

Interactive Image Synthesis with Panoptic Layout Generation

Bo Wang Tao Wu Minfeng Zhu Peng Du*
 Huawei Technologies

{wangbo341, zhuminfeng, dupeng25}@hisilicon.com, taowu1@huawei.com

Abstract

Interactive image synthesis from user-guided input is a challenging task when users wish to control the scene structure of a generated image with ease. Although remarkable progress has been made on layout-based image synthesis approaches, existing methods require high-precision inputs such as accurately placed bounding boxes, which might be constantly violated in an interactive setting. When placement of bounding boxes is subject to perturbation, layout-based models suffer from “missing regions” in the constructed semantic layouts and hence undesirable artifacts in the generated images. In this work, we propose Panoptic Layout Generative Adversarial Network (PLGAN) to address this challenge. The PLGAN employs panoptic theory which distinguishes object categories between “stuff” with amorphous boundaries and “things” with well-defined shapes, such that stuff and instance layouts are constructed through separate branches and later fused into panoptic layouts. In particular, the stuff layouts can take amorphous shapes and fill up the missing regions left out by the instance layouts. We experimentally compare our PLGAN with state-of-the-art layout-based models on the COCO-Stuff, Visual Genome, and Landscape datasets. The advantages of PLGAN are not only visually demonstrated but quantitatively verified in terms of inception score, Fréchet inception distance, classification accuracy score, and coverage. The code is available at <https://github.com/wb-fnalking/PLGAN>.

1. Introduction

Tremendous progress has been made on conditional image synthesis for creative design. Among different formats of conditional inputs are categories [2, 14, 25–27, 41], source images [15, 22, 44], text description [31, 38, 42], scene graphs [1, 12, 40] and semantic layouts [28, 36, 45]. To date, *interactive image synthesis* via conditional generative models remains a contemporary challenge. Text-to-image mod-

*Corresponding author.

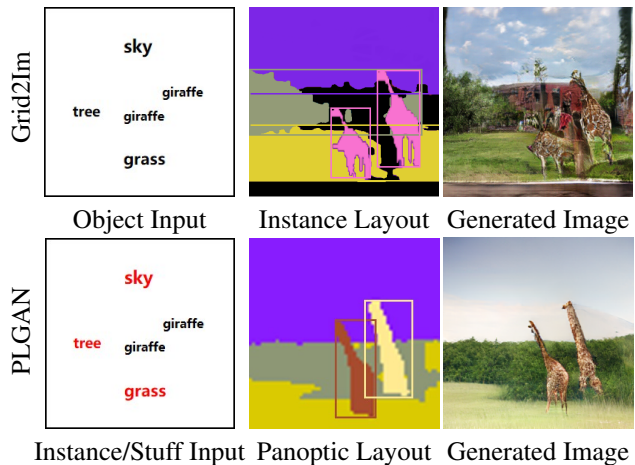


Figure 1. Scene-to-image synthesis by Grid2Im [1] vs. PLGAN. Unlike the instance layout utilized by Grid2Im that treats all objects as instances (things), the panoptic layout by PLGAN distinguishes objects between **instances** and **stuff**, thus eliminating artifacts in the *missing regions* (marked black in the layout).

els usually suffers from reasoning object locations and relations [42]. Image synthesis from semantic layouts provides an alternative way for computer-human interaction [9, 21] and yields aesthetically pleasing results [28, 36, 45]. However, high-quality semantic layouts demand professional skills in free-hand drawing from scratch by users, which prevent them from being used as a drag-and-drop GUI for novice users. In this respect, scene graph has attracted much attention in recent years [1, 12], as it requires only multiple object placements on the artboard and allows user-friendly manipulation with individual objects.

A milestone in scene-graph-to-image synthesis is made by Grid2Im [1], which roughly consists of two stages: layout construction and image generation. First, the input conditions are passed to construct an instance layout with per-object masks and bounding boxes. Secondly, conditioned on the instance layout a photo-realistic image is synthesized as the outcome of the image generation stage. Whilst Grid2Im [1] requires ground truth segmentation maps as su-

pervision signals, LostGAN [34,35] can learn the intermediate instance layout in a weakly-supervised way.

Besides [1, 34, 35], instance layout-based generative models such as [9, 12, 21, 43] have also driven progress on many cross-domain image synthesis tasks. A common caveat among aforementioned methods consists in that they are sensitive to spatial perturbation of scene objects and suffer from the *region missing* problem, especially in interactive scenarios. They may predict intermediate layouts with empty areas where pixels may not have correct category information. During the training phase, ground-truth bounding boxes and masks usually covers the whole image lattice. However, in interactive scenarios users can place objects with bounding boxes arbitrarily. In addition, the predicted per-object masks will not fill up the corresponding bounding boxes. Therefore, the intermediate layout may not be covered completely by object masks, yielding the region missing problem. An illustration of region missing by Grid2Im [1] is displayed in Fig. 1. Unsurprisingly, an imperfect semantic layout containing empty areas induces undesirable artifacts in the generated image.

In this work, we propose *Panoptic Layout Generative Adversarial Network* (PLGAN) for interactive image synthesis. Different from prior works that treat all objects invariably as instances (or things), we employ the panoptic segmentation theory [17] that splits objects into uncountable *stuff* (which refers to amorphous background such as grass, sky or sea) and countable *things* (which are foreground objects with well-defined shapes such as people, animals or vehicles). Furthermore, we develop the panoptic layout generation (PLG) module, which utilizes in parallel a stuff branch for stuff layout construction and an instance branch for instance layout construction. The instance branch predicts per-instance bounding boxes and masks as in [1,34]. The stuff branch generates pixel-wise masks for all stuff objects that cover the whole image lattice. Then the instance layout and the stuff layout are combined into a panoptic layout through an Instance- and Stuff-Aware Normalization (ISA-Norm) module. Images synthesis from such a panoptic layout successfully eliminates missing regions and behaves more robust to perturbation of object locations; see Fig. 1 for a visualized example.

Our contributions are summarized as follows:

- We leverage panoptic layouts in interactive image synthesis to resolve the region missing problem inherited by current instance layout-based approaches.
- Regarding model architecture, we propose to separate treatment of stuff and thing objects during layout constructions and later fuse the constructed instance and stuff layouts into a panoptic layout via Instance- and Stuff-Aware Normalization (ISA-Norm).
- Our experiments show qualitative and quantitative

comparisons on the COCO-Stuff dataset, Visual Genome, and Landscape datasets, and demonstrate of the merits of our PLGAN over state-of-the-art layout-based approaches.

2. Related Work

Even since the work of Generative Adversarial Networks (GAN) [5] has photo-realistic image synthesis attracted considerable attention. Further milestones are achieved by BigGAN [2] and StyleGAN [13] regarding network architecture and training strategy that bring impressive image quality in high resolution. While unconditional synthesis models like GAN take random noise as input, conditional synthesis [15, 25, 31, 42, 44] takes additional conditions (e.g., category, scene, layout) as input to control contents and styles of generated images. Among these, image synthesis from layout is an effective way for computer-human interaction [9, 21, 34], which requires to assign class label to each pixel. Layout-to-image models [28, 36, 45] can produce aesthetically pleasing results with multiple objects by leveraging locations and shapes of objects directly.

Instance Layout. Instance layout based methods consider each object as an instance attached with a bounding box and a shape independently, and assign every pixel with a category label and an instance ID. For instance, text-to-image model [9] takes text descriptions as inputs and adopts LSTM [8] to predict instance bounding boxes and masks. The configurable scene layout [1, 12, 34, 43] is a more user-friendly instrument, which consists of a set of objects with labels, locations and their interactions. SG2Im [12] and Grid2Im [1] take advantage of graph convolution networks [33] to extract information from scenes and build instance embedding features to predict bounding boxes and masks for layout construction. LostGAN [34] receives bounding boxes and classes as inputs for image synthesis. Compared to other condition formats, the scene layout provides a similar spatial structure to the target image and is easier to be constructed. Though the instance-based methods can generate realistic instances with recognizable shapes, they tend to cause region missing problems especially in high-resolution and interactive scenarios where the image lattice cannot be covered entirely by instance masks.

Panoptic Layout. Our proposed PLGAN leverages panoptic theory from panoptic segmentation. Panoptic segmentation was first termed in [17] to unify instance segmentation and semantic segmentation for scene understanding missions. Panoptic segmentation models [4, 20, 23] typically utilize separate branches to generate region-based instance layout for *things* and dense-pixel semantic layout for *stuff*. Inspired by previous works, our PLGAN adopts panoptic layouts in the context of interactive image generation, allowing users to manipulate the entire scene and

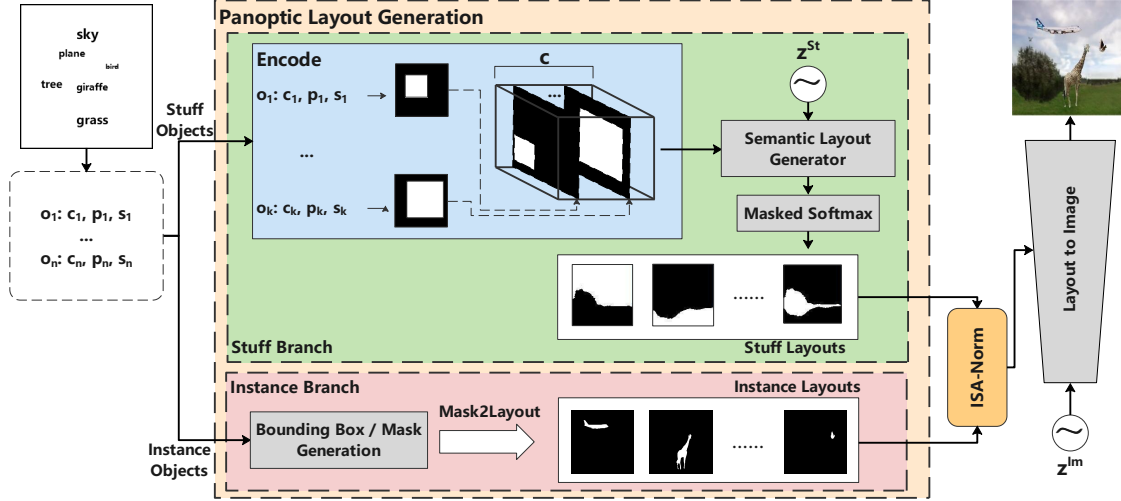


Figure 2. Overview of the PLGAN architecture. PLGAN is an interactive image synthesis model trained in an end-to-end manner. It consists of two stages: scene-to-layout generation and layout-to-image synthesis. Inspired by panoptic segmentation [17], panoptic layout generation (PLG) is proposed for scene-to-layout generation, which distinguishes between instances and stuff for generated objects. In particular, the stuff layout complements the instance layout, as the latter is prone to the region missing problem.

create photo-realistic images on the fly.

3. Method

The proposed PLGAN follows a two-stage procedure: **scene-to-layout generation** and **layout-to-image synthesis**; see Fig. 2 for an overview of its architecture. The input of PLGAN is a set of objects $O = \{o_1, o_2, \dots, o_n\}$ with each $o_i = (c_i, p_i, s_i)$, where $c_i \in \mathcal{C}$ is the object category (e.g., $|\mathcal{C}| = 171$ in the COCO-Stuff dataset), $p_i \in [0, 1]^2$ is the center position, and s_i is the object size (typically drawn from some predefined set, e.g., [1, 25]).

For scene-to-layout construction, we propose the **Panoptic Layout Generation (PLG)** module inspired from [17]. Specifically, PLG learns the mapping from scene to panoptic layout by embedding *stuff* and *thing* labels separately:

$$L^{St} = G^{St}(O^{St}, z^{St}), \quad (1)$$

$$L_i^{Th} = G^{Th}(o_i^{Th}, z_i^{Th}), \quad (2)$$

where $L^{St} \in \mathbb{R}^{H \times W \times |\mathcal{C}^{St}|}$ represents the *stuff layout*, $L_i^{Th} \in \mathbb{R}^{H \times W \times 1}$ is the *instance layout*, G^{St} and G^{Th} are two distinctive generators under stuff branch and instance branch respectively, O^{St} and $O^{Th} = \bigcup_i \{o_i^{Th}\}$ are sets of stuff objects and thing objects, and $z^{St}, z_i^{Th} \in \mathbb{R}^m$ are latent codes sampled from standard Gaussians.

Meanwhile, layout-to-image synthesis has been well explored in the recent literature. The PLGAN leverages the state-of-the-art models, such as Grid2Im [1], LostGAN-V1 [34], LostGAN-V2 [35] and CAL2I [6], for layout-to-image synthesis. In mathematical terms, a photorealistic image I^f

is produced by the generator G^{Im} :

$$I^f = G^{Im}(L^{St}, \{L_i^{Th}\}, z^{Im}), \quad (3)$$

from layouts $L^{St}, \{L_i^{Th}\}$ and a Gaussian latent code $z^{Im} \in \mathbb{R}^m$. Last but not least, we have integrated the **Instance- and Stuff-Aware Normalization (ISA-Norm)** into the layout-to-image stage, which is dedicated to the fusion of stuff layout and instance layouts.

3.1. Panoptic Layout Generation

In PLG we split object categories in \mathcal{C} into two disjoint subsets, **stuff** \mathcal{C}^{St} and **things** \mathcal{C}^{Th} , i.e., $\mathcal{C} = \mathcal{C}^{St} \cup \mathcal{C}^{Th}$. The stuff represents amorphous background regions of texture or material, such as grass, sky and road. In contrast, things are typically countable foreground objects with well-defined shapes, such as people, animals and vehicles. In order to eliminate missing regions from instance layout-based models [1, 43], we propose to split layout construction into **instance branch** and **stuff branch**, in analogy to panoptic segmentation [17, 19, 39]. Furthermore, we propose to fuse the stuff- and instance-layouts and then refine the panoptic layout conditioning on instance-to-instance and instance-to-stuff relations.

3.1.1 Instance Layout Branch

Similar to previous works [1, 12, 35], we generate the instance layout in two shots. First, we predict B(ounding) Box and mask for each instance object with conditional generative models. Given a thing object $o_i^{Th} = (c_i, p_i, s_i)$ with $c_i \in \mathcal{C}^{Th}$, both the mask- and BBox-generators take word

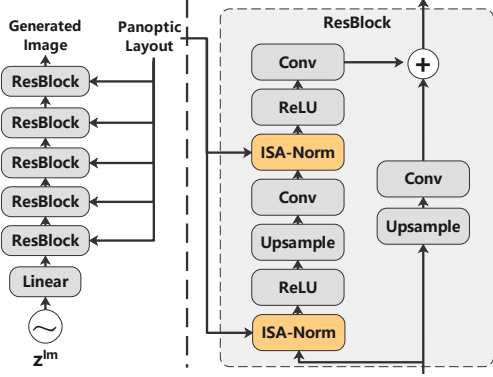


Figure 3. Illustration of the Layout-to-Image module.

embedding of object label, center position p_i and size s_i as the inputs. To simplify the model, we predict only the height and width of the BBox and then combine them with the input center position to obtain the final BBox. In the second step, all masks are resized into specific regions defined by BBoxes. Based on these BBoxes and masks, we further utilize the Mask2Layout module [35] to construct the instance layout $L^{Th} \in \mathbb{R}^{H \times W \times |O^{Th}|}$, where the slice $L_{\dots,i}^{Th}$ corresponds to the predicted mask of object o_i^{Th} .

3.1.2 Stuff Layout Branch

While instance layouts are generated independently from each other, stuff layouts are intercorrelated and ought to be generated jointly. More specifically, we first generate a coarse layout $L^{St,init} \in \mathbb{R}^{H \times W \times |C^{St}|}$, where each slice $L_{\dots,c}^{St,init}$ is a coarse mask of stuff object with label $c \in C^{St}$. Given a stuff object $o^{St} = (c, p, s)$, we generate a coarse mask by setting a square mask with s as its height and width and p as its center position (see Fig. 2). Then, we use the conditional generative model with four ResBlocks [28] to refine this coarse layout $L^{St,init}$ into a stuff layout \hat{L}^{St} . Finally, we normalize \hat{L}^{St} using a *masked softmax*:

$$L_{h,w,c}^{St} = \frac{e^{\hat{L}_{h,w,c}^{St}}}{\sum_{c' \in C^{in}} e^{\hat{L}_{h,w,c'}^{St}}}, \quad c \in C^{in}, \quad (4)$$

where $C^{in} \subset C^{St}$ contains the input stuff categories.

3.2. Layout-to-Image Synthesis

3.2.1 Conditional Image Synthesis

Synthesizing image from layout is a kind of conditional generative task that has recently embarked attention; see, e.g., Grid2Im [1], LostGAN [35] and CAL2I [6]. To accomplish layout-to-image synthesis in PLGAN, we follow the state-of-the-art method CAL2I [6] to construct our generative model. LostGAN [35] proposes ISLA-Norm which




Figure 4. Illustration of Instance- and Stuff-Aware Normalization.

takes instance layout as an input of its conditional generative model. To fully utilize the panoptic layout in the PLGAN setup, we propose ISA-Norm in place of ISLA-Norm.

Fig. 3 shows the design of the layout-to-image stage for 128×128 output resolution. The generator consists of one fully connected (FC) layer and five ResBlocks. The FC layer maps the image latent code z^{lm} of 128 dimensions to a $4 \times 4 \times 128$ tensor. Then five ResBlocks are employed to successively upsample this tensor to the final generated image of desired resolution, where each ResBlock blends the panoptic layout into upsampling via the ISA-Norm module.

3.2.2 Instance- and Stuff-Aware Normalization

Once instance and stuff layout templates are ready, we need to fuse them properly. In this respect, the ISLA-Norm module was used in LostGAN [35] (which only deals with instance layouts). ISLA-Norm accomplishes multi-object fusion by embedding instance layouts into affine transformations in BatchNorm layers. However, direct averaging over embedded instance objects is inappropriate for a panoptic layout. In a panoptic setup, the stuff layout will overspread the whole scene as background, and among stuff and instances there are widespread overlaps. To address this challenge, we propose *Instance- and Stuff-Aware Normalization* (ISA-Norm); see Fig. 4 for an illustration.

Let $X \in \mathbb{R}^{B \times H \times W \times C}$ be the 4D feature map resulting from the activation layer of the ResBlocks in Fig 3. The ISA-Norm transforms X as in a standard BatchNorm:

$$\mu_c = \frac{1}{BHW} \sum_{b,h,w} X_{b,h,w,c}, \quad (5)$$

$$\sigma_c = \sqrt{\frac{1}{BHW} \left(\sum_{b,h,w} X_{b,h,w,c} - \mu_c \right)^2 + \epsilon}, \quad (6)$$

$$\hat{X}_{b,h,w,c} = \frac{X_{b,h,w,c} - \mu_c}{\sigma_c} \cdot \gamma_{h,w,c} + \beta_{h,w,c}. \quad (7)$$

Here $\mu, \sigma \in \mathbb{R}^C$ are mean and standard deviation with respect to the batch dimension. The shift- and scale-parameters $\beta, \gamma \in \mathbb{R}^{H \times W \times C}$ are constructed as follows.

First, we use learnable matrices to embed the object classes and get $W_\gamma^{St} \in \mathbb{R}^{|C^{St}| \times C}$ and $W_\beta^{Th}, W_\gamma^{Th} \in \mathbb{R}^{|C^{Th}| \times C}$ for stuff and thing objects separately. Then we compute the foreground mask of thing objects:

$$M_{h,w} = \begin{cases} 1, & \text{if } \sum_c L_{h,w,c}^{Th} > \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where τ is a scalar value representing the foreground threshold ($\tau = 0.1$ in experiments). We further process the instance layouts conditioning on the current image feature via the Guided Filter (GF) [37], and project the instance and stuff layouts into semantic space using label embedding:

$$\tilde{L}_i^{Th} = \text{GF}(L_i^{Th}, X), \quad (9)$$

$$E_{h,w}^{Th,\star} = \frac{\left(\sum_i \tilde{L}_i^{Th} W_\star^{Th}\right)_{h,w}}{\left(\sum_i L_i^{Th}\right)_{h,w}}, \quad (10)$$

$$E^{St,\star} = L^{St} W_\star^{St}, \quad (11)$$

where $\star \in \{\gamma, \beta\}$ and $E^{Th,\star}, E^{St,\star} \in \mathbb{R}^{H \times W \times d}$. Details on GF are left to the Appendix. Finally, γ, β are fused from instance and stuff layout embedding:

$$\gamma_{h,w,c} = M_{h,w} E_{h,w,c}^{Th,\gamma} + (1 - M_{h,w}) E_{h,w,c}^{St,\gamma}, \quad (12)$$

$$\beta_{h,w,c} = M_{h,w} E_{h,w,c}^{Th,\beta} + (1 - M_{h,w}) E_{h,w,c}^{St,\beta}. \quad (13)$$

3.3. Model Objectives

The total loss used for PLGAN training is given by:

$$\begin{aligned} \mathcal{L} = & \lambda_1 \mathcal{L}_{box} + \lambda_2 \mathcal{L}_{img} + \lambda_3 \mathcal{L}_{obj} \\ & + \lambda_4 \mathcal{L}_{per} + \lambda_5 \mathcal{L}_{rec} + \lambda_6 \mathcal{L}_{app}. \end{aligned} \quad (14)$$

Here \mathcal{L}_{box} is the MSE between the predicted bounding boxes and the ground truth bounding boxes, \mathcal{L}_{img} and \mathcal{L}_{obj} are two *adversarial losses* for images and objects respectively, \mathcal{L}_{per} is the perceptual loss, \mathcal{L}_{rec} is the reconstruction loss, and \mathcal{L}_{app} is another adversarial loss for appearance. The balancing weights are set manually in the experiments, i.e., $\lambda_2 = 0.1$, $\lambda_1 = \lambda_3 = \lambda_4 = \lambda_5 = \lambda_6 = 1$.

Image and object losses. We use hinge loss [24] for \mathcal{L}_{img} :

$$\mathcal{L}_{img} = \mathbb{E}_{I^f \sim p_{fake}} [-D_{img}(I^f)], \quad (15)$$

where I^r is a real image drawn from training data and I^f is a fake image generated from Eq. (3). Notice that \mathcal{L}_{img} involves a discriminator D_{img} which is updated via minimizing the adversarial loss:

$$\begin{aligned} \mathcal{L}_{img}^{adv}(D_{img}) = & \mathbb{E}_{I^r \sim p_{real}} [\max(0, 1 - D_{img}(I^r))] \\ & + \mathbb{E}_{I^f \sim p_{fake}} [\max(0, 1 + D_{img}(I^f))], \end{aligned} \quad (16)$$

and similarly for the object loss \mathcal{L}_{obj} . The adversarial loss has been proven effective in generating realistic textures.

Reconstruction loss. The reconstruction loss measures the L_1 distance in pixels between the predicted images and the ground truth images, i.e.,

$$\mathcal{L}_{rec} = \mathbb{E}_{I^r \sim p_{real}, I^f \sim p_{fake}} [\|I^r - I^f\|_1]. \quad (17)$$

Perceptual loss. The perceptual loss encourages synthesized and real images to share similar feature representations, and is widely used in style transfer [11] and image synthesis [44]. With $\phi_j(\cdot)$ the activation of j -th layer from the VGG-19 network, the perceptual loss is defined as:

$$\mathcal{L}_{per} = \mathbb{E}_{I^r \sim p_{real}, I^f \sim p_{fake}} \left[\sum_j w_j \|\phi_j(I^r) - \phi_j(I^f)\|_1 \right], \quad (18)$$

with the feature balancing weights $\{w_j\}$. In our experiments, we compute \mathcal{L}_{per} using the activation of *conv1-1*, *conv2-1*, *conv3-1*, *conv4-1* and *conv5* layers, with the corresponding weights 1/32, 1/16, 1/8, 1/4 and 1.

Appearance Loss. Following CAL2I [6], we also introduce the appearance loss which penalizes the generator according to pixel-level misalignment:

$$\mathcal{L}_{app} = \mathbb{E}_{I^f \sim p_{fake}} [-D_{app}(A^f | I^f)], \quad (19)$$

The discriminator D_{app} is updated via minimizing the adversarial loss:

$$\begin{aligned} \mathcal{L}_{app}^{adv}(D_{app}) = & \mathbb{E}_{I^r \sim p_{real}} [\max(0, 1 - D_{app}(A^r | I^r))] \\ & + \mathbb{E}_{I^f \sim p_{fake}} [\max(0, 1 + D_{app}(A^f | I^f))]. \end{aligned} \quad (20)$$

In \mathcal{L}_{app} and \mathcal{L}_{app}^{adv} , we have used A^r and A^f which are the Gram matrices of object features in I^r and I^f respectively. The purpose of Gram matrices is to measure the spatial similarity between object features, hence better preserving location-sensitive information in synthesized images [6].

3.4. Implementation Details

The PLGAN model is trained using Pytorch [29] on a NVIDIA Tesla V100 GPU server. The training uses Adam optimizer [16] with learning rate 10^{-4} and batch size 128, and runs for 200 epochs on all tested datasets. Inference of the PLGAN model is tested on Huawei Atlas inference workstation equipped with Ascend AI Accelerator Card and Ascend Compute Architecture for Neural Networks (CANN).

4. Experiments

We evaluate the proposed PLGAN on three datasets: COCO-Stuff [3], Visual Genome [18], and our own Landscape dataset. The results from PLGAN are compared with

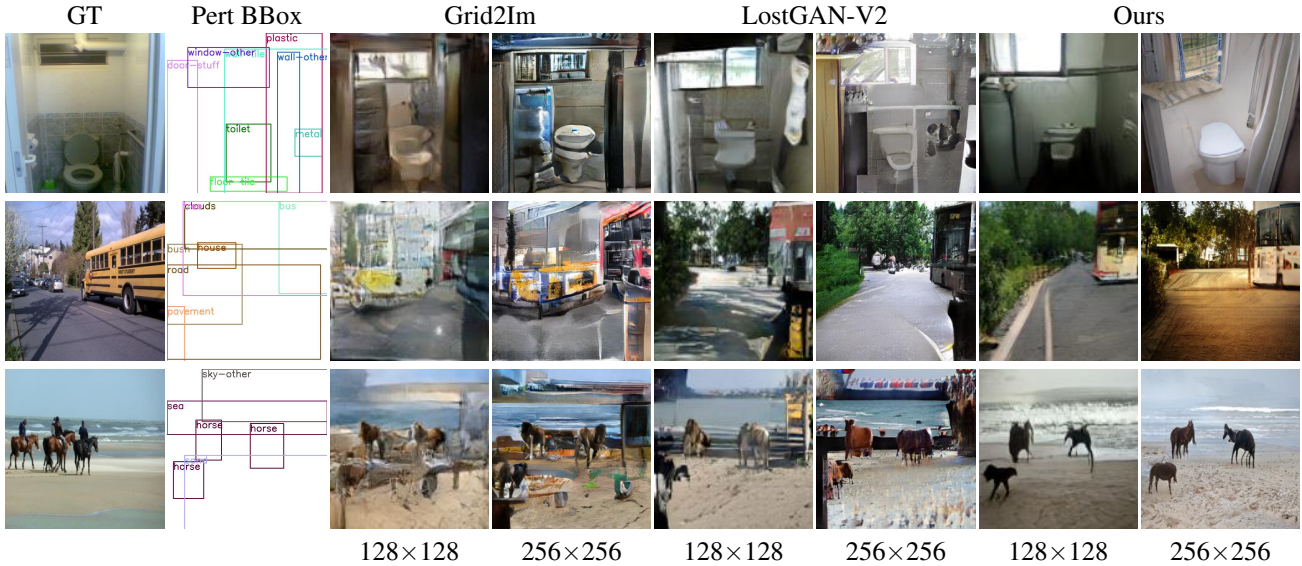


Figure 5. Visual comparison between sample images generated from perturbed BBoxes (Pert BBoxes) on the COCO-Stuff dataset.

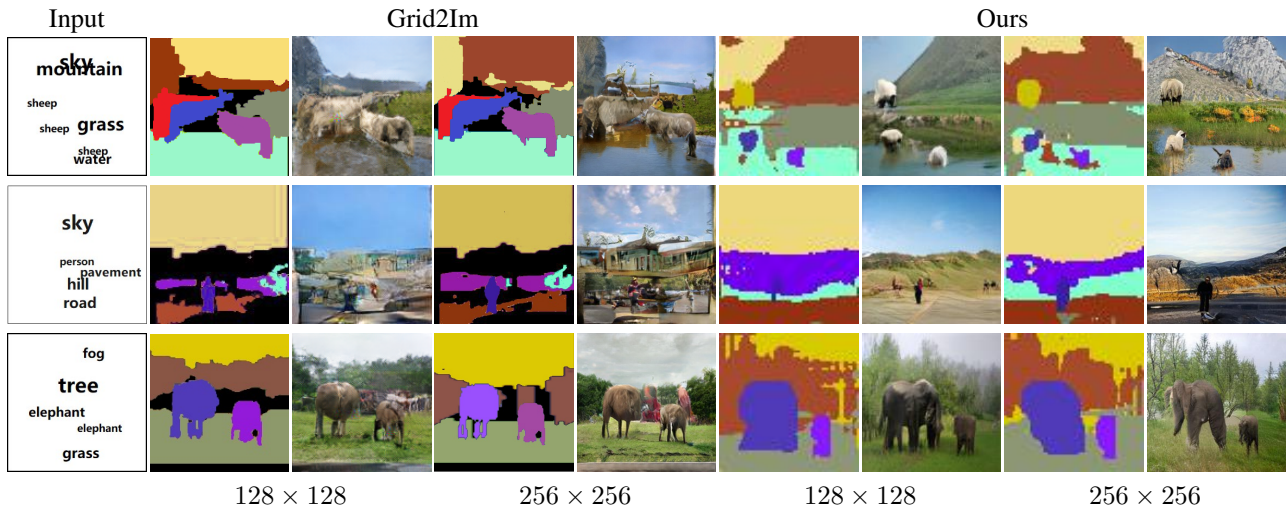


Figure 6. Visual comparison between instance layouts and panoptic layouts on the COCO-Stuff dataset.

state-of-the-art methods not only visually and but also quantitatively using widely adopted metrics. We also carry out ablation studies to evaluate effectiveness of the individual components of PLGAN. Due to space limitation, supplementary results are left to the Appendix.

Datasets. The **COCO-Stuff** dataset [3] annotates 40K training image and 5K validation images with bounding boxes and segmentation masks for 80 *thing* categories and 91 *stuff* categories. Following Ashual et al. [1], we choose images with three to eight objects and further filter images whose object coverage is less than 2%. We split categories of the **Visual Genome (VG)** dataset [18] into 92 *thing* categories and 87 *stuff* categories. We choose 62,565 training, 5,506 validation and 5,088 test images with 3 to 30 objects

in our experiments. We also created our own **Landscape** dataset by collecting 27k photos (25k train and 2k val) of 448² resolution from the Flickr website, in order to fully demonstrate the advantage of the stuff layout generation in PLGAN. The Landscape dataset contains only 23 *stuff* object classes (such as sky, sea and mountain) but no *thing* objects. We use a pre-trained UPSNet [39] to extract pixel-level segmentation masks for both *thing* and *stuff* objects.

Methods. We compare our PLGAN with state-of-the-art layout-to-image models: Layout2Im [43], Grid2Im [1], LostGAN-V1 [34], LostGAN-V2 [35] and CAL2I [6]. The results from these models are reproduced by the publicly released code.

Evaluation Metrics. Four metrics are adopted for

Table 1. Quantitative comparison with respect to Inception Score (higher is better), FID (lower is better) and CAS (higher is better) on the COCO-Stuff dataset. Pert1 BBox and Pert2 BBox are generated by perturbing GT BBox with different random biases on object center.

Methods	Resolution	IS \uparrow			FID \downarrow			CAS \uparrow
Real Images	64 \times 64	13.4 \pm 0.5			-			
Real Images	128 \times 128	22.3 \pm 0.4			-			51.04
Real Images	256 \times 256	30.4 \pm 0.6			-			
		GT BBox	Pert1 BBox	Pert2 BBox	GT BBox	Pert1 BBox	Pert2 BBox	GT BBox
Layout2Im [43]		9.1 \pm 0.1	7.7 \pm 0.2	7.0 \pm 0.2	37.53	44.57	50.58	27.32
Grid2Im [1]		10.3\pm0.1	-	-	48.7	-	-	-
LostGAN-V1 [34]	64 \times 64	9.8 \pm 0.2	-	-	34.31	-	-	28.81
Ours (CAL2I [6]+PLG)		10.3\pm0.1	9.2\pm0.1	8.2\pm0.1	21.85	28.01	34.52	29.50
Grid2Im [1]		11.2 \pm 0.3	7.4 \pm 0.1	6.4 \pm 0.2	63.2	77.76	87.89	25.89
LostGAN-V1 [34]		13.8 \pm 0.4	9.2 \pm 0.1	7.7 \pm 0.1	29.65	51.96	71.04	30.68
LostGAN-V2 [35]		14.2 \pm 0.4	9.9 \pm 0.1	8.1 \pm 0.2	24.76	43.82	59.34	31.98
CAL2I [6]		15.6\pm0.2	11.1 \pm 0.1	9.0 \pm 0.1	24.15	43.12	57.89	32.52
Grid2Im [1]+PLG	128 \times 128	12.7 \pm 0.1	11.0 \pm 0.2	9.5 \pm 0.2	45.83	51.53	60.24	26.74
LostGAN-V1 [34]+PLG		14.1 \pm 0.1	12.6 \pm 0.2	11.0 \pm 0.2	26.85	31.82	38.67	31.33
LostGAN-V2 [35]+PLG		14.6 \pm 0.2	12.8 \pm 0.1	11.4 \pm 0.1	25.43	30.80	36.75	32.86
Ours (CAL2I [6]+PLG)		15.6\pm0.3	13.2\pm0.2	11.7\pm0.2	22.70	29.03	35.40	33.86
Grid2Im [1]		15.2 \pm 0.1	7.7 \pm 0.4	4.4 \pm 0.1	65.95	147.85	253.59	20.54
LostGAN-V2 [35]	256 \times 256	18.2 \pm 0.2	12.2 \pm 0.2	9.5 \pm 0.2	30.82	56.67	77.56	30.33
Ours (CAL2I [6]+PLG)		18.9\pm0.3	15.8\pm0.2	14.2\pm0.2	29.10	40.14	46.89	32.33

the quantitative evaluation: Inception Score (IS) [32], Fréchet Inception Distance (FID) [7], Classification Accuracy Score (CAS), and Coverage (COV) [10]. In particular, we compute CAS [30, 35] by training a ResNet-101 model on the synthesized images to classify real images for the COCO-Stuff and VG datasets. Higher CAS is better (i.e., more identifiable objects). The Coverage (COV) measures the quality of intermediate semantic layouts, which is computed as the average percentage (ranged between 0 and 100) of empty area in the generated semantic layout. Higher COV is better (i.e., less empty artboard).

4.1. Qualitative Results

In Fig. 5 we show visual comparison between Grid2Im [1], LostGAN-V2 [35] and our PLGAN under 128² and 256² resolutions, all using perturbed BBoxes as input. It is clearly observable that Grid2Im and LostGAN-V2 which rely on instance-based layouts produce artifacts in regions between BBoxes and image borders. This is due to the fact that BBoxes of instances cannot occupy the whole image lattice. Even if there is no gap between BBoxes, mismatch between instance masks also causes the region missing problem in layouts by Grid2Im. With increasing resolution, this problem becomes even more apparent. Our PLGAN adopts panoptic-based layouts and processes stuff objects akin to semantic segmentation, hence the background

of a generated image naturally fills up the whole image lattice. We refer to Fig. 6 for visual difference between instance layouts and panoptic layouts.

4.2. Quantitative Results

In Tab. 1, we show quantitative comparisons on COCO-Stuff with respect to Inception Score, FID and CAS across different resolutions. “GT BBox” refers to the ground-truth annotations from the original dataset; “Pert1 BBox” and “Pert2 BBox” refer to randomly biased center positions of GT BBox in the range $[-0.3, 0.3]$ and $[-0.5, 0.5]$ respectively. Unsurprisingly, IS and FID under GT BBox are better than those under Pert1 BBox and Pert2 BBox. With GT BBox as input, our PLGAN (CAL2I+PLG) is the best in IS and FID but the difference to the competing models is small. However, the advantage of PLGAN in IS and FID is more apparent under Pert1 BBox and Pert2 BBox. This is due to that perfect placement of BBoxes potentially alleviates the region missing and overlapping problems that previous instance-layout based models are troubles with. To compare the CAS, we test the synthesized images with GT BBox as input only. Following LostGAN-V2 [35], we compute CAS on cropped and resized objects at 32² resolution from synthesized and real images. According to Tab. 1, our methods has higher CAS at different resolutions. This confirms that, with refined instance layouts by the ISA-Norm module, ob-

Table 2. Ablation study on the PLG and ISA-Norm module.

PLG	ISA-Norm		IS \uparrow			FID \downarrow			CAS \uparrow	Inference Time
	w/o GF	w/ GF	GT BBox	Pert1 BBox	Pert2 BBox	GT BBox	Pert1 BBox	Pert2 BBox	GT BBox	
			15.6 \pm 0.2	11.1 \pm 0.1	9.0 \pm 0.1	24.15	43.12	57.89	32.52	14 ms
✓			14.6 \pm 0.2	12.3 \pm 0.1	10.3 \pm 0.1	24.89	35.66	58.84	31.99	22 ms
✓	✓		15.0 \pm 0.3	12.5 \pm 0.1	10.8 \pm 0.1	23.65	33.09	52.73	31.09	22 ms
✓		✓	15.6 \pm 0.3	13.9 \pm 0.2	12.8 \pm 0.2	22.70	27.03	33.40	33.86	26 ms




Figure 7. IS, FID and Coverage curves with varying perturbation range on the COCO-Stuff dataset under resolution 128×128 .

Table 3. Ablation study on panoptic layout.

Layout	IS \uparrow	FID \downarrow	CAS \uparrow
Stuff Layout only	12.7 \pm 0.6	43.70	27.15
Instance Layout only	15.6 \pm 0.2	24.15	32.52
Panoptic Layout (Instance+Stuff)	15.6 \pm 0.3	22.70	33.86

jects in generated images will attain higher fidelity.

Notably, our PLGAN uses CAL2I as the layout-to-image generator, which can be replaced by other instance-layout based models. For this reason, we also include results in Tab. 1 that combine the PLG module and different layout-to-image generators. Noticeably PLG always improves instance-layout baselines.

We further elaborate the test on robustness against perturbation on GT BBoxes. By varying the perturbation range from 0 to 0.5, Fig. 7 plots IS, FID and COV curves for our PLGAN and other models from Tab. 1. It is observed that IS and FID deteriorate on all models as perturbation on GT BBoxes increases. However, our proposed PLG module always robustifies image synthesis compared to instance-layout baselines. The PLGAN which combines CAL2I and PLG is the best overall performer in terms of IS and FID. Regarding the coverage metric, we observe visible decay for all instance-layout baselines as the perturbation range increases. Meanwhile, our PLGAN maintains 100% COV thanks to the panoptic layout.

4.3. Ablation Studies

We now validate the effectiveness of individual components of the PLGAN. In Tab. 2, we use CAL2I [6] as

baseline and augment it with the PLG and ISA-Norm modules. We see that, although PLG resolves the region missing problem from the baseline, it sometimes degrades the image quality measured by the three metrics. The ISA-Norm is the right remedy to enhance the image quality generated from panoptic layouts. The PLGAN, with all three components combined, attains the highest score in all metrics and maintain the real time.

We also test the variants of PLGAN using either instance or stuff layout branch alone. In Tab. 3, we see that treating all objects as stuff or things alone yields inferior metrics to the panoptic layout approach.

5. Conclusion

This paper focuses on resolving the region missing problem and improving the robustness of scene-to-image synthesis in interactive scenarios. To this end, our PLGAN leverages panoptic theory and constructs instance and stuff layouts through separate branches. The resulting panoptic layouts eliminate missing regions and yield aesthetically pleasing images, even if perturbation on object locations is allowed. Extensive evidences on the COCO-Stuff, Visual Genome, and Landscape datasets advocate the superiority of PLGAN over the state-of-the-art methods.

6. Acknowledgement

We gratefully acknowledge the optimization support of Ascend CANN (Compute Architecture for Neural Networks) used for this research.

References

- [1] Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4561–4569, 2019. 1, 2, 3, 4, 6, 7
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 1, 2
- [3] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1209–1218, 2018. 5, 6
- [4] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12475–12485, 2020. 2
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, 2014. 2
- [6] Sen He, Wentong Liao, Michael Yang, Yongxin Yang, Yi-Zhe Song, Bodo Rosenhahn, and Tao Xiang. Context-aware layout to image generation with enhanced object appearance. In *CVPR*, 2021. 3, 4, 5, 6, 7, 8
- [7] 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 *Advances in neural information processing systems*, pages 6626–6637, 2017. 7
- [8] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [9] 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. 1, 2
- [10] Maor Ivgi, Yaniv Benny, Avichai Ben-David, Jonathan Berant, and Lior Wolf. Scene graph to image generation with contextualized object layout refinement. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 2428–2432. IEEE, 2021. 7
- [11] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 5
- [12] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018. 1, 2, 3
- [13] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2
- [14] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2020. 1
- [15] 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. 1, 2
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [17] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. 2, 3
- [18] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 5, 6
- [19] Jie Li, Allan Raventos, Arjun Bhargava, Takaaki Tagawa, and Adrien Gaidon. Learning to fuse things and stuff. *arXiv preprint arXiv:1812.01192*, 2018. 3
- [20] Qizhu Li, Xiaojuan Qi, and Philip HS Torr. Unifying training and inference for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13320–13328, 2020. 2
- [21] Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. Object-driven text-to-image synthesis via adversarial training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12174–12182, 2019. 1, 2
- [22] Xinyang Li, Shengchuan Zhang, Jie Hu, Liujuan Cao, Xiopeng Hong, Xudong Mao, Feiyue Huang, Yongjian Wu, and Rongrong Ji. Image-to-image translation via hierarchical style disentanglement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8639–8648, June 2021. 1
- [23] Yanwei Li, Xinze Chen, Zheng Zhu, Lingxi Xie, Guan Huang, Dalong Du, and Xingang Wang. Attention-guided unified network for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7026–7035, 2019. 2
- [24] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017. 5
- [25] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 1, 2
- [26] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. *arXiv preprint arXiv:1802.05637*, 2018. 1
- [27] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, pages 2642–2651, 2017. 1
- [28] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive nor-

- malization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. 1, 2, 4
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. 5
- [30] Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. *arXiv preprint arXiv:1905.10887*, 2019. 7
- [31] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*, 2016. 1, 2
- [32] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016. 7
- [33] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008. 2
- [34] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10531–10540, 2019. 2, 3, 6, 7
- [35] Wei Sun and Tianfu Wu. Learning layout and style reconfigurable gans for controllable image synthesis. *arXiv preprint arXiv:2003.11571*, 2020. 2, 3, 4, 6, 7
- [36] 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*, pages 8798–8807, 2018. 1, 2
- [37] Huikai Wu, Shuai Zheng, Junge Zhang, and Kaiqi Huang. Fast end-to-end trainable guided filter. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1838–1847, 2018. 5
- [38] Weihao Xia, Yujiu Yang, Jing-Hao Xue, and Baoyuan Wu. Tedigan: Text-guided diverse face image generation and manipulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2256–2265, June 2021. 1
- [39] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8818–8826, 2019. 3, 6
- [40] Cheng-Fu Yang, Wan-Cyuan Fan, Fu-En Yang, and Yu-Chiang Frank Wang. Layouttransformer: Scene layout generation with conceptual and spatial diversity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3732–3741, June 2021. 1
- [41] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, pages 7354–7363. PMLR, 2019. 1
- [42] 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 *Proceedings of the IEEE international conference on computer vision*, pages 5907–5915, 2017. 1, 2
- [43] Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. Image generation from layout. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8584–8593, 2019. 2, 3, 6, 7
- [44] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 1, 2, 5
- [45] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5104–5113, 2020. 1, 2