

Example-Guided Image Synthesis using Masked Spatial-Channel Attention and Self-Supervision

Haitian Zheng Haofu Liao Lele Chen
Wei Xiong Tianlang Chen Jiebo Luo

University of Rochester,
{hzheng15, hliao6, lchen63, wxiong5, tchen45, jluo}@cs.rochester.edu

Abstract. Example-guided image synthesis has recently been attempted to synthesize an image from a semantic label map and an exemplary image. In the task, the additional exemplar image provides the style guidance that controls the appearance of the synthesized output. Despite the controllability advantage, the existing models are designed on datasets with specific and roughly aligned objects. In this paper, we tackle a more challenging and general task, where the exemplar is a scene image that is semantically different from the given label map. To this end, we first propose a Masked Spatial-Channel Attention (MSCA) module which models the correspondence between two scenes via efficient decoupled attention. Next, we propose an end-to-end network for joint global and local feature alignment and synthesis. Finally, we propose a novel self-supervision task to enable training. Experiments on the large-scale and more diverse COCO-stuff dataset show significant improvements over the existing methods. Moreover, our approach provides interpretability and can be readily extended to other content manipulation tasks including style and spatial interpolation or extrapolation.

Keywords: Example-guided image synthesis · Self-supervised learning · Correspondence modeling · Efficient attention

1 Introduction

Conditional generative adversarial network (cGAN) [35] has recently made substantial progress in realistic image synthesis. In cGAN, a generator $\hat{x} = G(c, z)$ aims to output a realistic image \hat{x} with a constraint implicitly encoded by c . Conversely, a discriminator $D(x, c)$ learns such a constraint from ground-truth pairs $\langle x, c \rangle$ by predicting if $\langle \hat{x}, c \rangle$ is real or generated.

The current cGAN models [37,44,21] for semantic image synthesis aim to solve the *structural consistency* constraint where the output image $\hat{x} = G(c)$ is required to be aligned to a semantic label map c . However, for such a model, the style of \hat{x} is inherently determined by the model and thus cannot be controlled by the user. To provide desired controllability over the generated styles, previous studies [28,42] impose additional constraints and allow more inputs to the generator: $\hat{x}_{2 \rightarrow 1} = G(c_1, x_2, z)$, where x_2 is an exemplar image that guides

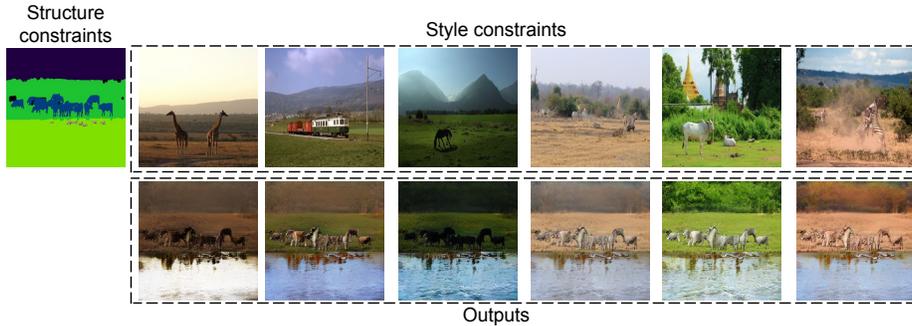


Fig. 1. Our task aims to synthesize style-consistent images from a semantic label map (row 1, column 1) and a **structurally and semantically different** exemplar image (row 1, columns 2-7). In spite of the differences between the two scenes, our model can synthesize high-quality images of consistent styles with the reference images.

the style of c_1 . However, previous studies are designed on specified datasets such as face [31,38], dancing [42] or street view [47], where the exemplar images and semantic label map usually contain similar semantics and spatial structures.

Different from the previous studies, we address a more challenging example-guided synthesis task that transfers styles *across semantically very different scenes*. As shown in Fig. 1, given a semantic label map c_1 (column 1) and an arbitrary scene image x_2 (row 1, column 2-7), the task aims to generate a new scene image $\hat{x}_{2 \rightarrow 1}$ (row 2) that matches the semantic structure of c_1 and the scene style of x_2 . The challenge is that scene images have complex semantic structures as well as diversified scene styles, and more importantly, the inputs c_1 and x_2 can be *structurally unaligned* and *semantically different*. Therefore, a mechanism is required to better match the structures and semantics for coherent synthesis.

In this paper, we propose a novel Masked Spatial-Channel Attention (MSCA) module (Section 3.2) to propagate features across unstructured scenes. Our module is inspired by a recent work [7] for attention-based object recognition, but instead, we propose a new cross-attention mechanism to model the semantic correspondence for image synthesis. Moreover, our method is based on the novel design of spatial-channel decoupling that allows efficient computation. To facilitate example-guided synthesis, we further improve the module by including: i) feature masking for semantic outlier filtering, ii) multi-scaling for global-local feature processing, and iii) resolution extending for image synthesis. As a result, our module provides both clear physical meaning and interpretability for the example-guided synthesis task.

We formulate the proposed approach under a unified synthesis network for joint feature extraction, alignment and image synthesis. We achieve this by applying MSCA modules to the extracted features for multi-scale feature domain alignment. Next, we apply a recent feature normalization technique, SPADE [37] on the aligned features to allow spatially-controllable synthesis. To facilitate the

learning of this network, we propose a novel self-supervision task. As opposed to [42], our scheme requires only semantically parsed images for training and does not rely on video data. We show that a model trained with this approach generalizes across different scene semantics (See Fig. 1).

Our main contributions include the following:

- A novel masked spatial-channel attention (MSCA) module to propagate features between two semantically different scenes.
- A unified example-guided synthesis network for joint feature extraction, alignment and image synthesis.
- A novel self-supervision scheme that only requires semantically annotated images for training but not at the testing (image synthesis) stage.
- Significant improvements over the existing methods on the COCO-stuff [3] dataset, as well as interpretability and easy extensions to other content manipulation tasks.

2 Related work

Generative Adversarial Networks Recent years have witnessed the progress of generative adversarial networks (GANs) [11] for image synthesis. A GAN model consists of a generator and a discriminator where the generator serves to produce realistic images that cannot be distinguished from the real ones by the discriminator. Recent techniques for realistic image synthesis include modified losses [1,34,39], model regularization [36], self-attention [49,2], feature normalization [25] and progressive synthesis [24].

Image-to-Image translation (I2I) I2I translation aims to translate images from a source domain to a target domain. The initial work of Isola *et al.* [21] proposes a conditional GAN framework to learn I2I translation with paired images. Wang *et al.* [44] improve the conditional GAN for high-resolution synthesis and content manipulation. To enable I2I translation without using paired data, a few works [51,30,18,26,4] apply the cycle consistency constraint in training. Recent works on photo-realistic image synthesis take semantic label maps as inputs for image synthesis. Specifically, Wang *et al.* [44] extend the conditional GAN for high-resolution synthesis, Chen *et al.* [6] propose a cascade refinement pipeline. More recently, Park *et al.* [37] propose spatial-adaptive normalization for realistic image generation.

Example-Guided Style Transfer and Synthesis Example guided style transfer [13,8] aims to transfer the style of an example image to a target image. Recent works [10,17,32,23,27,12,4,16,46] utilize deep neural network features to model and transfer styles. Several frameworks [18,19,33] perform style transfer via image domain style and content disentanglement. In addition, domain adaptation [4] applies a cycle consistency loss to perform cross-domain style transformation.

More recently, example-guided synthesis [28,42] is proposed to transfer the style of an example image to a target condition, e.g. a semantic label map. Specifically, Lin *et al.* [28] apply dual learning to disentangle the style for guided

synthesis, Wang *et al.* [42] extract style-consistent data pairs from videos for model training. In addition, Park *et al.* [37] adopt an I2I network under the auto-encoding framework for example-guided image synthesis. Different from [28,42,37], we address style alignment issue between arbitrary scenes for *region and semantic aware* style integration. Furthermore, our self-supervised learning scheme does not require video data and is a generalize and more challenging auto-encoding task.

Correspondence Matching for Synthesis Finding correspondence is critical for many synthesis tasks. For instance, Siarohin *et al.* [40] apply the affine transformation on reference person images to improve pose-guided person image synthesis, Wang *et al.* [43] use optical flow to align frames for coherent video synthesis. However, the affine transformation and optical flow cannot adequately model the correspondences between two structurally very different scenes.

Efficient Attention Modeling The self-attention [45,49] can capture general pair-wise correspondences. However, it is computationally intensive at high-resolution. To enable fast attention computation, GCNL [48] and CCCA [20] respectively apply Taylor series expansion and criss-cross attention to approximate self-attention. Alternatively, A^2 -Nets [7] factorize self-attention to solve video classification tasks. Inspired by [7], we propose an attention-based module named MSCA. It is worth noting MSCA is based on cross-attention and feature masking for modeling image correspondence.

3 Method

The proposed approach aims to generate scene images that align with given semantic maps. Differ from conventional semantic image synthesis methods [21,44,37], our model takes an exemplary scene as an extra input to provide more controllability over the generated scene image. Unlike existing example-based approaches [28,42], our model addresses a more challenging case where the exemplary inputs are structurally and semantically unaligned with the given semantic map.

Our method takes a semantic label map c_1 , a reference image x_2 and its corresponding parsed semantic label map \tilde{c}_2 as inputs and synthesizes an image $\hat{x}_{2 \rightarrow 1}$ which matches the style of x_2 and structure of x_1 using a generator G , $\hat{x}_{2 \rightarrow 1} = G(c_1, x_2, \tilde{c}_2)$. As shown in Fig. 2 left, the generator G consists of three parts, namely i) feature extraction ii) feature alignment and iii) image synthesis. In Sec. 3.1, we describe the first part that extracts features from inputs of both scenes. In Sec. 3.2, we propose a masked spatial-channel attention (MSCA) module to distill features and discover relations between two arbitrarily structured scenes. Unlike the affine-transformation [22] and flow-base warping [43], MSCA provides better interpretability to the scene alignment task. In Sec. 3.3, we introduce how to use the aligned features for image synthesis. Finally, in Sec. 3.4, we propose a self-supervised scheme to facilitate learning.

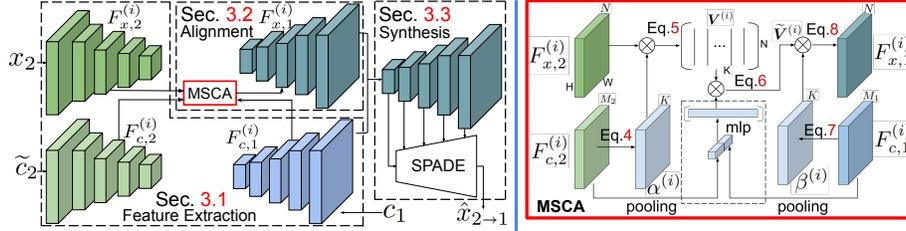


Fig. 2. **Left:** our generator consists of three steps, namely *feature extraction*, *feature alignment*, and *image synthesis*. We describe each step in its corresponding section, respectively. **Right:** The MSCA module for feature alignment (at scale i). Our module takes image feature map $F_{x,2}^{(i)}$ and segmentation feature map $F_{c,1}^{(i)}, F_{c,2}^{(i)}$ as inputs to output a new image feature map $F_{x,1}^{(i)}$ that is aligned to condition c_1 .

3.1 Feature Extraction

Taking an image x_2 and label maps c_1, \tilde{c}_2 as inputs, the feature extraction module extracts multi-scale feature maps for each input. Specifically, the feature map $F_{x,2}^{(i)}$ of image x_2 at scale i is computed by:

$$F_{x,2}^{(i)} = W_x^{(i)} * F_{\text{vgg}}^{(i)}(x_2), \quad \text{for } i \in \{0, \dots, L\}, \quad (1)$$

where $*$ denotes the convolution operation, $F_{\text{vgg}}^{(i)}$ denotes the feature map extracted by VGG-19 [41] at scale i , and $W_x^{(i)}$ denotes a 1×1 convolutional kernel for feature compression. L is the number scales and we set $L = 4$ in this paper.

For label map c_1 , its feature $F_{c,1}^{(i)}$ is computed by:

$$F_{c,1}^{(i)} = \begin{cases} \text{LReLU}(W_c^{(i)} * c_1^{(i)}) & \text{for } i = L, \\ \text{LReLU}(W_c^{(i)} * [\uparrow(F_{c,1}^{(i+1)}), c_1^{(i)}]) & \text{otherwise,} \end{cases} \quad (2)$$

where $\uparrow(\cdot)$ denotes $\times 2$ bilinear interpolation, $c_1^{(i)}$ denotes the resized label map, $W_c^{(i)}$ denotes a 1×1 convolutional kernel for feature extraction, and operation $[\cdot, \cdot]$ denotes channel-wise concatenation. Note that as scale i decreases from L down to 0, the feature resolutions in Eq. 2 are progressively increased to match a finer label maps $c_1^{(i)}$.

Similarly, applying Eq. 2 with the same weights to label map \tilde{c}_2 , we can extract its features $F_{c,2}^{(i)}$:

$$F_{c,2}^{(i)} = \begin{cases} \text{LReLU}(W_c^{(i)} * \tilde{c}_2^{(i)}) & \text{for } i = L \\ \text{LReLU}(W_c^{(i)} * [\uparrow(F_{c,2}^{(i+1)}), \tilde{c}_2^{(i)}]) & \text{otherwise} \end{cases} \quad (3)$$

3.2 Masked Spatial-channel Attention Module

As shown in Fig. 2 right, taking the image features $F_{x,2}^{(i)}$ and the label map features $F_{c,1}^{(i)}, F_{c,2}^{(i)}$ as inputs¹, the MSCA module generates a new image feature map $F_{x,1}^{(i)}$ that has the content of $F_{x,2}^{(i)}$ but is aligned with $F_{c,1}^{(i)}$. We elaborate the detailed procedures as follows:

Spatial Attention. Given feature maps $F_{x,2}^{(i)}, F_{c,2}^{(i)}$ of the exemplar scene, the module first computes a spatial attention tensor $\alpha^{(i)} \in [0, 1]^{K \cdot H \cdot W}$:

$$\alpha^{(i)} = \text{softmax}_{2,3}(\phi^{(i)} * [F_{x,2}^{(i)}, F_{c,2}^{(i)}]), \quad (4)$$

with $\phi^{(i)} \in \mathbb{R}^{(N+M_2) \cdot K}$ denoting a 1×1 convolutional filter and $\text{softmax}_{2,3}$ denoting a 2D softmax function on spatial dimensions $\{2, 3\}$. The output tensor contains K attention maps of resolution $H \times W$, which serve to attend K different spatial regions on image feature $F_{x,2}^{(i)}$.

Spatial Aggregation. Then, the module aggregates K feature vectors from $F_{x,2}^{(i)}$ using the K spatial attention maps of $\alpha^{(i)}$ from Eq. 4. Specifically, a matrix dot product is performed:

$$\mathbf{V}^{(i)} = \mathbf{F}_{x,2}^{(i)}(\boldsymbol{\alpha}^{(i)})^\top, \quad (5)$$

with $\boldsymbol{\alpha}^{(i)} \in [0, 1]^{K \cdot HW}$ and $\mathbf{F}_{x,2}^{(i)} \in \mathbb{R}^{N \cdot HW}$ denoting the reshaped versions of $\alpha^{(i)}$ and $F_{x,2}^{(i)}$, respectively. The output $\mathbf{V}^{(i)} \in \mathbb{R}^{N \cdot K}$ stores feature vectors spatially aggregated from the K independent regions of $F_{x,2}^{(i)}$.

Feature Masking. The exemplar scene x_2 may contain irrelevant semantics to the label map c_1 , and conversely, c_1 may contain semantics that are unrelated to x_2 . To address this issue, we apply feature masking on the output of Eq. 5 by multiplying $\mathbf{V}^{(i)}$ with a length- K gating vector at each row:

$$\tilde{\mathbf{V}}^{(i)} = (\mathbf{V}^{(i)})^T \circ \text{mlp}([\text{gap}(F_{c,1}^{(i)}), \text{gap}(F_{c,2}^{(i)})]), \quad (6)$$

where $\text{mlp}(\cdot)$ denotes a 2-layer MLP followed by a sigmoid function, gap denotes a global average pooling layer, \circ denotes broadcast element-wise multiplication, and $\tilde{\mathbf{V}}^{(i)}$ denotes the masked features. The design of feature masking in Eq. 6 resembles to Squeeze-and-Excitation [15]. Using the integration of global information from label maps c_1 and \tilde{c}_2 , features are filtered.

Channel Attention. Given feature $F_{c,1}^{(i)}$ of label map c_1 , a channel attention tensor $\beta^{(i)} \in [0, 1]^{K \cdot H \cdot W}$ is generated as follows:

$$\beta^{(i)} = \text{softmax}_1(\psi^{(i)} * F_{c,1}^{(i)}), \quad (7)$$

with $\psi^{(i)} \in \mathbb{R}^{M_1 \cdot K}$ denoting a 1×1 convolutional filter and softmax_1 denoting a softmax function on channel dimension. The output $\beta^{(i)}$ serves to dynamically reuse features from $\tilde{\mathbf{V}}^{(i)}$.

¹ We assume spatial resolution at scale i being $H \times W$ and channel size of $F_{x,2}^{(i)}, F_{c,1}^{(i)}, F_{c,2}^{(i)}$ being N, M_1, M_2 , respectively.

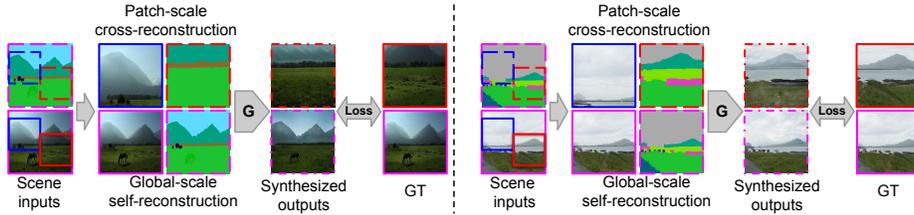


Fig. 3. Our self-supervision scheme performs cross-reconstruction at the patch scale (top row) and self-reconstruction at the global scale (bottom row). The solid, dashed and dotted bounding boxes respective represent images, semantic label maps, and synthesized outputs. Boxes with the same color are cropped from the same position.

Channel Aggregation. With channel attention $\beta^{(i)}$ computed in Eq. 7, feature vectors at HW spatial locations are aggregated again from $\tilde{\mathbf{V}}^{(i)}$ via matrix dot product:

$$\mathbf{F}_{x,1}^{(i)} = \tilde{\mathbf{V}}^{(i)}(\beta^{(i)})^\top, \quad (8)$$

where $\beta^{(i)} \in \mathbb{R}^{K \cdot HW}$ denotes the reshaped version of $\beta^{(i)}$. The output $\mathbf{F}_{x,1}^{(i)} \in \mathbb{R}^{N \cdot HW}$ represents the aggregated features at HW locations. The output feature map $F_{x,1}^{(i)}$ is generated by reshaping $\mathbf{F}_{x,1}^{(i)}$ to size $N \times H \times W$.

Remarks. Spatial attention (Eq. 4) and aggregation (Eq. 5) attend to K independent regions from feature $F_{x,2}^{(i)}$, then store the K features into $\mathbf{V}^{(i)}$. After feature masking, given a new label map c_1 , channel attention (Eq. 4) and aggregation (Eq. 8) combine $\tilde{\mathbf{V}}^{(i)}$ at each location to compute an output feature map. As a result, each output location finds its correspondent regional features or ignored via feature masking. In this way, the feature of example scene is aligned. Note that when $K = 1$ and $\alpha^{(i)}$ is constant, the above operations is essentially a global average pooling. We show in the experiment that $K = 8$ is sufficient to dynamically capture visually significant scene regions for alignment.

Multi-scaling. Both global color tone and local appearances are informative for the style-constraint synthesis. Therefore, we apply MSCA modules at all scales $i \in \{0, \dots, L\}$ to generate global and local features $F_{x,1}^{(i)}$.

3.3 Image Synthesis

The extracted features $F_{c,1}^{(i)}$ in Sec. 3.1 capture the semantic structure of c_1 , whereas the aligned features $F_{x,1}^{(i)}$ in Sec. 3.2 capture the appearance style of the example scene. In this section, we leverage $F_{c,1}^{(i)}$ and $F_{x,1}^{(i)}$ as control signals to generate output images with desired structures and styles.

Specifically, we adopt a recent synthesis model, SPADE [37], and feed the concatenation of $F_{x,1}^{(i)}$ and $F_{c,1}^{(i)}$ to the spatially-adaptive denormalization layer of SPADE at each scale. By taking the style and structure signal as inputs, spatially-controllable image synthesis is achieved. We refer readers to appendix for more network details of the synthesis module.

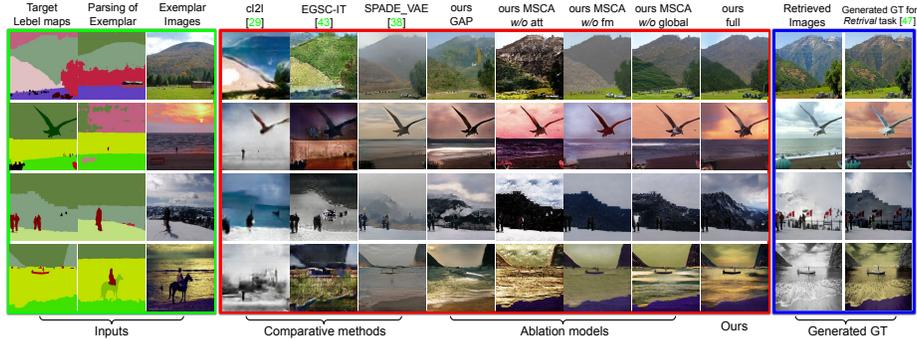


Fig. 4. Green box from left to right: the inputs for example-guided synthesis, i.e. target label maps, exemplar label parsing from Deeplab-v2 [5], and exemplar images. Red box from left to right: visual comparisons with ci2I [28], EGSC-IT [33], SPADE_VAE [37], four ablation models, and our full model. Blue box from left to right: the retrieved ground-truth before and after color correction [46]. Our full model generates the most style-consistent results with the exemplar images.

3.4 Self-Supervised Training

Training an example-guided synthesis model that can transfer styles across semantically different scenes is challenging. First, style-consistent scene images are hard to acquire. A previous work [42] generates style-consistent pairs from videos. However, collecting scene videos can be more labor intensive. Second, even with ground truth style-consistent pairs, the trained model is not guaranteed to generalize to a new arbitrary scene.

We propose a novel self-supervised scheme to enable style-transfer between two structurally and semantically different scenes. Our solution is motivated by the fact that the style of a scene image is stationary, meaning that patches cropped from the same scene share largely the same style. Moreover, non-overlapping patches from the same scene may contain new structures and semantic labeling, which is essential for the learned model to generalize better.

We first design a *cross-reconstruction* task at the patch scale: given patches x_p and x_q cropped from the same scene image x , the generator is asked to reconstruct x_p using x_q . Formally,

$$\hat{x}_p = G(c_p, x_q, \tilde{c}_q). \quad (9)$$

Note that c_p and \tilde{c}_q contain different semantic labeling. Therefore, the generator are required to infer the correlation between different semantic labeling for coherent style transfer. An illustrative example is shown in Fig. 3. More details on patch sampling is included in the appendix.

The cross-reconstruction task is designed at the patch scale and may not generalize well to the global scale. In fact, the generator trained with the patch-level task alone tends to generate repetitive local textures (in Sec. 4). Therefore, we further design a *self-reconstruction* task at the global scale, which reconstructs a global image x from itself:

$$\hat{x} = G(c, x, \tilde{c}). \quad (10)$$

Our training objective for generator G and discriminator D is formulated as:

$$\begin{aligned} \mathcal{L}(G, D) = & \log D(x_p, c_p, x_q, \tilde{c}_q) + \log(1 - D(\hat{x}_p, c_p, x_q, \tilde{c}_q)) + \mathcal{L}_{spade}(\hat{x}_p, x_p) \\ & + \lambda \{ \log D(x, c, x, \tilde{c}) + \log(1 - D(\hat{x}, c, x, \tilde{c})) + \mathcal{L}_{spade}(\hat{x}, x) \} \end{aligned} \quad (11)$$

where \mathcal{L}_{spade} refers to the VGG and GAN feature matching losses defined in [37] and λ is a parameter that controls the importance of the two self-supervised tasks. We set $\lambda = 1$ in our experiments. Our full objective for self-supervised training is:

$$G^* = \arg \min_G \max_D \mathcal{L}(G, D) \quad (12)$$

4 Experiments

Dataset Our model is trained on the *COCO-stuff* dataset [3]. It contains densely annotated images captured from various scenes. We remove images with indoor scenes and large objects from the dataset, resulting in 34,698/499 scene images for training/testing, respectively. The COCO-stuff dataset does not provide ground-truth for the example-guided synthesis task, i.e. two images with the exact same styles. To qualitatively evaluate the performances, we designed three tasks where the ground-truth image can be obtained: i) *duplicating task* self-reconstructs an image using itself as exemplar and its semantic label map as layout condition, ii) *mirroring task* reconstructs an image using its mirrored version as exemplar and the semantic label map as layout condition, iii) *retrieving task*: requires a model to reconstruct a ground-truth (GT) image using its semantic label map and a retrieved image from an image pool. To retrieve an image that best matches GT in styles, we first select 20 candidate images from the image pool that have the greatest label histogram intersections with the GT image. Afterwards, the best-matched image is selected out of candidates using SIFT Flow [29]. Finally, since the color of GT is different from the retrieved image, we apply color correction [46] on GT to eliminate color discrepancy. Examples of GT before and after color correction are shown in the blue box in Fig. 4.

Implementation Details We use a COCO-stuff pretrained Deeplab-v2 [5] model to generate semantic label maps from exemplar images. During training, we resize images to 512×512 then crop two non-overlapping patches of size 256×256 to facilitate patch-based cross-reconstruction. After 20 epochs, we increase the patch size to 384×384 for cross-patch reconstruction in order to improve generalization to global scenes. Details of the patch sampling procedure are provided in the appendix.

For the MSCA modules from scale 0 to 4, the number of attention maps K are respectively set to 8, 16, 16, 16, 16. The learning rate is set to 0.0002 for the generator and the discriminator. The weights of generator are updated every 5 iterations. Our synthesis model and all comparative models based on SPADE backbone are trained for 40 epochs to generate the results in the experiments.

task	measures	cI2I [28]	EGSC-IT [42]	SPADE-VAE [37]	ours GAP	ours MSCA w/o att	ours MSCA w/o fm	ours MSCA w/o global	ours full
retrieving	PSNR \uparrow	9.50	12.57	15.77	15.85	11.96	16.24	15.98	16.65
	LPIPS \downarrow	0.757	0.581	0.483	0.457	0.522	0.451	0.446	0.437
	FID \downarrow	228.63	163.23	102.68	101.74	112.83	100.01	96.66	91.91
	$\mathcal{L}_{style}\downarrow$	3.53e-3	1.69e-3	1.07e-3	6.40e-4	7.10e-3	7.62e-4	7.21e-4	5.34e-4
mirroring	PSNR \uparrow	9.46	12.44	15.37	15.80	11.95	16.02	16.58	17.03
	LPIPS \downarrow	0.759	0.602	0.477	0.438	0.510	0.437	0.421	0.397
	FID \downarrow	242.73	190.01	90.99	89.15	102.41	90.52	85.92	76.75
	$\mathcal{L}_{style}\downarrow$	4.14e-3	2.03e-3	1.67e-3	7.45e-4	7.99e-3	8.76e-4	6.69e-4	3.96e-4
duplicating	PSNR \uparrow	9.46	12.45	15.43	15.80	11.94	16.02	16.62	17.03
	LPIPS \downarrow	0.759	0.602	0.476	0.438	0.510	0.438	0.421	0.397
	FID \downarrow	242.81	190.02	90.97	89.22	103.23	90.56	86.20	76.64
	$\mathcal{L}_{style}\downarrow$	4.15e-3	2.03e-3	1.66e-3	7.41e-4	7.94e-3	8.81e-4	6.53e-4	3.97e-4

Table 1. Quantitative comparisons of different methods and ablation models in terms of PSNR, LPIPS [50], Fréchet Inception Distance (FID) [14] and style loss (\mathcal{L}_{style}) [9]. Higher scores are better for metrics with uparrow (\uparrow), and vice versa.



Fig. 5. With a slight modification and no further training, our model can perform style interpolation between exemplar inputs. Note how our model interpolates styles for new semantics, e.g. “river” in row 3.

Before training, we pretrain the spatial-channel attention with a lightweight feature decoder to improve training efficiency. Specifically, at each scale, the concatenation of $F_{x,1}^{(i)}$ and $F_{c,1}^{(i)}$ in Sec. 3.3 at each scale is fed into a 1×1 convolutional layer to reconstruct the ground-truth VGG feature at the corresponding scale. More details of the pretraining procedure is provided in the appendix.

Comparative Methods We compare our approach with an example-guided synthesis approach: variational autoencoding SPADE (SPADE.VAE) [37] which is based on a self-reconstruction loss for training. We also trained cI2I [28], EGSC-IT [33] and SCGAN [42] on COCO-stuff dataset. cI2I and EGSC-IT are originally designed for exemplar-guided image-to-image translation. As a result, we observed that cI2I and EGSC-IT have difficulty generating images from one-hot encoded semantic label maps. However, these models can synthesize reasonable images from color-encoded semantic label maps. Finally, we note that SCGAN is not directly applicable to COCO-stuff dataset, as its positive pairs are sampled from video data. We attempted to modify SCGAN such that its positive pairs can be generated from our self-supervision task. However, we could not achieve reasonable image outputs. We speculate that the negative sampling and semantic

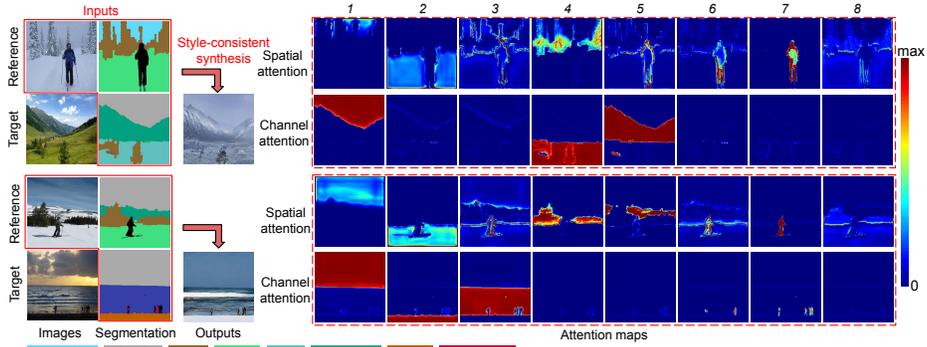


Fig. 6. Left: inputs and outputs of our model. Right: the $K = 8$ learned spatial and channel attention that attends and transfer feature between individual exemplar and target regions. By examining the semantics label maps, we observe the following transformation patterns: `sky_other` \rightarrow `clouds`, `tree` \rightarrow `{tree, hill}` for the 1st sample, and `clouds` \rightarrow `clouds`, `snow` \rightarrow `sand`, `other` \rightarrow `{surfboard, other}` for the 2nd sample.

consistency loss of `SCGAN` is not optimal for `COCO-stuff` dataset, as `COCO-stuff` dataset contains much larger variations for negative pairs. Finally, four ablation models are evaluated (see Ablation Study).

Quantitative Evaluation For quantitative evaluation, we apply PSNR as the low-level metric. Furthermore, perceptual-level metrics including Perceptual Image Patch Similarity Distance (LPIPS) [50], Fréchet Inception Distance (FID) [14] and style loss (\mathcal{L}_{style}) of [9] are evaluated on different methods. The linearly calibrated VGG model is used to compute LPIPS distance.

Among the four competitive methods (`cI2I`, `EGSC-IT`, `SPADE_VAE` and `ours full`) in Table 1, our method clearly outperforms the remaining methods both in low-level and perceptual-level measurements, suggesting that our model can better preserve color and texture appearances. Also, we observe that without further modification, the off-the-shelf example-guided image translation approaches cannot perform well on image synthesis tasks (`cI2I`, `EGSC-IT`). It suggests that example-guided image-synthesis task can be more challenging. Finally, a simple synthesis model (`ours GAP`) outperforms `SPADE_VAE`, suggesting that the self-supervised task in Sec. 3.4 is beneficial to the example-guided synthesis task (see Ablation Study for more details).

Qualitative Evaluation Fig. 4 qualitatively compares our approach against the remaining approach on four scenes. We observe that our full model generates more style-consistent results with the exemplar images. In comparisons, `SPADE_VAE` tends to generate results with low color contrast, as it lacks the mechanism and supervision to perform region-aware style transformation. In addition, the existing example-guided image-to-image approaches (`cI2I`, `EGSC-IT`) cannot generalize well to the image synthesis tasks.

Ablation Study To evaluate the effectiveness of our design, we separately train four variants of our model: i) `our GAP` that replaces the `MSCA` module

with global average pooling, ii) **ours MSCA w/o att** that keeps MSCA modules but replaces spatial and channel attention with one-hot label maps from source and target domains, respectively. In such a way, alignment is performed only for regions with the same semantic labeling, iii) **ours MSCA w/o fm** that keeps MSCA modules but removes the feature masking procedures, and iv) **ours MSCA w/o global** that is trained without using global-level self-reconstruction (Eq. 12) or increased patchsize.

In Table 1, our full model clearly achieves the best qualitative results. In Fig. 4, **ours GAP** tends to produce images with deviated colors since it averages the style features from all exemplar regions. In contrast, our model dynamically transfers appearance for individual regions. We observe that **ours w/o att** is less stable in training and cannot generate plausible results. We suspect that the label-level alignment generates more misaligned and noisier feature maps, thus hurting training. **ours MSCA w/o fm** tends to generate inconsistent colors for new semantic labels, for instance, the “hill” and “sky” regions in rows 1 and 2 of Fig. 4. In contrast, our model can eliminate the undesired influence of exemplar inputs on new semantic labels. **ours MSCA w/o global** performs reasonably well but it tends to generate repetitive local textures, while the self-reconstruction scheme helps our model generalize better at the global scale.

User Study We conduct a user study to qualitatively evaluate our method. Specifically, we retrieve an exemplar image for each testing label map, and ask 20 subjects to choose the most style-consistent results generated by our method and two competitive baselines (**SPADE_VAE** and **ours GAP**). To generate samples for the user study, we first rank the label histogram intersections with each target scene for all images in the image pool, and use the top 20 percentile images as exemplars². The subjects are given unlimited time to make their selections. For each subject, we randomly generate 100 questions from the dataset. Table 4 shows the evaluation results. First, all subjects strongly favor our results. Second, **ours GAP** is favored more than twice over **SPADE_VAE** [37], further suggesting that the proposed self-supervision scheme is effective since **ours GAP** is also trained with self-supervision.

Methods	SPADE_VAE [37]	ours GAP	ours full
Choose rate	15.6	29.3	55.0

Table 2. User preference study. The numbers indicate the percentage of user who favors the result generated by different methods. Two com

Effect of Attention To understand the effect of spatial-channel attention, we visualize the learned spatial and channel attention in Fig. 6. We observe that: a) spatial attention can attend to multiple regions of the reference image. For each reference region, channel attention finds the corresponding target region.

² This differs from the typical retrieve task that uses the top-1 image since the top 20 percentile images tend to be more semantically different from the target label maps.

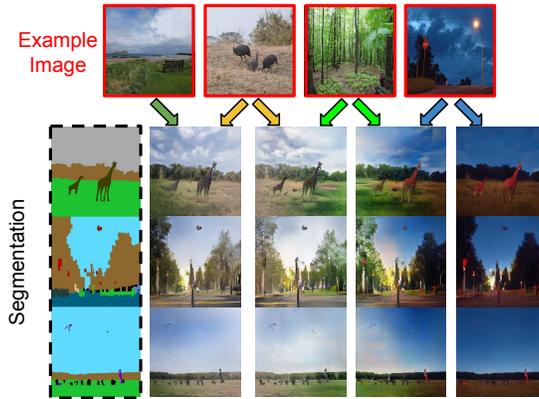


Fig. 7. With a slight modification and no further training, our model can perform spatial style interpolation. In this figure, we demonstrate a horizontal gradient style change on the output image. Please refer to Interpolation, Sec. 4 for more details.



Fig. 8. Given an exemplar patch at the center and the global semantic label map, our trained model can perform example-guided scene image extrapolation, i.e. generating style-consistent beyond-the-border images with semantic maps guidance.

b) spatial-channel attention can detect and utilize the similarities of semantic labels to facilitate style features transfer. In the first sample of Fig. 6, attention in channels 1, 4 respectively perform transformations: `sky_other` \rightarrow `clouds`, `tree` \rightarrow `{tree, hill}`. In the second sample, attention in channels 1, 2, 7 respectively perform transformations: `clouds` \rightarrow `clouds`, `snow` \rightarrow `sand` and `other` \rightarrow `{surfboard, other}`. We provide more analysis on the effect of attention in the appendix.

Interpolation We can easily control the synthesized styles in the test stage by manipulating spatial and channel attentions. First, by manipulating the spatial attention of two exemplar inputs, our trained model can perform style interpolation between the two exemplar. The results are shown in Fig. 5. Next, by manipulating the channel attention, our trained model can perform spatial style interpolation. Fig. 7 shows our model can interpolate between two images and generate horizontally gradient style changes. More details are included in the appendix.

Extrapolation Given a scene patch at the center our model can achieve scene extrapolation, i.e. generating beyond-the-border image content according to the semantic map guidance. A 512×512 extrapolated image is generated by weighted combining synthesized 256×256 patches at 4 corners and 10 other random locations. As shown in Fig. 8, our model generates visually plausible

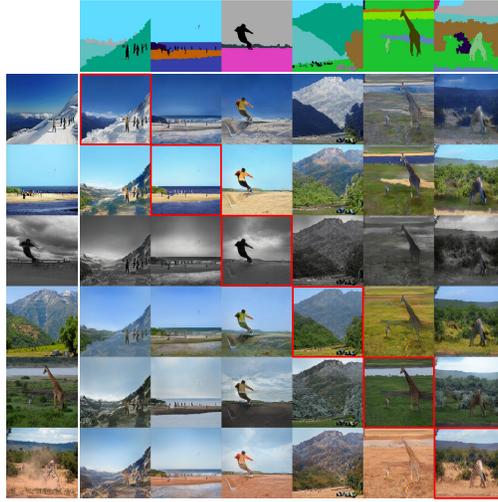


Fig. 9. Style-structure swapping on 6 arbitrary scenes at resolution 256×256 . Our model can generalize across recognizably different scenes of different semantics. Note that the images along the diagonal (red boxes) are *self-reconstruction*.

extrapolated images, showing the promise of our proposed framework for guided scene panorama generation.

Style Swapping Fig. 9 shows reference-guided style swapping between six arbitrary scenes. Our model can generalize across recognizably different scenes semantics and appearances, including snow, mountain, seashore, grassland, desert, artistic effect, and synthesize image with reasonable and consistent styles. More results and comparisons to other approaches are included in the appendix.

5 Conclusion

We propose to address a challenging example-guided image synthesis task between semantically very different scenes. To propagate information between two structurally unaligned and semantically different scenes, we propose an MSCA module that leverages decoupled cross-attention for adaptive correspondence modeling. With MSCA, we propose a unified model for joint global-local alignment and image synthesis. We further propose a patch-based self-supervision scheme that enables training. Experiments on the COCO-stuff dataset show significant improvements over the existing methods. Furthermore, our approach provides interpretability and can be extended to other content manipulation tasks.

Acknowledgment

This work is supported in part by NSF awards IIS-1704337, IIS-1722847, and IIS-1813709, as well as our corporate sponsors.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein gan. arXiv preprint arXiv:1701.07875 (2017)
2. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096 (2018)
3. Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1209–1218 (2018)
4. Chang, H., Lu, J., Yu, F., Finkelstein, A.: Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 40–48 (2018)
5. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* **40**(4), 834–848 (2017)
6. Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1511–1520 (2017)
7. Chen, Y., Kalantidis, Y., Li, J., Yan, S., Feng, J.: A²-nets: Double attention networks. In: Advances in Neural Information Processing Systems. pp. 352–361 (2018)
8. Efros, A.A., Freeman, W.T.: Image quilting for texture synthesis and transfer. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 341–346. ACM (2001)
9. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576 (2015)
10. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2414–2423 (2016)
11. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014)
12. Gu, S., Chen, C., Liao, J., Yuan, L.: Arbitrary style transfer with deep feature reshuffle. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8222–8231 (2018)
13. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 327–340. ACM (2001)
14. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances in Neural Information Processing Systems. pp. 6626–6637 (2017)
15. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)
16. Huang, H., Wang, H., Luo, W., Ma, L., Jiang, W., Zhu, X., Li, Z., Liu, W.: Real-time neural style transfer for videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 783–791 (2017)
17. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1501–1510 (2017)

18. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 172–189 (2018)
19. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 172–189 (2018)
20. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 603–612 (2019)
21. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
22. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in neural information processing systems. pp. 2017–2025 (2015)
23. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016)
24. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017)
25. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4401–4410 (2019)
26. Lee, H.Y., Tseng, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Diverse image-to-image translation via disentangled representations. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 35–51 (2018)
27. Liao, J., Yao, Y., Yuan, L., Hua, G., Kang, S.B.: Visual attribute transfer through deep image analogy. arXiv preprint arXiv:1705.01088 (2017)
28. Lin, J., Xia, Y., Qin, T., Chen, Z., Liu, T.Y.: Conditional image-to-image translation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5524–5532 (2018)
29. Liu, C., Yuen, J., Torralba, A.: Sift flow: Dense correspondence across scenes and its applications. *IEEE transactions on pattern analysis and machine intelligence* **33**(5), 978–994 (2010)
30. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: Advances in neural information processing systems. pp. 700–708 (2017)
31. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of the IEEE international conference on computer vision. pp. 3730–3738 (2015)
32. Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4990–4998 (2017)
33. Ma, L., Jia, X., Georgoulis, S., Tuytelaars, T., Van Gool, L.: Exemplar guided unsupervised image-to-image translation with semantic consistency. arXiv preprint arXiv:1805.11145 (2018)
34. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2794–2802 (2017)
35. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
36. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. arXiv preprint arXiv:1802.05957 (2018)

37. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
38. Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Nießner, M.: Face-forensics: A large-scale video dataset for forgery detection in human faces. arXiv preprint arXiv:1803.09179 (2018)
39. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: Advances in neural information processing systems. pp. 2234–2242 (2016)
40. Siarohin, A., Sangineto, E., Lathuilière, S., Sebe, N.: Deformable gans for pose-based human image generation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3408–3416 (2018)
41. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
42. Wang, M., Yang, G.Y., Li, R., Liang, R.Z., Zhang, S.H., Hall, P., Hu, S.M., et al.: Example-guided style consistent image synthesis from semantic labeling. arXiv preprint arXiv:1906.01314 (2019)
43. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. arXiv preprint arXiv:1808.06601 (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: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)
45. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7794–7803 (2018)
46. Yoo, J., Uh, Y., Chun, S., Kang, B., Ha, J.W.: Photorealistic style transfer via wavelet transforms. arXiv preprint arXiv:1903.09760 (2019)
47. Yu, F., Xian, W., Chen, Y., Liu, F., Liao, M., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving video database with scalable annotation tooling. arXiv preprint arXiv:1805.04687 (2018)
48. Yue, K., Sun, M., Yuan, Y., Zhou, F., Ding, E., Xu, F.: Compact generalized non-local network. In: Advances in Neural Information Processing Systems. pp. 6510–6519 (2018)
49. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. arXiv preprint arXiv:1805.08318 (2018)
50. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 586–595 (2018)
51. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)