

StEP: Style-based Encoder Pre-training for Multi-modal Image Synthesis

Moustafa Meshry Yixuan Ren Larry S. Davis Abhinav Shrivastava
University of Maryland, College Park

Abstract

We propose a novel approach for multi-modal Image-to-image (I2I) translation. To tackle the one-to-many relationship between input and output domains, previous works use complex training objectives to learn a latent embedding, jointly with the generator, that models the variability of the output domain. In contrast, we directly model the style variability of images, independent of the image synthesis task. Specifically, we pre-train a generic style encoder using a novel proxy task to learn an embedding of images, from arbitrary domains, into a low-dimensional style latent space. The learned latent space introduces several advantages over previous traditional approaches to multi-modal I2I translation. First, it is not dependent on the target dataset, and generalizes well across multiple domains. Second, it learns a more powerful and expressive latent space, which improves the fidelity of style capture and transfer. The proposed style pre-training also simplifies the training objective and speeds up the training significantly. Furthermore, we provide a detailed study of the contribution of different loss terms to the task of multi-modal I2I translation, and propose a simple alternative to VAEs to enable sampling from unconstrained latent spaces. Finally, we achieve state-of-the-art results on six challenging benchmarks with a simple training objective that includes only a GAN loss and a reconstruction loss.

1. Introduction

Image-to-Image (I2I) translation is the task of transforming images from one domain to another (e.g., semantic maps \rightarrow scenes, sketches \rightarrow photo-realistic images, etc.). Many problems in computer vision and graphics can be cast as I2I translation, such as photo-realistic image synthesis [1–3], super-resolution [4], colorization [5, 6], and inpainting [7]. Therefore, I2I translation has recently received significant attention in the literature. One main challenge in I2I translation is the multi-modal nature for many such tasks – the relation between an input domain A and an output domain B is often times one-to-many, where a single input image $I_i^A \in A$ can be mapped to different output images from domain B . For example, a sketch of a shoe or a handbag can be mapped to corresponding objects with different colors or

styles, or a semantic map of a scene can be mapped to many scenes with different appearance, lighting and/or weather conditions. Since I2I translation networks typically learn one-to-one mappings due to their deterministic nature, an extra input is required to specify an output mode to which an input image will be translated. Simply injecting extra random noise as input proved to be ineffective as shown in [2, 8], where the generator network just learns to ignore the extra noise and collapses to a single or few modes (which is one form of the mode collapse problem). To overcome this problem, Zhu *et al.* [8] proposed *BicycleGAN*, which trains an encoder network E , jointly with the I2I translation network, to encode the distribution of different possible outputs into a latent vector z , and then learns a deterministic mapping $G : (A, z) \rightarrow B$. So, depending on the latent vector z , a single input $I_i^A \in A$ can be mapped to multiple outputs in B . While *BicycleGAN* requires paired training data, several works, like MUNIT [9] and DRIT [10], extended it to the unsupervised case, where images in domains A and B are not in correspondence (‘unpaired’). One main component of unpaired I2I is a cross-cycle consistency constraint, where the network generates an intermediate output by swapping the styles of a pair of images, then swaps the style between the intermediate output again to reconstruct the original images. This enforces that the latent vector z preserves the encoded style information when translated from an image i to another image j and back to image i again. This constraint can also be applied to paired training data, where it encourages style/attribute transfer between images. However, training *BicycleGAN* [8] or its unsupervised counterparts [9, 10] is not trivial. For example, *BicycleGAN* combines the objectives of both conditional Variational Auto-Encoders (cVAEs) [11] and a conditional version of Latent Regressor GANs (cLR-GANs) [12, 13] to train their network. The training objective of [9, 10] is even more involved to handle the unsupervised setup.

In this work, we propose a novel weakly-supervised pre-training strategy to learn an expressive latent space for the task of multi-modal I2I translation. While end-to-end training of the encoder network E with the I2I translation network poses a convenience, we show that it can be advantageous to break down the training into proxy tasks. In specific, we

show both quantitatively and qualitatively that the proposed pre-training yields the following **advantages**:

- It learns a more powerful and expressive latent space. Specifically, we show that: (1) Our pre-trained latent space captures uncommon styles that are not well represented in the training set, while baselines like BicycleGAN [8] and MUNIT [9] fail to do so and instead tend to simplify such styles/appearances to the nearest common style in the train set. (2) Pre-training yields more faithful style capture and transfer. (3) Finally, the better expressiveness of the pre-trained latent space leads to more complex style interpolations compared to the baselines.
- The learned style embedding is not dependent on the target dataset and generalizes well across many domains, which can be useful especially when having limited training data.
- Style pre-training simplifies the training objective by requiring fewer losses, which also speeds up the training.
- Our approach improves the training stability and the overall output quality and diversity.

We note that our proposed style pre-training is weakly-supervised and doesn't require any manual labeling. Instead, it relies on a pre-trained VGG network [14] to provide training supervision. Our approach is inspired by and extends the work of Meshry *et al.* [15] which utilizes a staged training strategy to re-render scenes under different lighting, time of day, and weather conditions. Our work is also inspired by the standard training paradigm in visual recognition of first pre-training on a proxy task, either large supervised datasets (e.g., ImageNet) [16–18] or unsupervised tasks (e.g., [19, 20]), and then finetuning (transfer learning) on the desired task. Similarly, we propose to pre-train the encoder using a proxy task that encourages capturing style into a latent space. Our goal is to highlight the benefits of encoder pre-training and demonstrate its effectiveness for multi-modal image synthesis. In particular, we make the following **contributions**:

- We propose to pre-train an encoder to learn a low-dimensional projection of Gram matrices (from Neural Style Transfer) and show that the pre-trained embedding is effective for multi-modal I2I translation, and that it simplifies and stabilizes the training.
- We show that the pre-trained latent embedding is not dependent on the target domain and generalizes well to other domains (transfer learning).
- We provide a study of the importance of different loss terms for multi-modal I2I translation network.
- We propose an alternative to enable sampling from a latent space instead of enforcing a prior as done in VAE training.
- We achieve state-of-the-art results on six benchmarks in terms of style capture and transfer, and diversity of results.

2. Related work

Deep generative models. There has been incredible progress in the field of image synthesis using deep neural

networks. In its unconditional setting, a decoder network learns to map random values drawn from a prior distribution (typically Gaussian) to output images. Variational Auto-Encoders (VAEs) [21] assume a bijection mapping between output images and some latent distribution and learn to map the latent distribution to a unit Gaussian using the reparameterization trick. Alternatively, Generative Adversarial Networks (GANs) [22] directly map random values sampled from a unit Gaussian to images, while using a discriminator network to enforce that the distribution of generated images resembles that of real images. Recent works proposed improvements to stabilize the training [23–26] and improve the quality and diversity of the output [27, 28]. Other works combine both VAEs and GANs into a hybrid VAE-GAN model [29, 30].

Conditional image synthesis. Instead of generating images from input noise, the generator can be augmented with side information in the form of extra conditional inputs. For example, Sohn *et al.* [11] extended VAEs to their conditional setup (cVAEs). Also, GANs can be conditioned on different information, like class labels [31–33], language description [34, 35], or an image from another domain [1, 2]. The latter is called Image-to-Image translation.

Image-to-Image (I2I) translation. I2I translation is the task of transforming an image from one domain, such as a sketch, into another domain, such as photo-realistic images. While there are regression-based approaches to this problem [1, 36], significant successes in this field are based on GANs and the influential work of pix2pix [2]. Following the success of pix2pix [2], I2I translation has since been utilized in a large number of tasks, like inpainting [7], colorization [5, 6], super-resolution [4], rendering [15, 37, 38], and many more [39–41]. There has also been works to extend this task to the unsupervised setting [36, 42–46], to multiple domains [47, 48], and to videos [49, 50].

Multi-modal I2I translation. Image translation networks are typically deterministic function approximators that learn a one-to-one mapping between inputs and outputs. To extend I2I translation to the case of diverse multi-modal outputs, Zhu *et al.* [8] proposed the BicycleGAN framework that learns a latent distribution that encodes the variability of the output domain and conditions the generator on this extra latent vector for multi-modal image synthesis. Wang *et al.* [3, 50] learn instance-wise latent features for different objects in a target image, which are clustered after training to find f fixed modes for different semantic classes. At test time, they sample one of the feature clusters for each object to achieve multi-modal synthesis. Other works extended the multi-modal I2I framework to the unpaired setup, where images from the input and output domains are not in correspondence [9, 10, 51], by augmenting BicycleGAN with different forms of a cross-cycle consistency constraint between two unpaired image pairs. In our work, we focus on

the supervised setting of multi-modal I2I translation. We propose a pre-training strategy to learn a latent distribution that encodes the variability of the output domain. The learned distribution can be easily adapted to new unseen datasets with simple finetuning, instead of training from scratch.

3. Approach

Current multi-modal image translation networks require an extra input z that allows for modelling the one-to-many relation between an input domain A and an output domain B as a one-to-one relation from a pair of inputs $(A, z) \rightarrow B$. In previous approaches, there has been a trade-off between simplicity and effectiveness for providing the input z . On one hand, providing random noise as the extra input z maintains a simple training objective (same as in pix2pix [2]). However, [2, 8] showed that the generator has little incentive to utilize the input vector z since it only encodes random information, and therefore the generator ends up ignoring z and collapsing to one or few modes. On the other hand, *BicycleGAN* [8] combines the objectives of both conditional Variational Auto-Encoder GANs (cVAE-GAN) and conditional Latent Regressor GANs (cLR-GAN) to learn a latent embedding z simultaneously with the generator G . Their training enforces two cycle consistencies: $B \rightarrow z \rightarrow \hat{B}$ and $z \rightarrow \hat{B} \rightarrow \hat{z}$. This proved to be very effective, but the training objective is more involved, which makes the training slower. Also, since the latent embedding is being trained simultaneously with the the generator, hyper-parameter tuning becomes more critical and sensitive. The training objective of more recent works (e.g., [9, 10]) is even more complicated. We aim to combine the best of both worlds: an effective training of a latent embedding that models the distribution of possible outputs, while retaining a simple training objective. This would allow for faster and more efficient training, as well as less sensitivity to hyper-parameters. We observe that the variability in many target domains can be represented by the style diversity of images in the target domain B , where the style is defined in terms of the Gram matrices used in the Neural Style Transfer literature [52]. However, using Gram matrices directly to represent styles is not feasible due to its very high dimensionality. So, instead we learn an embedding by separately training an encoder network E on an auxiliary task to optimize for $z = E(I^B)$ capturing the style of an image I^B . Visualizing the pre-trained latent space shows that our pre-trained encoder models different modes of the output distribution (e.g., different colors, lighting and weather conditions, ... etc.) as clusters of images with similar styles as shown in §4.7. Then, to synthesize an image $\hat{I}^B = G(I^A, z)$, the input latent can be used to clearly distinguish the style cluster to which the output belongs. This makes for an effective and more stable training of the generator G , since G is just required to discover the correlation between output images and their corresponding style em-

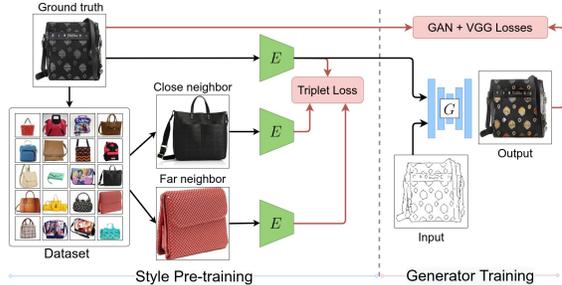


Figure 1: Overview of our training pipeline. Stage 1: pre-training the style encoder E using a triplet loss. Stages 2, 3: training the generator G , and finetuning both G, E together using GAN and reconstruction losses.

bedding z . Moreover, experimental evaluation shows that the proposed style-based pre-training yields better results in terms of more faithful style capture and transfer, as well as better output quality and diversity.

To incorporate this into BicycleGAN [8], we replace the simultaneous training of the encoder E and the generator G with a staged training (Figure 1) as follows:

- **Stage 1:** Pre-train E on a proxy task that optimizes an embedding of images in the output domain B into a low-dimensional style latent space, such that images with similar styles lie closely in that space (i.e., clustered).
- **Stage 2:** Train the generator network G while fixing the encoder E , so that G learns to associate the style of output images to their deterministic style embedding $z = E(I^B)$.
- **Stage 3:** Finetune both the E and G networks together, allowing for the style embedding to be further adapted to best suit the image synthesis task for the target domain.

Next, we explain how to pre-train the style encoder network E in §3.1, and how to train the generator G using the pre-learned embeddings (§3.2). Finally, we demonstrate the generalization of pre-training the style encoder E in §3.3.

3.1. Weakly-supervised encoder pre-training

The goal of pre-training the encoder network E is to learn a deterministic mapping from the style of a target image $I_i^B \in B$ to a latent style code $z_i = E(I_i^B)$. Ideally, images with similar styles should be close in the style embedding space, while images with different styles should be far apart. To supervise training such an embedding, we utilize the style loss [52] as a distance metric to measure the style similarity between any two given images. The style encoder network E is then trained using a triplet loss [53], where the input is a triplet of images (I_a, I_p, I_n) , where (I_a, I_p) have similar style, while (I_a, I_n) have different style, as measured by the style loss metric. The training objective for E is given by:

$$\mathcal{L}^{\text{tri}}(I_a, I_p, I_n) = \max([\|z_a - z_p\|^2 - \|z_a - z_n\|^2 + \alpha], 0) + \lambda \mathcal{L}^{\text{reg}}(z_a, z_p, z_n) \quad (1)$$



Figure 2: Qualitative comparison with baselines. Ours better matches the ground truth (GT) style.



Figure 3: Style transfer for different datasets. We show output for applying different styles to each input image.

where α is a separation margin, λ is a relative weighting parameter between the main triplet objective and an optional regularization term $\mathcal{L}^{\text{reg}}(\cdot)$ which is an $L2$ regularization to encourage learning a compact latent space.

Triplet selection. To generate triplets for pre-training the encoder E , we compute the set of k_c closest and k_f furthest neighbors for each anchor image I_a as measured by the style loss. Then, for each anchor image I_a , we randomly sample a positive image I_p and a negative image I_n from the set of closest and furthest neighbors respectively. We found that, for large datasets, it is sufficient to generate triplets for a subset of the training images. One challenge is the set of images with an outlier style. Such images will be furthest neighbors to most images, and can mislead the training by just projecting outlier images to separate clusters. To deal with this, we sample the negative style image I_n from a larger set of furthest neighbors; while the positive image I_p is sampled from a small set of closest neighbors so that it would have reasonable style similarity to the anchor image.

3.2. Generator training

After pre-training the style encoder E (stage 1), we have established a mapping from images in the output domain, $I^B \in B$, to their style-embedding $z = E(I^B)$. Feeding the style embedding as input to the generator during training, the generator has good incentive to associate the style of output images to their corresponding style embedding instead of learning to hallucinate the style. It’s important to retain the deterministic correspondence between images and their style

codes to facilitate the job of the generator to discover this correlation. This is why, during stage 2, we keep the weights of the style encoder, E , fixed. The forward pass reconstructs a training image I_i^B as $\hat{I}_i^B = G(I_i^A, z_i)$, where $z_i = E(I_i^B)$. The training objective is similar to that of pix2pix [2]:

$$\mathcal{L}^{\text{img}}(I_i^B, \hat{I}_i^B) = \mathcal{L}_{\text{cGAN}}(I_i^B, \hat{I}_i^B) + \lambda_{\text{rec}} \mathcal{L}_{\text{rec}}(I_i^B, \hat{I}_i^B) \quad (2)$$

where we use the Least Square GAN loss (LSGAN) [25] for the $\mathcal{L}_{\text{cGAN}}$ term, and a VGG-based perceptual loss [14] for the reconstruction term \mathcal{L}_{rec} . Once the generator has learned to associate the output style with the input style embedding, stage 3 finetunes both the generator, G , and the style encoder, E , together using the same objective (2).

Style sampling. To perform multimodal synthesis on a given input at test time, we can capture the latent vector z from any existing image and transfer the style to the generated image. However, if we wish to sample styles directly from the latent distribution, one option is to enforce a prior on the latent distribution. For example, we found it effective to add an $L2$ regularization on the latent vectors to enforce zero-mean embeddings and limit the variance of the latent space. We then compute an empirical standard deviation for sampling. Another alternative to enable sampling is to train a mapper network \mathcal{M} to map the unit Gaussian to the latent distribution. This can be done as a post-processing step after the style encoder has been trained and finetuned. Specifically, we propose to train a mapper network \mathcal{M} using the nearest-neighbor based Implicit Maximum Likelihood Estimation (IMLE) training [54, 55]. The training objective

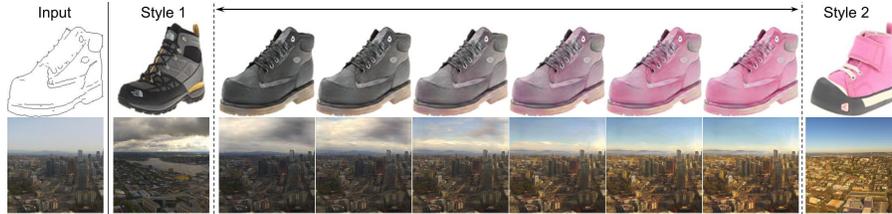


Figure 4: Style interpolation. Left column is the input to the generator G , second and last columns are input style images to the style encoder, and middle images are linear interpolation in the embedding space (figure better seen in zoom).

is given by:

$$\mathcal{M} = \arg \min_{\mathcal{C}} \sum_i \|z_i - \tilde{\mathcal{M}}(e_i)\|_2^2, e_i = \arg \min_{r_j} \|z_i - \tilde{\mathcal{M}}(r_j)\|_2^2 \quad (3)$$

where $\{r_j\}$ is a set of random samples from the unit Gaussian prior, and for each latent code z_i , we select e_i that generates the nearest neighbor $\mathcal{M}(e_i)$ to z_i .

3.3. Generalizing the pre-training stage

The use of Gram matrices for Neural Style Transfer proved to be very effective and it reliably captures the style of arbitrary input images. This implies that Gram matrices can reliably encode styles from a wide range of domains, and they are not specific to a certain domain. Therefore, we hypothesize that encoder pre-training using a style-based triplet loss would learn a generic style embedding that can generalize across multiple domains and be effective for multi-modal I2I translation. This would allow for performing the pre-training stage only once using auxiliary training data. The finetuning stage eventually tweaks the embedding to better suit the specific target domain B . We validate our hypothesis experimentally in §4, and show that pre-training the style encoder on datasets other than the target domain B doesn't degrade the performance. It can even improve the performance if the target dataset is small, in which case pre-training on an auxiliary dataset helps with the generalization of the overall model.

4. Experimental evaluation

Datasets. We evaluate our approach on five standard I2I translation benchmarks used in [2, 8]; Architectural labels \rightarrow photo, aerial \rightarrow map, edges \rightarrow shoes/handbags and night \rightarrow day. In addition, we use the Space Needle timelapse dataset [56], which consists of 2068 paired images with a 8280×1080 resolution, where the input domain includes images with temporally smoothed appearance, and the output domain contains real images spanning different lighting and weather conditions.

Baselines. While we report numbers for retrained models using the official code released with BicycleGAN (BicycleGAN v0) for completeness, we mainly compare to two stronger baselines:

- **BicycleGAN v1:** we implement BicycleGAN using the same network architecture as used in our approach to have

a fair comparison (see supp. material for implementation details).

- **MUNIT-p:** We train MUNIT [9] in a *paired* setup by applying its cross-cycle consistency constraint as follows: the input is a pair of training examples $(I_1^A, I_1^B), (I_2^A, I_2^B)$ for which we obtain their respective style embeddings $z_1 = E(I_1^B), z_2 = E(I_2^B)$. We then apply a 2-step cyclic reconstruction of I_1^B, I_2^B ; in the first step, we generate both images with a swapped style $u = G(I_1^A, z_2), v = G(I_2^A, z_1)$. In the second step, we re-capture the latent style vectors $\hat{z}_2 = E(u), \hat{z}_1 = E(v)$ and generate the original images I_1^B, I_2^B by swapping the style again: $\hat{I}_1^B = G(I_1^A, \hat{z}_1), \hat{I}_2^B = G(I_2^A, \hat{z}_2)$. We add a cyclic reconstruction loss for \hat{I}_1^B, \hat{I}_2^B .

4.1. Image reconstruction

We report the reconstruction quality of validation set images, using both *PSNR* and *AlexNet*-based *LPIPS* [57] metrics, in Table 1. Note that our results with the pre-trained embeddings without finetuning (stage 2) are on-par-with the baselines. This verifies the validity of our approach and that style-based encoder pre-training successfully learns to distinguish different modes in the output domain, which proves effective for training multi-modal I2I networks. Finetuning (stage 3) further improves our results compared to the baselines. See supp. material for more quantitative comparison including Inception Scores (*IS*). Figure 2 shows qualitatively how our approach reconstructs the target style more faithfully. Our approach not only matches the ground truth colors better, but also texture (e.g., left column, first and third rows). While the baselines rely on VAEs to provide the style sampling property, we observe that, for a low dimensional latent space, the noise robustness added by VAEs reduces the expressiveness of the latent space, since slight changes to one style code would still be mapped to the same point in the latent space. This explains why our approach, which doesn't use VAEs, achieves more faithful style capture and reconstruction. We verified this by studying the effect of removing the VAE component from the baselines, which improved their performance as shown in §4.5.

4.2. Style transfer and sampling

Figure 3 shows style transfer to images from the validation set of different datasets. Note how the style trans-



Figure 5: Style sampling for different datasets using our approach after full training (e.g., phase 3). We sample either from $N(\mu, \sigma)$, where μ, σ are computed from the train set (middle), or using the mapper network \mathcal{M} (right).

Table 1: Validation set reconstruction quality, as measured by $PSNR$ (higher is better) and $LPIPS$ [57] (lower is better), for various datasets. We compare between retraining BicycleGAN [8] authors’ released code (Bicycle v0), the baselines (BicycleGAN v1 and MUNIT-p) described in §4, and our approach both before finetuning (ours - stage 2), and after finetuning (ours - stage 3).

	edges2handbags		edges2shoes		labels2facades		night2day		maps		space needle	
	PSNR \uparrow	LPIPS \downarrow										
Bicycle v0	17.08	0.255	20.24	0.177	12.64	0.431	13.25	0.520	14.32	0.396	–	–
Bicycle v1	18.52	0.198	21.84	0.124	13.08	0.378	13.88	0.491	14.67	0.359	19.72	0.233
MUNIT-p	19.23	0.192	22.51	0.132	13.36	0.375	14.48	0.480	16.17	0.407	19.84	0.238
Ours - stage 2	18.01	0.209	21.40	0.140	13.44	0.383	14.34	0.476	15.08	0.392	21.39	0.227
Ours - stage 3	18.91	0.177	22.68	0.117	13.44	0.370	15.05	0.452	15.15	0.349	22.11	0.187

fer copies the weather conditions in the Space Needle and Night2day datasets. For example, in the Space Needle dataset, we show sunset, sunny, foggy and cloudy weather. Also, the Night2day examples exhibit variation in lighting conditions including transferring whether the surface is sunlit or not, as well as different cloud patterns and clear skies. Comparison to the baselines’ style transfer results in the supplementary material further highlights the improvements of our approach. We can also sample random styles directly from the latent distribution as described in §3.2. Figure 5 shows results for both adhoc sampling from the assumed $N(\mu, \sigma)$ empirical distribution, as well as formally sampling from a unit Gaussian using the mapper network \mathcal{M} . Note that the diversity of sampled styles doesn’t stem from simple color changes; for example, sampled styles for the Space Needle dataset show clear weather changes, such as cloudy/sunny, different cloud patterns and even sampling foggy weather which was present in some images in the training set. In the Maps dataset, the existence and/or density of bushes clearly varies between different sampled styles. Also, in the Edges2handbags dataset, the texture of the bag varies

between smooth and rough leather (better seen in zoom). While the results of both sampling methods look good, we note that the assumption for adhoc sampling is not explicitly enforced, and thus could lead to sampling bad style codes outside the distribution (see supp. material for examples).

4.3. Style interpolation

Figure 4 shows style interpolation by linearly interpolating between two latent vectors. For example, note the smooth change in lighting and cloud patterns when going from cloudy to sunny in the Space Needle dataset. More interpolation results on other datasets, and comparison with interpolation results of the baselines can be found in the supplementary material.

4.4. Pre-training generalization

Since the notion of style, as defined in the Neural Style Transfer literature, is universal and not specific to a certain domain, we hypothesized that style-based encoder pre-training would learn a generic style embedding that can generalize across multiple domains and be effective for

Table 2: Ablation study of the effect of different components and loss terms using the edges2handbags dataset. We study direct and cyclic reconstructions on ground truth images (dir_recon, cyc_recon), discriminator loss on direct reconstructions and on generated images with a randomly sampled style (D_dir, D_rand_z), latent reconstruction (z_recon), L2 and KL regularization on the latent vector z (z_L2, z_KL), and finally the use of VAE vs. just an auto-encoder.

Approach	Loss setup								IS \uparrow	PSNR \uparrow	LPIPS \downarrow
	dir_recon	cyc_recon	D_dir	D_rand_z	z_recon	z_L2	z_KL	VAE			
Bicycle v1	✓	–	✓	✓	✓	–	✓	✓	2.31 \pm 0.05	18.28 \pm 0.30	0.201 \pm 0.003
MUNIT-p	✓	✓	✓	✓	✓	–	✓	✓	2.45 \pm 0.07	18.96 \pm 0.30	0.192 \pm 0.002
Bicycle v2	✓	–	✓	✓	✓	–	✓	–	2.36 \pm 0.12	19.02 \pm 0.10	0.175 \pm 0.001
MUNIT-p v2	✓	✓	✓	✓	✓	–	✓	–	2.44 \pm 0.06	19.34 \pm 0.07	0.176 \pm 0.002
Bicycle v3	✓	–	✓	✓	✓	✓	–	–	2.34 \pm 0.08	19.21 \pm 0.06	0.177 \pm 0.002
MUNIT-p v3	✓	✓	✓	✓	✓	✓	–	–	2.33 \pm 0.04	19.24 \pm 0.09	0.180 \pm 0.001
Ours v1	✓	✓	✓	✓	✓	✓	–	–	2.41 \pm 0.07	18.97 \pm 0.13	0.189 \pm 0.004
Ours v2	✓	–	✓	✓	✓	✓	–	–	2.43 \pm 0.03	18.94 \pm 0.10	0.183 \pm 0.002
Ours v3	✓	–	✓	–	–	✓	–	–	2.42 \pm 0.03	18.94 \pm 0.05	0.176 \pm 0.001
Ours v4	✓	–	✓	–	–	–	–	–	2.46 \pm 0.03	18.94 \pm 0.02	0.177 \pm 0.001

Table 3: Generalization of a pretrained style encoder E . We report validation set reconstruction for the edges2handbags and night2day datasets when pretraining with different datasets. Stages 2, 3 show results before/after finetuning E respectively.

Dataset	pretrain dataset	Stage 2		Stage 3	
		PSNR \uparrow	LPIPS \downarrow	PSNR \uparrow	LPIPS \downarrow
edges2handbags	edges2handbags	18.01	0.209	18.91	0.177
	edges2shoes	17.89	0.215	18.96	0.176
	space_needle	17.86	0.221	19.02	0.175
night2day	night2day	13.75	0.489	15.15	0.454
	space_needle	14.34	0.476	15.05	0.452
	edges2handbags	13.91	0.492	15.03	0.461

multi-modal image I2I translation. Here, we experimentally verify our hypothesis in Table 3. For a target dataset, we train the generator G three times, each with different pre-training of the style encoder E : (1) same dataset pre-training: pre-train E using the output domain B of the target dataset. (2) similar-domain pre-training: pre-train on a different dataset, but whose output domain bears resemblance to the output domain of the target dataset (e.g., edges2shoes and edges2handbags, or day images from night2day and the Space Needle timelapse dataset). (3) different-domain pre-training: pre-train on a different dataset whose output domain has different styles from that of the target dataset (e.g., edges2handbags and the Space Needle timelapse datasets, or night2day and edges2handbags datasets). Table 3 shows that without finetuning (i.e., stage 2), the edges2handbags dataset shows a slight performance degradation when going from pre-training on the same dataset, to pre-training on a similar-domain dataset, and finally pre-training on a different-domain dataset. On the other hand, the night2day dataset has only ~ 100 unique scenes for training. So, pre-training on another dataset such as Space Needle generalizes better to new scenes in the validation set, since it helps avoid

Table 4: Diversity score is the average LPIPS distance [57]. User preference score is the percentage a method is preferred over ‘Ours v4’, on the edges2shoes dataset.

Approach	LPIPS \uparrow (transfer)	LPIPS \uparrow (sampling)	User preference \uparrow
Bicycle v1	0.102	0.119	30.0%
MUNIT-p	0.138	0.132	37.7%
Ours v1	0.153	0.148	46.5%
Ours v2	0.171	0.140	41.1%
Ours v3	0.149	0.165	50.4%
Ours v4	0.154	0.132	50%

overfitting the small number of unique scenes in the training set. After finetuning, performance differences further reduce to be insignificant. We also investigate the generalization of the proposed encoder pre-training to the case of using non-style distance metrics in the supplementary material.

4.5. Ablative study

We investigate the role of different loss terms as we transition from the loss setup of the baselines to that of our training approach. We first remove the variational part in both BicycleGAN v1 and MUNIT-p baselines resulting in Bicycle v2, MUNIT-p v2. We further remove the Gaussian prior and replace the KL loss with an L2 regularization in Bicycle v3, MUNIT-p v3. To maintain random latent vector sampling during training without a prior, we sample a random training image, and use its style code. We define different versions of our approach (v1, v2, v3, and v4) based on different loss setup during training as follows: we start with ‘Ours v1’, which has the same setup as MUNIT-p v3, except that it uses pre-trained embeddings as described in §3.1. We then remove cyclic reconstruction, random z sampling, and L2 regularization terms resulting in ‘Ours v2’, ‘v3’, and ‘v4’ re-

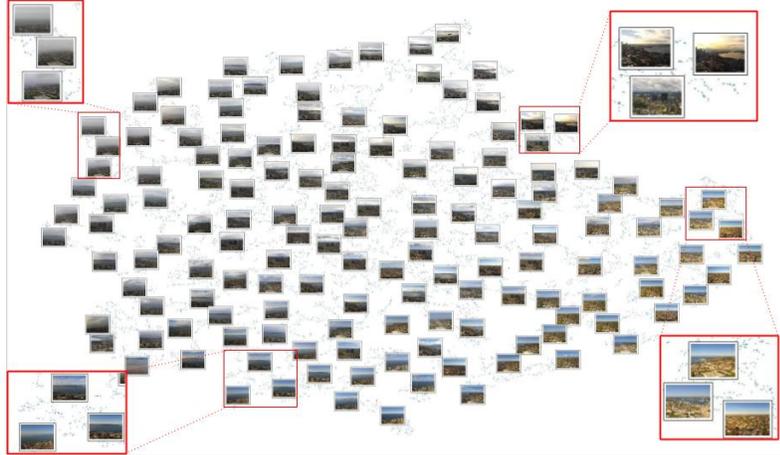


Figure 6: t-SNE plot for the pre-trained latent space of the Space Needle timelapse dataset. Images with similar styles (e.g., weather conditions and lighting) lie closely in the latent space. (figure best seen in zoom)

spectively. We run each setup on the edges2handbags dataset. In order to draw more reliable conclusions, we repeat each experiment 3 times and report the mean and standard deviation in Table 2. We notice that removing the variational part in VAEs is enough to improve the reconstruction results. While VAEs in general are robust to noise in the input latent, we observe that this comes at the expense of the expressiveness of the latent space (e.g., less faithful style capture and transfer), especially for low dimensional latents. We also observe that our approach generally performs better with less constraints (loss terms). For example, “Ours v1, v2” have lower results than their “Bicycle v3”, “MUNIT-p v3” counterparts. This shows that the main benefit of pre-trained embeddings is when the network is less constrained.

4.6. Diversity and user study

We evaluate diversity by computing the average LPIPS distance over 1600 outputs. We measure diversity on two setups: we sample 100 validation images, and (1) apply style transfer from 16 randomly sampled images, or (2) we sample 16 random codes using the mapper network \mathcal{M} to obtain 1600 outputs. We also measure the realism and faithfulness of style transfer through a user study, where 30 participants are shown an input shoe sketch, an input style image and two style transfer outputs. They are asked to choose which output looks more realistic, and if both are realistic, then which transfers the style more faithfully. We fix ‘Ours v4’ approach as anchor and compare other methods to it. Table 4 shows that the baselines achieve lower diversity and user preference compared to our approach, specially in the style transfer setup. Different variations of our method, except for ‘Ours v2’ yield similar diversity and user preference scores. We observe that ‘Ours v2’ shows artifacts in some outputs, leading to higher diversity but lower user preference. Our diversity results for the style sampling setup have some

variation and are sensitive to the mapper network training, but are still either on-par or better than the baselines.

4.7. Visualizing pre-trained embeddings

Figure 6 visualizes the pre-trained latent space learned by the style encoder E . The visualization shows meaningful clusters of similar styles (weather conditions for the Space Needle timelapse dataset). Refer to supp. material for visualization of the latent space after finetuning, as well as the latent space learned by the baselines.

More analysis and discussion: Please refer to the supplementary material for convergence analysis, training time comparison, and more quantitative and qualitative results.

5. Conclusion

We investigated the effectiveness of Style-based Encoder Pre-training (StEP) for the task of multi-modal I2I translation. The proposed pre-training can be done once on auxiliary data and generalizes well to several domains. This allows for a faster training of I2I translation networks with fewer losses and achieves more faithful style capture and transfer. Furthermore, we studied the contribution of different loss terms to the task of multi-modal I2I translation, where we discovered that noise added by a variational auto-encoder can limit the expressiveness of low-dimensional latent spaces. We proposed two simple alternatives to VAEs to provide latent code sampling. Finally, we achieved state-of-the-art results on several benchmarks.

Acknowledgements. We would like to thank Ricardo Martin-Brualla for helping with initial concepts and reviewing drafts. We also thank Kamal Gupta and Alex Hanson for suggesting the acronym for our approach, and the members of the Perception and Intelligence (PI) Lab for their helpful feedback. The project was partially funded by DARPA MediFor (FA87501620191) and DARPA SAIL-ON (W911NF2020009) programs.

References

- [1] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *iccv*, pages 1511–1520, 2017. [1](#), [2](#)
- [2] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. [1](#), [2](#), [3](#), [4](#), [5](#)
- [3] 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 *CVPR*, 2018. [1](#), [2](#), [11](#)
- [4] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. [1](#), [2](#)
- [5] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *eccv*, pages 649–666. Springer, 2016. [1](#), [2](#)
- [6] Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. *tog*, 36(4):119, 2017. [1](#), [2](#)
- [7] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. [1](#), [2](#)
- [8] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *NeurIPS*, 2017. [1](#), [2](#), [3](#), [5](#), [6](#), [11](#), [13](#)
- [9] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. [1](#), [2](#), [3](#), [5](#)
- [10] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Kumar Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *ECCV*, 2018. [1](#), [2](#), [3](#), [11](#)
- [11] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *nips*, pages 3483–3491, 2015. [1](#), [2](#)
- [12] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *iclr*, 2016. [1](#)
- [13] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastrogiuseppe, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *iclr*, 2016. [1](#)
- [14] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. [2](#), [4](#), [11](#)
- [15] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rendering in the wild. *cvpr*, 2019. [2](#)
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. [2](#)
- [17] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *ECCV*, pages 181–196, 2018.
- [18] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*, pages 843–852, 2017. [2](#)
- [19] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, pages 1422–1430, 2015. [2](#)
- [20] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, pages 69–84. Springer, 2016. [2](#)
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014. [2](#)
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *nips*, pages 2672–2680, 2014. [2](#)
- [23] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *nips*, pages 5767–5777, 2017. [2](#)
- [24] Animesh Karnewar and Raghu Sessa Iyengar. Msggan: Multi-scale gradients gan for more stable and synchronized multi-scale image synthesis. *arXiv preprint arXiv:1903.06048*, 2019.
- [25] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. [4](#), [11](#)
- [26] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *iclr*, 2016. [2](#)
- [27] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018. [2](#)
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *cvpr*, 2019. [2](#)
- [29] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. *icml*, 2016. [2](#)
- [30] Mihaela Rosca, Balaji Lakshminarayanan, and David Warde-Farley Shakir Mohamed. Variational approaches for auto-encoding generative adversarial networks. *stat*, 1050:15, 2017. [2](#)
- [31] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. [2](#)
- [32] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *icml*, pages 2642–2651. JMLR. org, 2017.
- [33] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *nips*, pages 4790–4798, 2016. [2](#)
- [34] Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Generating images from captions with attention. *iclr*, 2016. [2](#)
- [35] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *icml*, 2016. [2](#)
- [36] Yedid Hoshen and Lior Wolf. Nam: Non-adversarial unsupervised domain mapping. In *eccv*, pages 436–451, 2018. [2](#)
- [37] Ricardo Martin-Brualla, Rohit Pandey, Shuoran Yang, Pavel

- Pidlypenskyi, Jonathan Taylor, Julien Valentin, Sameh Khamis, Philip Davidson, Anastasia Tkach, Peter Lincoln, Adarsh Kowdle, Christoph Rhemann, Dan B Goldman, Cem Keskin, Steve Seitz, Shahram Izadi, and Sean Fanello. Looking-Good: Enhancing performance capture with real-time neural re-rendering. In *Proc. SIGGRAPH Asia*, 2018. 2
- [38] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *tog*, 2019. 2
- [39] Hao Dong, Simiao Yu, Chao Wu, and Yike Guo. Semantic image synthesis via adversarial learning. In *ICCV*, 2017. 2
- [40] Xiaolong Wang and Abhinav Gupta. Generative image modeling using style and structure adversarial networks. In *ECCV*, 2016.
- [41] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, 2017. 2
- [42] Taeksoo Kim, Moonsoo Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *ICML*, 2017. 2
- [43] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017.
- [44] Liqian Ma, Xu Jia, Stamatios Georgoulis, Tinne Tuytelaars, and Luc Van Gool. Exemplar guided unsupervised image-to-image translation. *iclr*, 2019.
- [45] Amélie Royer, Konstantinos Bousmalis, Stephan Gouws, Fred Bertsch, Inbar Mosseri, Forrester Cole, and Kevin Murphy. Xgan: Unsupervised image-to-image translation for many-to-many mappings. *arXiv preprint arXiv:1711.05139*, 2017.
- [46] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 2
- [47] 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 *cvpr*, pages 8789–8797, 2018. 2
- [48] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, pages 8188–8197, 2020. 2
- [49] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. *arXiv preprint arXiv:1808.07371*, 2018. 2
- [50] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Nikolai Yakovenko, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *NeurIPS*, 2018. 2
- [51] Amjad Almahairi, Sai Rajeshwar, Alessandro Sordani, Philip Bachman, and Aaron Courville. Augmented CycleGAN: Learning many-to-many mappings from unpaired data. In *ICML*, 2018. 2
- [52] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 3
- [53] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *cvpr*, pages 815–823, 2015. 3
- [54] Yedid Hoshen, Ke Li, and Jitendra Malik. Non-adversarial image synthesis with generative latent nearest neighbors. In *cvpr*, pages 5811–5819, 2019. 4
- [55] Ke Li and Jitendra Malik. Implicit maximum likelihood estimation. *arXiv preprint arXiv:1809.09087*, 2018. 4
- [56] Ricardo Martin-Brualla. Space needle timelapse. <https://bit.ly/2RMKZ3B>, 2007. Accessed: 2019-09-22. 5
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *cvpr*, pages 586–595, 2018. 5, 6, 7
- [58] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, 2014. 11
- [59] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, et al. Deep face recognition. In *bmvc*, volume 1, page 6, 2015. 13
- [60] Wu Jie. Facial expression recognition. <https://github.com/WuJie1010/Facial-Expression-Recognition.Pytorch>, 2018. Accessed: 2019-09-22. 15
- [61] KDEF. Karolinska directed emotional faces (kdef) dataset. <http://kdef.se>, 2017. Accessed: 2019-09-22. 15

A. Appendix

A.1. Implementation details

Triplet selection requires computing the set of nearest and furthest neighbors to each anchor image. When pre-training using a large dataset, we found it sufficient to randomly sample a subset of 8000 images and sample triplets from this subset. This number was chosen to ensure fast nearest-neighbor computation.

The generator network G has a symmetric encoder-decoder architecture based on [3], with extra skip connections by concatenating feature maps of the encoder and decoder. We use a multiscale-patchGAN discriminator [3] with 3 scales and employ a LSGAN [25] loss. The mapper network \mathcal{M} is a multi-layer perceptron (MLP) with three 128-dimensional hidden layers and a \tanh activation function. For the reconstruction loss, we use the perceptual loss [14] evaluated at $\text{conv}_{i,2}$ for $i \in [1, 5]$ of VGG [58] with linear weights of $w_i = 1/2^{6-i}$ for $i \in [1, 5]$. The architecture of the style encoder E is adopted from [10], and we use a latent style vector $z \in \mathbb{R}^8$. Our optimizers setup is similar to that in [8]. We use three *Adam* optimizers: one for the generator G and encoder E , another for the discriminator D , and another optimizer for the generator G alone with $\beta_1 = 0, \beta_2 = 0.99$ for the three optimizers, and learning rates of 0.001, 0.001 and 0.0001 respectively. We use a separate *Adam* optimizer for the mapper network \mathcal{M} with $\beta_1 = 0.5, \beta_2 = 0.99$, and a learning rate of 0.01 with a decay rate of 0.7 applied every 50 steps. Relative weights for the loss terms are $\lambda_{\text{cGAN}} = 1, \lambda_{\text{rec}} = 0.02$ and $\lambda_{L2} = 0.01$ for the GAN loss, reconstruction loss, and $L2$ latent vector regularization respectively. When sampling triplets for any anchor image I_c , we use $k_c = 5, k_f = 13$ for the size of the set of close and far neighbors respectively.

A.2. Training time

Simplifying the training objective allows for faster training, as well as a larger batch size due to lower memory usage. Table 5 shows the processing time per 1000 training images for the baselines as well as different variations of our approach as defined in Table 2 in the main text.

Table 5: Training time (in seconds) per 1000 images for the baselines, as well as different versions of our approach (defined in Table 2 in the main text).

Approach	Batch size	time/kimg↓ (sec)	Max batch size	time/kimg↓ (sec)
Bicycle v1	8	93.11	12	85.36
MUNIT-p	8	155.72	8	155.72
Ours v1	8	145.50	8	145.50
Ours v2	8	98.55	12	93.04
Ours v3	8	64.92	16	53.80

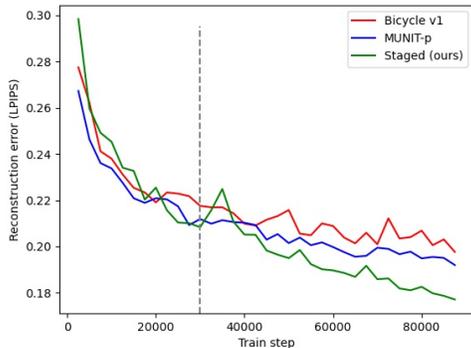


Figure 7: Convergence comparison between the proposed staged training (ours - v3) and the BicycleGAN baselines measured by the reconstruction error (LPIPS) of the validation set of the edges2handbags dataset. Dotted line shows the transition between stages 2 and 3 of our training (i.e, switching from a fixed E to finetuning both G and E together).

A.3. Convergence analysis

Figure 7 compares the convergence of our staged training compared to the BicycleGAN baselines. The dotted line in the graph marks the transition between stages 2 and 3 of our training (i.e, switching from a fixed pre-trained encoder E to finetuning both G and E together). We measure the reconstruction error (LPIPS) of the validation set of the edges2handbags dataset as the training progresses. Results show that with a fixed pre-trained encoder, our staged training starts with higher error than the baselines, but quickly drops to show similar performance as the baselines, and even beats the baselines before switching to stage 3 (marked by a dotted line). When starting to finetune the encoder E , we get a spike in the reconstruction error as the network adapts to the shift in the pre-trained embeddings, but then our staged training steadily widens the performance gap with the baselines. This shows the importance of the finetuning stage to tweak the pre-trained embeddings to better serve the image synthesis task for the target domain.

A.4. More quantitative comparison

We report the Inception Score (IS) computed over the validation set of various datasets in Table 6. Surprisingly, results after finetuning (“ours - stage 3”) are slightly worse than those before finetuning (“ours - stage 2”), but both are still better than the baselines except for the maps dataset. We also note that Inception Score is not very suited to image-to-image translation tasks, since it prefers output diversity with respect to ImageNet classes, not within-class diversity as in our case.

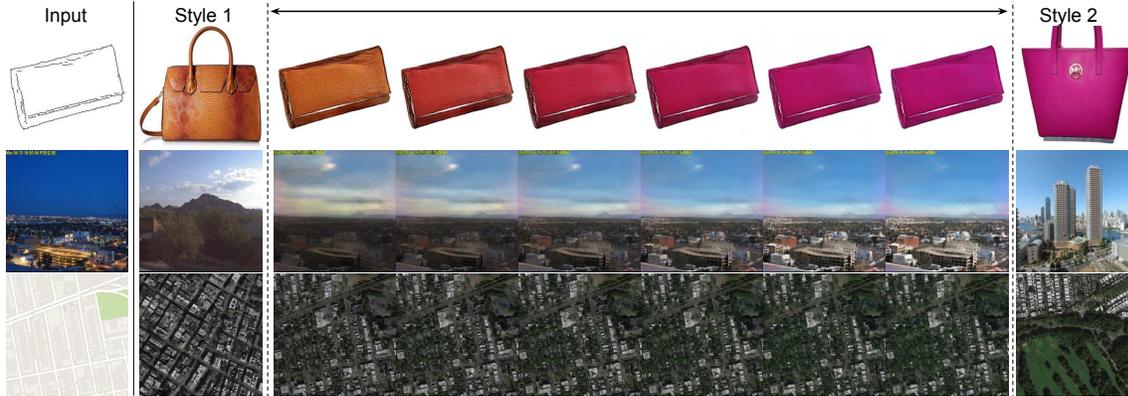


Figure 8: Style interpolation. Left column is the input to the generator G , second and last columns are input style images to the style encoder, and middle images are linear interpolation in the embedding space (figure better seen in zoom).

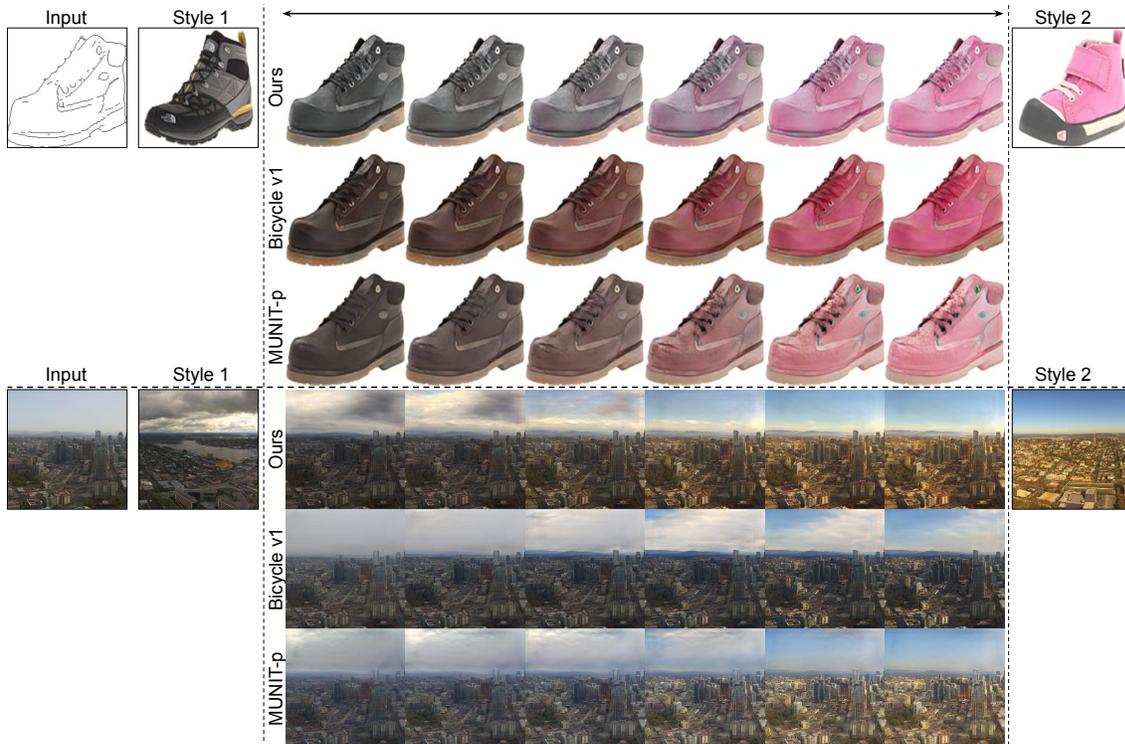


Figure 9: Style interpolation. Left column is the input to the generator G , second and last columns are input style images to the style encoder, and middle images are linear interpolation in the embedding space (figure better seen in zoom).

Table 6: Inception score comparison (higher is better) for different datasets.

	handbags	shoes	facades	night2day	maps	space needle
Bicycle v1	2.13	2.83	1.41	1.65	3.26	1.82
MUNIT-p	2.07	2.64	1.45	1.74	3.57	1.77
Ours - stage 2	2.22	2.75	1.61	1.76	3.32	1.90
Ours - stage 3	2.15	2.85	1.56	1.84	3.28	1.89

A.5. More style interpolations

Figure 8 shows style interpolation on more datasets. Notice that, in the edges2handbags results, not only the color is transferred, but also the texture varies from non-smooth to smooth leather. Also, in the maps dataset, the density of bushes varies smoothly. Figure 9 further compares our interpolation results with the baselines. Our results show more complex interpolations, as evidenced by the change in lighting and cloud patterns, as well as more faithful style transfer compared to the baselines.



Figure 10: Style transfer comparison on different datasets. For each dataset, we apply different styles to the same input image and show the output of different methods.

A.6. Style transfer comparison

We compare style transfer performance of our approach against that of the baselines in Figure 10. Our approach faithfully captures and transfers colors and weather conditions (including sky and surface lighting) compared to the baselines. We attribute the inferior results of the baselines to the reliance on VAEs to train the latent space. This is because noise added by VAEs means that slight changes to one style would still be mapped to the same point in the latent space, which limits the capacity of low dimensional latent space. On the other hand, our pre-trained embeddings don't rely on VAEs and hence, can discriminate between more styles.

A.7. Latent space visualization

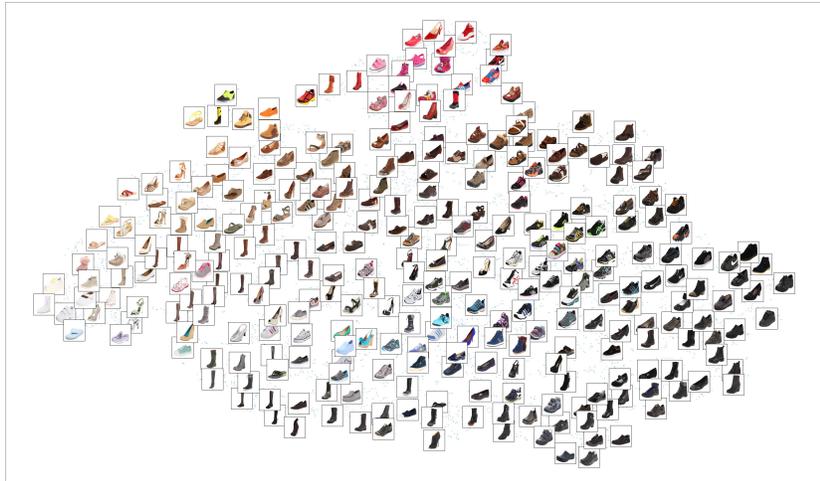
Figure 11a visualizes the latent space learned by the style encoder E after pretraining and before finetuning (a), after finetuning (b), and the latent space learned by BicycleGAN [8] (c). The embedding learned through pre-training (i.e. before training the generator G) shows meaningful clusters, which verifies the validity of the proposed style-based pre-training. Finetuning smooths the style clusters and brings the latent space closer to that of BicycleGAN.

A.8. Encoder pre-training with non-style metrics

Pre-training the encoder using a style-based triplet loss showed to be successful for multi-modal image translation tasks where the variability in the target domain is mainly color-based. This is shown in the results obtained on several benchmarks, even before the finetuning stage (“ours - stage 2” in Table 1 of the main text). We note though that the usage of style-loss as a distance metric for triplet sampling is just one choice and can be replaced with other distance metrics depending on the target application. Triplet sampling with style distance results in learning an embedding space where images with similar colors/styles lie closely in that space as shown in Section A.7. If, for example, we sample triplets instead based on the distance between VGG-Face [59] embeddings, the encoder will learn a latent space which is clustered by identity. In this section, we aim to validate that the proposed pre-training strategy can be extended to multi-modal image-to-image translation tasks with non-style variability. We inspect the task of manipulating facial expressions, where the input is a neutral face, and the output can have other emotions or facial expressions. For this task, similar emotions should be embedded closely in the latent space. We therefore use an off-the-shelf facial expression recognition system to compute the emotion sim-



(a) Our approach: after style pretraining.



(b) Our approach: after finetuning.



(c) BicycleGAN v1 baseline.

Figure 11: t-SNE plots for the latent style space learned by the style encoder E (a) after style pretraining, (b) after finetuning, and (c) using the BicycleGAN v1 baseline.

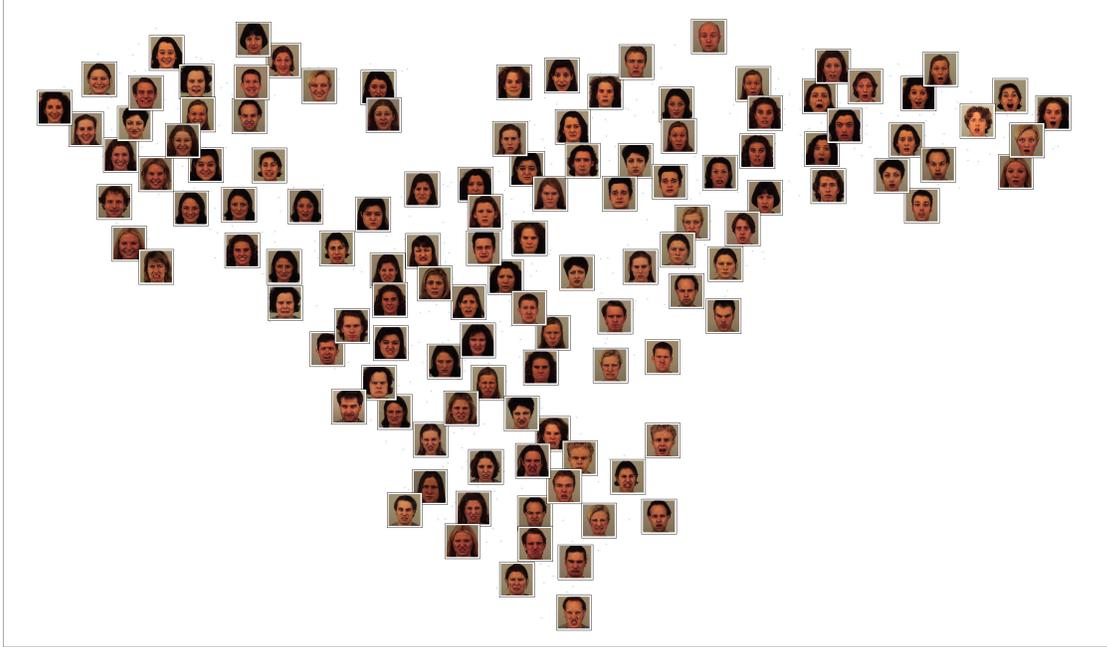


Figure 12: t-SNE plot for the pre-trained latent space learned for facial expressions on a subset of the KDEF dataset.

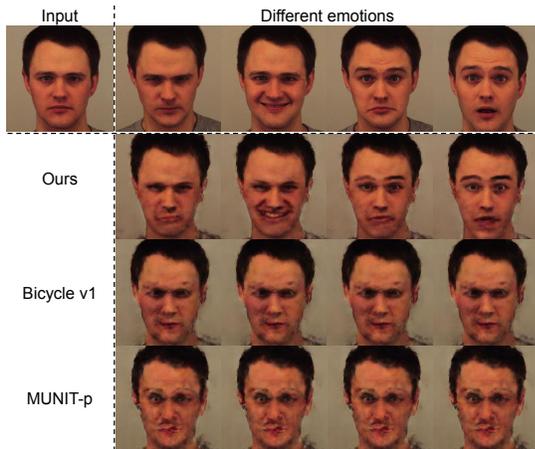


Figure 13: Emotion translation results. First row shows the input image, as well as the ground truth images from which we encode the latent emotion vector for reconstruction. Our staged training approach is able to achieve multi-modal synthesis, while the baselines collapse to a single mode.

ilarity/distance between any pair of images. Specifically, we compute the emotion distance as the euclidean distance between the 512-dimensional feature map of the last layer of a pretrained classification network (e.g., [60]). We visualize the learned latent space in Figure 12, which shows clusters with similar emotions or facial expressions. We also show example translation results on a holdout set of the

front-view images of the KDEF dataset [61] in Figure 13. We note that the generator successfully learns to manipulate facial expressions based solely on the pre-trained embeddings (without the finetuning stage). On the other hand, the BicycleGAN-based baselines collapsed to a single mode (over 3 different runs). This shows that our staged-training approach is stable and not sensitive to hyper-parameters, unlike the BicycleGAN baselines which will require careful hyper-parameter tuning to work properly on this task. We also point out that the poor output quality is mainly due to using a pixel-wise reconstruction loss for the generator training, while the input-output pairs in this dataset are not aligned. We didn't investigate improving the generator training as this is orthogonal to verifying the generalization of encoder pre-training.

A.9. Style sampling comparison

Figure 14 compares style sampling using the mapper network \mathcal{M} vs adhoc sampling from the assumed $N(\mu, \sigma)$ of an $L2$ -regularized latent space, where μ, σ are empirically computed from the training set. Note that adhoc sampling can sometimes sample bad style codes outside the distribution (e.g. third image in first row, and first image in third row in the right side of Figure 14), since the assumption that a $L2$ -regularized space would yield normally distributed latents with zero mean and low standard deviation is not explicitly enforced.



Figure 14: Sixteen randomly sampled styles using both the mapper network \mathcal{M} (left), as well as adhoc sampling from the empirically computed $N(\mu, \sigma)$ distribution of a $L2$ -regularized latent space (right). Adhoc sampling could sample bad style codes outside the latent distribution (marked in red).