This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Freestyle Layout-to-Image Synthesis

Han Xue^{1,2} Zhiwu Huang^{2,3} Qianru Sun² Li Song^{1,4*} Wenjun Zhang¹ ¹School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University ²Singapore Management University ³University of Southampton ⁴MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University {xue_han, song_li, zhangwenjun}@sjtu.edu.cn gianrusun@smu.edu.sg zhiwu.huang@soton.ac.uk



Figure 1. Freestyle Layout-to-Image Synthesis (FLIS) results generated by using our model. Each has two kinds of inputs: a layout of semantic masks (on the 1st column), and a text (on the top of each result). For each layout, we show three example results with edited texts (3rd-5th columns). They validate that our model is able to introduce new attributes (3rd column), styles (4th column), and objects (5th column), which are all unseen during training, in the synthesized images. The generated hornless unicorn is due to the layout constraint.

Abstract

Typical layout-to-image synthesis (LIS) models generate images for a closed set of semantic classes, e.g., 182 common objects in COCO-Stuff. In this work, we explore the freestyle capability of the model, i.e., how far can it generate unseen semantics (e.g., classes, attributes, and styles) onto a given layout, and call the task Freestyle LIS (FLIS). Thanks to the development of large-scale pretrained language-image models, a number of discriminative models (e.g., image classification and object detection) trained on limited base classes are empowered with the ability of unseen class prediction. Inspired by this, we opt to leverage large-scale pre-trained text-to-image diffusion models to achieve the generation of unseen semantics. The key challenge of FLIS is how to enable the diffusion model to synthesize images from a specific layout which very likely violates its pre-learned knowledge, e.g., the model never sees "a unicorn sitting on a bench" during its pre-training. To this end, we introduce a new module called Rectified Cross-Attention (RCA) that can be conveniently plugged in the diffusion model to integrate semantic masks. This "plug-in" is applied in each cross-attention layer of the model to rectify the attention maps between image and text tokens. The key idea of RCA is to enforce each text token to act on the pixels in a specified region, allowing us to freely

^{*}Corresponding author. Part of this work was done when Han Xue served as a visiting Ph.D. student at Singapore Management University.

put a wide variety of semantics from pre-trained knowledge (which is general) onto the given layout (which is specific). Extensive experiments show that the proposed diffusion network produces realistic and freestyle layout-to-image generation results with diverse text inputs, which has a high potential to spawn a bunch of interesting applications. Code is available at https://github.com/essunny310/FreestyleNet.

1. Introduction

Layout-to-image synthesis (LIS) is one of the prevailing topics in the research of conditional image generation. It aims to generate complex scenes where a user requires fine controls over the objects appearing in a scene. There are different types of layouts including bboxes+classes [13], semantic masks+classes [29, 37], key points+attributes [33], skeletons [26], to name a few. In this work, we focus on using semantic masks to control object shapes and positions, and using texts to control classes, attributes, and styles. We aim to offer the user the most fine-grained controls on output images. For example, in Figure 1 (2nd column), the train and the cat are generated onto the right locations defined by the input semantic layouts, and they should also be aligned well with their surroundings.

However, most existing LIS models [24, 29, 37, 49, 52] are trained on a specific dataset from scratch. Therefore, the generated images are limited to a closed set of semantic classes, e.g., 182 objects on COCO-Stuff [3]. While in this paper, we explore the *freestyle* capability of the LIS models for the generation of unseen classes (including their attributes and styles) onto a given layout. For instance, in Figure 1 (5th column), our model stretches and squeezes a "warehouse" into a mask of "train", and the "warehouse" is never seen in the training dataset of LIS. In addition, it is also able to introduce novel attributes and styles into the synthesized images (in the 3rd and 4th columns of Figure 1). The *freestyle* property aims at breaking the in-distribution limit of LIS, and it has a high potential to enlarge the applicability of LIS results to out-of-distribution domains such as data augmentation in open-set or long-tailed semantic segmentation tasks.

To overcome the in-distribution limit, large-scale pretrained language-image models (*e.g.*, CLIP [31]) have shown the ability of learning general knowledge on an enormous amount of image-text pairs. The learned knowledge enables a multitude of discriminative methods trained with limited base classes to predict unseen classes. For instance, the trained models in [31,51,60] based on pre-trained CLIP are able to do zero-shot or few-shot classification, which is achieved by finding the most relevant text prompt (*e.g.*, *a photo of a cat*) given an image. However, it is non-trivial to apply similar ideas to generative models. In classification, simple texts (*e.g.*, *a photo of a cat*) can be conveniently mapped to labels (*e.g.*, *cat*). It is, however, not intuitive to achieve this mapping from compositional texts (*e.g.*, *a train is running on the railroad with grass on both sides*) to synthesized images. Thanks to the development of pre-training in image generation tasks, we opt to leverage large-scale pre-trained language-image models to the downstream generation tasks.

To this end, we propose to leverage the large-scale pretrained text-to-image diffusion model to achieve the LIS of unseen semantics. These diffusion models are powerful in generating images of high quality that are in line with the input texts. This can provide us with a generative prior with a wide range of semantics. However, it is challenging to enable the diffusion model to synthesize images from a specific layout which very likely violates its pre-trained knowledge, e.g., the model never sees "a unicorn sitting on a bench" during its pre-training. Textual semantics should be correctly arranged in space without impairing their expressiveness. A seminal attempt, PITI [48], learns an encoder to map the layout to the textual space of a pre-trained text-toimage model. Although it demonstrates improved quality of the generated images, PITI has two main drawbacks. First, it abandons the text encoder of the pre-trained model, losing the ability to freely control the generation results using different texts. Second, its introduced layout encoder implicitly matches the space of the pre-trained text encoder and hence incurs inferior spatial alignment with input layouts.

In this paper, we propose a simple yet effective framework for the suggested freestyle layout-to-image synthesis (FLIS) task. We introduce a new module called Rectified Cross-Attention (RCA) and plug it into a pre-trained textto-image diffusion model. This "plug-in" is applied in each cross-attention layer of the diffusion model to integrate the input layout into the generation process. In particular, to ensure the semantics (described by text) appear in a region (specified by layout), we find that image tokens in the region should primarily aggregate information from the corresponding text tokens. In light of this, the key idea of RCA is to utilize the layout to rectify the attention maps computed between image and text tokens, allowing us to put desired semantics from the text onto a layout. In addition, RCA does not introduce any additional parameters into the pre-trained model, making our framework an effective solution to freestyle layout-to-image synthesis.

As shown in Figure 1, the proposed diffusion network (FreestyleNet) allows the freestyle synthesis of high-fidelity images with novel semantics, including but not limited to: synthesizing unseen objects, binding new attributes to an object, and rendering the image with various styles. Our main contributions can be summarized as follows:

• We propose a novel LIS task called FLIS. In this task, we exploit large-scale pre-trained text-to-image diffusion models with layout and text as conditions.

- We introduce a parameter-free RCA module to plug in the diffusion model, allowing us to generate images from a specified layout effectively while taking full advantage of the model's generative priors.
- Extensive experiments demonstrate the capability of our approach to translate a layout into high-fidelity images with a wide range of novel semantics, which is beyond the reach of existing LIS models.

2. Related Work

Deep generative models like Generative Adversarial Networks (GANs) [2, 10, 20] and diffusion models [5, 14, 40–43] have achieved impressive results for image synthesis. Compared to unconditional synthesis which yields an image from a random noise, conditional image synthesis models [7, 58] allow more controllability over the generation. Among them, semantic image synthesis and text-based image synthesis are the most relevant to this paper.

Semantic Image Synthesis. This task aims to generate images based on the given semantic layouts [1, 4, 19, 24, 27, 29, 30, 39, 44–47, 49, 50, 52, 61–63]. A pioneer work, Pix2Pix [19], utilizes an encoder-decoder generator and a PatchGAN discriminator to perform semantic image synthesis. Pix2PixHD [49] enhances Pix2Pix with coarse-to-fine and multi-scale network architectures for high-resolution image synthesis. SPADE [29] suggests modulating the activations in normalization layers using the affine parameters predicted from the input semantic maps. Instead of modifying the features, CC-FPSE [24] and SC-GAN [52] learn to produce convolutional kernels and semantic vectors from the semantic maps to condition the generation. OASIS [44] designs a segmentation-based discriminator to synthesize images with a better alignment to the input label maps. Despite their impressive performance, these methods are limited in generating fully controlled images with diverse semantics. While in this paper, we propose a freestyle layout-to-image synthesis framework that is armed with much more controllability, enabling 1) the generation of semantics beyond the pre-defined semantic categories in the training dataset, and 2) the separate modulation of each class in the layout with text. A recent work, PITI [48], leverages a pre-trained diffusion model to facilitate imageto-image translation, achieving impressive results. However, its generation is restricted by single-modality condition such as mask.

Text-based Image Synthesis. Generating images from texts is another popular task of image synthesis. Significant progress has been made by exploiting progressive generation [56], attention mechanisms [53], and cross-modal contrastive approaches [55]. Recent text-to-image synthesis methods such as GLIDE [28], DALL-E2 [32], Imagen [36], Parti [54], CogView2 [6], and Stable Diffusion [34] are built

upon either auto-regressive or diffusion models. Trained on large-scale image-text pairs, these methods demonstrate an astonishing ability to generate high-fidelity images from Nevertheless, with only text as input, they fall texts. short of fine-grained control over the synthesized results (e.g., specifying the layout). To advance this, several approaches [13, 16, 22, 33] attempt to introduce the location of objects and generate images in a hierarchical manner, but the controls are limited to ambiguous locations (indicated by bboxes). In contrast, our method manages to generate semantics with fine-grained (pixel-level) control. PoE-GAN [17] and Make-A-Scene [8] accept both text and layouts, but they support only in-distribution generation. In addition, PoE-GAN [17] merely performs global text-based manipulations. Our proposed method supports text-guided and out-of-distribution LIS.

In addition, some recent works succeed in applying the pre-trained text-to-image diffusion models to various applications including text-guided image editing [11, 21], image inpainting [25], and personalized generation [9, 35], but none of them can fulfill our goal. In this paper, we excavate the general knowledge learned in a pre-trained text-to-image diffusion model to enable the layout-to-image synthesis to be performed in a freestyle manner.

3. Problem Definition

The objective of general layout-to-image synthesis (LIS) using masks is to learn a mapping function f(l) = x that generates a realistic image x from the input mask l. This mapping function is expected to generate semantics onto the given mask, and is usually trained in a specific dataset. The generated semantics are thus highly constrained by the dataset. To achieve the generation of freestyle semantics, we define a new task called freestyle layout-to-image synthesis (FLIS), and introduce a feasible solution for it. Specifically, we introduce to condition the model with not only masks but also texts. Masks control the layout of semantics, while texts elaborate what specific semantics to "put onto" the layout. Therefore, FLIS mapping function f can be formulated as f(l, y) = x, where l and y are the input masks and texts respectively.

An intuitive approach to initializing the input text y is to use the labels of semantic masks in l, e.g., "train bush grass railroad" in the first example of Figure 1. In any specific image generation dataset, the vocabulary of these labels is fixed, e.g., 182 objects on COCO-Stuff [3]. To go beyond this, we opt to adopt the large-scale diffusion models such as Stable Diffusion [34] that have seen far more labels during their pre-training on billions of open data (e.g., LAION-5B [38]). We expect to leverage the general knowledge of such models to enable the generation of open-set semantics in the task of FLIS.

The main challenge is how to synthesize images from



Figure 2. An overview of the proposed FreestyleNet. By applying Rectified Cross-Attention (RCA) in the pre-trained textto-image diffusion model, Stable Diffusion [34], our method allows for freestyle layout-to-image synthesis, where a layout can be translated into realistic and creative images with a wide range of semantics specified by the input texts. Note that the base text shown here is simply derived from the semantic layout in order to train our network on image-layout pairs. During inference, we can freely edit the text to generate diverse images. Here the text encoder and the diffusion model are adopted from Stable Diffusion and we omit the autoencoder of Stable Diffusion for simplicity.

a specific layout which very likely violates the pre-trained knowledge of such models, *e.g.*, the model has seen the images of "unicorn" but never seen "the unicorn sitting on a bench" (with the exact relative location of "unicorn" and "bench" being specified) during its pre-training. To solve the issue, we propose a framework with Rectified Cross-Attention (RCA) to explicitly "inject" a specific layout (a sequence of masks) into the generation process of pre-trained diffusion models. The key idea is to ensure that the cross-attention of each text token is only performed in the region defined by its corresponding mask. The proposed framework will be elaborated in the next section.

4. Proposed Method

We propose a framework of FLIS (FreestyleNet) by leveraging the pre-trained text-to-image diffusion network: Stable Diffusion [34]. An overview of our method is presented in Figure 2. The original Stable Diffusion adopts a text encoder of CLIP [31] to extract text embeddings from the input text and then injects them into a diffusion model through cross-attention to perform generation. Starting from a noise x_T , the reverse diffusion is performed using a denoising U-Net inside the model for T time steps to generate the final result x_0 . To empower Stable Diffusion with the ability of synthesizing images on a specified layout, we propose a Rectified Cross-Attention (RCA) to integrate the layout in each cross-attention layer of the U-Net. In the following, we will elaborate on 1) how to rectify the pre-trained text-to-image diffusion model in the context of FLIS, and 2) how to further exploit the freestyle capability of the rectified diffusion model for FLIS.

4.1. Rectifying Diffusion Model

We propose to rectify the pre-trained diffusion model to facilitate FLIS. The rectification is performed by: 1) representing semantics by using text embeddings, 2) injecting semantics to layout through the proposed RCA, and 3) finetuning the diffusion model on layout-based training data.

Representing Semantics. The rich semantics learned in the pre-trained text-to-image model lie in the textual space. To deploy them for FLIS, we need to represent the desired semantics by texts. Here we provide a simple yet effective solution. We first describe each semantic (*e.g.*, an object class from COCO-Stuff [3]) with a word (*e.g.*, "cat") or several words (*e.g.*, "teddy bear"), which we call a *concept*. Then, we stack all *concepts* corresponding to the classes appeared in the input semantic layout ("sky mountain snow person" in Figure 2), and input them into the text encoder. The resulting text embeddings serve as the semantics to perform the following generation.

Injecting Semantics to Layout. Given the text embeddings (*i.e.*, semantics), the next step is to inject the semantics into the input layout. However, it is not intuitive how to achieve that on a pre-trained text-to-image diffusion model.

Before diving into our solution, we elaborate the crossattention layer in this pre-trained model. The crossattention layer takes two inputs: text embeddings and image features. It calculates an attention map between them, and outputs a weighted sum on the text values using the weights on the attention map. Each channel of the attention map is related to a text embedding. For instance, we can find the channel corresponding to a text embedding extracted from the word "cat". The spatial layout of this cat in the generated image is similar to this channel, as concluded in [11,18]. Therefore, by modifying this channel, we can change the shape and location of this cat in the generated image. Inspired by this, we propose a Rectified Cross-Attention (RCA) to integrate the input layout into the image generation process. The idea of RCA is to rectify the attention maps of each text embedding using the corresponding mask. In this way, it determines the spatial layout of the generated image.

As illustrated in Figure 3, the input to RCA consists of three parts: input image feature, text embeddings, and layout. RCA computes image queries Q, text keys K, and text values V through three projection layers, respectively. After that, attention score maps \mathcal{M} are calculated as:

$$\mathcal{M} = \frac{QK^T}{\sqrt{d}} \in \mathbb{R}^{C \times H \times W},\tag{1}$$

where d is the scaling factor that is set as the dimension of the queries and keys, and C, H, W are the channel number, height, and weight of \mathcal{M} , respectively. Next, we utilize the layout to constrain the spatial distribution of the attention maps by rectifying the attention.



Figure 3. An illustration of Rectified Cross-Attention (RCA). It utilizes the layout to rectify the attention score maps calculated between the text and image tokens. By forcing each text token to affect only pixels in the region specified by the layout, the spatial alignment between the generated image and the given layout is guaranteed. Q, K, V are image queries, text keys, and text values, which are computed through projection layers respectively.

We translate the input one-channel semantic layout into a *C*-channel layout, where each channel is a binary map for an individual *concept*. For the text embedding corresponding to the *concept*, we use its associated semantic mask as the binary map. For other embeddings, *e.g.*, paddings, we set the values on the binary map to 1. To match the spatial size of the attention score maps \mathcal{M} , we resize the *C*-channel layout to obtain $L \in \mathbb{R}^{C \times H \times W}$. \mathcal{M} is then rectified by:

$$\widehat{\mathcal{M}}_{i,j}^k = \begin{cases} \mathcal{M}_{i,j}^k, & L_{i,j}^k = 1, \\ -inf, & L_{i,j}^k = 0, \end{cases}$$
(2)

where $\mathcal{M}_{i,j}^k$ and $L_{i,j}^k$ denote the values at position (i, j) in the *k*-th channels of \mathcal{M} and *L*, respectively. $\widehat{\mathcal{M}}$ represents the rectified attention score maps. Thereafter, we can calculate the output image feature \mathcal{O} of this RCA layer by:

$$\mathcal{O} = \operatorname{softmax}(\widehat{\mathcal{M}})V.$$
(3)

After applying softmax, the rectified attention map becomes a binary map with the spatial distribution similar to the input layout. Using this map, the model forces each text embedding to affect only pixels in the region specified by the input layout. In other words, the semantics from the text are injected to the layout.

Model Fine-tuning. Equipped with the proposed RCA, the pre-trained diffusion model can be fine-tuned on any sets of layout-based training data. The objective for fine-tuning is the same as that for pre-training (image denoising) [34]:

$$\mathcal{L}(\theta) = \mathbb{E}_{z,l,y,\epsilon \sim \mathcal{N}(0,1),t} \Big[\|\epsilon - \epsilon_{\theta}(z_t, t, l, c_{\phi}(y))\|_2^2 \Big], \quad (4)$$

where z is the latent code extracted from the input image, l is the input layout, y is the input text, ϵ is a noise term, t is the time step, ϵ_{θ} is the denoising U-Net, z_t is the noisy version of z at time t, and c_{ϕ} is the text encoder. We merely fine-tune the denoising U-Net while keeping the text encoder and the autoencoder of Stable Diffusion frozen. We refer readers to the original paper of Stable Diffusion [34] for the training details.

4.2. Generating Freestyle Images

The fine-tuned diffusion model is endowed with the ability to synthesize images given both layouts and texts. By describing the desired semantics in the input text, we generalize the model by excavating its pre-trained knowledge to perform freestyle layout-to-image synthesis. Here we showcase three capabilities of the proposed model for FLIS as well as the mechanisms behind them.

Binding New Attributes. Having a layout that contains an object of interest (*e.g.*, cat), one may further want to generate a specific species of cat (*e.g.*, tabby cat). In this scenario, we first replace "cat" with "tabby cat" in the text input. Then, following Eq. (2), we adopt the binary map of the input layout L (corresponding to the region you want to place this tabby cat) to rectify the two attention score maps, which are related to "tabby" and "cat" respectively.

Specifying the Style. We can render a given image in the desired style by adding a description (*e.g.*, "an ink painting of") to the input text. Since we aim to specify the global style of the generated image, the attention score maps corresponding to these added words are not rectified.

Generating Unseen Objects. Benefiting from the general



Figure 4. Example results and capabilities from our FreestyleNet in performing freestyle layout-to-image synthesis (FLIS). Zoom in for a better view.

knowledge learned from the pre-training phase, our model is able to generate objects from unseen classes. This is achieved by simply describing the object in the text input and rectifying the corresponding attention map through Eq. (2) with the layout related to this object.

5. Experiments

5.1. Experimental Settings

Datasets. We conduct experiments on two challenging datasets: COCO-Stuff [3] and ADE20K [59]. COCO-Stuff consists of 118,287 training and 5,000 validation images, which are annotated with 182 semantic classes. ADE20K has 20,210 images for training and 2,000 for validation, including 150 semantic classes.

Implementation Details. We adopt the original training settings of Stable Diffusion [34] to optimize our model, unless otherwise noted. The base learning rate is set to 10^{-6} . We use 50 PLMS [23] sampling steps with a classifier-free guidance [15] scale of 2. All the experiments are conducted on a single NVIDIA A100 GPU.

5.2. Qualitative Evaluation on FLIS

In order to demonstrate the viability of our FreestyleNet for freestyle layout-to-image synthesis (FLIS), we showcase some synthesis results covering a variety of scenarios. As presented in Figures 1 and 4, the proposed model is capable of but not limited to 1) binding new attributes to the objects, 2) specifying the styles for the synthesized images, and 3) generating unseen objects. The model we use to generate these images is fine-tuned on COCO-Stuff. For each sample here, we conduct two kinds of layout-to-image synthesis, *i.e.*, in-distribution and out-of-distribution. For the former, the text input only contains the semantics that are defined in the 182 classes of COCO-Stuff. While for the out-of-distribution synthesis (depicted with a red box), we apply some words with a wide range of new semantics in the text input. Our method consistently produces high-fidelity images that faithfully reveal the novel semantics described in the text while conforming to the given layouts. Some impressive features enabled by our approach such as generating a "Bugatti Veyron" or hornless "unicorn" and putting



Figure 5. Qualitative comparison results with the state-of-the-art layout-to-image synthesis methods on COCO-Stuff (top 2 rows) and ADE20K (bottom 2 rows). Note that our FreestyleNet takes text and layout as input while other methods adopt only the layout.

the scene "in Minecraft" break the in-distribution limit, allowing the users to create the images in a *freestyle* way^{*}.

5.3. Comparison with LIS Baselines

We compare our method against the state-of-the-art LIS baselines including Pix2PixHD [49], SPADE [29], CC-FPSE [24], OASIS [44], SC-GAN [52], and PITI [48]. Notably, since generating unseen semantics is beyond the reach of these methods, we conduct the comparison under the in-distribution setting. While other methods use only one input modality (layout), the proposed model adopts two (text and layout). We follow the standard test settings to evaluate all methods on COCO-Stuff and ADE20K. Fréchet Inception Distance (FID) [12] and mean Intersection-over-Union (mIoU) are adopted to assess the realism of the generated images and the semantic alignment between the generated images and the input layouts, respectively.

Quantitative comparison results are reported in Table 1. The proposed model outperforms all the competing methods in terms of the FID metric, indicating the high visual quality of the images synthesized by our method. The most recent method, PITI, also exhibits superior generation quality by leveraging a pre-trained diffusion model. However, it implicitly learns to map the input layouts into the space of the pre-trained text encoder, resulting in clearly worse

Method	COCO-Stuff		ADE20K	
	FID↓	mIoU↑	FID↓	mIoU↑
Pix2PixHD [49]	111.5	14.6	81.8	20.3
SPADE [29]	22.6	37.4	33.9	38.5
CC-FPSE [24]	19.2	41.6	31.7	43.7
OASIS [44]	17.0	44.1	28.3	48.8
SC-GAN [52]	18.1	42.0	29.3	45.2
PITI [48]	16.1	34.1	27.9	29.4
FreestyleNet (ours)	14.4	40.7	25.0	41.9

Table 1. Quantitative comparison results with the state-of-the-art layout-to-image synthesis methods. Note that our model takes text and layout as input while other methods adopt only the layout.

spatial alignment (reflected in low mIoU values) than ours. Although our method does not suffer from this problem by explicitly constraining the spatial layout of image features through RCA, it obtains less preferred results in mIoU compared to some of the other LIS methods. We argue that mIoU is not fully suitable to evaluate the proposed FLIS method. While the text input does not contain unseen semantics under the in-distribution setting, our method may sometimes generate more general semantics that stem from the pre-trained knowledge (*e.g.*, the wall behind the switch in the 3rd row of Figure 5), leading to mispredictions of class labels by the segmentation model used to calculate mIoU. In addition, our goal is to synthesize images that are

 $^{^{*}\}mathrm{A}$ comparison to a concurrent work ControlNet [57] is provided on our GitHub page.



Figure 6. Effectiveness of the proposed RCA. AM represents attention maps in each RCA. By sharing/exchanging the attention maps rectified by RCA between different semantics, our method succeeds in local style assignment/appearance swapping.

spatially, not semantically, aligned with the input layout. As later shown in the visual comparison, our generated images have a strong spatial alignment with the input layouts.

Qualitative comparison results are shown in Figure 5. Compared to the other baselines, our method consistently produces sharper images with fine details, which is consistent with the quantitative evaluation results.

For a fair comparison, we replace PITI's backbone GLIDE [28] with Stable Diffusion. The comparison results (see Figure S7 and Table S1 in the supplementary) indicate that the clearly improved spatial alignment are due to RCA, and the image quality is also enhanced by our framework.

5.4. Effectiveness of RCA

To further validate the effectiveness of the proposed RCA, we show two representative examples in Figure 6. In the first row, if we share the rectified attention maps corresponding to "sky" with those associated with "drawn by Van Gogh" when specifying the image style, only the sky will be rendered in the style of Van Gogh. The second row showcases that our model is also capable of swapping the appearance of two objects by exchanging their corresponding rectified attention maps. Both of these two cases demonstrate that RCA can force the semantics described by the text to appear in the region specified by the layout.

5.5. Limitations

We find several failure cases of our method. As presented in Figure 7, our model cannot produce satisfactory results (distortion in kite and failed appearance swap) when the specified semantics or scenes are counterfactual (*e.g.*,



Figure 7. Failure cases. AM represents attention maps in each RCA. Our method may struggle with generating some rare semantics. More results are provided in the supplementary material.

marshmallow in the sky and horse riding a person). One possible reason is that the diffusion model itself does not learn sufficient knowledge of these rarely appeared or unreasonable semantics. Another reason could be the fine-tuning on downstream datasets, likely resulting in some damage to pre-trained knowledge. Our model cannot work as intended if given contradictory layout and prompt (*e.g.*, the unicorn in Figure 1 is hornless). A feasible solution is to relax the layout constraints, *e.g.*, by increasing the rectified attention score when $L_{i,j}^k = 0$ (indicating regions outside the *k*-th semantic) in Eq. (2). Besides, our method needs to have a user-defined set of category names, which could be expensive to collect in case of long-tailed datasets.

6. Conclusions

In this paper, we introduce a new task called freestyle layout-to-image synthesis (FLIS), which aims to generate unseen semantics described by text onto a given layout. To that end, we propose to leverage the generative prior from pre-trained text-to-image diffusion models. After inserting the proposed Rectified Cross-Attention (RCA) module, the pre-trained model is armed with the FLIS capability, breaking the in-distribution limit that hinders the applicability of previous LIS models. We demonstrate a plethora of applications by translating a single layout into high-fidelity images with a variety of novel semantics using our FreestyleNet.

Acknowledgments. This work was funded by the Singapore MOE Academic Research Fund Tier 1 grant (MSS21C002). It was also supported by "the Fundamental Research Funds for the Central Universities", 111 Project, China under Grants B07022 and (Sheitc) 150633, and Shanghai Key Laboratory of Digital Media Processing and Transmissions, China. It was also funded by the A*STAR under its AME YIRG Grant (Project No.A20E6c0101).

References

- Stephan Alaniz, Thomas Hummel, and Zeynep Akata. Semantic image synthesis with semantically coupled vq-model. *arXiv preprint arXiv:2209.02536*, 2022. 3
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096, 2018. 3
- [3] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocostuff: Thing and stuff classes in context. In *CVPR*, pages 1209–1218, 2018. 2, 3, 4, 6
- [4] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, pages 1511–1520, 2017. 3
- [5] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, pages 8780–8794, 2021. 3
- [6] Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: Faster and better text-to-image generation via hierarchical transformers. *arXiv preprint arXiv:2204.14217*, 2022. 3
- [7] Wan-Cyuan Fan, Yen-Chun Chen, Dongdong Chen, Yu Cheng, Lu Yuan, and Yu-Chiang Frank Wang. Frido: Feature pyramid diffusion for complex scene image synthesis. In AAAI, 2023. 3
- [8] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scenebased text-to-image generation with human priors. In *ECCV*, pages 89–106, 2022. 3
- [9] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-toimage generation using textual inversion. arXiv preprint arXiv:2208.01618, 2022. 3
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014. 3
- [11] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626, 2022. 3, 4
- [12] 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 *NeurIPS*, 2017. 7
- [13] Tobias Hinz, Stefan Heinrich, and Stefan Wermter. Generating multiple objects at spatially distinct locations. arXiv preprint arXiv:1901.00686, 2019. 2, 3
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, pages 6840–6851, 2020. 3
- [15] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598, 2022. 6
- [16] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. Inferring semantic layout for hierarchical textto-image synthesis. In *CVPR*, pages 7986–7994, 2018. 3

- [17] Xun Huang, Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Multimodal conditional image synthesis with productof-experts gans. In *ECCV*, pages 91–109, 2022. 3
- [18] Drew A Hudson and Larry Zitnick. Generative adversarial transformers. In *ICML*, pages 4487–4499. PMLR, 2021. 4
- [19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 1125–1134, 2017. 3
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. 3
- [21] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. *arXiv* preprint arXiv:2210.09276, 2022. 3
- [22] Jiadong Liang, Wenjie Pei, and Feng Lu. Layout-bridging text-to-image synthesis. arXiv preprint arXiv:2208.06162, 2022. 3
- [23] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. arXiv preprint arXiv:2202.09778, 2022. 6
- [24] Xihui Liu, Guojun Yin, Jing Shao, Xiaogang Wang, et al. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In *NeurIPS*, 2019. 2, 3, 7
- [25] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In CVPR, pages 11461–11471, 2022. 3
- [26] Zhengyao Lv, Xiaoming Li, Xin Li, Fu Li, Tianwei Lin, Dongliang He, and Wangmeng Zuo. Learning semantic person image generation by region-adaptive normalization. In *CVPR*, pages 10806–10815, 2021. 2
- [27] Zhengyao Lv, Xiaoming Li, Zhenxing Niu, Bing Cao, and Wangmeng Zuo. Semantic-shape adaptive feature modulation for semantic image synthesis. In *CVPR*, pages 11214– 11223, 2022. 3
- [28] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 3, 8
- [29] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, pages 2337–2346, 2019. 2, 3, 7
- [30] Xiaojuan Qi, Qifeng Chen, Jiaya Jia, and Vladlen Koltun. Semi-parametric image synthesis. In CVPR, pages 8808– 8816, 2018. 3
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 2, 4
- [32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3

- [33] Scott E Reed, Zeynep Akata, Santosh Mohan, Samuel Tenka, Bernt Schiele, and Honglak Lee. Learning what and where to draw. In *NeurIPS*, 2016. 2, 3
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684– 10695, 2022. 3, 4, 5, 6
- [35] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. arXiv preprint arXiv:2208.12242, 2022. 3
- [36] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. arXiv preprint arXiv:2205.11487, 2022. 3
- [37] Edgar Schönfeld, Vadim Sushko, Dan Zhang, Juergen Gall, Bernt Schiele, and Anna Khoreva. You only need adversarial supervision for semantic image synthesis. In *ICLR*, 2021. 2
- [38] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. arXiv preprint arXiv:2210.08402, 2022. 3
- [39] Yupeng Shi, Xiao Liu, Yuxiang Wei, Zhongqin Wu, and Wangmeng Zuo. Retrieval-based spatially adaptive normalization for semantic image synthesis. In *CVPR*, pages 11224–11233, 2022. 3
- [40] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256– 2265. PMLR, 2015. 3
- [41] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502, 2020. 3
- [42] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *NeurIPS*, pages 12438–12448, 2020. 3
- [43] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456, 2020. 3
- [44] Vadim Sushko, Edgar Schönfeld, Dan Zhang, Juergen Gall, Bernt Schiele, and Anna Khoreva. You only need adversarial supervision for semantic image synthesis. arXiv preprint arXiv:2012.04781, 2020. 3, 7
- [45] Zhentao Tan, Dongdong Chen, Qi Chu, Menglei Chai, Jing Liao, Mingming He, Lu Yuan, Gang Hua, and Nenghai Yu. Efficient semantic image synthesis via class-adaptive normalization. *TPAMI*, 2021. 3
- [46] Zhentao Tan, Qi Chu, Menglei Chai, Dongdong Chen, Jing Liao, Qiankun Liu, Bin Liu, Gang Hua, and Nenghai Yu. Semantic probability distribution modeling for diverse semantic image synthesis. *TPAMI*, 2022. 3

- [47] Hao Tang, Song Bai, and Nicu Sebe. Dual attention gans for semantic image synthesis. In ACMMM, pages 1994–2002, 2020. 3
- [48] Tengfei Wang, Ting Zhang, Bo Zhang, Hao Ouyang, Dong Chen, Qifeng Chen, and Fang Wen. Pretraining is all you need for image-to-image translation. arXiv preprint arXiv:2205.12952, 2022. 2, 3, 7
- [49] 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*, pages 8798–8807, 2018. 2, 3, 7
- [50] Weilun Wang, Jianmin Bao, Wengang Zhou, Dongdong Chen, Dong Chen, Lu Yuan, and Houqiang Li. Semantic image synthesis via diffusion models. arXiv preprint arXiv:2207.00050, 2022. 3
- [51] Yabin Wang, Zhiwu Huang, and Xiaopeng Hong. S-prompts learning with pre-trained transformers: An occam's razor for domain incremental learning. In *NeurIPS*, 2022. 2
- [52] Yi Wang, Lu Qi, Ying-Cong Chen, Xiangyu Zhang, and Jiaya Jia. Image synthesis via semantic composition. In *ICCV*, pages 13749–13758, 2021. 2, 3, 7
- [53] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Finegrained text to image generation with attentional generative adversarial networks. In CVPR, pages 1316–1324, 2018. 3
- [54] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. arXiv preprint arXiv:2206.10789, 2022. 3
- [55] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-toimage generation. In *CVPR*, pages 833–842, 2021. 3
- [56] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, pages 5907–5915, 2017. 3
- [57] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. arXiv preprint arXiv:2302.05543, 2023. 7
- [58] Zhu Zhang, Jianxin Ma, Chang Zhou, Rui Men, Zhikang Li, Ming Ding, Jie Tang, Jingren Zhou, and Hongxia Yang. Ufcbert: Unifying multi-modal controls for conditional image synthesis. In *NeurIPS*, pages 27196–27208, 2021. 3
- [59] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, pages 633–641, 2017. 6
- [60] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 130(9):2337–2348, 2022. 2
- [61] Junchen Zhu, Lianli Gao, Jingkuan Song, Yuan-Fang Li, Feng Zheng, Xuelong Li, and Heng Tao Shen. Label-guided generative adversarial network for realistic image synthesis. *TPAMI*, 2022. 3
- [62] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In CVPR, pages 5104–5113, 2020. 3

[63] Zhen Zhu, Zhiliang Xu, Ansheng You, and Xiang Bai. Semantically multi-modal image synthesis. In CVPR, pages 5467–5476, 2020. 3