} \sim p_g} [\log(1 - D(\hat{y}))] \\ \mathcal{L}_S &= \mathcal{L}_c \\ \mathcal{L}_G &= \mathbb{E}_{\hat{y} \sim p_g} [\log(1 - D(\hat{y}))] + \lambda d(S_x, S_{\hat{y}}) \end{aligned} \quad (4)$$




Figure 5. **Error map visualization.** Our LSeSim has small errors on the left where ground truth paired data is provided, while having large errors on the right for unpaired data.

where  $\mathcal{L}_D$  is the adversarial loss for the discriminator  $D(\cdot)$ ,  $\hat{y}$  is the translated image, and  $\mathcal{L}_S$  is the contrastive loss for the structure representation network  $f(\cdot)$ .  $\mathcal{L}_G$  is the loss for the generation (translation) network  $G(\cdot)$ , which consists of the style loss term and the structure loss term.  $\lambda$  is a hyper-parameter to trade off between style and content.

### 3.4. Analysis

Readers may wonder why the proposed F/LSeSim losses would perform better than existing feature-level losses [25, 34, 39]. An intuitive interpretation is that *self-similarity deals only with spatial relationships of co-occurring signals, rather than their original absolute values.*

To provide further clarity, we consider a scenario where given a semantic map  $x$  (Fig. 5), the task is to translate it to a photorealistic image  $y$ . We consider an ideal result (the paired ground truth  $y_{align}$ ) and a wrong result (another image  $y_{unalign}$ ), respectively. Under such a setting, a good structure loss should penalize the wrong result, while supporting the ideal result. To visualize the error maps, for each corresponding pair of query patches in  $x$  and  $y$  we computed the error at that patch location for different losses. As can be seen, pixel-level loss [46] is naturally unsuitable when there are large domain gaps, and while Perceptual loss [25] will report significant errors for both aligned and unaligned results. PatchNCE [39] mitigates the problem by calculating the cosine distance of features, but it can be seen the loss map still retains high errors in many regions within the aligned result, due to extracted features consisting of appearance attributes, such as color and texture.

In contrast, appearance attributes are ignored in LSeSim by representing scene structure as a spatially-correlative map. Fig. 5 shows that our LSeSim leads to low errors for the aligned image (left), even when they are in quite different domains, but large errors for the non-aligned image (right). Even for  $y_{unalign}$ , LSeSim differentiates between related structures (e.g. roads) and unrelated structures (trees vs windows), with lower errors for the former. Hence LSeSim can better help preserve scene structure even across large domain gaps.



Figure 6. **Comparing results under different content losses.** All results are reported following the same setting of CycleGAN [56], except using different content losses. Our model generates much better visual results with only loss modification.

In Fig. 6, we report a qualitative comparison of various losses that be applied to a same translation network architecture. All methods following the setting in CycleGAN [57], except that the content loss is changed. Cycle-consistency is achieved using the auxiliary generator and discriminator, and all other methods are one-sided translation. We find that our method produces results with much better visual quality.

**Discussion.** Similar to conventional feature-level losses [25, 34], our F/LSeSim is computed in a deep feature space. However, we represent the structure as multiple spatially-correlative maps. So rather than directly at feature level which is not free from domain-specific attributes, our comparison is done at a more abstract level that is intended to transcend domain specificity.

While attention maps have been used in previous image translation works [8, 1], it is fundamentally different from our F/LSeSim in concept — their attention maps effectively function as saliency maps to guide the translation, but content preservation is primarily still dependent on cycle-consistency loss. In our case, the multiple spatially-correlative maps are used to encode and determine invariance in scene structure. Our F/LSeSim also differs from the content loss used in [30], in which the self-similarity was calculated at random positions without a clear purpose. Our F/LSeSim is on the other hand organized at a local patch level to explicitly represent the scene structure. As shown in § 4.2, our local structure representation is better than just using random spatial relationships. Furthermore, our LSeSim is a metric learned from the infoNCE loss, which generalizes well robustly on various tasks. While PatchNCE loss [39] can also learn feature similarity using contrastive loss, it directly compares features in two domains.

## 4. Experiments

To demonstrate the generality of our method, we instantiated F/LSeSim in multiple frameworks on various I2I translation tasks, including *single-modal*, *multi-modal*, and even *single-image* translation. For each task, we used a suitable baseline architecture, but replaced their content losses with our F/LSeSim loss. In addition, we are only interested in scenarios where scene structure is preserved during the

Method	Cityscapes		Horse→Zebra	
	pixAcc↑	FID↓	FID↓	Mem↓
CycleGAN [56]	57.2	76.3	77.2	4.81
MUNIT [22]	58.4	91.4	98.0	9.43
DRIT++ [32]	60.3	96.2	88.5	11.2
Distance [5]	47.2	75.9	67.2	2.72
GcGAN [13]	65.5	57.4	86.7	4.68
CUT [39]	68.8	56.4	45.5	3.33
FSeSim	69.4	53.6	40.4	<b>2.65</b>
LSeSim	<b>73.2</b>	<b>49.7</b>	<b>38.0</b>	2.92

Table 1. **Quantitative comparison on single-modal image translation.** FID [18] measures the distance between distributions of generated images and real images. “Mem” denotes the memory cost during training.

translation [23, 56, 57], rather investigating translations incorporating shape modification [9, 10, 37, 28, 4].

### 4.1. Single-Modal Unpaired Image Translation

We first evaluated our loss on the classical single-modal unpaired I2I translation task.

**Implementation details.** In this task, we chose CycleGAN [56] as the reference architecture, but only used half of their pipeline and replaced the cycle-consistency loss with our F/LSeSim loss. Specifically, we used the Resnet-based generator with PatchGAN discriminator [23]. Details can be found on their [website](#).

Our FSeSim is based on the ImageNet-pretrained VGG16 [47], where we used features from layers `relu3_1` and `relu4_1`. While the LSeSim employs the same structure as FSeSim, the weights are not fixed and additionally two convolution layers, implemented as  $1 \times 1$  kernels, are included to select better features. As for the selection of patches to build the contrastive loss, we found that random sampling the patch locations performed much better than uniform sampling on a grid, leading to better convergence when training the structure representation network. We set  $\lambda = 10$  in FSeSim and  $\tau = 0.07$  in LSeSim.

**Metrics.** Our evaluation protocols are adopted from previous work [18, 40, 39]. We first used the popular Fréchet




Figure 7. **Qualitative comparison on single-modal image translation.** Here, we show results for *horse*→*zebra* and *label*→*image*.

Inception Distance (FID) [18]<sup>1</sup> to assess the visual quality of generated images by comparing the distance between distributions of generated and real images in a deep feature domain. For *semantic image synthesis*, we further applied semantic segmentation to the generated images to estimate how well the predicted masks match the ground truth segmentation masks as in [6, 49, 40, 39]. Following [39, 24], we used the pre-trained DRN [54].

**Results.** In Table 1, we reported either published results or our reproductions with publicly-available code, choosing the better. Our simple, inexpensive losses substantially outperformed state-of-the-art methods, including two-sided frameworks with multiple cycle-consistency losses [57, 22, 32], and one-sided frameworks using self-distance [5], geometry consistency [13] and contrastive loss [39].

When compared to CycleGAN [56] and CUT [39], although we used the same settings for the generator and discriminator, our method led to significant improvement. Unlike CUT [39] that depends on an identity pass for good performance, our results were achieved by training with only one pass using F/LSeSim and GAN losses. This suggests that once we explicitly decouple scene structure and appearance, it is easier for the model to modify the visual appearance correctly. As our model belongs to one-sided image translation that does *not* require additional generators and discriminators, our model is also memory-efficient.

Qualitative results are shown in Figs. 6 and 7. In Fig. 6, despite keeping the same settings and only comparing different content losses, our method translated the zebra appearance more cleanly. We also compared results using the same examples as [39] in Fig. 7, where our method achieved better visual results.

<sup>1</sup>As claimed in StyleGANv2 [27], ImageNet-pretrained classifiers tend to evaluate the distribution on texture than shape, while humans focus on shape. The best FID score does *not* ensure the best image quality for translated images. As such, for a fair comparison, we reported the best FID score from all trained epochs for all methods, rather than the score in the last epoch.

## 4.2. Multi-Modal Unpaired Image Translation

Our F/LSeSim is also naturally suited for multi-modal image translation, since the use of our spatial-correlative maps imposes only structural consistency and not appearance constraints. We performed a thorough comparison of F/LSeSim to state-of-the-art methods, along with comprehensive ablation experiments.

**Implementation details.** Our multi-modal setting is based on MUNIT [22, 32], except our model uses only one generator and one discriminator of MUNIT [22] without requiring the auxiliary generators and discriminators for multiple cycle training. Specifically, we used the Resnet-based generator with Instance Normalization (IN) [48] in the encoder and Adaptive Instance Normalization (AdaIN) [21, 26] in the decoder, plus multi-scale discriminators [49]. The details of the architecture can be found on their [website](#). The F/LSeSim used is identical to that used in § 4.1.

**Metrics.** Besides using FID to measure quality, we also used the average LPIPS distance [55] to evaluate the diversity of generated results. The LPIPS distance is calculated by comparing the features of two images. Following [22, 57], we computed the distances between 1900 pairs, sampling 100 images 19 times. We also report the latest metrics of Density and Coverage (D&C) [36], which separately evaluate the diversity and fidelity of generated results. Likewise, we used the 1900 sampled pairs to compute D&C scores. Higher scores here indicate larger diversity and better coverage to the ground-truth domain, respectively.

**Results.** We compared our F/LSeSim to state-of-the-art methods in multi-modal image translation in Table 2. Our method outperformed the two baselines, MUNIT [22] and DRIT++ [32], although we deployed the same network architecture. BicycleGAN [57] achieved larger diversity on all tasks by adding noise to all decoders through the U\_Net [41], but the tradeoffs are worse visual results, due to



Figure 8. **Qualitative comparison on multi-modal image translation.** Here, we show the examples of *winter*→*summer* and *night*→*day*.

Method	Winter→Summer			Night→Day		
	LPIPS ↑	FID ↓	D & C ↑	LPIPS ↑	FID ↓	D & C ↑
Real images	0.770	44.1	0.997 / 0.986	0.684	146.1	0.977 / 0.962
BicycleGAN [57]	<b>0.285</b>	99.2 ± 3.2	<b>0.536</b> / 0.667	<b>0.349</b>	290.9 ± 6.5	<b>0.375</b> / 0.515
MUNIT [22]	0.160	97.4 ± 2.2	0.439 / 0.707	0.152	267.1 ± 2.7	0.271 / 0.548
DRIT++ [32]	0.186	93.1 ± 2.0	0.494 / 0.753	0.167	258.5 ± 2.3	0.298 / 0.631
<b>FSeSim</b>	0.216	90.5 ± 1.9	0.501 / 0.779	0.203	234.3 ± 2.8	0.332 / 0.638
<b>LSeSim</b>	0.232	<b>89.4 ± 1.9</b>	0.516 / <b>0.793</b>	0.215	<b>224.9 ± 2.0</b>	0.347 / <b>0.652</b>

Table 2. **Quantitative evaluation on multi-modal image translation task.** LPIPS distance [55] measures the diversity of generated images by comparing the features of two images, while (D&C) [36] evaluates the diversity and fidelity by matching whole features in the generated and real datasets.

Configuration		Horse → Zebra		Night → Day		
		FID ↓	Mem(GB) ↓	FID ↓	LPIPS ↑	D & C ↑
A	STROTSS [30] (random SeSim)	70.1	2.68	262.7 ± 3.6	0.162	0.289 / 0.554
B	Baseline (global SeSim on single layer)	53.7	2.97	<b>173.2 ± 2.2</b>	0.168	0.303 / <b>0.664</b>
C	(B): Global → Patch (32 × 32)	45.8	<b>2.61</b>	231.3 ± 2.5	0.181	0.317 / 0.634
D	(C): Single → Multi (relu3_1, relu4_1)	43.3	2.65	229.4 ± 2.1	0.177	0.311 / 0.646
E	(D): l1loss → 1 - cosine	40.4	2.65	234.3 ± 2.8	0.203	0.332 / 0.638
F	Ours LSeSim	<b>38.0</b>	2.92	224.9 ± 2.0	<b>0.215</b>	<b>0.347</b> / 0.652

Table 3. **Ablation study on both single- and multi-modal image translation.** Refer to ablation experiments in main text for details.

the larger noise being directly added to the last generative layer. In contrast, following in MUNIT [22], we only added noise to the middle layers of generation, through AdaIN.

In Fig. 8, we show qualitative comparisons of our method to MUNIT [22] on *winter* → *summer*, and *night* → *day* tasks. As can be seen, our model not only generated higher quality translated results, but also produced more diverse solutions for these multi-modal tasks. We believe this is because the formulation of our F/LSeSim will only maintain structural fidelity, and does not impose penalties on appropriate appearance modifications in the target domain.

**Ablation Experiments.** To understand the influence of different components for the proposed spatially-correlative

loss, we ran a number of ablations. The quantitative results are reported in Table 3 for both single- and multi-modal image translation. In this table, **row A** shows the performance of the baseline method [30] which utilizes self-similarity as content loss. However, it calculates the similarity using random sampled features and does not have an explicit connection to spatial structure. **Row B** is a global attention map. While this version performed well and ran faster by avoiding sampling, it has two main limitations. First, the original global attention module is memory intensive and *cannot* be applied to multiple scales nor to large feature spaces. Second, as evident from Fig. 10, the spatially distant correlation is essentially noise (as is also the case for the Random baseline of [30]), which is detrimental to the results. Compared




Figure 9. **High-resolution painting to photorealistic image** on single-image translation. Zoom in to see the details.




Figure 10. **Ablation study on self-similarity maps**. A, B, D and F correspond to the settings in Table 3, respectively.

to the global version, **row C** largely improved the performance as clearer shapes are captured in the local patches. In **row D**, we applied local attention to multiple layers with a fixed path size. This results in the spatially-correlative maps having different receptive fields, which further improves the performance. **Row E** replaces the  $l_1$  distance with cosine distance. While the improvement in image quality is not obvious, the diversity scores increased substantially. This is due to the cosine similarity supporting only the correlation between the two spatially-correlative maps without encouraging the maps to be fully same. **Row E** shows the performance of the full model (same as in Tables 1 and 2), where LSeSim of **row F** improved on many metrics, and had better visual results.

### 4.3. Single-Image Unpaired Image Translation

To further test the generalization ability, we applied the FSeSim to a high-resolution single-image translation task. Here, only one source and one target image are provided for training, but they are unpaired. This task is conceptually similar to the style transfer [14, 25, 33], except that here we trained a SinGAN-like [42, 45] model that captures the distribution of a single image through the adversarial learning, rather than using a fixed style loss [25].

**Implementation details.** The single-image translation setting is based on the latest CUT [39] method, except that the PatchNCE loss is replaced by our FSeSim loss. In de-

tail, the StyleGAN2-based generator and discriminator [27] with the gradient penalty [35, 26] are used. To further increase simplicity, we removed the identity loss in CUT [39], and only used a GAN loss in conjunction with the proposed FSeSim loss to assess the appearance and structure separately. As  $64 \times 64$  crops have to be taken from a high-resolution image for training here, it becomes less useful to further subsample “positive” and “negative” patches. Therefore, we only use our FSeSim to train the model, without using LSeSim with contrastive loss.

**Results.** In Fig. 9, we show qualitative results from the CUT [39] paper on the *painting*→*photo* task. As evident in the highlighted regions, our model generated not only higher quality results, but they were also closer to the target image style than existing methods, including classical style transfer models, such as WCT<sup>2</sup> [53] and STRORSS [30], as well as the latest single-image translation CUT [39] model.

## 5. Conclusion

We introduced F/LSeSim, a new structure consistency loss that focuses only on spatially-correlative relationships, without regard to visual appearances. Our F/LSeSim is naturally suitable for tasks that require structure consistency, and can be easily applied to existing architectures. We demonstrated its generality to various unpaired I2I translation tasks, where a simple replacement of the existing content losses with F/LSeSim led to solid performance improvements.

**Acknowledgements** This research was supported by the National Research Foundation, Singapore under its International Research Centres in Singapore Funding Initiative. This research was also supported by the Monash FIT Start-up Grant. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.



## References

- [1] Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, and Kwang In Kim. Unsupervised attention-guided image-to-image translation. *Advances in Neural Information Processing Systems*, 31:3693–3703, 2018. [5](#)
- [2] Matthew Amodio and Smita Krishnaswamy. Travelgan: Image-to-image translation by transformation vector learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8983–8992, 2019. [2](#)
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15535–15545, 2019. [3](#)
- [4] Kyungiune Baek, Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Hyunjung Shim. Rethinking the truly unsupervised image-to-image translation. *arXiv preprint arXiv:2006.06500*, 2020. [5](#)
- [5] Sagie Benaim and Lior Wolf. One-sided unsupervised domain mapping. In *Advances in neural information processing systems*, pages 752–762, 2017. [2](#), [5](#), [6](#), [12](#)
- [6] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1511–1520, 2017. [1](#), [2](#), [6](#)
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. [3](#)
- [8] Xinyuan Chen, Chang Xu, Xiaokang Yang, and Dacheng Tao. Attention-gan for object transfiguration in wild images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 164–180, 2018. [5](#)
- [9] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018. [2](#), [5](#)
- [10] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020. [5](#)
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [1](#)
- [12] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in neural information processing systems*, pages 658–666, 2016. [1](#), [2](#)
- [13] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, Kun Zhang, and Dacheng Tao. Geometry-consistent generative adversarial networks for one-sided unsupervised domain mapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2427–2436, 2019. [2](#), [5](#), [6](#), [12](#)
- [14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. [2](#), [8](#)
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [1](#)
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. [3](#)
- [17] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. [3](#)
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017. [5](#), [6](#), [15](#)
- [19] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2018. [3](#)
- [20] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018. [2](#)
- [21] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. [6](#)
- [22] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018. [2](#), [5](#), [6](#), [7](#), [12](#), [14](#), [15](#)
- [23] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. [1](#), [5](#)
- [24] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. Tsit: A simple and versatile framework for image-to-image translation. In *ECCV*, 2020. [2](#), [6](#)
- [25] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In

- Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2, 6, 8
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 2, 6, 8
- [28] Sunnie S. Y. Kim, Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Deformable style transfer. In *European Conference on Computer Vision (ECCV)*, 2020. 5
- [29] Taeksoo Kim, Moonsoo Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017. 1, 2
- [30] Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Style transfer by relaxed optimal transport and self-similarity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10051–10060, 2019. 5, 7, 8
- [31] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51, 2018. 2
- [32] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision*, pages 1–16, 2020. 2, 5, 6, 7
- [33] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4990–4998, 2017. 8
- [34] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 768–783, 2018. 2, 3, 4, 5
- [35] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International Conference on Machine learning (ICML)*, pages 3481–3490. PMLR, 2018. 8
- [36] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *International Conference on Machine Learning*, 2020. 6, 7, 15
- [37] Ori Nizan and Ayellet Tal. Breaking the cycle-colleagues are all you need. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7860–7869, 2020. 5
- [38] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [39] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, 2020. 1, 2, 3, 4, 5, 6, 8, 12, 14, 15
- [40] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. 2, 5, 6, 15
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 6
- [42] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4570–4580, 2019. 8
- [43] Eli Shechtman and Michal Irani. Matching local self-similarities across images and videos. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. 1
- [44] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000. 1
- [45] Assaf Shocher, Shai Bagon, Phillip Isola, and Michal Irani. Ingan: Capturing and retargeting the "dna" of a natural image. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 8
- [46] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–2116, 2017. 1, 2, 4
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 2, 3, 5, 14
- [48] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017. 2, 6
- [49] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 2, 6
- [50] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018. 3
- [51] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015. 1
- [52] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *Proceedings of the IEEE international conference on computer vision*, pages 2849–2857, 2017. 1, 2

- [53] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9036–9045, 2019. [8](#)
- [54] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 472–480, 2017. [6](#)
- [55] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [2](#), [6](#), [7](#), [15](#)
- [56] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [1](#), [2](#), [5](#), [6](#), [12](#), [14](#)
- [57] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Advances in neural information processing systems*, pages 465–476, 2017. [2](#), [5](#), [6](#), [7](#), [15](#)