

Attribute-Guided Image Generation from Layout

Ke Ma¹
mark1123@cs.ubc.ca

Bo Zhao⁴
zhaobo.cs@gmail.com

Leonid Sigal^{1,2,3}
lsigal@cs.ubc.ca

¹ University of British Columbia
Vancouver, Canada

² Vector Institute for AI
Toronto, Canada

³ CIFAR AI Chair

⁴ Bank of Montreal
Toronto, Canada

Abstract

Recent approaches have achieved great successes in image generation from structured inputs, e.g., semantic segmentation, scene graph or layout. Although these methods allow specification of objects and their locations at image-level, they lack the fidelity and semantic control to specify visual appearance of these objects at an instance-level. To address this limitation, we propose a new image generation method that enables instance-level attribute control. Specifically, the input to our attribute-guided generative model is a tuple that contains: (1) object bounding boxes, (2) object categories and (3) an (optional) set of attributes for each object. The output is a generated image where the requested objects are in the desired locations and have prescribed attributes. Several losses work collaboratively to encourage accurate, consistent and diverse image generation. Experiments on Visual Genome [15] dataset demonstrate our model’s capacity to control object-level attributes in generated images, and validate plausibility of disentangled object-attribute representation in the image generation from layout task. Also, the generated images from our model have higher resolution, object classification accuracy and consistency, as compared to the previous state-of-the-art.

1 Introduction

Controlled image generation methods have achieved great successes in recent years, driven by the advances in conditional Generative Adversarial Networks (GANs) [4, 14, 25, 27, 29, 37, 38, 42, 43] and disentangled representations [13, 44]. The goal of these methods is to generate high-fidelity images from various user specified guidelines (conditions), such as textual descriptions [7, 23, 32, 36, 37], attributes [8, 11, 18, 22, 26, 41], scene graphs [10, 11, 19], layout [61, 40] and semantic segmentation [2, 8, 9, 11, 21, 28, 34, 42, 43]. The high-level nature of most of these specifications is desirable from ease of use and control point of views, but severely impoverished in terms of pixel-level spatial and appearance information, leading to a challenging image generation problem.

In this paper, we specifically focus on image generation from layout, where the input is a coarse spatial layout of the scene (e.g., bounding boxes and corresponding object class

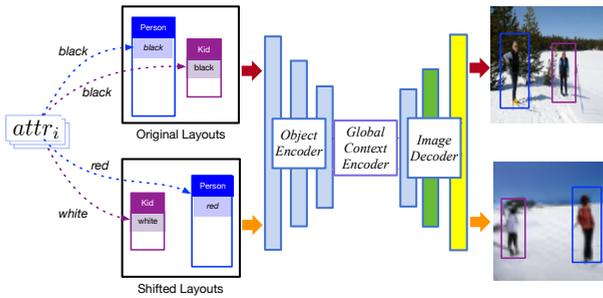


Figure 1: **Attribute-guided Image Generation from Layout.** Unlike prior layout-based image generation architectures, our model allows for instance-level granular semantic attribute control over individual objects (e.g., specifying that a `person` should be wearing something `black` (top) or `red` (bottom)); it also ensures appearance consistency when bounding boxes in layout undergo translation.

labels) and the output is an image consistent with the specified layout. Compared to text-to-image [29, 37] and scene-graph-to-image [0, 10, 19] generation paradigms, layout-to-image provides an easy, spatially aware and interactive abstraction for specifying desired content. This makes this paradigm compelling and effective for users across the spectrum of potential artistic skill sets; from children and amateurs to designers. Image generation from a layout is a relatively new problem, with earlier methods using layout only as an intermediate representation [0, 10], but not a core abstraction or specification exposed to the user.

Layout2Im [40] was the first model proposed for image generation from layout, followed by more recent LostGAN [51], which improved on the performance in terms of overall image quality. However, all current image generation from layout frameworks [51, 40] are limited in a couple of fundamental ways. First, they *lack ability to semantically control individual object instances*. While both Layout2Im and LostGAN model distributions over appearances of objects through appearance [40] or style [51] latent codes, neither is able to control these variations semantically. One can imagine using encoded sample patches depicting desired objects as an implicit *control* mechanism (i.e., generate an instance of a `tree` or `sky` that resembles an example in a given image patch), however, this is in the very least awkward and time consuming from the user perspective. Second, they generally *lack consistency* – are not spatially equivariant. Intuitively, shifting a location (bounding box) of an object in the layout specification, while keeping appearance/style latent code fixed, should result in the object simply shifting by the relatively same amount in the output image (property known as equivariance). However, current models fail to achieve this intuitive consistency. Finally, they are *limited to low-resolution* output images, typically of size 64×64 .

In this paper, we address these challenges by proposing a new framework for attribute-guided image generation from layout, building on, and substantially extending, the backbone of [40]. In particular, we show that a series of simple and intuitive architectural changes: incorporating (optional) attribute information, adopting a global context encoder, and adding additional image generation path where object locations are shifted – leads to the instance-level fine-grained control over the generation of objects, while increasing the image quality and resolution. We call this model attribute-guided layout2im (see Figure 1).

Contributions: Our contributions are three fold: (1) our attribute-guided layout2im architecture allows (but does not require) direct instance-level granular semantic attribute control over individual objects; (2) is directly optimized to be consistent, or equivariant, with re-

spect to spatial shifts of object bounding boxes in the input layout; and (3) allows for higher-resolution output images of size up to 128×128 by utilizing global context encoder and progressive upsampling. Both qualitatively and quantitatively we show state-of-the-art generative performance on Visual Genome [15] benchmark dataset, while benefiting from the desirable control properties, unavailable in other models. The code and pretrained models will be made available ¹.

2 Related Work

Image Generation from Scene Graph: Scene graph is a convenient directed graph structure designed to represent a scene, where nodes correspond to objects and edges to the relationships between objects. Recently, scene graphs have been used in many image generation tasks due to their flexibility in explicitly encoding object-object relationships and interactions [10, 11, 19]. The typical generation process involves two phases: (1) a graph convolutional neural network [6] is applied to the scene graph to predict the scene layout (*i.e.*, bounding boxes and segmentation masks); and (2) the layout is then decoded into an image. Different from methods that generate image as a whole, Li *et al.* [19] propose a semi-parametric approach and crop refining network to reconcile the isolated crops into an integrated image. Unlike in our approach, the scene layout in these models is used as an intermediate semantic representation to bridge abstract scene graph input and an image output.

Image Generation from Layout: Image layout, comprising bounding box locations, often serves as an intermediate step for image generation (see above). Zhao *et al.* [20] proposed image generation from layout as a task in its own right, where the image is generated from bounding boxes and corresponding object categories. To combine multiple objects, [10] sequentially fuse object feature maps using a convolutional LSTM (cLSTM) network; the resulting fused feature map is then decoded to an image. Turkoglu *et al.* [33], on the other hand, divide the generation into multiple stages where objects are added to the canvas one by one. To better bridge the gap between layouts and images, Li *et al.* [17] uses a shape generator to outline the object shape and provide the model fine-grained information from text using object-driven attention. Similarly, [31] learns object-level fine-grained mask maps that outline the object shape. In addition, [12, 31] show that incorporating layout information into normalization layer yields better results: adopting instance normalization technique [12] in generator realize multi-object style control [31], whereas spatially-adaptive normalization [28] based on segmentation mask modulates activation in upsampling layers. Taking the inspiration from [12, 28, 31], we apply spatially-adaptive normalization in our generator to compute layout-aware affine transformations in normalization layers.

Semantic Image Synthesis: Semantic image synthesis is an image-to-image translation task. While most methods use conditional adversarial training [9, 24], such as pix2pix [9], pix2pixHD [24], cVAE-GAN and cLR-GAN [43], others such as Cascaded Refinement Networks [2] also yields plausible results. To preserve semantic information, normalization techniques like SPADE [28] have recently been deployed to modulate the activations in normalization layers through a spatially adaptive and learned transformation. Semantic image synthesis can also serve as an intermediate step for image modeling [35]. In addition, some image-to-image translation tasks are achieved in unsupervised manner [8, 21, 42], but most existing semantic image synthesis models require paired data.

¹<https://github.com/UBC-Computer-Vision-Group/attribute-guided-image-generation-from-layout>

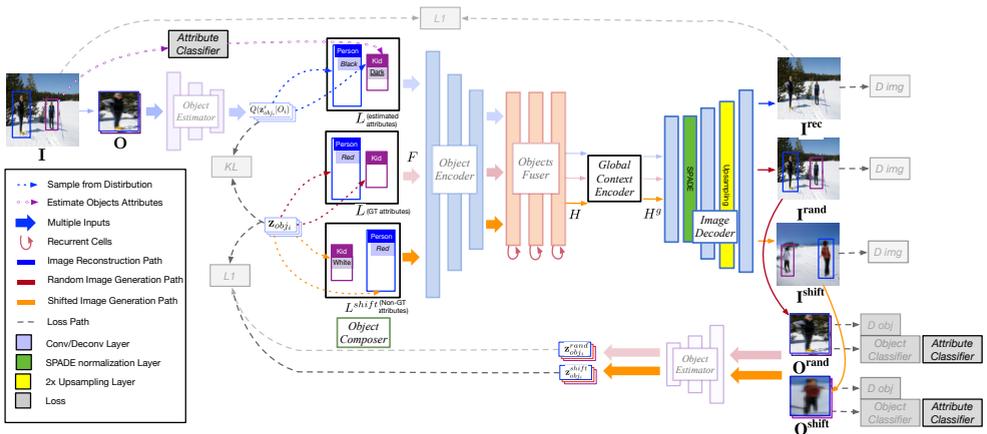


Figure 2: **Overview of our Attribute-Guided Layout2Im training pipeline.** Our architecture has three generation paths: Image reconstruction path (top/blue), Image generation path (middle/red) and Layout reconfiguration path (bottom/orange). Attribute classifier is used in reconstruction path to estimate attributes of objects that do not have any attribute annotations. Non-GT (sampled) attributes are used in image generation path and layout reconfiguration path to disentangle attribute information from class (\mathbf{w}_i) and appearance (\mathbf{z}_{obj_j}).

Attribute-guided Image Generation: In image generation, various attempts have been made to specify the attributes of generated images and objects. For example, [9, 18, 22, 26] aim to edit the attributes of the generated image using natural language descriptions. In [18, 41] authors embed a visual attribute vector (e.g., blonde hair) for attribute-guided face generation. Li *et al.* [19] also incorporates object-level appearance information in the input, but it relies on external memory to sample objects. Another line of the literature concentrates on editing images by providing reference images (e.g., [1]) to transfer style. Different from prior approaches, our model allows direct attribute control over individual instances of objects in complex generative scenes.

3 Approach

Let us start by formally defining the problem. Let Λ be an image canvas (e.g., of size 128×128) and let $L = \{(\ell_i, \mathbf{attr}_i, \mathbf{bbox}_i)_{i=1}^m\}$ be a layout consisting of m labeled object instances with class labels $\ell_i \in \mathcal{C}$, attributes $\mathbf{attr}_i = \{a_{ij}\}_{j=1}^{n_i}$, and bounding boxes defined by top-left and bottom-right coordinates on the canvas, $\mathbf{bbox}_i \subseteq \Lambda$, where $|\mathcal{C}|$ is the total number of object categories and $a_{ij} \in \mathcal{A}$ is the j -th attribute of i -th object instance; \mathcal{A} is the attribute set. Note that each object might have more than one attribute. Let \mathbf{z}_{obj_j} be the latent code for object instance $(\ell_i, \mathbf{attr}_i, \mathbf{bbox}_i)$, modeling class- and attribute-unconstrained appearance. We denote the set of all object instance latent codes in the layout L by $Z_{obj} = \{\mathbf{z}_{obj_j}\}_{j=1}^m$.

Attribute-guided image generation from a layout can then be formulated as learning a generator function G which maps given input (L, Z_{obj}) to the output image I conforming to specifications:

$$I = G(L, Z_{obj}; \Theta_G), \quad (1)$$

where Θ_G parameterizes the generator function G and corresponds to the set of parameters that need to be learned.

Different from previous layout to image generation methods is explicit, but optional (\mathbf{att} ; can be \emptyset or sampled from the prior), inclusion of the attributes. Further, we specifically aim to learn G which is, at least to some extent, equivariant with respect to location of objects \mathbf{bbox}_i in the layout. Our attribute-guided layout2im formulation builds on and improves [40], as such it shares some of the basic architecture design principles outlined in Zhao *et al.* [40].

Training: The overall training pipeline of the proposed approach is illustrated in Figure 2. Given the input image I and its layout $L = \{(\ell_i, \mathbf{att}_i, \mathbf{bbox}_i)_{i=1}^m\}$, our model first creates a word embedding \mathbf{w}_i for each object label ℓ_i and multi-hot attribute embedding $\mathbf{e}_i \in \{0, 1\}^{|A|}$ for object attribute(s)² \mathbf{att}_i , and form a joint object-attribute embedding $M(\mathbf{w}_i \oplus \mathbf{e}_i)$ where \oplus is the vector concatenation and M is a MLP layer, composed of three fully connected layers that map the concatenated vector to a lower dimensional vector. A set of object latent codes $Z_{obj} = \{\mathbf{z}_{obj_i}\}_{i=1}^m$ are sampled from the standard prior normal distribution $\mathcal{N}(0, 1)$, and a new $L^{shift} = \{(\ell_i, \mathbf{att}_i, \mathbf{bbox}_i^{shift})_{i=1}^m\}$ is constructed, where \mathbf{bbox}_i^{shift} represents bounding boxes that are randomly shifted in the canvas Λ . The shifts are horizontal to maintain consistency. Then, our model estimates another set of latent codes $Z'_{obj} = \{\mathbf{z}'_{obj_i}\}_{i=1}^m$, where \mathbf{z}'_{obj_i} is sampled from the posterior $Q(\mathbf{z}'_{obj_i} | O_i)$ conditioned on CNN features of object O_i cropped from input image I . We effectively end up with three datasets:

- Set 1: (L, Z'_{obj}) for reconstruction of the original image. Mapping this input through generator G should result in an image I^{rec} , which is a reconstruction of the training image I serving as the source of the layout L ;
- Set 2: (L, Z_{obj}) for generation of a new image sharing the original layout. The output of G here should be image I^{rand} that shares the layout with above, but where appearance of each object instance is novel (sampled from the prior).
- Set 3: (L^{shift}, Z_{obj}) which is used to generate an image from reconfigured layout (*i.e.*, reconfiguration path, see Suppl. Sec. 1.1 for details). The output should be a corresponding shifted image I^{shift} , which shares latent appearance codes with those from Set 2.

The same pipeline is applied to all three input sets simultaneously: multiple object feature maps F_i are constructed based on the layout L_i and $(\mathbf{z}_{obj_i} \oplus M(\mathbf{w}_i \oplus \mathbf{e}_i))$, and then fed into the object encoder and the objects fuser sequentially, generating a fused hidden feature map H containing information from all specified objects. Lastly, we incorporate a global context embedding \mathbf{g} onto H to form a context-aware feature canvas H^g , and decode it back to an image with a combination of deconvolution, upsampling and SPADE normalization layers [23]. For generated object O_i in I^{rand} and I^{shift} , we make the object estimator regress the sampled latent codes \mathbf{z}_{obj_i} based on O_i to encourage O_i to be consistent with the latent code \mathbf{z}_{obj_i} , and use an auxiliary object classifier and attribute classifier to ensure O_i has the desired category and attributes. To train the model adversarially, we also introduce a pair of discriminators, D_{img} and D_{obj} , to classify results as being real/fake at image and object level.

Inference: At inference time, the model is able to synthesize a realistic image from a given (user specified) layout L and object latent codes Z_{obj} sampled from the prior normal object appearance distribution $\mathcal{N}(0, 1)$. The attributes can be specified by the user or sampled from prior class distributions of object-attribute co-occurrence counts, which we also estimate from data during training. In this way, attribute can be treated as “optional” at individual instance level; *i.e.*, one can specify sub-set of attributes for any sub-set of instances.

²Note exactly n_i elements of \mathbf{e}_i will be 1 and the rest are 0.

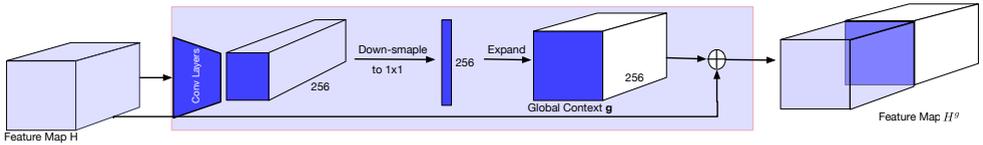


Figure 3: **Global Context Encoder.** The aggregated feature map H is fed into a set of conv layers; then it is down-sampled, spatially expanded and concatenated to H itself to form context-aware H^g feature map that can then be decoded into an image.

3.1 Attribute Encoder

Visual Genome [15] dataset provides attribute descriptions for some objects. For example, a car object might have attributes `red` and `parked`, and a person object might have attributes `smile` and `tall`. There are over 40,000 attributes in the datasets. We only keep the most common 106 attributes for simplicity. In other words, $|\mathcal{A}| = 106$. Each object might have more than one attribute, hence we adopt multi-hot embedding for the attribute encoder. If no attributes are specified for the object, the attribute embedding would be a vector of zeros, *i.e.*, $\mathbf{e}_i = \mathbf{0}$. We concatenate the multi-hot embedding with object word embedding and pass it through an MLP layer to obtain the final object-attribute embedding $M(\mathbf{w}_i \oplus \mathbf{e}_i)$, which is then concatenated with (prior sampled) latent code \mathbf{z}_{obj_i} to construct the object feature map F_i . The feature map F_i is therefore constructed by filling `bboxi` of the feature canvas with $M(\mathbf{w}_i \oplus \mathbf{e}_i) \oplus \mathbf{z}_{obj_i}$.

Attribute Disentanglement: For two novel image generation paths, mainly (L, Z_{obj}) and (L^{shift}, Z_{obj}) we further entice the model to use attribute embedding \mathbf{e}_i to determine appearance of corresponding objects. To explicitly disentangle attribute information from \mathbf{z}_{obj_i} and \mathbf{w}_i during training, we randomly choose a subset of training objects and replace their GT attributes with new attributes sampled from the attributes frequency distribution for the object class. By doing this, we force the generator to use the attributes code \mathbf{e}_i to generate objects with corresponding attributes, instead of encoding attribute information into \mathbf{z}_{obj_i} and/or \mathbf{w}_i .

3.2 Global Context Encoder

At the last stage of the generation process, the fused feature map H is decoded back to the output image by a stack of deconvolution layers. However, the generated image obtained from simply decoding H often contains objects that are not coherent within a scene. For example, it is observed that some generated objects and the background appear incoherent and exhibit sharp transitions (image patch pasting effect). Hence, it is desirable to explicitly incorporate global context information \mathbf{g} in each receptive field of the feature map H , so that, locally, object generation is more informed. Since H contains the information for all objects, it itself is a natural choice for encoding the global context \mathbf{g} (Figure 3). To encode \mathbf{g} , we feed the 8×8 feature map H through two convolution layers to downsample it to a 2×2 feature map, which is average pooled to a global context vector. We then expand the vector to the size of the fused feature map H . This concatenation, $H^g = (\mathbf{g} \oplus H)$, is the final feature map used to decode the image.

3.3 SPADE normalization

Spatially-adaptive (SPADE) normalization [28] is an improved normalization technique that prevents semantic information from being washed away by conventional normalization layers. In SPADE normalization, the learnable transformation (*i.e.*, scale and shift) is learned directly from the semantic layouts. In our model, the feature map H resembles the semantic layout because H encodes both spatial and semantic information. Hence, we add SPADE normalization layers between the deconvolution layers in our image decoder where H is used as the semantic layout. As we show later, in the ablation study (Table 3), the object accuracy of generated results improves when we adopt SPADE.

3.4 Loss Function

The structure of our discriminator D follows the discriminator proposed in layout2im [40] (see Supplemental 1.2 for details), but adds an additional term for attribute classifier cls^{att} , which predicts the attribute of cropped objects and is used to train the generator to synthesize objects with correct attributes. It is trained on real objects O and their attributes \mathcal{A} .

Our GAN model utilizes two adversarial losses: Image Adversarial Loss \mathcal{L}_{adv}^{img} and Object Adversarial Loss \mathcal{L}_{adv}^{obj} . Five more losses, including KL Loss \mathcal{L}_{KL} , Image Reconstruction Loss \mathcal{L}_1^{img} , Object Latent code Reconstruction \mathcal{L}_1^{latent} , Auxiliari Object Classification Loss \mathcal{L}_{AC}^{obj} and Auxiliari Attribute Classification Loss \mathcal{L}_{AC}^{att} , are added to facilitate the generation of realistic images. Due to lack of space we provide details of the loss terms in Supplemental Material (Section 1.3). As the result, the generator G minimizes:

$$\mathcal{L}_G = \lambda_1 \mathcal{L}_{adv}^{img} + \lambda_2 \mathcal{L}_{adv}^{obj} + \lambda_3 \mathcal{L}_{AC}^{obj} + \lambda_4 \mathcal{L}_{AC}^{att} + \lambda_5 \mathcal{L}_{KL} + \lambda_6 \mathcal{L}_1^{img} + \lambda_7 \mathcal{L}_1^{latent} \quad (2)$$

and the discriminator D minimizes:

$$\mathcal{L}_D = -\lambda_1 \mathcal{L}_{adv}^{img} - \lambda_2 \mathcal{L}_{adv}^{obj} + \lambda_3 \mathcal{L}_{AC}^{obj} + \lambda_4 \mathcal{L}_{AC}^{att} \quad (3)$$

where λ_i are weights for different loss terms.

4 Experiments

Datasets: We evaluate our proposed model on Visual Genome [15] datasets. We preprocess and split the dataset following the settings of [10, 31, 40]. In total, we have 62,565 training, 5,506 validation and 5,088 testing images. Each image contains 3 to 30 objects from 178 categories, and each object has 0 to 5 attributes from 106 attribute set. We are unable to train on COCO [20] because it does not provide attribute annotations.

Experimental setup: Our experiments use the pre-trained VGG-net [30] as the base model to compute the inception score for generated image. For object classification loss and the attribute classification loss, our experiments adopted the ResNet-50 model [5] and replace its last fc layer with the corresponding dimensions. Both object and attribute classifier are trained using the objects cropped from real training images. For attribute accuracy, we estimate the attributes of generated objects using a separately trained attribute classifier which consists of five residual blocks, and compute the recall and precision against the GT attribute annotations. Lastly, we generate two sets of test images and use LPIPS metric [39] to compute the diversity score. More specifically, we take the activation of conv1 to conv5 from AlexNet [16], and normalize them in the channel dimension and take the L2 distance. We

Method (64 × 64)	FID ↓	Inception ↑	Obj Acc ↑	Diversity	Attribute Score ↑		Consistency ↑	
					Recall	Precision	bg	fg
Real Images	-	13.9 ± 0.5	49.13	-	0.30	0.31	-	-
sg2im [10]	74.61	6.3 ± 0.2	40.29	0.15 ± 0.12	0.07	0.15	0.87	0.84
Itr. SG [11]	65.3	5.6 ± 0.5	28.23	0.18 ± 0.12	0.04	0.09	0.82	0.81
layout2im [12]	40.07	8.1 ± 0.1	48.09	0.17 ± 0.09	0.09	0.22	0.87	0.85
LostGAN [6]	34.75	8.7 ± 0.4	27.50	0.34 ± 0.10	0.17	0.06	0.63	0.61
Ours	33.09	8.1 ± 0.2	48.82	0.10 ± 0.02 0.20 ± 0.01	0.26	0.30	0.90	0.89

Table 1: **Performance of 64 × 64 image generation** on Visual Genome [13] dataset) For Diversity Score of our model, we have two versions of attribute use: GT attribute specification (top), and sampled attributes from prior class distributions of object-attribute co-occurrence counts (bottom). For Attribute Score, we predict the attributes of generated objects and calculate recall and precision against GT attributes. We trained Interactive SG without the GT object mask. ↑: higher is better; ↓: lower is better; bg: background, fg: foreground.

then average across spatial dimension and across all layers to get the LPIPS metric. 1-LPIPS metric is also used for consistency score, where we compute the foreground diversity between each object before and after it is shifted, and compute the background diversity for the rest of the image which did not undergo the shift. Higher consistency for both is better.

Baselines: We compare our model with Sg2Im [10], Interactive Scene Graph [11], Layout2im [12] and LostGAN [6].

4.1 Quantitative Results

Table 1 and 2 shows the image generation results when trained using different models. For 64 × 64 images, our attribute-guided image generation from layout outperforms all other models in terms of object accuracy, attribute score and consistency score. Notably, our attribute classification score (recall and precision) outperform other models with a substantial margin, demonstrating our model’s capability to control the attributes of generated objects. For consistency in layout reconfiguration, our consistency score is the highest for both background and foreground in the generated images, reflecting the effectiveness of the layout reconfiguration path. Note, as expected, specifying attributes limits the diversity of output objects (0.10 ± 0.02). However, sampled from prior class distributions of object-attribute co-occurrence counts leads to much higher diversity of generated images (0.20 ± 0.01).

We also conduct experiments at 128 × 128 resolution and compare with LostGAN [6]. Our model obtains better results on the object accuracy, attribute score and consistency score.

4.2 Qualitative Results

Figure 4 demonstrates our model’s ability to control the attributes of generated objects. For each image, we pick an object and replace its current attribute with a different one, while keeping the rest of the layout unchanged. It can be seen from Figure 4 that our model is able to change the attributes of the object of interest, and keep the rest of the image unmodified.

Figure 5 compares the results before and after some object bounding boxes in the canvas are horizontally shifted. For each images pair, the image on the left is generated from the GT layout, and the image on the right from the reconfigured layout. Our model shows better

Method (128 × 128)	FID ↓	Inception ↑	Obj Acc ↑	Diversity	Attribute Score ↑		Consistency ↑	
					Recall	Precision	bg	fg
Real Images	-	26.15 ± 0.23	41.92	-	0.27	0.27	-	-
LostGAN [50]	29.36	11.1 ± 0.6	25.89	0.43 ± 0.09	0.19	0.04	0.54	0.51
Ours	39.12	8.5 ± 0.1	31.02	0.15 ± 0.09	0.10	0.25	0.84	0.80

Table 2: **Performance of 128 × 128 images** on Visual Genome [15] dataset. We note that images generated by LostGAN [50] contains too many attributes signals, which explains its high recall and low precision. ↑: higher is better; ↓: lower is better.

Method (64 × 64)	Inception ↑	Accu. ↑	Diversity	Attribute Score ↑		Consistency ↑	
				Recall	Precision	bg	fg
w/o attribute specification	7.9 ± 0.05	48.01	0.19 ± 0.08	0.08	0.13	0.88	0.87
w/o location change	7.8 ± 0.1	48.96	0.12 ± 0.05	0.25	0.30	0.86	0.84
w/o SPADE [28]	7.9 ± 0.1	45.05	0.15 ± 0.07	0.23	0.29	0.89	0.88
w/o context encoder	7.7 ± 0.1	47.96	0.13 ± 0.15	0.24	0.30	0.89	0.87
full model	8.0 ± 0.2	48.82	0.10 ± 0.02	0.26	0.30	0.90	0.89

Table 3: **Ablation study of our model** on Visual Genome [15] dataset by removing different objectives. Inception is the inception score, Accu. is the object classification accuracy, and Diversity is the diversity score. ↑: higher is better; ↓: lower is better.

layout reconfigurability than other methods. For example, in Figure 5(b’) the boat is moved, in (d’) two human are moved, and in (e’) the horse is moved. In contrast, layout2im [40] and LostGAN [50] either change the theme of the image (see 5(f’)), or have missing objects (see 5(d’)) for reconfigured layouts. This is also reflected in their much lower consistency score.

Additional examples at 128 × 128 resolution are in Supplemental, Figure 1 and 2. Supplemental Figure 4 shows generated images obtained using different SoTA models. Our method consistently outperforms baselines in quality and consistency of generated images. LostGAN [50] fails to generate plausible human faces, and layout2im [40] does not generate realistic objects such as food.

4.3 Ablation Study

We demonstrate the necessity of our key components by comparing scores of several ablated models trained on Visual Genome [15] dataset. As shown in Table 3, removing any components is detrimental to the model’s performance. Not surprisingly, attribute specification is the key to the successful attribute classification. The absence of layout reconfiguration path decreases the inception score by 0.2, slightly increases the classification accuracy and, more importantly, reduces the consistency for reconfigured layouts. SPADE [28] is beneficial for object classification accuracy, and global context encoder improves inception score by 0.3.

5 Conclusions

This paper proposes attribute-guided image generation from layout, an effective approach to control the visual appearance of generated images in instance-level. We also showed that the model ensures visual consistence of generated images when bounding boxes in layout undergo translation. Qualitative and quantitative results on Visual Genome [15] datasets demonstrated our model’s ability to synthesize images with instance-level attribute control and improved level of visual consistence.



Figure 4: **Examples of 64×64 generated images with modified attributes on Visual Genome [15] datasets obtained by our proposed method.**

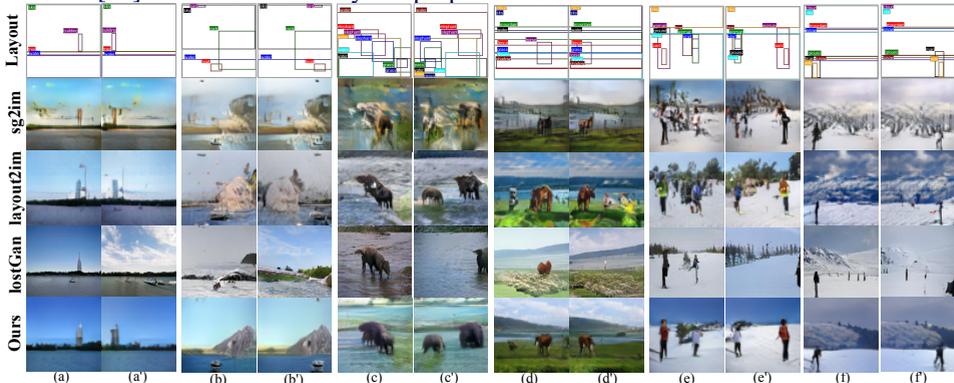


Figure 5: **Examples of 64×64 generated images with horizontally shifted bounding boxes on Visual Genome [15] datasets**

Acknowledgments

This work was funded, in part, by the Vector Institute for AI, Canada CIFAR AI Chair, NSERC CRC and the NSERC DG and Discovery Accelerator Grants. Hardware support was provided by JELF CFI grant and Compute Canada under the RAC award.

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 (ICCV)*, pages 4561–4569, 2019.
- [2] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1511–1520, 2017.
- [3] Hao Dong, Simiao Yu, Chao Wu, and Yike Guo. Semantic image synthesis via adversarial learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 5706–5714, 2017.
- [4] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

- [6] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- [7] 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.
- [8] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1125–1134, 2017.
- [10] 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 (CVPR)*, pages 1219–1228, 2018.
- [11] Levent Karacan, Zeynep Akata, Aykut Erdem, and Erkut Erdem. Learning to generate images of outdoor scenes from attributes and semantic layouts. *arXiv preprint arXiv:1612.00215*, 2016.
- [12] 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 (CVPR)*, pages 4401–4410, 2019.
- [13] Hadi Kazemi, Seyed Mehdi Iranmanesh, and Nasser Nasrabadi. Style and content disentanglement in generative adversarial networks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 848–856. IEEE, 2019.
- [14] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1857–1865, 2017.
- [15] 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.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [17] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12174–12182, 2019.
- [18] Xinyang Li, Jie Hu, Shengchuan Zhang, Xiaopeng Hong, Qixiang Ye, Chenglin Wu, and Rongrong Ji. Attribute guided unpaired image-to-image translation with semi-supervised learning. *arXiv preprint arXiv:1904.12428*, 2019.

- [19] Yikang Li, Tao Ma, Yeqi Bai, Nan Duan, Sining Wei, and Xiaogang Wang. Pastegan: A semi-parametric method to generate image from scene graph. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, pages 3950–3960. 2019.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [21] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017.
- [22] Yongyi Lu, Yu-Wing Tai, and Chi-Keung Tang. Attribute-guided face generation using conditional cycleGAN. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 282–297, 2018.
- [23] Elman Mansimov, Emilio Parisotto, Jimmy Ba, and Ruslan Salakhutdinov. Generating images from captions with attention. In *International Conference on Learning Representations*, 2016.
- [24] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [25] Takeru Miyato and Masanori Koyama. cGANs with projection discriminator. In *International Conference on Learning Representations*, 2018.
- [26] Seonghyeon Nam, Yunji Kim, and Seon Joo Kim. Text-adaptive generative adversarial networks: manipulating images with natural language. In *Advances in Neural Information Processing Systems*, pages 42–51, 2018.
- [27] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2642–2651, 2017.
- [28] 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*, pages 2337–2346, 2019.
- [29] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, pages 1060–1069, 2016.
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- [31] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 10531–10540, 2019.
- [32] Fuwen Tan, Song Feng, and Vicente Ordonez. Text2scene: Generating abstract scenes from textual descriptions. *arXiv preprint arXiv:1809.01110*, 2018.

- [33] Mehmet Ozgur Turkoglu, William Thong, Luuk Spreeuwers, and Berkay Kicanaoglu. A layer-based sequential framework for scene generation with gans. *arXiv preprint arXiv:1902.00671*, 2019.
- [34] 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 (CVPR)*, pages 8798–8807, 2018.
- [35] Xiaolong Wang and Abhinav Gupta. Generative image modeling using style and structure adversarial networks. In *European Conference on Computer Vision (ECCV)*, pages 318–335. Springer, 2016.
- [36] 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 (CVPR)*, pages 1316–1324, 2018.
- [37] 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 (ICCV)*, pages 5907–5915, 2017.
- [38] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*, 2018.
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018.
- [40] 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 (CVPR)*, pages 8584–8593, 2019.
- [41] Xingran Zhou, Siyu Huang, Bin Li, Yingming Li, Jiachen Li, and Zhongfei Zhang. Text guided person image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3663–3672, 2019.
- [42] 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 (ICCV)*, pages 2223–2232, 2017.
- [43] 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*, pages 465–476, 2017.
- [44] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, and Bill Freeman. Visual object networks: image generation with disentangled 3d representations. In *Advances in Neural Information Processing Systems*, pages 118–129, 2018.