SEAN: Image Synthesis with Semantic Region-Adaptive Normalization

Peihao Zhu¹ Rameen Abdal¹ Yipeng Qin² Peter Wonka¹

¹KAUST ²Cardiff University

Figure 1: Face image editing controlled via style images and segmentation masks. a) source images. b) reconstruction of the source image; segmentation mask shown as small inset. c - f) four separate edits; we show the image that provides new style information on top and show the part of the segmentation mask that gets edited as small inset. The results of the successive edits are shown in row two and three. The four edits change hair, mouth and eyes, skin tone, and background, respectively.

Abstract

We propose semantic region-adaptive normalization (SEAN), a simple but effective building block for Generative Adversarial Networks conditioned on segmentation masks that describe the semantic regions in the desired output image. Using SEAN normalization, we can build a network architecture that can control the style of each semantic region individually, e.g., we can specify one style reference image per region. SEAN is better suited to encode, transfer, and synthesize style than the best previous method in terms of reconstruction quality, variability, and visual quality. We evaluate SEAN on multiple datasets and report better quan-

titative metrics (e.g. FID, PSNR) than the current state of the art. SEAN also pushes the frontier of interactive image editing. We can interactively edit images by changing segmentation masks or the style for any given region. We can also interpolate styles from two reference images per region. Code: <https://github.com/ZPdesu/SEAN>.

1. Introduction

In this paper we tackle the problem of synthetic image generation using conditional generative adversarial networks (cGANs). Specifically, we would like to control the layout of the generated image using a segmentation mask

Figure 2: Editing sequence on the ADE20K dataset. (a) source image, (b) reconstruction of the source image, (c-f) various edits using style images shown in the top row. The regions affected by the edits are shown as small insets.

that has labels for each semantic region and “add” realistic styles to each region according to their labels. For example, a face generation application would use region labels like eyes, hair, nose, mouth, *etc.* and a landscape painting application would use labels like water, forest, sky, clouds, *etc.* While multiple very good frameworks exist to tackle this problem [22, 8, 39, 44], the currently best architecture is SPADE [38] (also called GauGAN). Therefore, we decided to use SPADE as starting point for our research. By analyzing the SPADE results, we found two shortcomings that we would like to improve upon in our work.

First, SPADE uses only one style code to control the entire style of an image, which is not sufficient for high quality synthesis or detailed control. For example, it is easily possible that the segmentation mask of the desired output image contains a labeled region that is not present in the segmentation mask of the input style image. In this case, the style of the missing region is undefined, which yields low quality results. Further, SPADE does not allow using a different style input image for each region in the segmentation mask. Our first main idea is therefore to control the style of each region individually, *i.e.*, our proposed architecture accepts one style image per region (or per region instance) as input.

Second, we believe that inserting style information only in the beginning of a network is not a good architecture choice. Recent architectures [25, 31, 2] have demonstrated that higher quality results can be obtained if style information is injected as normalization parameters in multiple layers in the network, *e.g.* using AdaIN [19]. However, none of these previous networks use style information to generate spatially varying normalization parameters. To alleviate this shortcoming, our second main idea is to design a normalization building block, called SEAN, that can use style

input images to create spatially varying normalization parameters per semantic region. An important aspect of this work is that the spatially varying normalization parameters are dependent on the segmentation mask as well as the style input images.

Empirically, we provide an extensive evaluation of our method on several challenging datasets: CelebAMask-HQ [28, 24, 32], CityScapes [10], ADE20K [50], and our own Façades dataset. Quantitatively, we evaluate our work on a wide range of metrics including FID, PSNR, RMSE and segmentation performance; qualitatively, we show examples of synthesized images that can be evaluated by visual inspection. Our experimental results demonstrate a large improvement over the current state-of-the-art methods. In summary, we introduce a new architecture building block *SEAN* that has the following advantages:

1. SEAN improves the quality of the synthesized images for conditional GANs. We compared to the state of the art methods SPADE and Pix2PixHD and achieve clear improvements in quantitative metrics (*e.g.* FID score) and visual inspection.
2. SEAN improves the per-region style encoding, so that reconstructed images can be made more similar to the input style images as measured by PSNR and visual inspection.
3. SEAN allows the user to select a different style input image for each semantic region. This enables image editing capabilities producing much higher quality and providing better control than the current state-of-the-art methods. Example image editing capabilities are interactive region by region style transfer and per-region style interpolation (See Figs. 1, 2, and 5).

2. Related Work

Generative Adversarial Networks. Since their introduction in 2014, Generative Adversarial Networks (GANs) [14] have been successfully applied to various image synthesis tasks, *e.g.* image inpainting [48, 11], image manipulation [52, 5, 1] and texture synthesis [29, 43, 12]. With continuous improvements on GAN architecture [40, 25, 38], loss function [33, 4] and regularization [16, 36, 34], the images synthesized by GANs are becoming more and more realistic. For example, the human face images generated by StyleGAN [25] are of very high quality and are almost indistinguishable from photographs by untrained viewers. A traditional GAN uses noise vectors as the input and thus provides little user control. This motivates the development of conditional GANs (cGANs) [35] where users can control the synthesis by feeding the generator with conditioning information. Examples include class labels [37, 34, 6], text [41, 18, 46] and images [22, 53, 30, 44, 44, 38]. Our work is built on the conditional GANs with image inputs, which aims to tackle image-to-image translation problems.

Image-to-Image Translation. Image-to-image translation is an umbrella concept that can be used to describe many problems in computer vision and computer graphics. As a milestone, Isola *et al.* [22] first showed that image-conditional GANs can be used as a general solution to various image-to-image translation problems. Since then, their method is extended by several following works to scenarios including: unsupervised learning [53, 30], few-shot learning [31], high resolution image synthesis [44], multi-modal image synthesis [54, 21] and multi-domain image synthesis [9]. Among various image-to-image translation problems, semantic image synthesis is a particularly useful genre as it enables easy user control by modifying the input semantic layout image [28, 5, 15, 38]. To date, the SPADE [38] model (also called GauGAN) generates the highest quality results. In this paper, we will improve SPADE by introducing per-region style encoding.

Style Encoding. Style control is a vital component for various image synthesis and manipulation applications [13, 29, 20, 25, 1]. *Style* is generally not manually designed by a user, but extracted from reference images. In most existing methods, styles are encoded in three places: i) statistics of image features [13, 42]; ii) neural network weights (*e.g.* fast style transfer [23, 51, 47]); iii) parameters of a network normalization layer [20, 27]). When applied to style control, the first encoding method is usually time-consuming as it requires a slow optimization process to match the statistics of image features extracted by image classification networks [13]. The second approach runs in real-time but each neural network only encodes the style of selected reference images [23]. Thus, a separate neural network is required

to be trained for each style image, which limits its application in practice. To date, the third approach is the best as it enables arbitrary style transfer in real-time [20] and it is used by high-quality networks such as StyleGAN [25], FUNIT [31], and SPADE [38]. Our per-region style encoding also builds on this approach. We will show that our method generates higher quality results and enables more-detailed user control.

3. Per Region Style Encoding and Control

Given an input style image and its corresponding segmentation mask, this section shows i) how to distill the per-region styles from the image according to the mask and ii) how to use the distilled per-region style codes to synthesize photo-realistic images.

3.1. How to Encode Style?

Per-Region Style Encoder. To extract per region styles, we propose a novel style encoder network to distill the corresponding style code from each semantic region of an input image simultaneously (See the subnetwork *Style Encoder* in Figure 4 (A)). The output of the style encoder is a $512 \times s$ dimensional style matrix $\mathbf{ST}$, where s is the number of semantic regions in the input image. Each column of the matrix corresponds to the style code of a semantic region. Unlike the standard encoder built on a simple downscaling convolutional neural network, our per-region style encoder employs a “bottleneck” structure to remove the information irrelevant to styles from the input image. Incorporating the prior knowledge that styles should be independent of the shapes of semantic regions, we pass the intermediate feature maps (512 channels) generated by the network block *TConv-Layers* through a region-wise average pooling layer and reduce them to a collection of 512 dimensional vectors. As implementation detail we would like to remark that we use s as the number of semantic labels in the data set and set columns corresponding to regions that do not exist in an input image to 0. As a variation of this architecture, we can also extract style per region instance for datasets that have instance labels, *e.g.*, CityScapes.

3.2. How to Control Style?

With per-region style codes and a segmentation mask as inputs, we propose a new conditional normalization technique called Semantic Region-Adaptive Normalization (SEAN) to enable detailed control of styles for photo-realistic image synthesis. Similar to existing normalization techniques [20, 38], SEAN works by modulating the scales and biases of generator activations. In contrast to all existing methods, the modulation parameters learnt by SEAN are dependent on both the style codes and segmentation masks. In a SEAN block (Figure 3), a style map is first gen-

Figure 3: SEAN normalization. The input are style matrix $\mathbf{ST}$ and segmentation mask $\mathbf{M}$. In the upper part, the style codes in $\mathbf{ST}$ undergo a per style convolution and are then broadcast to their corresponding regions according to $\mathbf{M}$ to yield a style map. The style map is processed by conv layers to produce per pixel normalization values γ^s and β^s . The lower part (light blue layers) creates per pixel normalization values using only the region information similar to SPADE.

erated by broadcasting the style codes to their corresponding semantic regions according to the input segmentation mask. This style map along with the input segmentation mask are then passed through two separate convolutional neural networks to learn two sets of modulation parameters. The weighted sums of them are used as the final SEAN parameters to modulate the scales and biases of generator activations. The weight parameter is also learnable during training. The formal definition of SEAN is introduced as follows.

Semantic Region-Adaptive Normalization (SEAN). A SEAN block has two inputs: a style matrix $\mathbf{ST}$ encoding per-region style codes and a segmentation mask $\mathbf{M}$. Let $\mathbf{h}$ denote the input activation of the current SEAN block in a deep convolutional network for a batch of N samples. Let H , W and C be the height, width and the number of channels in the activation map. The modulated activation value at site $(n \in N, c \in C, y \in H, x \in W)$ is given by

$$\gamma_{c,y,x}(\mathbf{ST}, \mathbf{M}) \frac{h_{n,c,y,x} - \mu_c}{\sigma_c} + \beta_{c,y,x}(\mathbf{ST}, \mathbf{M}) \quad (1)$$

where $h_{n,c,y,x}$ is the activation at the site before normalization, the modulation parameters $\gamma_{c,y,x}$ and $\beta_{c,y,x}$ are weighted sums of $\gamma_{c,y,x}^s$, $\gamma_{c,y,x}^o$ and $\beta_{c,y,x}^s$, $\beta_{c,y,x}^o$ respectively (See Fig. 3 for definition of γ and β variables):

$$\begin{aligned} \gamma_{c,y,x}(\mathbf{ST}, \mathbf{M}) &= \alpha_\gamma \gamma_{c,y,x}^s(\mathbf{ST}) + (1 - \alpha_\gamma) \gamma_{c,y,x}^o(\mathbf{M}) \\ \beta_{c,y,x}(\mathbf{ST}, \mathbf{M}) &= \alpha_\beta \beta_{c,y,x}^s(\mathbf{ST}) + (1 - \alpha_\beta) \beta_{c,y,x}^o(\mathbf{M}) \end{aligned} \quad (2)$$

μ_c and σ_c are the mean and standard deviation of the activation in channel c :

variation in channel c :

$$\mu_c = \frac{1}{NHW} \sum_{n,y,x} h_{n,c,y,x} \quad (3)$$

$$\sigma_c = \sqrt{\frac{1}{NHW} \left(\sum_{n,y,x} h_{n,c,y,x}^2 \right) - \mu_c^2} \quad (4)$$

4. Experimental Setup

4.1. Network Architecture

Figure 4 (A) shows an overview of our generator network, which is building on that of SPADE [38]. Similar to [38], we employ a generator consisting of several SEAN ResNet blocks (SEAN ResBlk) with upsampling layers.

SEAN ResBlk. Figure 4 (B) shows the structure of our SEAN ResBlk, which consists of three convolutional layers whose scales and biases are modulated by three SEAN blocks respectively. Each SEAN block takes two inputs: a set of per-region style codes $\mathbf{ST}$ and a semantic mask $\mathbf{M}$. Note that both inputs are adjusted in the beginning: the input segmentation mask is downsampled to the same height and width of the feature maps in a layer; the input style codes from $\mathbf{ST}$ are transformed per region using a 1×1 conv layer A_{ij} . We observed that the initial transformations are indispensable components of the architecture because they transform the style codes according to the different roles of each neural network layer. For example, early layers might control the hair styles (e.g. wavy, straight) of human face images while later layers might refer to the lighting and color. In addition, we observed that adding noise to the input of SEAN can improve the quality of synthesized images. The scale of such noise is adjusted by a per-channel scaling factor B that is learnt during training, similar to StyleGAN [25].

4.2. Training and Inference

Training We formulate the training as an image reconstruction problem. That is, the style encoder is trained to distill per-region style codes from the input images according to their corresponding segmentation masks. The generator network is trained to reconstruct the input images with the extracted per-region style codes and the corresponding segmentation masks as inputs. Following SPADE [38] and Pix2PixHD [44], the difference between input images and reconstructed images are measured by an overall loss function consisting of three loss terms: conditional adversarial loss, feature matching loss [44] and perceptual loss [23]. Details of the loss function are included in the supplementary material.

Inference During inference, we take an arbitrary segmentation mask as the mask input and implement per-region style

Figure 4: SEAN generator. (A) On the left, the style encoder takes an input image and outputs a style matrix ST . The generator on the right consists of interleaved SEAN ResBlocks and Upsampling layers. (B) A detailed view of a SEAN ResBlock used in (A).

Method	CelebAMask-HQ			CityScapes			ADE20K			Façades		
	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR
Pix2PixHD [44]	0.68	0.15	17.14	0.69	0.13	18.32	0.51	0.22	13.81	0.53	0.16	16.30
SPADE [38]	0.63	0.21	14.30	0.64	0.18	15.77	0.45	0.28	11.52	0.44	0.22	13.87
Ours	0.73	0.12	18.74	0.70	0.13	18.61	0.58	0.17	16.16	0.58	0.14	17.14

Table 1: Quantitative comparison of reconstruction quality. Our method outperforms current leading methods using similarity metrics SSIM, RMSE, and PSNR on all the datasets. For SSIM and PSNR, higher is better. For RMSE, lower is better.

control by selecting a separate 512 dimensional style code for each semantic region as the style input. This enables a variety of high quality image synthesis applications, which will be introduced in the following section.

5. Results

In the following we discuss quantitative and qualitative results of our framework.

Implementation details. Following SPADE [38], we apply Spectral Norm [36] in both the generator and discriminator. Additional normalization is performed by SEAN in the generator. We set the learning rates to 0.0001 and 0.0004 for the generator and discriminator, respectively [17]. For the optimizer, we choose ADAM [26] with $\beta_1 = 0, \beta_2 = 0.999$. All the experiments are trained on 4 Tesla v100 GPUs. To get better performance, we use a synchronized version of batch normalization [49] in the SEAN normalization

blocks.

Datasets. We use the following datasets in our experiments: 1) CelebAMask-HQ [28, 24, 32] containing 30,000 segmentation masks for the CelebAHQ face image dataset. There are 19 different region categories. 2) ADE20K [50] contains 22,210 images annotated with 150 different region labels. 3) Cityscapes [10] contains 3500 images annotated with 35 different region labels. 4) Façades dataset. we use a dataset of 30,000 façade images collected from Google Street View [3]. Instead of manual annotation, the segmentation masks are automatically calculated by a pre-trained DeepLabV3+ Network [7].

We follow the recommended splits for each dataset and note that all images in the paper are generated from the test set only. Neither the segmentation masks nor the style images have been seen during training.

Metrics. We employ the following established metrics to compare our results to state of the art: 1) segmentation ac-

Figure 5: Style interpolation. We take a mask from a source image and reconstruct with two different style images (Style1 and Style2) that are very different from the source image. We then show interpolated results of the per-region style codes

Method	CelebAMask-HQ			CityScapes			ADE20K			Façades
	mIoU	accu	FID	mIoU	accu	FID	mIoU	accu	FID	FID
Ground Truth	73.14	94.38	9.41	66.21	93.69	32.34	39.38	78.76	14.51	14.40
Pix2PixHD [44]	76.12	95.76	23.69	50.35	92.09	83.24	22.78	73.32	43.0	22.34
SPADE [38]	77.01	95.93	22.43	56.01	93.13	60.51	35.37	79.37	34.65	24.04
Ours	75.69	95.69	17.66	57.88	93.59	50.38	34.59	77.16	24.84	19.82

Table 2: Quantitative comparison using semantic segmentation performance measured by mIoU and accu and generation performance measured by FID. Our method outperforms current leading methods in FID on all the datasets.

curacy measured by mean Intersection-over-Union (mIoU) and pixel accuracy (accu), 2) FID [17], 3) peak signal-to-noise ratio (PSNR), 4) structural similarity (SSIM) [45], 5) root mean square error (RMSE).

Quantitative comparisons. In order to facilitate a fair comparison to SPADE, we report reconstruction perfor-

Method	SSIM	RMSE	PSNR	FID
Pix2PixHD [44]	0.68	0.15	17.14	23.69
SPADE [38]	0.63	0.21	14.30	22.43
SPADE++	0.67	0.16	16.71	20.80
SEAN-level encoders	0.74	0.11	19.70	18.17
ResBlk-level encoders	0.72	0.12	18.86	17.98
unified encoder	0.73	0.12	18.74	17.66
w/ downsampling	0.73	0.12	18.74	17.66
w/o downsampling	0.72	0.13	18.32	18.67
w/ noises	0.73	0.12	18.74	17.66
w/o noises	0.72	0.13	18.49	18.05

Table 3: Ablation study on CelebAMask-HQ dataset. See supplementary materials for further details.

mance where only a single style image is employed. We train one network for each of the four datasets and report the previously described metrics in Table 1 and 2. We choose SPADE [38] as the currently best state-of-the-art method and Pix2PixHD [44] as the second best method in the comparisons. Based on visual inspection of the results, we found that the FID score is most indicative of visual quality of the results and that a lower FID score often (but not always) corresponds to better visual results. Generally, we can observe that our SEAN framework clearly beats the state of the art on all datasets.

Qualitative results. We show visual comparisons for the four datasets in Fig. 6. We can observe that the quality difference between our work and the state of the art can be significant. For example, we notice that SPADE and Pix2Pix cannot handle more extreme poses and older people. We conjecture that our network is more suitable to learn more of the variability in the data because of our improved style encoding. Since a visual inspection of the results is very important for generative modeling, we would like to refer the reader to the supplementary materials and the accompanying video for more results. Our per-region style encoding

Figure 6: Visual comparison of semantic image synthesis results on the CelebAMask-HQ, ADE20K, CityScapes and Façades dataset. We compare Pix2PixHD, SPADE, and our method. See supplementary materials for further details.

Figure 7: Style crossover. In addition to style interpolation (bottom row), we can perform crossover by selecting different styles per ResBlk. We show two transitions in the top two rows. The blue / orange bars on top of the images indicate which styles are used by the six ResBlks. We can observe that earlier layers are responsible for larger features and later layers mainly determine the color scheme.

also enables new image editing operations: iterative image editing with per-region style control (See Figs. 1 and 2), style interpolation and style crossover (See Fig. 5 and 7).

Variations of SEAN generator (Ablation studies). In Table 3, we compare many variations of our architecture with previous work on the CelebAMask-HQ dataset.

According to our analysis, Pix2PixHD has a better style extraction subnetwork than SPADE (because the PSNR reconstruction values are better), but SPADE has a better generator subnetwork (because the FID scores are better). We therefore build another network that combines the style encoder from Pix2PixHD with the generator from SPADE and call this architecture SPADE++. We can observe that SPADE++ improves SPADE by a bit, but all our architecture variations still have a better FID and PSNR score.

To evaluate our design choices, we report the performance of variations of our encoder. First we compare three variants to encode style: SEAN-level encoder, ResBlk-level encoder, and unified encoder. A fundamental design choice is how to split the style computation into two parts: a shared computation in the style encoder and a per layer computation in each SEAN block. The first two variants perform all style extraction per-layer, while the unified encoder is the architecture described in the main part of the paper. While the other two encoders have better reconstruction performance, they lead to lower FID scores. Based on visual inspection, we also confirmed that the unified encoder leads to better visual quality. This is especially noticeable for difficult inputs (see supplementary materials). Second, we evaluate if downsampling in the style encoder is a good idea and evaluate a style encoder with bottleneck (with downsampling) with another encoder that uses less downsampling. This test confirms that introducing a bottleneck in the style

encoder is a good idea. Lastly, we test the effect of the learnable noise. We found the learnable noise can help both similarity (PSNR) and FID performance. More details of the alternative architectures are provided in supplementary materials.

6. Conclusion

We propose semantic region-adaptive normalization (SEAN), a simple but effective building block for Generative Adversarial Networks (GANs) conditioned on segmentation masks that describe the semantic regions in the desired output image. Our main idea is to extend SPADE, the currently best network, to control the style of each semantic region individually, e.g. we can specify one style reference image per semantic region. We introduce a new building block, SEAN normalization, that can extract style from a given style image for a semantic region and processes the style information to bring it in the form of spatially-varying normalization parameters. We evaluate SEAN on multiple datasets and report better quantitative metrics (e.g. FID, PSNR) than the current state of the art. While SEAN is designed to specify style for each region independently, we also achieve big improvements in the basic setting where only one style image is provided for all regions. In summary, SEAN is better suited to encode, transfer, and synthesize style than the best previous method in terms of reconstruction quality, variability, and visual quality. SEAN also pushes the frontier of interactive image editing. In future work, we plan to extend SEAN to processing meshes, point clouds, and textures on surfaces.

Acknowledgement This work was supported by the KAUST Office of Sponsored Research (OSR) under Award No. OSR-CRG2018-3730.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4432–4441, 2019. 3
- [2] Yazeed Alharbi, Neil Smith, and Peter Wonka. Latent filter scaling for multimodal unsupervised image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1458–1466, 2019. 2
- [3] Dragomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stéphane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, and Josh Weaver. Google street view: Capturing the world at street level. *Computer*, 43(6):32–38, 2010. 5
- [4] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017. 3
- [5] David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. Semantic photo manipulation with a generative image prior. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 38(4), 2019. 3
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis, 2018. 3
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation, 2018. 5
- [8] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks, 2017. 2
- [9] Yunje 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, 2017. 3
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5
- [11] Ugur Demir and Gozde Unal. Patch-based image inpainting with generative adversarial networks. *arXiv preprint arXiv:1803.07422*, 2018. 3
- [12] Anna Frühstück, Ibraheem Alhashim, and Peter Wonka. Tilegan. *ACM Transactions on Graphics*, 38(4):1–11, Jul 2019. 3
- [13] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. A neural algorithm of artistic style. *arXiv*, Aug 2015. 3
- [14] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. 3
- [15] Shuyang Gu, Jianmin Bao, Hao Yang, Dong Chen, Fang Wen, and Lu Yuan. Mask-guided portrait editing with conditional gans. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- [16] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017. 3
- [17] 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, 2017. 5, 6
- [18] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7986–7994, 2018. 3
- [19] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization, 2017. 2
- [20] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 3
- [21] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation, 2018. 3
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks, 2016. 2, 3
- [23] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution, 2016. 3, 4
- [24] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation, 2017. 2, 5
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2018. 2, 3, 4
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. 5
- [27] Dmytro Kotovenko, Artsiom Sanakoyeu, Sabine Lang, and Bjorn Ommer. Content and style disentanglement for artistic style transfer. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 3
- [28] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. *arXiv preprint arXiv:1907.11922*, 2019. 2, 3, 5
- [29] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III*, 2016. 3
- [30] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks, 2017. 3
- [31] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation, 2019. 2, 3
- [32] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 2, 5
- [33] Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares gen-

- erative adversarial networks. *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 3
- [34] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge?, 2018. 3
- [35] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 3
- [36] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks, 2018. 3, 5
- [37] Takeru Miyato and Masanori Koyama. cgans with projection discriminator, 2018. 3
- [38] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2, 3, 4, 5, 6, 7
- [39] Xiaojuan Qi, Qifeng Chen, Jiaya Jia, and Vladlen Koltun. Semi-parametric image synthesis, 2018. 2
- [40] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 3
- [41] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*, 2016. 3
- [42] Omry Sendik and Daniel Cohen-Or. Deep correlations for texture synthesis. *ACM Transactions on Graphics (TOG)*, 36(5):161, 2017. 3
- [43] Ron Slossberg, Gil Shamaï, and Ron Kimmel. High quality facial surface and texture synthesis via generative adversarial networks. In *European Conference on Computer Vision*, pages 498–513. Springer, 2018. 3
- [44] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 3, 4, 5, 6, 7
- [45] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE TRANSACTIONS ON IMAGE PROCESSING*, 13(4):600–612, 2004. 6
- [46] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1316–1324, 2018. 3
- [47] Shuai Yang, Zhangyang Wang, Zhaowen Wang, Ning Xu, Jiaying Liu, and Zongming Guo. Controllable artistic text style transfer via shape-matching gan. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 3
- [48] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5505–5514, 2018. 3
- [49] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 5
- [50] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 5
- [51] Yang Zhou, Zhen Zhu, Xiang Bai, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Non-stationary texture synthesis by adversarial expansion. *arXiv preprint arXiv:1805.04487*, 2018. 3
- [52] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative visual manipulation on the natural image manifold. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016. 3
- [53] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks, 2017. 3
- [54] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, 2017. 3