

Example-Guided Style-Consistent Image Synthesis from Semantic Labeling

Miao Wang¹ Guo-Ye Yang² Ruilong Li² Run-Ze Liang²
 Song-Hai Zhang² Peter M. Hall³ Shi-Min Hu^{2,1}

¹ State Key Laboratory of Virtual Reality Technology and Systems, Beihang University
² Department of Computer Science and Technology, Tsinghua University, Beijing
³ University of Bath

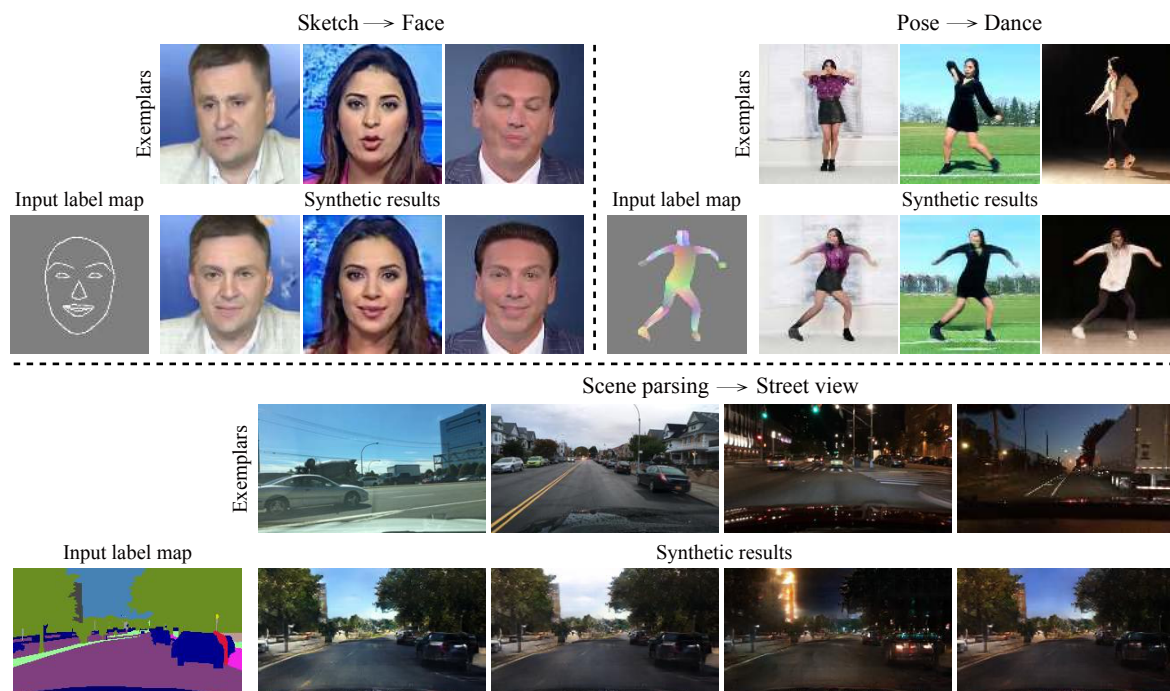


Figure 1: We present a generative adversarial framework for synthesizing images from semantic label maps as well as image exemplars. Our synthetic results are photorealistic, semantically consistent to the label maps (facial expression, pose or scene segmentation map) and style-consistent with corresponding exemplars.

Abstract

Example-guided image synthesis aims to synthesize an image from a semantic label map and an exemplary image indicating style. We use the term “style” in this problem to refer to implicit characteristics of images, for example: in portraits “style” includes gender, racial identity, age, hairstyle; in full body pictures it includes clothing; in street scenes it refers to weather and time of day and such like. A semantic label map in these cases indicates facial expression, full body pose, or scene segmentation. We propose a solution to the example-guided image synthesis problem using conditional generative adversarial networks with style consistency. Our key contributions are (i) a novel style con-

sistency discriminator to determine whether a pair of images are consistent in style; (ii) an adaptive semantic consistency loss; and (iii) a training data sampling strategy, for synthesizing style-consistent results to the exemplar. We demonstrate the efficiency of our method on face, dance and street view synthesis tasks.

1. Introduction

In image-to-image translation tasks, mappings between two visual domains are learnt. Various computer vision and graphics problems are addressed and formulated using the image-to-image translation framework, including super-resolution [30, 28], colorization [29, 49], inpainting [38, 21], style transfer [23, 34] and photorealistic image

synthesis [22, 6, 46]. In the photorealistic image synthesis problem, images are generated from abstract semantic label maps such as pixel-wise segmentation maps or sparse landmarks. In this paper, we study the problem of example-guided image synthesis. Given an input semantic label map x and a guidance image I , the goal is to synthesize a photorealistic image, y , which is semantically consistent with the label map x , while being style-consistent with the exemplar I , so $(x, I) \mapsto y$. Style consistency is automatically determined: in portraits, style consistency refers to the fact that we want our synthetic output to be plausibly of the same genetic type as an input exemplar; in full body images style consistency means the same clothing; and in street scenes it includes such things as the same weather and time of day. Representative applications are shown in Figure 1.

Example-based image synthesis cannot be solved with a straightforward combination of photorealistic image synthesis based on pix2pixHD [22, 46] and style transfer [34]; the style of the input exemplar is not well kept in the synthetic result, see Figure 5. Recently, example-guided image-to-image translation frameworks [20, 31, 2] are proposed using a disentangled model to represent content and style or identity and attributes, however they fail to synthesize photorealistic results from abstract semantic label maps. The challenges are multi-fold: first, the ground truth photorealistic result for each label map given an arbitrary exemplar is not available for training; second, the synthetic results should be photorealistic while semantically consistent with the source label maps; last but not least, the synthetic result should be stylistically consistent with the corresponding image exemplar.

We present a method for this example-guided image synthesis problem with conditional generative adversarial networks. We build on the recent pix2pixHD [46] for image synthesis to ensure photorealism, with the crucial contributions of:

- a novel style consistency discriminator to enforce style consistency of a pair of images (see Section 3.2.2);
- an adaptive semantic consistency loss to ensure quality (see Section 3.2.3);
- a data sampling strategy that ensures we need only a weakly supervised approach for training (see Section 3.3).

2. Related Work

Generative Adversarial Networks. In recent years, generative adversarial networks (GANs) [11, 1] for image generation have progressed rapidly [22, 46]. Driven by adversarial losses, generators and discriminators compete with each other: discriminators aim to distinguish the generated fake images from the target domain; generators try

to fool discriminators. Technologies to improve GANs include: progressive GANs [19, 48, 24], training objective and process designs [42, 1, 37, 43], *etc.* In this paper, we use GANs for example-guided image generation with style consistency awareness.

Image-to-Image Translation and Photorealistic Image Synthesis. The goal of image-to-image translation is to translate images from a source domain to a target domain. Isola *et al.* [22] proposed the conditional GAN framework for various image-to-image translation tasks with paired images for supervision. Wang *et al.* [46] extended this work for high-resolution image synthesis and interactive manipulation. Recently, researchers proposed to solve the unsupervised image-to-image translation problem with cycle consistency to overcome the lack of unpaired training data [51, 25, 33, 52, 20, 31, 5]. Photorealistic image synthesis [6, 39, 46] is a specific application of image-to-image translation, where images are synthesized semantically from abstract label maps. Chen *et al.* [6] proposed a cascade framework to synthesize high-resolution images from pixel-wise labeling maps. Wang *et al.* [46] proposed a framework for instance-level image synthesis with conditional GANs.

Very recently, a few works [16, 20, 31, 35] have been proposed to transfer the style or attributes of an exemplar to the source image, where the images belong to photorealistic domains (*aka* domain adaptation). Our goal differs from these works by aiming at synthesizing photos from an abstract semantic label domain rather than a photorealistic image domain. Zheng *et al.* [50] proposed a clothes changing system to change the clothing of a person in image. Chan *et al.* [4] presented a network to synthesize a dance video from a target dance video and an source exemplar video. Different from our model, it was trained for every input exemplar video. Ma *et al.* [36] proposed to synthesize person images from pose keypoints. We show in Section 4 that our method outperforms the state-of-the-art methods.

Style Transfer. Style transfer is a long-standing problem in computer vision and graphics, which aims to transfer the style of a source image to a target image or target domain. Some approaches [14, 10, 23, 34, 18, 32, 12, 5, 17] transfer style based on single exemplar, where others learn a general style of a target domain with a holistic sense [51, 20, 31, 7]. Similar to our model, the PairedCycleGAN model [5] uses a style discriminator to distinguish whether a pair of facial images wear the same make-up in the making-up application. However, in their discriminator, the input image pair must be accurately aligned via warping; a generator is learned for each facial component. Our style consistency discriminator, in contrast, provides a general solution for image synthesis from both sparse labels (*e.g.* sketch and pose) and pixel-wise dense labels (*e.g.* scene parsing).

3. Example-guided Image Synthesis

In this section, we first review the baseline model pix2pixHD [46], then describe our method, a conditional generative adversarial network for synthesizing photorealistic images from semantic label maps given specific exemplars. Finally we show how to appropriately prepare training data for our framework.

3.1. The pix2pixHD Baseline

The pix2pixHD [46] is a powerful image synthesis and interactive manipulation framework based on the pioneering conditional image-to-image translation method pix2pix [22]. Let x be a label map from a semantic label domain \mathcal{X} , the goal of pix2pixHD is to synthesize an image y , from x : $x \mapsto y$. It consists of a hierarchically integrated generator G and multi-scale discriminators $\mathcal{D} = \{D_1, D_2, D_3\}$ to handle high-resolution synthesis tasks. The goal of the generator G is to translate semantic label maps to photorealistic images, and the objective of the discriminators is to distinguish generated fake images from real ones at different resolution. The training dataset $\{(x_i, y_i)\}$ consists of pairs of label map x_i and corresponding real image y_i .

The pix2pixHD optimizes a multi-task problem with a standard GAN loss \mathcal{L}_{GAN} and feature matching loss \mathcal{L}_{FM} :

$$\min_G (\max_D \sum_{k=1,2,3} \mathcal{L}_{\text{GAN}}(G, D_k) + \lambda \mathcal{L}_{\text{FM}}(G, D_k)), \quad (1)$$

where $\mathcal{L}_{\text{GAN}}(G, D_k)$ is the standard GAN loss given by:

$$\mathbb{E}_{(x,y)} [\log D_k(x, y)] + \mathbb{E}_x [\log(1 - D_k(x, G(x))),] \quad (2)$$

and $\mathcal{L}_{\text{FM}}(G, D_k)$ is the feature matching loss given by:

$$\mathbb{E}_{(x,y)} \sum_{i=1}^T \frac{1}{N_i} [||D_k^i(x, y) - D_k^i(x, G(x))||_1], \quad (3)$$

where T is the layer size and N_i is the feature size in corresponding discriminator layer. An optional perceptual loss is introduced as the L_1 loss between pre-trained VGG network [44] features.

One appealing feature of pix2pixHD is the instance-level image manipulation with a feature embedding technique. Given an instance-level segmentation map, pix2pixHD is able to synthesize an image with a specific appearance from an instance exemplar in the same object category. We will show that without the input instance-level pixel-wise segmentation map as a constraint, our model is still able to synthesize images with styles automatically transferred from exemplar images.

3.2. Our Model

Let I be a guidance image from a natural image domain \mathcal{Y} . Our goal is to synthesize an image y , from a semantic label map x and an image I : $(x, I) \mapsto y$. The role of

I is to provide a style constraint to image synthesis: the output image y must be style-consistent with the exemplar I . Our problem is more difficult than that solved by pix2pixHD. One particular challenge we face is that given an input label map x , the ground truth images $\{y\}$ for arbitrary guided style exemplars $\{I\}$ are missing. To solve this weakly-supervised problem, we learn style consistency between pairs of images: they could be style-consistent image pairs $\{(y_1^S, y_2^S)\}$ or style-inconsistent image pairs $\{(y_1^N, y_2^N)\}$ (see Section 3.3).

An overview of our method is illustrated in Figure 2. It builds upon a single-scale version of pix2pixHD, and contains: (i) a generator G , with semantic map x , style example I and its corresponding label $F(I)$ as input and output a synthetic image; (ii) a standard discriminator D_R to distinguish real images from fake ones given conditional inputs; and (iii) we introduce a *style consistency* discriminator D_{SC} to detect whether the synthetic image and the guidance image I are style-compatible, which operates on image pairs from domain \mathcal{Y} . Here, $F(\cdot)$ is an operator which, given an image produces a set of semantic labels that represent the image (choices of $F(\cdot)$ are given in Section 4.2); for convenience $F(I)$ can be visualized as an image, provided the viewer recalls that the image contains semantic labels. Our objective function contains three losses: a *standard adversarial loss*; a novel *adversarial style consistency loss*; and a novel *adaptive semantic consistency loss*.

3.2.1 Standard Adversarial Loss

We apply standard adversarial losses via the standard discriminator D_R as:

$$\mathcal{L}_{\text{StdAdv}}(G, D_R) = \mathbb{E}_{(x,y)} [\log D_R(x, y)] + \mathbb{E}_{(x,I)} [\log(1 - D_R(x, G(x, I, F(I))))], \quad (4)$$

where the G tries to synthesize images that look similar to real images from image domain \mathcal{Y} regardless of specific styles, while given an image conditioned with the corresponding label map, the D_R aims to determine the image is real or fake.

3.2.2 Adversarial Style Consistency Loss

With the standard adversarial loss, the generator G is able to synthesize images matching the data distribution of domain \mathcal{Y} , however the synthetic results are not guaranteed to be style-consistent with the corresponding guidance I . We introduce the style consistency loss $\mathcal{L}_{\text{SCAdv}}$ using a discriminator D_{SC} associated with a pair of images — either both real, or one real and one synthetic:

$$\begin{aligned} \mathcal{L}_{\text{SCAdv}}(G, D_{SC}) &= \mathbb{E}_{(y_1^S, y_2^S)} [\log D_{SC}(y_1^S, y_2^S)] \\ &+ \mathbb{E}_{(y_1^N, y_2^N)} [\log(1 - D_{SC}(y_1^N, y_2^N))] \\ &+ \mathbb{E}_{(x,I)} [\log(1 - D_{SC}(I, G(x, I, F(I))))], \end{aligned} \quad (5)$$

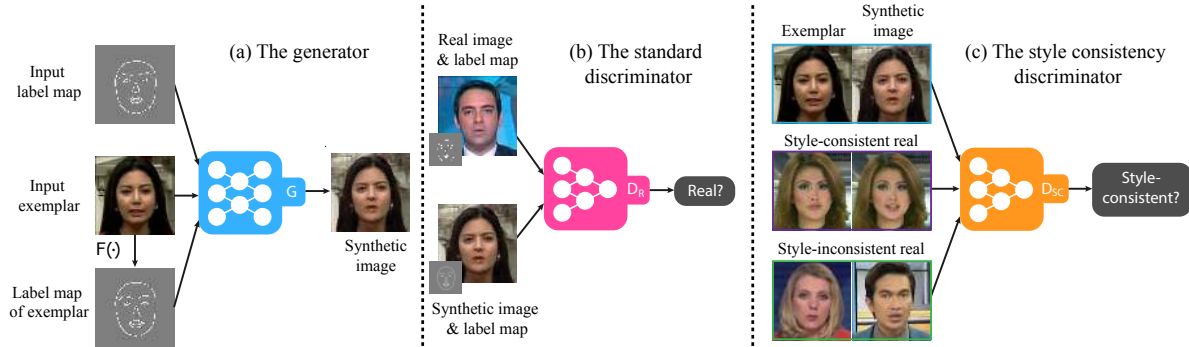


Figure 2: Overview of our framework consisting of a generator G and two discriminators D_R and D_{SC} . (a) Given an input label map, a guided example and its labels generated by a known function $F(\cdot)$, the generator G tries to synthesize an image semantically consistent to the labels, while being style-consistent to the exemplar. (b) The standard discriminator D_R learns to distinguish between real and synthetic images on conditional input. (c) The style consistency discriminator D_{SC} aims to distinguish between style-consistent image pairs and style-inconsistent image pairs.

where y_1^S and y_2^S are a pair of sampled real images from domain \mathcal{Y} with the same style, y_1^N and y_2^N are a pair of sampled real images from domain \mathcal{Y} with different styles. We introduce the data sampling strategy in Section 3.3.

With the proposed adversarial style consistency loss \mathcal{L}_{SCAdv} , the discriminator D_{SC} tries to learn awareness of style consistency between a pair of images, while the generator G tries to fool D_{SC} by generating an image with the same style to exemplar I .

3.2.3 Adaptive Semantic Consistency Loss

The semantic consistency loss is introduced to reconstruct an image from a label map in the semantic sense of *e.g.* sketch. It may appear we could use the error between the input labels x , and the predicted labels from the synthetic image, $G(x, I, F(I))$, for example $\|x - F(G(x, I, F(I)))\|_1$ or some variant thereof. However, different applications give distinct meanings to the semantic label maps, with the consequence that the gradient of the loss will, in general, vary between applications. This would mean selecting hyper-parameters λ to combine losses on a per-application basis.

We avoid this problem by always computing semantic consistency losses between images: the synthetic image $G(x, I, F(I))$ and specifically an image z which is a priori known to be consistent with a given semantic map x . Typically the image z is drawn from the training dataset and we have $x = F(z)$. A particular issue with our adopted scheme is that such losses will try to converge the network output on the image z , which by choice is photorealistic and is semantically consistent with x . Such behavior would work perfectly when z and I are sampled from images $\{(y_1^S, y_2^S)\}$ with the same style, but could force the output away from the desired style when z and I are “style-wise” different, *i.e.* $(z, I) \in \{(y_1^N, y_2^N)\}$.

Our solution, is to use a novel *adaptive* VGG loss com-

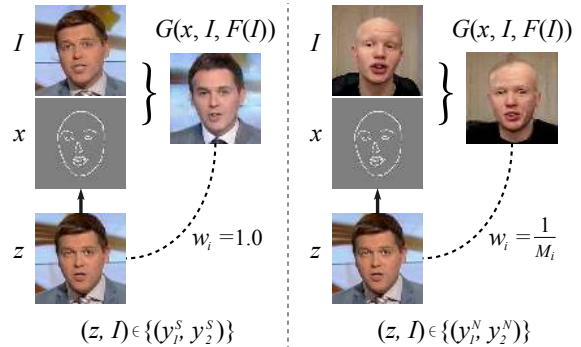


Figure 3: Adaptive weight for semantic consistency loss.

puted via a pre-trained model [23] between the synthetic image $G(x, I, F(I))$ and the real image z of label map x . An adaptive weighting scheme is proposed for per-layer VGG loss computation, to ensure the semantic consistency of the synthetic image to x :

$$\mathcal{L}_{SC}(G) = \sum_{i=1}^N w_i \|L^{(i)}(z) - L^{(i)}(G(x, I, F(I)))\|_1, \quad (6)$$

where $L^{(i)}$ represents the i -th layer feature extractor of the VGG network, and w_i is the adaptive weight for the i -th layer. We set $w_i = 1.0$ to gain the impact of details from shallow layers when z and I are from style-consistent sampled pairs $\{(y_1^S, y_2^S)\}$, and $w_i = \frac{1}{M_i}$ to suppress the impact of detail matching for style-inconsistent pair $(z, I) \in \{(y_1^N, y_2^N)\}$; M_i is the number of elements in the i -th feature layer. The adaptive weighting scheme is illustrated in Figure 3.

Full Objective. The final loss is formulated as:

$$\begin{aligned} \mathcal{L}(G, D_R, D_{SC}) = & \mathcal{L}_{StdAdv}(G, D_R) \\ & + \lambda_1 \mathcal{L}_{SCAdv}(G, D_{SC}) \\ & + \lambda_2 \mathcal{L}_{SC}(G), \end{aligned} \quad (7)$$



Figure 4: Representative sampled data for training networks using FaceForensics [41], YouTube Dances and BDD100K [47] datasets. Each row shows pairs of sampled images from the above three datasets.

where λ_1 and λ_2 control the relative importance of the terms, our full objective is given by:

$$G^* = \arg \min_G \max_{D_R, D_{SC}} \mathcal{L}(G, D_R, D_{SC}). \quad (8)$$

3.3. Sampling Strategy for Style-consistent and Style-inconsistent Image Pairs

So far, we have introduced the core techniques of our network. However one prerequisite to our method is to obtain style-consistent image pairs $\{(y_1^S, y_2^S)\}$ and style-inconsistent image pairs $\{(y_1^N, y_2^N)\}$. Thus the datasets for prior image-to-image translation works [22, 46, 51, 31, 20] are not feasible for our training.

A key idea for training data acquisition is to collect image pairs from videos. In face and dance synthesis tasks, we observed that: (i) within a short temporal period of a video, the style of frame contents are ensured to be the same, and (ii) frames from different videos probably have different styles (e.g. different gender, hairstyles, skin colors and make-up in the face image synthesis application). We thus randomly sample pairs of frames within $T = 10$ frames from a video and regard them as style-consistent ones $\{(y_1^S, y_2^S)\}$. For style-inconsistent pairs $\{(y_1^N, y_2^N)\}$, we firstly randomly sample pairs of frames from different videos, then manually label whether images from each sampled pair are style-consistent or not.

In the street view synthesis task, as large scale street view videos with different styles are not easy to collect, we use images from the BDD100K dataset [47]. In BDD100K, street view images and the *weather*, *time of day* attributes are provided. We coarsely categorize the images into 13 style groups based on the attributes, then sample style-consistent image pairs inside each group and sample style-inconsistent image pairs between groups. Figure 4 shows representative sampled pairs of images.

4. Experiments

4.1. Implementation Details

We implement our model based on the single-scale pix2pixHD framework and experiment with images with size 256×256 (256×144 for street view synthesis). The generator G contains several Convolution-InstanceNorm-ReLU-Stride-2 layers to encode deep features, then 9 residual blocks [13] and finally some Convolution-InstanceNorm-ReLU-Stride-0.5 layers to synthesize images. For both discriminators D_R and D_{SC} , we use PatchGANs [22] with several Convolution-InstanceNorm-LeakyReLU-Stride-2 layers with the exception that InstanceNorm is not applied in the first layer. The slope for LeakyReLU is set as 0.2. For all the experiments, we set $\lambda_1 = 10$ and $\lambda_2 = 10$ in Equation 7. All the networks are trained from scratch on an NVIDIA GTX 1080 Ti GPU using the Adam solver [27] with a batch size of 1. The learning rate is initially fixed as 0.0002 for the first 500K iterations and linearly decayed to zero over the next 500K iterations. We use LSGANs [37] for stable training. For more details, please refer to the supplementary material.

4.2. Datasets

We evaluate our method on face, dance and street view image synthesis tasks, using the following datasets:

Sketch→Face. We use the real videos in the FaceForensics dataset [41], which contains 854 videos of reporters broadcasting news. We use the image sampling strategy described in Section 3.3 to acquire training image pairs from video, then apply face alignment algorithm [26] to localize facial landmarks, crop facial regions and resize them to size 256×256 . The detected facial landmarks are connected to create face sketches as function $F(\cdot)$.

Pose→Dance. We download 150 solo dance videos from YouTube, crop out the central body regions and resize them to 256×256 . As the number of videos is small, we evenly split each video into the first part and the second part along the time-line, then sample training data only from the first parts and sample testing data only from the second parts of all the videos. The function $F(\cdot)$ is implemented using concatenated pre-trained DensePose [40] and OpenPose [3] pose detection results to provide pose labels.

Scene parsing→Street view. We use the BDD100k dataset [47] to synthesize street view images from pixel-wise semantic labels (*i.e.* scene parsing maps). We use the state-of-the-art scene parsing network DANet [9] as the function $F(\cdot)$. Please find more details in our supplementary material.

4.3. Baselines

We compare our method with the following algorithms:



Figure 5: Example-based face image synthesis on the FaceForensics dataset. The first column shows the input labels, the second column shows the input style example, next columns show the results from our method and our ablation studies, pix2pixHD, pix2pixHD with DPST, MUNIT and PairedMUNIT.

pix2pixHD and **pix2pixHD [46] with DPST [34]**. pix2pixHD is the image-to-image translation baseline. A default image could be synthesized using pix2pixHD with its style then transferred to the guided example using Deep Photo Style Transfer (DPST) method.

MUNIT [20] and PairedMUNIT. MUNIT is the state-of-the-art unsupervised image-to-image translation method with disentangled content and style representations that are able to translate images to given exemplars. We modify MUNIT by integrating pairwise style information to the original model and adaptively computing losses with style (denoted as PairedMUNIT).

Ours without \mathcal{L}_{SC} , \mathcal{L}_{SCAdv} or adaptive weights for ablation studies. All of the methods are trained on the datasets introduced in Section 4.2.

4.4. Evaluation Metrics

Photorealism and Semantic Consistency. We use the Fréchet Inception Distance [15] to evaluate the realism and faithfulness of the synthetic results. This metric is widely used for implicit generative models, because it correlates with the visual quality of generated samples. A smaller FID is often favored by the human subjects. We further evaluate semantic consistency by translating the synthetic images back to the label domain and comparing the accuracy to the input labels. For tasks Sketch→Face and Pose→Dance, we use the labeling endpoint error (LEPE) between the input label map x and the labels generated by $F(\cdot)$ to compute the label accuracy. For task Scene parsing→Street view, we use scene parsing score (SPS) [9] on synthetic street view images to measure the segmentation accuracy.

| | Sketch→Face | Pose→Dance | Parsing→Street |
|-------------|--------------|--------------|----------------|
| pix2pixHD | 39.86 | 92.33 | 157.46 |
| MUNIT | 148.57 | 158.47 | 235.84 |
| PairedMUNIT | 142.08 | 161.22 | 259.95 |
| Ours | 31.26 | 33.39 | 96.23 |

Table 1: Photorealism comparison measured by Fréchet Inception Distance (FID) [15].

| Method | Sketch→Face | Pose→Dance |
|-------------|---------------|---------------|
| pix2pixHD | 0.0050 | 0.0163 |
| MUNIT | 0.0107 | 0.0958 |
| PairedMUNIT | 0.0080 | 0.0502 |
| Ours | 0.0085 | 0.0186 |

Table 2: Semantic consistency measured by normalized label endpoint error for different methods in face and dance image synthesis tasks.

Style Consistency. We perform a human perceptual study to compare style consistency from human point of view. We show pairs of our result and the result from baseline methods to invited subjects and ask which one they see as being closer to the guidances’ style.

4.5. Results

Main Results. In Figure 5, we show our results (column 3) and the results from baseline methods in the Sketch→Face synthesis application on the test set. While the pix2pixHD is able to generate photorealistic images consistent with the input semantic labels, it is not able to keep the style (e.g. gender, hair, skin color) from input exemplars in the synthetic results, even enhanced by the deep

| | pix2pixHD | pix2pixHD+DPST | PairedMUNIT |
|------|-----------|----------------|-------------|
| Ours | 89.12% | 80.67% | 90.88% |

Table 3: Style consistency evaluation by human option study on Sketch→Face synthesis. Each cell lists the percentage where our result is preferred over the other method.

photo style transfer effect (column 7 and 8). The unsupervised method MUNIT and its improvement PairedMUNIT fail to generate photorealistic results from semantic maps in this application (column 9 and 10). The possible reason for their failures is that they assume that the input and output domains share the same content space, which is not true in image synthesis applications from semantic label maps.

Table 1 gives the quantitative evaluation of the photorealism measured by FID in various image synthesis tasks, where our method performs the best. The semantic consistency of synthetic results to the input labels is given by LEPE in Table 2. It can be seen that the pix2pixHD obtains the best semantic consistency to the input labels, because it does not lose semantic accuracy by totally ignoring style consistency. Our method outperforms MUNIT and PairedMUNIT.

For style consistency evaluation, we conduct a human perception study commonly used in image-to-image translation works [22, 51, 6, 46, 8]. The input exemplars and pairwise synthetic results sampled from our method and a baseline method are shown to the subjects with unlimited watching time. Then the subjects were asked “Which image is closer to the exemplar in terms of style?” Images for user study were randomly sampled from the test set; each pair was shown in random order and guaranteed to be examined by at least 30 subjects. The ratios of votes our method got over baseline methods are given in Table 3. Our method won more user preferences in pairwise comparison. The quantitative results shown that our results are more photorealistic and more style-consistent with the exemplars.

We conducted ablation studies to verify our model. As can be seen in Figure 5, without the adaptive weight scheme in \mathcal{L}_{SC} , the quality of results is slightly reduced; without the semantic loss \mathcal{L}_{SC} , the semantic consistency would lose; without the style consistency adversarial loss \mathcal{L}_{SCAdv} , the target style is not maintained. Quantitative photorealism statistics reported in Table 4 validated the above observation. We further extract 50×50 eye patches from synthetic images and exemplars and compute the VGG feature distance between them. Table 5 indicates that the weight adaptation makes a quantitative improvement of style consistency.

Figure 6 shows the in-the-wild synthesis results from our model using Internet images. The results indicate that the model generalizes well for “unseen” cases. We provide more results in the supplementary material.

Pose→Dance Synthesis. Figure 7 shows a visual com-

| Method | FID | Method | FID |
|-----------------------------|-------|--------------------------------|--------------|
| Ours w/o adaptive weights | 35.59 | Ours w/o \mathcal{L}_{SCAdv} | 58.08 |
| Ours w/o \mathcal{L}_{SC} | 76.59 | Ours | 31.26 |

Table 4: Ablation study: Fréchet Inception Distance (FID) of our results and alternatives on the Sketch→Face synthesis task.

| | Ours | Ours w/o adapt. weights | w/o \mathcal{L}_{SC} | w/o \mathcal{L}_{SCAdv} |
|-----------|--------------|-------------------------|------------------------|---------------------------|
| VGG Dist. | 0.643 | 0.654 | 0.898 | 0.703 |

Table 5: VGG feature distance of eye patches between synthetic image and exemplar.

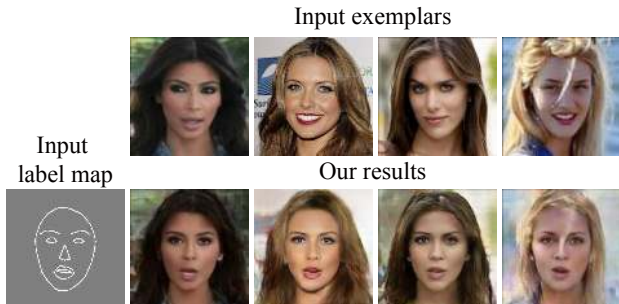


Figure 6: In-the-wild Sketch→Face synthesis.



Figure 7: Dance synthesis from pose maps.




Figure 8: Pose→Dance comparison with Ma *et al.* [36].

parison of our method and baselines in the Pose→Dance synthesis application. The semantic consistency of synthetic results to the input labels measured using LEPE are given in Table 2. Although the facial regions of our results

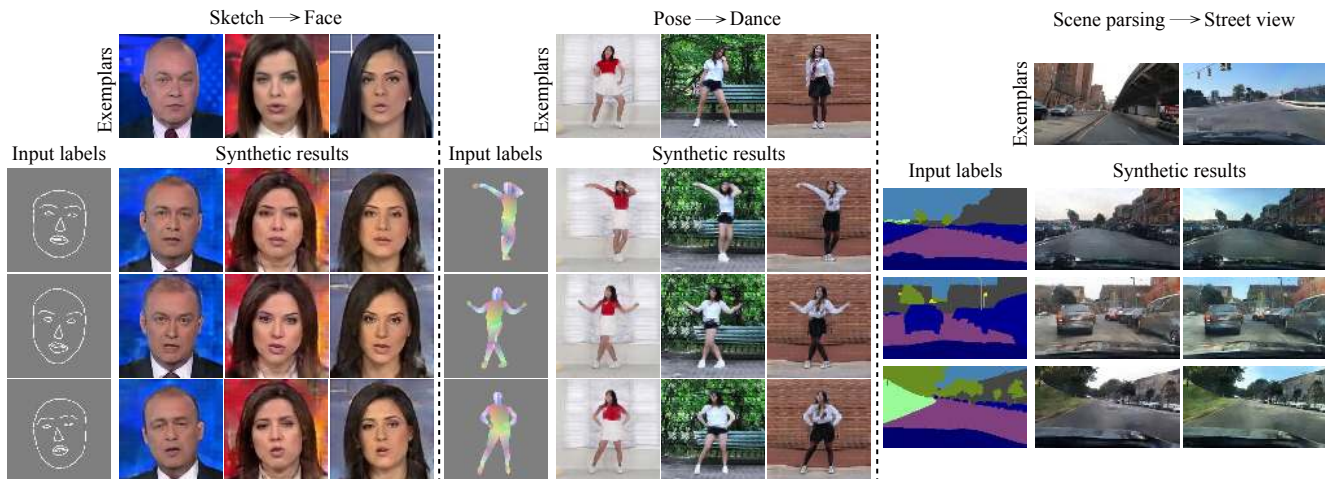


Figure 9: More results of example-based image synthesis on face, dance and street view synthesis tasks.



Figure 10: Street view synthesis from scene parsing maps and corresponding exemplars.

are blurry without including facial landmarks in the input pose labels, our model still produces images that are style-consistent with the guidance images while consistent with the semantic labels. Figure 8 shows the visual comparison with Ma *et al.* [36] on the dancing dataset. The generated poses and clothes in our results are visually better.

Scene parsing → Street view Synthesis. A comparison of our method and baselines in the Scene parsing → Street view task is given in Figure 10. The semantic consistency of synthetic results to the input labels measured using SPS are given in Table 6. Although the scene in the guidance images are not quite the same as the semantics of the input label maps, our model is able to produce images that are semantically consistent with the segmentation map and style-consistent with the guidance image.

Figure 9 shows more results. Our network can faithfully synthesize images from various semantic labels and exemplars. Please find more results in the supplementary file.

5. Conclusions

In this paper, we present a novel method for example-guided image synthesis with style-consistency from general-form semantic labels. During network training, we propose to sample style-consistent and style-inconsistent image pairs from video to provide style awareness to the

| Method | Per-pixel acc. | Per-class acc. | Class IOU |
|----------------|----------------|----------------|--------------|
| pix2pixHD | 83.85 | 36.17 | 0.310 |
| MUNIT | 58.58 | 18.99 | 0.139 |
| PairedMUNIT | 62.96 | 22.41 | 0.160 |
| Ours | 84.71 | 39.44 | 0.333 |
| Original image | 86.74 | 52.25 | 0.452 |

Table 6: Semantic consistency measured by scene parsing score [9] for different methods on the street view image synthesis task.

model. Beyond that, we introduce the style consistency adversarial losses and the style consistency discriminator, as well as the semantic consistency loss with adaptive weights, to produce plausible results. Qualitative and quantitative results in different applications show that the proposed model produces realistic and style-consistent images better than those from prior arts.

Limitations and Future Work. Our network is mainly trained on cropped video data whose resolution is limited (*e.g.* 256×256), we did not use the multi-scale architecture as pix2pixHD did for high-resolution image synthesis (*e.g.* 512×512 resolution or more). Moreover, the synthetic background in face and dance image synthesis tasks may be blurry, because the semantic labels do not specify any background scenes. Lastly, we have demonstrated the efficiency of our method in several synthesis applications, however the results in other applications could be effected by the performance of the state-of-the-art semantic labeling function $F(\cdot)$. In the future, we plan to extend this framework to video domain [45] and synthesize style-consistent videos to given exemplars.

Acknowledgements. We thank the anonymous reviewers for the valuable discussions. This work was supported by the Natural Science Foundation of China (Project Number: 61521002, 61561146393). Shi-Min Hu is the corresponding author.

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017. [2](#)
- [2] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. Towards open-set identity preserving face synthesis. In *CVPR*, 2018. [2](#)
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. [5](#)
- [4] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. *arXiv preprint arXiv:1808.07371*, 2018. [2](#)
- [5] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In *CVPR*, 2018. [2](#)
- [6] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, 2017. [2, 7](#)
- [7] Yang Chen, Yu-Kun Lai, and Yong-Jin Liu. Cartoongan: Generative adversarial networks for photo cartoonization. In *CVPR*, pages 9465–9474, 2018. [2](#)
- [8] 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*, 2018. [7](#)
- [9] Jun Fu, Jing Liu, Haijie Tian, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. *arXiv preprint arXiv:1809.02983*, 2018. [5, 6, 8](#)
- [10] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. [2](#)
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. [2](#)
- [12] Shuyang Gu, Congliang Chen, Jing Liao, and Lu Yuan. Arbitrary style transfer with deep feature reshuffle. In *CVPR*, 2018. [2](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [5](#)
- [14] Aaron Hertzmann, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. Image analogies. In *SIGGRAPH*, pages 327–340, 2001. [2](#)
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017. [6](#)
- [16] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. [2](#)
- [17] Haozhi Huang, Hao Wang, Wenhan Luo, Lin Ma, Wenhao Jiang, Xiaolong Zhu, Zhifeng Li, and Wei Liu. Real-time neural style transfer for videos. In *CVPR*, July 2017. [2](#)
- [18] Xun Huang and Serge J Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. [2](#)
- [19] Xun Huang, Yixuan Li, Omid Poursaeed, John Hopcroft, and Serge Belongie. Stacked generative adversarial networks. In *CVPR*, 2017. [2](#)
- [20] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. [2, 5, 6](#)
- [21] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and Locally Consistent Image Completion. *ACM Trans. Graph.*, 36(4):107:1–107:14, 2017. [1](#)
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arxiv*, 2016. [2, 3, 5, 7](#)
- [23] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. [1, 2, 4](#)
- [24] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. [2](#)
- [25] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. *arXiv preprint arXiv:1703.05192*, 2017. [2](#)
- [26] Davis E King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10(Jul):1755–1758, 2009. [5](#)
- [27] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [28] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *CVPR*, 2017. [1](#)
- [29] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *ECCV*, 2016. [1](#)
- [30] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. [1](#)
- [31] 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. [2, 5](#)
- [32] Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. Visual attribute transfer through deep image analogy. *ACM Trans. Graph.*, 36(4), 2017. [2](#)
- [33] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NIPS*, pages 700–708, 2017. [2](#)
- [34] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. *arXiv preprint arXiv:1703.07511*, 2017. [1, 2, 6](#)
- [35] Liqian Ma, Xu Jia, Stamatios Georgoulis, Tinne Tuytelaars, and Luc Van Gool. Exemplar guided unsupervised image-to-image translation. *arXiv preprint arXiv:1805.11145*, 2018. [2](#)

- [36] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, pages 405–415, 2017. [2](#), [7](#), [8](#)
- [37] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. [2](#), [5](#)
- [38] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. [1](#)
- [39] Xiaojuan Qi, Qifeng Chen, Jiaya Jia, and Vladlen Koltun. Semi-parametric image synthesis. In *CVPR*, 2018. [2](#)
- [40] Iasonas Kokkinos Riza Alp Güler, Natalia Neverova. Densepose: Dense human pose estimation in the wild. *arXiv*, 2018. [5](#)
- [41] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv*, 2018. [5](#)
- [42] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016. [2](#)
- [43] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *CVPR*, 2017. [2](#)
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [3](#)
- [45] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *NIPS*, 2018. [8](#)
- [46] 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. [2](#), [3](#), [5](#), [6](#), [7](#)
- [47] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018. [5](#)
- [48] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaoqi Huang, Xiaogang Wang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. *arXiv preprint*, 2017. [2](#)
- [49] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. [1](#)
- [50] Zhao-Heng Zheng, Hao-Tian Zhang, Fang-Lue Zhang, and Tai-Jiang Mu. Image-based clothes changing system. *Computational Visual Media*, 3(4):337–347, 2017. [2](#)
- [51] 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](#), [5](#), [7](#)
- [52] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *NIPS*, 2017. [2](#)