



# COCO-FUNIT: Few-Shot Unsupervised Image Translation with a Content Conditioned Style Encoder

Kuniaki Saito<sup>1,2</sup>(✉) , Kate Saenko<sup>1</sup> , and Ming-Yu Liu<sup>2</sup> 

<sup>1</sup> Boston University, Boston, USA  
{keisaito,saenko}@bu.edu

<sup>2</sup> NVIDIA, Santa Clara, USA  
mingyul@nvidia.com

**Abstract.** Unsupervised image-to-image translation intends to learn a mapping of an image in a given domain to an analogous image in a different domain, without explicit supervision of the mapping. Few-shot unsupervised image-to-image translation further attempts to generalize the model to an unseen domain by leveraging example images of the unseen domain provided at inference time. While remarkably successful, existing few-shot image-to-image translation models find it difficult to preserve the structure of the input image while emulating the appearance of the unseen domain, which we refer to as the *content loss* problem. This is particularly severe when the poses of the objects in the input and example images are very different. To address the issue, we propose a new few-shot image translation model, COCO-FUNIT, which computes the style embedding of the example images conditioned on the input image and a new module called the constant style bias. Through extensive experimental validations with comparison to the state-of-the-art, our model shows effectiveness in addressing the *content loss* problem. Code and pretrained models are available at <https://nvlabs.github.io/COCO-FUNIT/>.

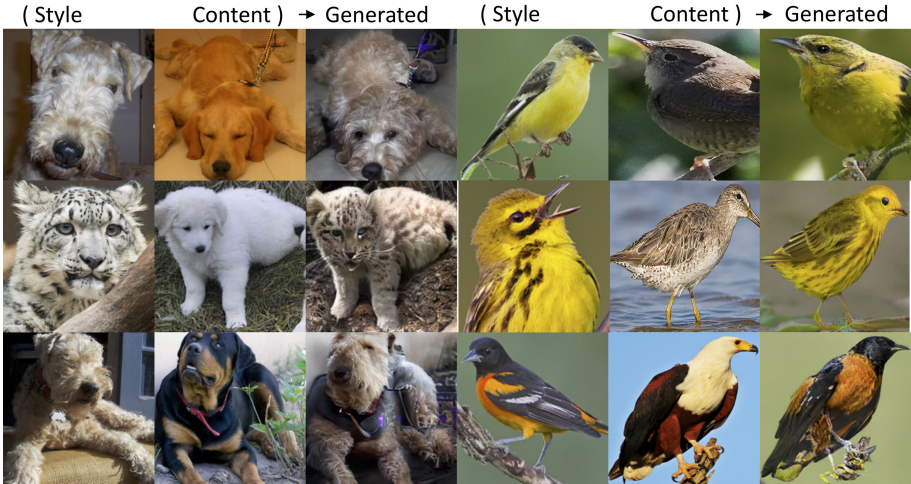
**Keywords:** Image-to-image translation · Generative Adversarial Networks

## 1 Introduction

Image-to-Image translation [18, 44] concerns learning a mapping that can translate an input image in one domain into an analogous image in a different domain. Unsupervised image-to-image translation [5, 21, 25, 26, 28, 39, 47] attempts to learn such a mapping without paired data. Thanks to the introduction of novel

---

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-58580-8\\_23](https://doi.org/10.1007/978-3-030-58580-8_23)) contains supplementary material, which is available to authorized users.




**Fig. 1.** Given as few as one style example image from an object class unseen during training, our model can generate a photorealistic translation of the input content image in the unseen domain.

network architectures and learning objective terms, the state-of-the-art has advanced significantly in the past few years. However, while existing unsupervised image-to-image translation models can generate realistic translations, they still have several drawbacks. First, they require a large amount of images from the source and target domains for training. Second, they cannot be used to generate images in unseen domains. These limitations are addressed in the few-shot *unsupervised* image-to-image translation framework [27]. By leveraging example-guided episodic training, the few-shot image translation framework [27] learns to extract the domain-specific style information from a few example images in the unseen domain during test time, mixes it with the domain-invariant content information extracted from the input image, and generates a few-shot translation output as illustrated in Fig. 2 (Fig. 1).

However, despite showing encouraging results on relatively simple tasks such as animal face and flower translation, the few-shot translation framework [27] frequently generates unsatisfactory translation outputs when the model is applied to objects with diverse appearance, such as animals with very different body poses. Often, the translation output is not well-aligned with the input image. The domain invariant content that is supposed to remain unchanged disappears after translation, as shown in Fig. 3. We will call this issue the *content loss* problem. We hypothesize that solving the content loss problem would produce more faithful and photorealistic few-shot image translation results.

But why does the content loss problem occur? To learn the translation in an unsupervised manner, Liu *et al.* [27] rely on inductive bias injected by the network design and adversarial training [10] to transfer the appearance from the example images in the unseen domain to the input image. However, as there is no supervision, it is difficult to control what to be transferred precisely. Ideally, the



**Fig. 2.** Few-shot image-to-image translation. **Training.** The training set consists of many domains. We train a model to translate images between these domains. **Deployment.** We apply the trained model to perform few-shot image translation. Given a few examples from a test domain, we aim to translate a content image into an image analogous to the test class.

transferred appearance should contain just the style. In reality, it often contains other information, such as the object pose.

In this paper, we propose a novel network architecture to counter the content loss problem. We design a style encoder called the *content-conditioned style encoder*, to hinder the transmission of task-irrelevant appearance information to the image translation process. In contrast to the existing style encoders, our style code is computed by conditioning on the input content image. We use a new architecture design to limit the variance of the style code. We conduct an extensive experimental validation with a comparison to the state-of-the-art method using several newly collected and challenging few-shot image translation datasets. Experimental results, including both automatic performance metrics and user studies, verify the effectiveness of the proposed method in dealing with the content loss problem.

## 2 Related Works

**Image-to-Image Translation.** Most of the existing models are based on the Generative Adversarial Network (GAN) [10] framework. Unlike unconditional GANs [10, 12, 19, 20, 30], which learn to map random vectors to images, existing image-to-image translation models are mostly based on conditional GANs where they learn to generate a corresponding image in the target domain conditioned on the input image in the source domain. Depending on the availability of paired input and output images as supervision in the training dataset, image-to-image translation models can be divided into supervised [4, 18, 29, 33, 35, 40, 41, 43, 44, 46, 48, 49] or unsupervised [1, 3, 5, 9, 16, 21, 23, 25, 26, 28, 34, 37, 39, 47]. Our work falls in the category of unsupervised image-to-image translation. However, instead of learning a mapping between two specific domains, we aim at learning a flexible mapping that can be used to generate images in many unseen domains. Specifically, the mapping is only determined at test time, via example images. When using example images from a different unseen domain, the same model can generate images in the new unseen domain.



**Fig. 3.** Illustration of the *content loss* problem. The images generated by the baseline [27] fail to preserve domain invariant appearance information in the content image. The animals’ bodies are sometimes merged with the background (column 3, & 4), scales of the generated body parts are sometimes inconsistent with the input (column 5), and some body parts absent in the content image show up (column 1 & 2). Our proposed method solves this “content loss” problem.

**Multi-domain Image Translation.** Several works [2, 5, 6, 17] extend the unsupervised image translation to multiple domains. They learn a mapping between multiple seen domains, simultaneously. Our work differs from the multi-domain image translation works in that we aim to translate images to *unseen* domains.

**Few-Shot Image Translation.** Several few-shot methods are proposed to generate human images [13, 38, 41, 42], scenes [41], or human faces [11, 42, 45] given a few instances and semantic layouts in a test time. These methods operate in the supervised setting. During training, they assume access to paired input (layout) and output data. Our work is most akin to the FUNIT work [27] as we aim to learn to generalize the translation to unseen domain without paired input and output data. We build on top of the FUNIT work where we first identify the *content loss* problem and then address it with a novel content-conditioned style encoder architecture.

**Example-Guided Image Translation** refers to methods that generate a translation of an input conditioning on some example images. Existing works

in this space [16, 27, 33] use a style encoder to extract style information from the example images. Our work is also an example-guided image translation method. However, unlike the prior works where the style code is computed independent of the input image, our style code is computed by conditioning on the input image, where we normalize the style code using the content to prevent over-transmission of the style information to the output.

**Neural Style Transfer** studies approaches to transfer textures from a painting to a real photo. While existing neural style transfer methods [8, 15, 24] can generalize to unseen textures, they cannot generalize to unseen shapes, necessary for image-to-image translation. Our work is inspired by these works, but we focus on generalizing the generation of both unseen shapes and textures, which is essential to few-shot unsupervised image-to-image translation.

### 3 Method

In this section, we start with a brief explanation of the problem setup, introduce the basic architecture, and then describe our proposed architecture. Throughout the paper, the two words, “class” and “domain”, are used interchangeably since we treat each object class as a domain.

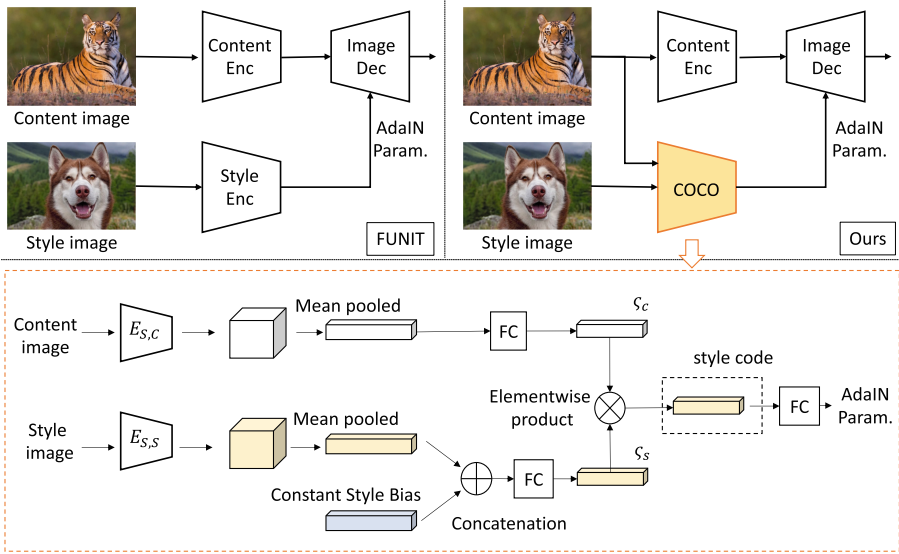
**Problem Setting.** Figure 2 provides an overview of the few-shot image translation problem [27]. Let  $X$  be a training set consists of images from  $K$  different domains. For each image in  $X$ , the class label is known. Note that we operate in the unsupervised setting where corresponding images between domains are *unavailable*. The few-shot image-to-image translation model learns to map a “content” image in one domain to an analogous image in the domain of the input “style” examples. In the test phase, the model sees a few example images from an unseen domain and performs the translation.

During training, a pair of content and style images  $x_c, x_k$  is randomly sampled. Let  $x_k$  denote a style image in domain  $k$ . The content image  $x_c$  can be from any domains in  $K$ . The generator  $G$  translates  $x_c$  into an image of class  $k$  ( $\bar{x}_k$ ) while preserving the content information of  $x_c$ .

$$\bar{x}_k = G(x_c, x_k) \tag{1}$$

In the test phase, the generator takes style images from a domain unseen during training, which we call the target domain. The target domain can be any related domain, not included in  $K$ .

**FUNIT Baseline.** FUNIT uses an example-guided conditional generator architecture as illustrated in the top-left of Fig. 4. It consists of three modules, 1) content encoder  $E_c$ , 2) style encoder  $E_s$ , and 3) image decoder  $F$ .  $E_c$  takes content image  $x_c$  as input and outputs content embedding  $z_c$ .  $E_s$  takes style image  $x_s$  as input and output style embedding  $z_s$ . Then,  $F$  generates an image using  $z_c$  and  $z_s$ , where  $z_s$  is used to generate the mean and scale parameters of adaptive instance normalization (AdaIN) layers [15] in  $F$ . The AdaIN design is based on



**Fig. 4. Top.** The FUNIT baseline [27] vs. our COCO-FUNIT. To highlight, we use a novel style encoder called the content-conditioned style encoder where the content image is also used in computing the style code for few-shot unsupervised image-to-image translation. **Bottom.** Detail design of the content-conditioned style encoder. Please refer to the main text for more details.

the assumption that the domain-specific information can be governed by the first and second order statistics of the activation and has been used in several GAN frameworks [16, 20, 27]. We further note that when multiple example/style images are present. FUNIT extracts a style code from each image and uses the average style code as the final input to  $F$ . To sum up, in FUNIT the image translation is formalized as follows,

$$z_c = E_c(x_c), \quad z_s = E_s(x_s), \quad \bar{x} = F(z_c, z_s). \quad (2)$$

**Content Loss.** As illustrated in Fig. 3, the FUNIT method suffers from the content loss problem—the translation result is not well-aligned with the input image. While a direct theoretical analysis is likely elusive, we conduct an empirical study, aiming at identify the cause of the content loss problem. As shown in Fig. 5, we compute different translation results of a content image based on a different style image where each of the style images is cropped from the same original style image. In the plot, we show variations of the deviation of the extracted style code due to different crops. Ideally, the plot should be constant as long as the crop covers sufficient appearance signature of the target class since that should be all required to generate a translation in the unseen domain. However, the FUNIT style encoder produces very different style codes as using different crops. Clearly, the style code contains other information about the style

image such as the object pose. We hypothesize this is the cause of the content loss problem and revisit the translator network design for addressing it.

**Content-Conditioned Style Encoder (COCO).** We hypothesize that the content loss problem can be mitigated if the style embedding is more robust to small variations in the style image. To this end, we design a new style encoder architecture, called the COntent-COnditioned style encoder (COCO). There are several distinctive features in COCO. The most obvious one is the conditioning in the content image as illustrated in the top-right of Fig. 4. Unlike the style encoder in FUNIT, COCO takes *both* content and style image as input. With this content-conditioning scheme, we create a *direct* feedback path during learning to let the content image influence how the style code is computed. It also helps reduce the direct influence of the style image to the extract style code.

The bottom part of Fig. 4 details the COCO architecture. First, the content image is fed into encoder  $E_{S,C}$  to compute a spatial feature map. This content feature map is then mean-pooled and mapped to a vector  $\zeta_c$ . Similarly, the style image is fed into encoder  $E_{S,S}$  to compute a spatial feature map. The style feature map is then mean-pooled and concatenated with an input-independent bias vector, which we refer to as the constant style bias (CSB). Note that while the regular bias in deep networks is added to the activations, in CSB, the bias is concatenated with the activations. The CSB provides a fixed input to the style encoder, which helps compute a style code that is less sensitive to the variations in the style image. In the experiment section, we show that the CSB can also be used to control the type of appearance information that is transmitted from the style image. When the CSB is activated, mostly texture-based appearance information is transferred. Note that the dimension of the CSB is set to 1024 through the paper.

The concatenation of the style vector and the CSB is mapped to a vector  $\zeta_s$  via a fully connected layer. We then perform an element-wise product operation to  $\zeta_c$  and  $\zeta_s$ , which is our final style code. The style code is then mapped to produce the AdaIN parameters for generating the translation. Through this element-wise product operation, the resulting style code is heavily influenced by the content image. One way to look at this mechanism is that it produces a customized style code for the input content image.

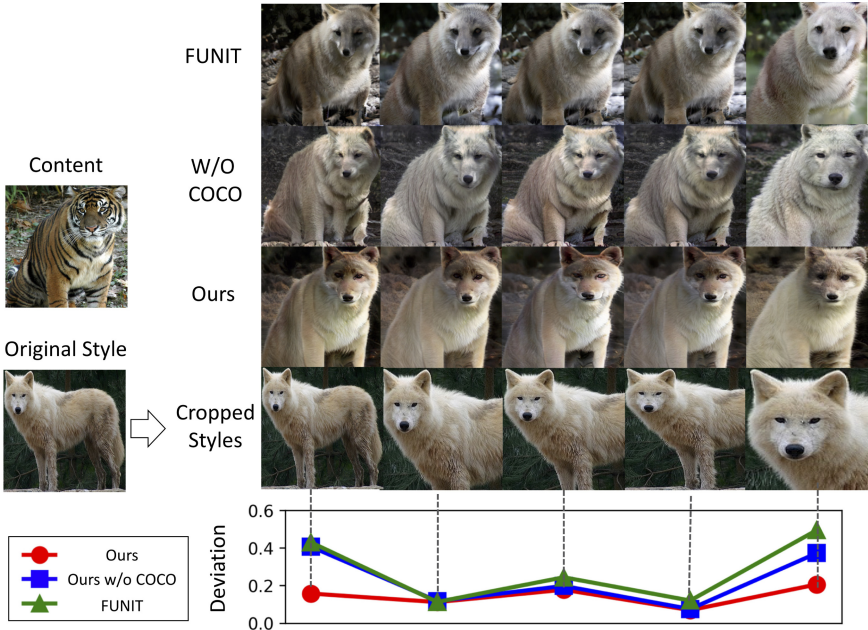
We use the COCO as a drop-in replacement for the style encoder in FUNIT. Let  $\phi$  denote the COCO mapping. The translation output is then computed via

$$z_c = E_c(x_c), z_s = \phi(E_{s,s}(x_s), E_{s,c}(x_c)), \bar{x} = F(z_c, z_s). \quad (3)$$

As shown in Fig. 5, the style code extracted by the COCO is more robust to variations in the style image. Note that we set  $E_{S,C} \equiv E_C$  to keep the number of parameters in our model similar to that in FUNIT.

We note that the proposed COCO architecture shows only one way to generate the style code conditioned on the content and to utilize the CSB. Certainly, there exist other design choices that could potentially lead to better translation performance. However, since this is the first time these two components are used for the few-shot image-to-image translation task, we focus on analyzing their

contribution in one specific design, i.e., our design. An exhaustive exploration is beyond the scope of the paper and is left for future work.



**Fig. 5.** We compare variations of the computed style codes due to variations in the style images for different methods. Note that for a fair comparison, in addition to the original FUNIT baseline [27], we create an improved FUNIT method by using our improved design for the content encoder, image decoder, and discriminator, which is termed “Ours w/o COCO”. “Ours” is our full algorithm where we use COCO as a drop-in replacement for the style encoder in the FUNIT framework. In the bottom part of the figure, we plot the variations of the style code due to using different crops of a style image. Specifically, the style code for each style image is first extracted for each method. We then compute the mean of the style codes for each method. The magnitudes of the deviations from the mean style code are then plotted. Note that to calibrate the network weights in different methods, all the style codes are first normalized by the mean extracted from 500 style images for each method. As shown in the figure, “Ours” produces more consistent translation outputs, which is a direct consequence of a more consistent style code extraction mechanism.

In addition to the COCO, we also improve the design of the content encoder, image decoder, and discriminator in the FUNIT work [27]. For the content encoder and image decoder, we find that replacing the vanilla convolutional layers in the original design with residual blocks [14] improves the performance so does replacing the multi-task adversarial discriminator with the project-based discriminator [32]. In Appendix D of our full technical report [36], we report their individual contribution to the few-shot image translation performance.



**Learning.** We train our model using three objective terms. We use the GAN loss ( $\mathcal{L}_{GAN}(D, G)$ ) to ensure the realism of the generated images given the class of the style images. We use the image reconstruction loss ( $\mathcal{L}_R(G)$ ) to encourage the model to reconstruct images when both the content and the style are from the same domain. We use the discriminator feature matching loss ( $\mathcal{L}_{FM}(G)$ ) to minimize the feature distance between real and fake samples in the discriminator feature space, which has the effect of stabilizing the adversarial training and contributes to generating better translation outputs as shown in the FUNIT work. In Appendix B of our full technical report [36], we detail the computation of each loss. Overall the objective is

$$\min_D \max_G \mathcal{L}_{GAN}(D, G) + \lambda_R \mathcal{L}_R(G) + \lambda_F \mathcal{L}_{FM}(G), \tag{4}$$



**Fig. 6.** Results on one-shot image-to-image translation. Column 1 & 2 are from the Carnivores dataset. Column 3 & 4 are from the Birds dataset. Column 5 & 6 are from the Mammals dataset. Column 7 & 8 are from the Motorbikes dataset.

**Table 1.** Results on the benchmark datasets.

Dataset	Method	mFID ↓	PAcc ↑	mIoU ↑	User style Preference ↑	User content Preference ↑
Carnivores	FUNIT	147.8	59.8	44.6	16.5	11.9
	Ours	<b>107.8</b>	<b>66.5</b>	<b>52.1</b>	<b>83.5</b>	<b>88.1</b>
Mammals	FUNIT	245.8	35.3	23.3	23.6	27.8
	Ours	<b>109.3</b>	<b>48.8</b>	<b>35.5</b>	<b>76.4</b>	<b>72.2</b>
Birds	FUNIT	89.2	52.4	37.2	38.5	37.5
	Ours	<b>74.6</b>	<b>53.3</b>	<b>38.3</b>	<b>61.5</b>	<b>62.5</b>
Motorbikes	FUNIT	275.0	85.6	73.8	17.8	17.4
	Ours	<b>56.2</b>	<b>94.6</b>	<b>90.3</b>	<b>82.2</b>	<b>82.6</b>

where  $\lambda_R$  and  $\lambda_F$  denote trade-off parameters for two losses. We set  $\lambda_R$  0.1 and  $\lambda_F$  1.0 in all of the experiments.



Fig. 7. Two-shot image translation results on the Carnivores dataset.



Fig. 8. Two-shot image translation results on the Birds dataset.

## 4 Experiments

We evaluate our method on challenging datasets that contain large pose variations, part variations, and category variations. Unlike existing few-shot image-to-image translation works, which focus on translations between reasonably-aligned images or simple objects, our interest is in the translations between likely misaligned images of highly articulate objects. Throughout the experiments, we use  $256 \times 256$  as our default image resolution for both inputs and outputs.




Fig. 9. Two-shot image translation results on the Mammals dataset.

**Implementation.** We use Adam [22] with  $lr = 0.0001$ ,  $\beta_1 = 0.0$ , and  $\beta_2 = 0.999$  for all methods. Spectral normalization [31] is applied to the discriminator. The final generator is a historical average version of the intermediate generators [19] where the update weight is 0.001. We train the model for 150,000 iterations in total. For every competing model, we compute the scores every 10,000 iterations and report the scores of the iteration that achieves the smallest mFID. Each training batch consists of 64 content images, which are evenly distributed on a DGX machine with 8 V100 GPUs, each with 32 GB RAM.

**Datasets.** We benchmark our method using 4 datasets. Each of the dataset contains objects with diverse poses, parts, and appearances.

- *Carnivores.* We build the dataset using images from the ImageNet dataset[7]. We pick up images from the 149 carnivorous animals and used 119 as the source/seen classes and 30 as the target/unseen classes.

**Table 2.** Ablation study on the Carnivores and Birds dataset. “Ours w/o CC” represents a baseline where the content conditioning part in COCO is removed. “Ours w/o CSB” represents a baseline where the CSB is removed. Detailed architecture of these baselines are given in Appendix A of our full technical report [36].

Method	Carnivores			Birds		
	mFID↓	PAcc↑	mIou ↑	mFID↓	PAcc↑	mIou ↑
Ours w/o COCO	<b>99.6</b>	62.5	47.8	<b>68.8</b>	52.8	37.9
Ours w/o CSB	107.1	61.8	46.9	74.1	52.5	37.7
Ours w/o CC	110.0	<b>66.7</b>	<b>52.1</b>	75.3	52.8	37.9
Ours	107.8	66.5	<b>52.1</b>	74.6	<b>53.3</b>	<b>38.3</b>

- *Mammals.* We collect 152 classes of herbivore animal images using Google image search and combine them with the Carnivores dataset to build the Mammals dataset. Out of the 301 classes, 236 classes are used for the source/seen and the rest is used for the target/unseen.
- *Birds.* We collect 205 classes of bird images using Google image search. 172 classes are used for training and the rest is used for the testing.
- *Motorbikes.* We also collected 109 classes of motorbike images in the same way. 92 classes are used as the source and the rest is used for the target.

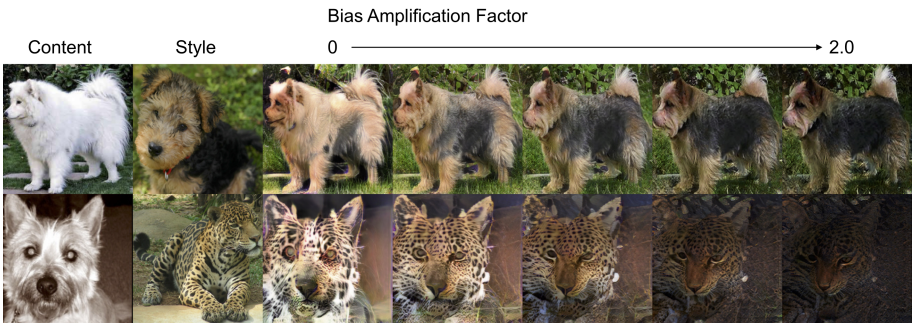
**Evaluation Protocol.** For each dataset, we train a model using the source classes mentioned above and test the performance on the target classes for each competing methods. In the test phase, we randomly sample 25,000 content images and pair each of them with a few style images from a target class to compute the translation. Unless specified otherwise, we use the one-shot setting for performance evaluation as it is the most challenging few-shot setting. We evaluate the quality of the translated images using various metrics as explained below.

**Performance Metrics.** Ideally, a translated image should keep the structure of the input content image, such as the pose or scale of body parts, unchanged when emulating the appearances of the unseen domain. Existing work mainly focused on the style transfer evaluation because the experiments are performed on well-aligned images or images of simple objects. To consider both the style translation and content preservation, we employ the following metrics. First, we evaluate the style transfer by measuring distance between the distribution of the translated images and the distribution of the real images in the unseen domain. Second, the content preservation is evaluated by measuring correspondence between a content and a translated image. Third, we conduct a user study to compute human preference scores on both the style transfer and content preservation of the translation results. The details of the performance metrics are given in Appendix C of our technical report [36].

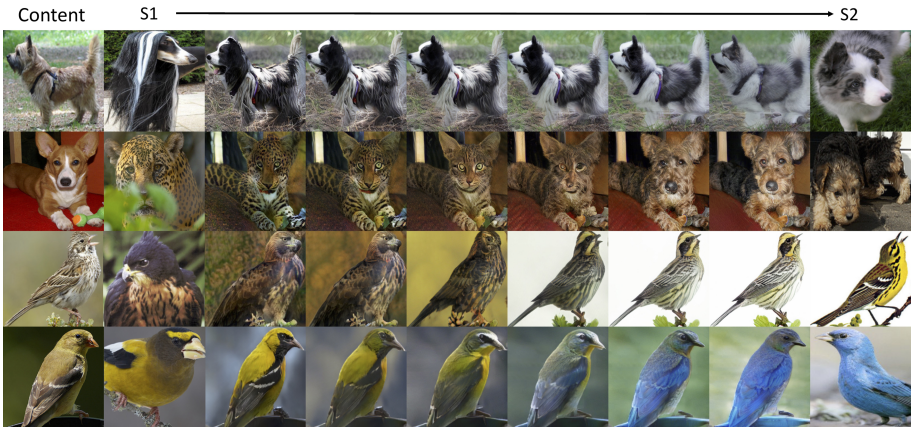
**Baseline.** We compare our method with the FUNIT method because it outperforms many baselines with a large margin as described in Liu *et al.* [27].

Therefore, a direct comparison with this baseline can verify the effectiveness of the proposed method for the few-shot image-to-image translation task.

**Main Results.** The comparison results is summarized in Table 1. As shown, our method outperforms FUNIT by a large margin in all the datasets on both automatic metrics and human preference scores. This validates the effectiveness of our method for few-shot unsupervised image-to-image translation. Figure 6 and 3 compare the one-shot translation results computed by the FUNIT method and our approach. We find images generated by the FUNIT method contain many artifacts while our method can generate photorealistic and faithful translation outputs. In Fig. 7, 8, and 9, we further visualize two-shot translation results. More visualization results are provided in the supplementary materials (Fig. 10).



**Fig. 10.** By changing the amplification factor  $\lambda$  of the CSB, our model generates different translation outputs for the same pair of content and style images.



**Fig. 11.** We interpolate the style codes from two example images from two different unseen domains. Our model can generate photorealistic results using these interpolated style codes. More results are in the supplementary materials.

**Ablation Study.** In Table 2, we ablate modules in our architecture and measure their impact on the few-shot translation performance using the Carnivores and Birds datasets. Now, let us walk through the results. First, we find using the CSB improve content preservation scores (“Ours w/o CSB” vs “Ours”), reflected by the better PAcc and mIoU scores achieved. Second, using content conditioning improves style transferring (“Ours w/o CC” vs “Ours”), reflected by the better mFID scores achieved. We also note that despite “Ours w/o COCO” achieves a better mFID, it is in the expense of large content loss.

**Effect of the CSB.** We conduct an experiment to understand how the CSB designed added to our COCO influences the translation results. Specifically, during testing, we multiply the CSB with a scalar  $\lambda$ . We then change the  $\lambda$  value to visualize its effect as shown in Fig. reffig:csbmanipulation. Interestingly, different values of  $\lambda$  generate different translation results. When the value is small, the model mostly changes the texture of the content image. With a large  $\lambda$  value, both the shape and texture are changed.

**Unseen Style Blending.** Here, we show an application where we combine two style images from two unseen domains to create a new unseen domain. Specifically, we first extract two style codes from two images from two different unseen domains. We then mix their styles by linear interpolating the style codes. The results are shown in Fig. 11 where the leftmost image is the content and row indicated by s1 and s2 are the two style images. We find the intermediate style codes render plausible translation results.

**Failure Cases.** While our approach effectively addresses the content loss problem, it still have several failure modes. We discuss these failure modes in the supplementary materials.

## 5 Conclusion

We introduced the COCO-FUNIT architecture, a new style encoder for few-shot image-to-image translation that extracts the style code from the example images from the unseen domain conditioning on the input content image and uses a constant style bias design. We showed that the COCO-FUNIT can effectively address the content loss problem, proven challenging for few-shot image-to-image-translation.

**Acknowledgements.** We would like to thank Jan Kautz for his insightful feedback on our project. Kuniaki Saito and Kate Saenko were supported by Honda, DARPA and NSF Award No. 1535797. Kuniaki Saito contributed to this work during his internship at NVIDIA.

## References

1. AlBahar, B., Huang, J.B.: Guided image-to-image translation with bi-directional feature transformation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)

2. Anoosheh, A., Agustsson, E., Timofte, R., Van Gool, L.: ComboGAN: unrestrained scalability for image domain translation. arXiv preprint [arXiv:1712.06909](https://arxiv.org/abs/1712.06909) (2017)
3. Benaim, S., Wolf, L.: One-shot unsupervised cross domain translation. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2018)
4. Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: *IEEE International Conference on Computer Vision (ICCV)* (2017)
5. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: StarGAN: unified generative adversarial networks for multi-domain image-to-image translation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
6. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: StarGAN v2: diverse image synthesis for multiple domains. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009)
8. Gatys, L.A., Ecker, A.S., Bethge, M.: Texture synthesis using convolutional neural networks. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2015)
9. Gokaslan, A., Ramanujan, V., Ritchie, D., Kim, K.I., Tompkin, J.: Improving shape deformation in unsupervised image-to-image translation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *ECCV 2018*. LNCS, vol. 11216, pp. 662–678. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01258-8\\_40](https://doi.org/10.1007/978-3-030-01258-8_40)
10. Goodfellow, I., et al: Generative adversarial networks. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2014)
11. Gu, Q., Wang, G., Chiu, M.T., Tai, Y.W., Tang, C.K.: LADN: local adversarial disentangling network for facial makeup and de-makeup. In: *IEEE International Conference on Computer Vision (ICCV)* (2019)
12. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
13. Han, X., Wu, Z., Wu, Z., Yu, R., Davis, L.S.: VITON: an image-based virtual try-on network. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
15. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: *IEEE International Conference on Computer Vision (ICCV)* (2017)
16. Huang, X., Liu, M.-Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *ECCV 2018*. LNCS, vol. 11207, pp. 179–196. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01219-9\\_11](https://doi.org/10.1007/978-3-030-01219-9_11)
17. Hui, L., Li, X., Chen, J., He, H., Yang, J.: Unsupervised multi-domain image translation with domain-specific encoders/decoders. arXiv preprint [arXiv:1712.02050](https://arxiv.org/abs/1712.02050) (2017)
18. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
19. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: *International Conference on Learning Representations (ICLR)* (2018)

20. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
21. Kim, T., Cha, M., Kim, H., Lee, J.K., Kim, J.: Learning to discover cross-domain relations with generative adversarial networks. In: International Conference on Machine Learning (ICML) (2017)
22. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR) (2015)
23. Lee, H.-Y., Tseng, H.-Y., Huang, J.-B., Singh, M., Yang, M.-H.: Diverse image-to-image translation via disentangled representations. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11205, pp. 36–52. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01246-5\\_3](https://doi.org/10.1007/978-3-030-01246-5_3)
24. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
25. Liang, X., Lee, L., Dai, W., Xing, E.P.: Dual motion GAN for future-flow embedded video prediction. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
26. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
27. Liu, M.Y., Huang, X., Mallya, A., Karras, T., Aila, T., Lehtinen, J., Kautz, J.: Few-shot unsupervised image-to-image translation. In: IEEE International Conference on Computer Vision (ICCV) (2019)
28. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2016)
29. Liu, X., Yin, G., Shao, J., Wang, X., et al.: Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2019)
30. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Smolley, S.P.: Least squares generative adversarial networks. In: IEEE International Conference on Computer Vision (ICCV) (2017)
31. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: International Conference on Learning Representations (ICLR) (2018)
32. Miyato, T., Koyama, M.: cGANs with projection discriminator. In: International Conference on Learning Representations (ICLR) (2018)
33. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
34. Pumarola, A., Agudo, A., Martinez, A.M., Sanfeliu, A., Moreno-Noguer, F.: GANimation: anatomically-aware facial animation from a single image. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11214, pp. 835–851. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01249-6\\_50](https://doi.org/10.1007/978-3-030-01249-6_50)
35. Qi, X., Chen, Q., Jia, J., Koltun, V.: Semi-parametric image synthesis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
36. Saito, K., Saenko, K., Liu, M.Y.: COCO-FUNIT: few-shot unsupervised image translation with a content conditioned style encoder. arXiv preprint [arXiv:2007.07431](https://arxiv.org/abs/2007.07431) (2020)
37. Shen, Z., Huang, M., Shi, J., Xue, X., Huang, T.S.: Towards instance-level image-to-image translation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)



38. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: Animating arbitrary objects via deep motion transfer. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
39. Taigman, Y., Polyak, A., Wolf, L.: Unsupervised cross-domain image generation. In: International Conference on Learning Representations (ICLR) (2017)
40. Wang, C., Zheng, H., Yu, Z., Zheng, Z., Gu, Z., Zheng, B.: Discriminative region proposal adversarial networks for high-quality image-to-image translation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11205, pp. 796–812. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01246-5\\_47](https://doi.org/10.1007/978-3-030-01246-5_47)
41. Wang, M., Yang, G.Y., Li, R., Liang, R.Z., Zhang, S.H., Hall, P.M., Hu, S.M.: Example-guided style-consistent image synthesis from semantic labeling. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
42. Wang, T.C., Liu, M.Y., Tao, A., Liu, G., Kautz, J., Catanzaro, B.: Few-shot video-to-video synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2019)
43. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2018)
44. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional GANs. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
45. Zakharov, E., Shysheya, A., Burkov, E., Lempitsky, V.: Few-shot adversarial learning of realistic neural talking head models. In: IEEE International Conference on Computer Vision (ICCV) (2019)
46. Zhao, B., Meng, L., Yin, W., Sigal, L.: Image generation from layout. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
47. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: IEEE International Conference on Computer Vision (ICCV) (2017)
48. Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Toward multimodal image-to-image translation. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
49. Zhu, S., Urtasun, R., Fidler, S., Lin, D., Change Loy, C.: Be your own prada: fashion synthesis with structural coherence. In: IEEE International Conference on Computer Vision (ICCV) (2017)