

SceneGenie: Scene Graph Guided Diffusion Models for Image Synthesis

Azade Farshad^{1,2} *Yousef Yeganeh^{1,2} *Yu Chi¹ *Chengzhi Shen¹Björn Ommer³Nassir Navab¹¹Technical University of Munich²MCML³Ludwig Maximilian University of Munich

Prompt: a sheep by another sheep standing on the grass with sky above and the ocean by a tree and a boat on the grass



Figure 1: Synthesized images from SG2Im [28], GLIDE [46], DALL.E 2 [53], and SceneGenie. For SG2Im, and SceneGenie, the sentence is first converted to a scene graph before feeding into the model. While the sentence describes two sheep and a boat in a specified scene, text-to-image generators like GLIDE, and even larger models like DALL.E 2 generate considerably inaccurate results; while SceneGenie accurately represents the scene defined by the prompt.

Abstract

Text-conditioned image generation has made significant progress in recent years with generative adversarial networks and more recently, diffusion models. While diffusion models conditioned on text prompts have produced impressive and high-quality images, accurately representing complex text prompts such as the number of instances of a specific object remains challenging.

To address this limitation, we propose a novel guidance approach for the sampling process in the diffusion model that leverages bounding box and segmentation map information at inference time without additional training data. Through a novel loss in the sampling process, our approach guides the model with semantic features from CLIP embeddings and enforces geometric constraints, leading to high-resolution images that accurately represent the scene. To obtain bounding box and segmentation map information, we structure the text prompt as a scene graph and enrich the nodes with CLIP embeddings. Our proposed model

achieves state-of-the-art performance on two public benchmarks for image generation from scene graphs, surpassing both scene graph to image and text-based diffusion models in various metrics. Our results demonstrate the effectiveness of incorporating bounding box and segmentation map guidance in the diffusion model sampling process for more accurate text-to-image generation. Project Page: scenegenie.github.io/SceneGenie/

1. Introduction

Image generation using deep neural networks is a rapidly evolving field in computer vision, with the objective of creating models that have a deep understanding of the objects and scenes they are creating. In recent years, significant progress has been made in text-to-image synthesis using Recurrent Neural Networks (RNNs) [75] and Generative Adversarial Networks (GANs) [49, 60], which can generate high-quality, photorealistic images from textual descriptions. Lately, diffusion models, a class of generative mod-

*The first three authors contributed equally to this work

els, excelled GAN models [9] and became the prominent method in the image generation task. However, most of these methods often struggle with creating complex scenes from long, natural language descriptions. This is because sentences are linear structures that may not efficiently describe complex scenes.

To tackle this problem, we propose SceneGenie, which is a novel layout-based approach for guiding the sampling process of a diffusion model. Our method leverages bounding box and segmentation information as a guidance in the reverse sampling process. The bounding box and segmentation map information are predicted by a Graph Neural Network (GNN) after structuring the text prompt in the form of a scene graph. We propose using scene graphs as they are powerful structured representations of objects and their relationships in both the image and language domains.

Our proposed guidance is similar to classifier guidance. We compute the classifier gradients for each object based on the distance between the CLIP image embedding in the region of interest (RoI) for that specific object and the corresponding CLIP text embedding for that object in the form of a *Photo of an obj*. To compute object-wise gradients in the RoI, we inject gaussian noise outside the RoI and then compute the total gradient as the weighted sum of the gradients for different objects in the scene. For segmentation guidance, we take advantage of the first-stage autoencoder of the diffusion model to measure how semantically close the segmentation map and the image are. Using such guidance for the diffusion model sampling results in higher quality and more accurate images that better represent the input prompt.

Recently, there have been approaches such as Make a Scene [14, 66] that condition the diffusion model directly on the segmentation map or scene layout, or methods such as SDEdit [40] that use the segmentation map as the initialization in the sampling process. However, these methods either require additional training for the input condition or need different architectural designs for different conditions.

Our proposed method differs from these other works in that it directly optimizes the sampling process using additional information and does not necessarily need to be paired with the images in the training dataset due to its usage at inference time. This allows us to create more complex and accurate scenes while still maintaining the high quality of generated images. We demonstrate the effectiveness of our approach through experiments on public benchmarks, showing that our method outperforms existing text-to-image diffusion models as well as state-of-the-art scene graph to image approaches without any additional training.

In summary, our work makes several key contributions:

- We propose a novel approach for guiding the sampling process in a diffusion model that places greater emphasis on the regions of interest (RoI) by incorporating the

gradients computed from predicted bounding box and segmentation maps.

- Our proposed guidance is applied during the reverse sampling process. Therefore, it does not require any additional training and can be applied to any diffusion model architecture.
- To enable the use of bounding box guidance in the sampling process, we propose a novel method of noise injection outside the RoI. For the segmentation guidance, we take advantage of the first-stage autoencoder of the diffusion model. Therefore, we effectively leverage the bounding box and segmentation map information and improve the accuracy of generated images.
- Our method achieves notably higher image generation performance compared to scene graph to image models in high resolution image generation, and outperforms the state-of-the-art in text-to-image generation.
- Finally, we demonstrate that incorporating CLIP embeddings as node features in the scene graph improves the accuracy of bounding box and segmentation predictions.

2. Related Work

The high dimensionality of images poses a challenge for image generation based on deep learning. Recent advances in generative models, in particular, Generative Adversarial Networks (GANs) [16], have boosted the quality and diversity of generated images. A line of works explore generative models for unconditional image generation [31, 30]. Conditional image generation models have also been explored [41] with a diverse set of priors such as semantic segmentation maps [5, 64, 49], natural language descriptions [72, 35] or translating from one image domain to another using paired [25] or unpaired data [74]. Conditional image generation models also enable the possibility of interactive image manipulation by partial image generation using hand-crafted part replacement [24] or by incorporating a user interface for specifying the locations that need to be inpainted [33, 3]. The inpainting process of the specified regions [50, 38, 68] can also be guided by semantic information [69, 23, 47] or edges [70, 44]. For instance, in GLIDE [46], the model is capable of replacing original content with the guidance from the CLIP [52] embeddings by extra fine-tuning.

Diffusion Models A recent and impressive improvement in image generation is achieved by diffusion models [22, 9]. Diffusion models [22, 9] are generative models that produce images by successively denoising images. Unconditional image generation with diffusion models initiated with the




Figure 2: **Overview of SceneGenie.** Our pipeline starts by getting a text prompt as input. The triplets are structured in form of a scene graph. The graph is then processed by a GCN which outputs object embeddings per node in the graph. The embeddings are used to predict bounding box coordinates and pseudo-segmentation maps corresponding to each object in the scene, as well as a final segmentation map from the whole image. We use a diffusion model to generate images conditioned on the text prompt by guiding the sampling process through the bounding box and segmentation map information.

denoising approach with the work of Sohl *et al.* [57] that formalized diffusion models as multi-scale convolutions, and then the same work was extended by Ho *et al.* [22]; later, Song *et al.* [58] proposed a non-Markovian method for the forward process. Conditional image generation by diffusion models based on classifier guidance was proposed by Dhariwal *et al.* [9], which uses a classifier to guide the diffusion model during the sampling process. In many applications, the conditional diffusion models are utilized; for instance, in Palette [56], the diffusion model is conditioned on a low-resolution image to generate a high-resolution image, or in SDEdit [40], the diffusion model is conditioned with a low-quality image to sharpen and enhance colors and textures. Recently, image generation conditioned on text has captured a lot of attention, where CLIP [52] guidance is often utilized.

Contrastive Language-Image Pre-Training By successfully incorporating text-and-image pairs through contrastive learning, Contrastive Language-Image Pre-Training (CLIP) [52] has been extensively applied in object detection [63, 17], image captioning [43, 20], and text-to-image generation [46]. CLIP [52] is able to capture similar representation beyond modality by minimizing the distance between text and picture embeddings from the same pair while maximizing the distance between those from dissimilar pairs.

Scene Graph to Image Scene graphs [29] are graphs that represent a scene by defining the objects in the scene as nodes in the graph and the relationships between them as edges. Scene graphs gained more attention recently due to the rise of large-scale scene graph datasets such as Visual Genome [32], MOMA [39], and Action Genome [27]. A broad line of works [37, 65, 45, 19, 51, 71, 62, 59] explore the generation of scene graphs from images, while Johnson *et al.* proposed SG2Im [28], which is the first pipeline that attempts to generate images from scene graphs. Typically, the task can be divided into two parts: first, convert a scene graph to an intermediate layout, then use the layout as an input to conditional GANs [5, 49, 60, 12] for image synthesis. Some works [73, 60, 61] focus on generating images directly from scene layouts. Recent works either focus on enhancing the model’s capacity for graph understanding [15] or improving intermediate layout quality [67]. Herzig *et al.* [18] deals with semantic equivalence in large complex scene graphs, while [26] intends to reduce blurry and overlapping objects in the scene layout in a coarse-to-fine manner. Modified GCNs are engaged in a cascaded refinement network in [5] to reduce high-dimensional embeddings. However, the literature has widely studied that deep Graph Neural Networks suffer from over-smoothing issues [48, 34, 1] that average all information in the graph, causing semantic ambiguities. Recently, there has been models [8, 42, 2, 13, 11] proposed that focus on image manipulation, where users are able to control the synthesis results by modifying the scene

graph interactively. More recently video generation from scene graphs [6] has also been introduced.

3. Background

Diffusion models use a Markov chain that gradually adds Gaussian noise to an image x_0 to get the approximate posterior $q(x_{1:T}|x_0)$ with x_1, \dots, x_T being the noisy versions of x_0 . If T is large enough x_T is approximated by $\mathcal{N}(0, \mathcal{I})$. By learning the reverse process $p_\theta(x_{t-1}|x_t) := \mathcal{N}(\mu_\theta(x_t), \Sigma_\theta(x_t))$ of this Markov chain, one can generate new images $x_0 \sim p_\theta(x_0)$ from pure noise $x_T \sim \mathcal{N}(0, \mathcal{I})$ by gradually denoising in a sequence of steps $x_{T-1}, x_{T-2}, \dots, x_0$. Such a model is obtained by generating noisy samples $x_t \sim q(x_t|x_0)$ and training a model θ (typically a U-Net) to predict the added noise using an MSE loss:

$$L_{DM} = E_{x_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(x_t)\|^2] \quad (1)$$

This model can then successively generate images by denoising images step by step starting from pure noise.

3.1. Guided Diffusion

Denoising diffusion probabilistic models (DDPM) [22] have shown exceptional performance in unconditional image generation. Yet, generating images with desired semantics is still challenging due to DDPM’s nature as a stochastic generation process. Therefore previous work has focused on classifier guiding [9], perturbing the mean $\mu_\theta(x_t|y)$ and variance $\Sigma_\mu(x_t|y)$ of the diffusion model by a classifier gradient. The perturbed mean $\hat{\mu}_\theta(x_t|y)$ is given by

$$\hat{\mu}_\theta(x_t|y) = \mu_\theta(x_t|y) + \alpha \cdot \Sigma_\theta(x_t|y) \nabla_{x_t} \log p_\phi(y|x_t) \quad (2)$$

Where α is a hyperparameter called guidance scale that controls sample quality vs. sample diversity [9].

3.2. Latent Diffusion Models

Latent Diffusion Models (LDMs) [54] are trained to apply the diffusion process and reverse sampling process on image latent space, which significantly reduces the computational complexity compared to the diffusion models trained on the image space. Latent embeddings of the images are coded by KI-autoencoder [54] or VQGAN [10]. The diffusion and denoising processes of LDM [54] can be derived as:

$$q(z_{1:T}|z_0) = \prod_{t=1}^T q(z_t|z_{t-1}) \quad (3)$$

$$p_\theta(z_{0:T}) = p(z_T) \prod_{t=1}^T p_\theta(z_{t-1}|z_t) \quad (4)$$

4. Methodology

Our method consists of two steps: 1) Training a model for the prediction of bounding boxes and segmentation

maps from a scene graph obtained from the text prompt, 2) Generating the image guided by the bounding box coordinates, the segmentation map, and the text embedding. We focus on the extraction of bounding boxes and segmentation maps from scene graphs in subsection 4.1, while subsection 4.2 concentrates on guiding the diffusion model in order to generate more accurate images.

We are given a dataset \mathcal{D} of images x , and text prompts τ , and bounding box coordinates c . Optionally, we can have segmentation maps s . The text prompts are split into triplets in the form of (object, predicate, subject), where the predicate r defines the relationship between the object and the subject. The scene graph \mathcal{G} is composed of the the object categories o and the relationships r , where object categories are the nodes in the graph and the edges are the relationships between them. More formally, a graph can be formed as $\mathcal{G} = (O, E)$ where $O = \{o_1, \dots, o_n\}$ are n objects in the graph, and $E = \{(o_i, r, o_j) | o_i, o_j \in O, r \in R\}$ with R as the relationship category between objects.

4.1. Scene Graph to Segmentation (SG2SEG)

Given the scene graph \mathcal{G} with objects (nodes) and relationships (edges), we aim to synthesize segmentation maps that transform the information from the text space to the image space. These segmentations ought to be realistic in terms of object shapes, as well as semantically consistent in terms of object relationships.

Firstly, we acquire object embeddings using CLIP [52] features in each node of the graph. The assumption is that CLIP [52] is able to generate the object features f_{obj} that are consistent with the text feature f_{text} describing each object. To form the input to CLIP’s [52] text encoder, we build a prompt for each $o_i \in O$, e.g., *a photo of an [obj]* with *[obj]* substituted by the corresponding object classes. The generated CLIP [52] features $f_{text} \in R^{n \times 512}$ are then fed to a Graph Neural Network to learn the object embedding $f_{emb} \in R^{n \times d}$, where d is a hyperparameter that controls the embedding dimensionality. As a trade-off between computation and model expressiveness, we select $d = 128$. Note that, we treat edge E , that stands for relationships, as learnable embeddings.

Secondly, we explicitly constrain the output of our SG2SEG on ground truth, instead of learning an intermediate representation, as in SG2Im [28]. Given the i -th object embedding $f_{emb}^i \in R^{1 \times d}$, we apply a mask regression network (Mask Net) and box regression network (B. Box Net) as in SG2Im [28], but remove the remaining parts of layout sampling and merging. The benefits are twofold: 1) By explicitly constraining the predicted mask and bounding boxes, we are able to achieve higher quality in both results, 2) Predictions are more reliable with even less computational power.

The SG2SEG network is optimized with three objective

functions:

1) Box loss $L_{box} = \sum_{i=1}^m ||c_i, \hat{c}_i||_1$ is the L1 difference between the 4 coordinate values of the predicted and ground truth bounding boxes.

2) Mask loss $L_{mask} = BCE(m_i, \hat{m}_i)$ is the binary cross entropy loss for each predicted object mask and ground truth.

3) Segmentation loss $L_{seg} = ||s, \hat{s}||_1$ is the L1 difference between the segmentation map from predicted masks and the ground truth. It helps the network to generate more consistent results when merging multiple object masks together.




Figure 3: Example of Gaussian Noise padded bounding box in the image. We compute the gradients (∂obj) for each object in the image based on the CLIP score and then aggregate the gradients for all the object for the backpropagation step.

4.2. Conditional Image Generation with Diffusion Models

In our work, we replace the gradient of the classifier in Equation 2 with combinations of gradients defined in subsection 4.2.1, subsection 4.2.2, and subsection 4.2.4. The main objective is to guide the image generation toward a correct scene layout and object realism.

4.2.1 CLIP Text Guidance

In order to generate an image which corresponds to a specific input text in the sampling process, we add guidance in the sampling process. Given an input text $text$ and input image img along with CLIP [18] image encoder E_i and CLIP [18] text encoder E_t , we calculate the gradient of the CLIP score with respect to the latent space z of input image of LDM [54]. The calculated gradient is used in the sampling process for the guidance as described in equation 2. It is formulated as:

$$L_{text} = E_i(img) * E_t(text) \quad (5)$$

For LDM, the gradient is formulated as:

$$\nabla L_{text} = -\frac{\partial L_{text}}{\partial z} \quad (6)$$

For other diffusion model whose diffusion process is on image space x , the gradient is computed by:

$$\nabla L_{text} = -\frac{\partial L_{text}}{\partial x} \quad (7)$$

4.2.2 CLIP Bounding Box Guidance

In addition to using an input prompt to generate the entire image, we want to make a certain region in image that corresponds to the input prompt. Here we propose CLIP [18] Bounding Box Guidance. From SG2SEG framework, the size and location of bounding box of a certain object is inferred.

For an object obj_k in a bounding box, we pad the bounding box with Gaussian Noise into the size of the original image as illustrated in Figure 3. Then we calculate the CLIP score between CLIP [52] image embedding of Gaussian Noise padded Bounding Box obj_k and CLIP text embedding of the category of object l_k in format starting with "A photo of" as follows:

$$L_{obj_k} = E_i(obj_k) * E_t(l_k) \quad (8)$$

where E_i and E_t are pretrained CLIP Image Encoder and Text Encoder.

Assuming there are N objects $obj_1, obj_2, \dots, obj_N$ in one image, we have the entire bounding box guidance score as weighted sum of the object bounding box guidance:

$$L_{boxg} = \sum_{i=1}^N w_i * L_{obj_i} \quad (9)$$

The weights w_i are normalized and proportional to the size of each bounding box. The gradient used for the guidance of LDM[54] will be:

$$\nabla L_{boxg} = -\frac{\partial L_{boxg}}{\partial z} \quad (10)$$

, where z is the image latent of LDM[54]. For other diffusion model whose diffusion process is on image space x , the gradient is formulated as

$$\nabla L_{boxg} = -\frac{\partial L_{boxg}}{\partial x} \quad (11)$$

4.2.3 Augmented CLIP Bounding Box Guidance

Based on CLIP Bounding Box guidance, we propose an Augmented CLIP Bounding Box Guidance by strengthening the guidance gradient with gradient of Gaussian noise.




Figure 4: **Some qualitative results on the comparison of SceneGenie against related work on the COCO stuff [4] test set.** As it can be seen, the images generated by SceneGenie represent the given prompt more accurately compared to previous work. SceneGenie, in addition to high quality image generation, correctly generates the number of given instances in the image and represents the scene more accurately overall.

The aim of Augmented CLIP Bounding Box Guidance is to increase the guidance in the region where the object should appear. Along with calculating the gradient of a Gaussian Noise padded image above, we also calculate the gradient of a pure Gaussian Noise γ with the same size.

$$\nabla L_{gauss} = -\frac{\partial L_{gauss}}{\partial \gamma} \quad (12)$$

The formulation for this Augmented CLIP Bounding Box Guidance $\nabla L_{aug,boxg}$ is:

$$\nabla L_{aug,boxg} = \lambda * (\nabla L_{boxg} - \nabla L_{gauss}) + \nabla L_{gauss} \quad (13)$$

where ∇L_{bbox} represents the gradient of Gaussian Noise padded bounding box, ∇L_{gauss} represents the gradient of pure Gaussian Noise with the same size, λ is a hyperparameter to control the intensity of guidance. If λ is set to 1, Augmented CLIP Bounding Box Guidance is exactly the same as the vanilla CLIP Bounding Box Guidance.

4.2.4 Segmentation Map Guidance

Given the first-stage autoencoder of the LDM [54], $T(\cdot)$, segmentation map s for the image and generated image x in the reverse process of the diffusion model, we calculate the score to measure how semantically close the segmentation

map and the generated image are. The score is formulated as:

$$L_{segg} = T(s) * T(x) \quad (14)$$

, the gradient of the score function for LDM [54] with respect to the latent space z will be

$$\nabla L_{segg} = -\frac{\partial L_{segg}}{\partial z} \quad (15)$$

Then, the total diffusion guidance gradient is computed as follows:

$$\nabla L_{diff} = \nabla L_{segg} + \nabla L_{text} + \nabla L_{aug,boxg} \quad (16)$$

5. Experiments

In this section, we present the implementation details of our method and the results of our experiments on two public benchmarks, which are commonly used for image generation from scene graphs, namely Visual Genome [32], and COCO stuff [4]. We evaluate our model both quantitatively and qualitatively on these datasets and compare them against the state-of-the-art in Text2Image and Scene Graph to image models.

5.1. Datasets

The Visual Genome [32] dataset consists of images and semantic scene graph annotations, along with the bounding box coordinates. The relationships in the scene graphs of VG dataset are purely semantic and only implicitly encode geometric information; while the COCO [4] dataset does not originally include scene graph annotations, the bounding box coordinates and the captions in this dataset are used to generate geometric scene graphs. The COCO dataset includes images with semantic segmentation, bounding box coordinates and captions as annotations.

5.2. Experimental Setup

For all the experiments unless specified, we use a pre-trained U-Net [55] as our diffusion model based on LDM [54], and perform image generation using our proposed guidance. The diffusion model is pre-trained on the ImageNet dataset [7]. Our model does not require any fine-tuning and is applicable to existing networks while the guidance happens during the inference time. We adopt the LDM-8 (KL) pre-trained model. The sampling process is done via DDIM [58] sampling with 100 sampling steps. For the model trained on 64×64 images from COCO, we combine our guidance with GLIDE [46].

We report the performance of our model using inception score (IS), Fréchet Inception Distance (FID) and Kernel Inception Distance (KID) which are common image quality metrics. In addition, we report the Semantic Object Accuracy [21] metrics (SOA-O and SOA-I) for our model and

the LDM [54] that checks whether a pre-trained object detection model recognizes the given objects. We use CLIP [52] as the text encoder. We empirically found 0.5 as the best value for scaling segmentation guidance in the total guidance. The architecture details will be provided in the supplementary material.

Since the VG dataset does not include semantic segmentation annotations, we omit the scene graph to segmentation step in our model and only predict the bounding box coordinates. For the same reason, our final model in VG only uses CLIP embeddings and bounding box guidance.

5.3. Results

We provide qualitative and quantitative results of our approach compared against the state-of-the-art. We present two variations of our model in Table 1 and Table 2, which are either with predicted or ground truth bounding box and segmentation map information. We present more qualitative results on COCO and VG in the supplementary material.

Comparison against SOTA The results of our model compared against the state-of-the-art on COCO [4] and Visual Genome (VG) [32] datasets are provided in Table 1 and Table 2 respectively. As it can be seen, our proposed model SceneGenie, outperforms the state-of-the-art diffusion model, LDM [54] as well as the scene graph to image [28] model on both datasets.

We also present some qualitative results on COCO in Figure 4. The qualitative results show that our proposed model generates more accurate images conditioned on the prompt. One main advantage of our model is in situations, where the number of object instances are defined. In such cases, the text guided image generation models fail in representing the scene correctly, while SceneGenie generates a more representative image.

Ablation Study We present an ablation study of the components of our model in Table 3. We analyze different values for scaling the augmented bounding box in the diffusion process, and we find the best value of 1.2 based on FID. The best overall performance is obtained by combining the bounding box and segmentation guidance with GT values. We also analyze the effect of incorporating CLIP embeddings in the graph nodes for the models with predicted bounding box and segmentation map and show its effectiveness in improving the image generation quality. In addition, we measure the bounding box prediction error with and without using CLIP embeddings for the nodes in the graph. The bbox prediction error is 0.736, and 0.749, with an without CLIP embeddings respectively.

Table 1: **Comparison against SOTA on COCO stuff [4].** We present the results on 64×64 , and 256×256 resolutions. We present the results of different methods with different generator architectures. The models identified by **Pred** use predicted bounding box, segmentation map, or scene layouts, while **GT** identifies experiments with ground truth information.

Method	Pred / GT	Resolution	IS \uparrow	FID \downarrow	KID ($\times 10^2$) \downarrow	SOA-C ($\times 10^2$) \uparrow	SOA-I ($\times 10^2$) \uparrow
SG2Im [28]	GT	64×64	5.30	113.61	$58_{\pm 0.3}$	-	-
SceneGenie (Ours)	GT	64×64	9.05	67.51	$7.86_{\pm 0.087}$	-	-
SG2Im [28]	GT	256×256	6.6	127.0	-	-	-
PasteGAN [36]	GT	256×256	11.0	70.2	-	-	-
Specifying [2]	GT	256×256	12.4	65.2	-	-	-
Canonical [18]	GT	256×256	19.5	64.65	$7.03_{\pm 0.177}$	33.94	48.55
LDM [54]	GT	256×256	22.24	63.83	$6.06_{\pm 0.114}$	45.38	57.22
SceneGenie (Ours)	GT	256×256	21.72	63.05	$5.54_{\pm 0.105}$	45.67	56.91
SceneGenie (Ours + Seg)	GT	256×256	21.50	62.38	$5.10_{\pm 0.095}$	45.80	57.39
Canonical [18]	Pred	256×256	9.03	113.30	$7.67_{\pm 0.173}$	34.78	50.93
SceneGenie (Ours)	Pred	256×256	22.16	63.27	$4.98_{\pm 0.101}$	43.80	56.61

Table 2: **Comparison against SOTA on Visual Genome [32].** The results are presented on images with 256×256 resolution. Our final model is the combination of predicted bounding box, with augmented bounding box guidance.

Method	IS \uparrow	FID \downarrow	KID ($\times 10^2$) \downarrow
Canonical [18]	16.5	45.7	-
LDM [54]	20.02	42.69	$8.63_{\pm 0.505}$
SceneGenie (Ours)	20.25	42.21	$8.43_{\pm 0.517}$

Table 3: **Ablation Study on COCO stuff [4].** We study the different components of our model. We analyze the effect of bounding box and segmentation guidance, and the different values for the bounding box guidance scale. **B**: Bounding Box, **S**: Segmentation Map.

Guidance	λ	IS \uparrow	FID \downarrow	KID ($\times 10^2$) \downarrow
GT				
-	-	$22.24_{\pm 1.778}$	63.83	$6.06_{\pm 0.114}$
B	1	$21.46_{\pm 1.49}$	63.14	$5.82_{\pm 0.114}$
B	1.1	$21.93_{\pm 1.44}$	63.14	$5.21_{\pm 0.103}$
B	1.2	$21.72_{\pm 1.45}$	<u>63.05</u>	$5.54_{\pm 0.105}$
B	1.3	$21.87_{\pm 1.60}$	64.19	$5.78_{\pm 0.109}$
B	1.4	$22.07_{\pm 1.83}$	63.76	$5.93_{\pm 0.111}$
B + S	-	$21.50_{\pm 1.31}$	62.38	$5.10_{\pm 0.095}$
Pred				
B	-	$21.73_{\pm 1.65}$	63.61	$5.88_{\pm 0.114}$
B + CLIP	-	$22.04_{\pm 2.19}$	63.37	$5.31_{\pm 0.112}$
B + S + CLIP	-	$22.16_{\pm 1.65}$	63.27	$4.98_{\pm 0.101}$

5.4. Discussion

The introduction of bounding box and segmentation map guidance in our approach enables the model to accurately represent the scene. As it can be seen in the qualitative re-

sults in Figure 4, the generated images by our model represent the input prompt more accurately. Specifically, when the input prompt defines the number of objects in the scene, previous works fail to correctly generate the specified number of objects (e.g. Generating an image of two skiers instead of one or an image of one animal instead of two). Our model can be used either with predicted bounding box and segmentation map information from a text prompt or with ground truth bounding box and segmentation map. Both variations outperform the state-of-the-art in text to image and scene graph to image generation.

5.5. Limitations

Despite the high capacity of our method in generation of accurate and high quality images, it still fails to generate high quality images of complex structures such as faces. This limitation is consistent in different models and has also been existing in previous work. We believe that, by fine-tuning the model on a more constrained dataset of, for example, faces, this issue can be solved. Another issue is the high time consumption for the generation of image in the reverse sampling process, which is common in diffusion models. One limitation of our method is that the guidance process requires predicted segmentation maps and bounding box information, which can be tackled by employing off-the-shelf semantic segmentation and object detection models.

6. Conclusion

In this work, we presented a novel guidance for the sampling process in a diffusion model. Our proposed guidance enforces geometric constraints in the sampling process using the bounding box and segmentation information predicted from a scene graph. To improve the prediction of bounding box and segmentation map from the scene graph,

we encode the nodes with CLIP embedding. Our proposed guidance, as well as the employment of CLIP embeddings in the graph nodes, facilitate the generation of higher quality and more accurate images. The proposed guidance does not require any training and is applicable during the sampling process. Our method achieves better performance compared to the models trained for conditional scene graph to image generation without any training on the target datasets and also outperforms the state-of-the-art in text to image generation.

References

- [1] Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. *arXiv preprint arXiv:2006.05205*, 2020. [3](#)
- [2] Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. In *ICCV*, 2019. [3](#), [8](#)
- [3] Mehdi Astaraki, Francesca De Benetti, Yousef Yeganeh, Iuliana Toma-Dasu, Örjan Smedby, Chunliang Wang, Nassir Navab, and Thomas Wendler. Autopaint: A self-inpainting method for unsupervised anomaly detection. *arXiv preprint arXiv:2305.12358*, 2023. [2](#)
- [4] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocosuff: Thing and stuff classes in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1209–1218, 2018. [6](#), [7](#), [8](#)
- [5] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, pages 1511–1520, 2017. [2](#), [3](#)
- [6] Yuren Cong, Jinhui Yi, Bodo Rosenhahn, and Michael Ying Yang. Ssgvs: Semantic scene graph-to-video synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2554–2564, 2023. [4](#)
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [7](#)
- [8] Helisa Dhamo, Azade Farshad, Iro Laina, Nassir Navab, Gregory D Hager, Federico Tombari, and Christian Rupprecht. Semantic image manipulation using scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5213–5222, 2020. [3](#)
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [2](#), [3](#), [4](#)
- [10] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. [4](#)
- [11] Azade Farshad. Representation learning for semantic scene understanding. In *HAI 2023: Augmenting Human Intellect*, pages 445–458. IOS Press, 2023. [3](#)
- [12] Azade Farshad, Sabrina Musatian, Helisa Dhamo, and Nassir Navab. Migs: Meta image generation from scene graphs. In *BMVC*, 2021. [3](#)
- [13] Azade Farshad, Yousef Yeganeh, Helisa Dhamo, Federico Tombari, and Nassir Navab. Dispositionet: Disentangled pose and identity in semantic image manipulation. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21–24, 2022*. BMVA Press, 2022. [3](#)
- [14] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XV*, pages 89–106. Springer, 2022. [2](#)
- [15] Sarthak Garg, Helisa Dhamo, Azade Farshad, Sabrina Musatian, Nassir Navab, and Federico Tombari. Unconditional scene graph generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16362–16371, 2021. [3](#)
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [2](#)
- [17] Wei Han, Pooya Khorrani, Tom Le Paine, Prajit Ramachandran, Mohammad Babaeizadeh, Honghui Shi, Jianan Li, Shuicheng Yan, and Thomas S Huang. Seq-nms for video object detection. *arXiv preprint arXiv:1602.08465*, 2016. [3](#)
- [18] Roei Herzig, Amir Bar, Huijuan Xu, Gal Chechik, Trevor Darrell, and Amir Globerson. Learning canonical representations for scene graph to image generation. In *European Conference on Computer Vision*, pages 210–227. Springer, 2020. [3](#), [5](#), [8](#)
- [19] Roei Herzig, Moshiko Raboh, Gal Chechik, Jonathan Berant, and Amir Globerson. Mapping images to scene graphs with permutation-invariant structured prediction. In *NeurIPS*, 2018. [3](#)
- [20] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. [3](#)
- [21] Tobias Hinz, Stefan Heinrich, and Stefan Wermter. Semantic object accuracy for generative text-to-image synthesis. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1552–1565, 2020. [7](#)
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [2](#), [3](#), [4](#)
- [23] Seunghoon Hong, Xinchen Yan, Thomas E Huang, and Honglak Lee. Learning hierarchical semantic image manipulation through structured representations. In *Advances in Neural Information Processing Systems*, pages 2713–2723, 2018. [2](#)
- [24] Shi-Min Hu, Fang-Lue Zhang, Miao Wang, Ralph R Martin, and Jue Wang. Patchnet: A patch-based image representation for interactive library-driven image editing. *ACM Transactions on Graphics (TOG)*, 32(6):1–12, 2013. [2](#)
- [25] 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*, pages 1125–1134, 2017. [2](#)
- [26] 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. 3
- [27] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action genome: Actions as compositions of spatio-temporal scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10236–10247, 2020. 3
- [28] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *CVPR*, 2018. 1, 3, 4, 6, 7, 8
- [29] J. Johnson, R. Krishna, M. Stark, L. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. In *CVPR*, 2015. 3
- [30] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *arXiv preprint arXiv:2106.12423*, 2021. 2
- [31] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 2
- [32] 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. *IJCV*, 123(1):32–73, 2017. 3, 7, 8
- [33] Lei Li, Kai Fan, and Chun Yuan. Cross-modal representation learning and relation reasoning for bidirectional adaptive manipulation. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 3222–3228. International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track. 2
- [34] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Thirty-Second AAAI conference on artificial intelligence*, 2018. 3
- [35] Yitong Li, Zhe Gan, Yelong Shen, Jingjing Liu, Yu Cheng, Yuexin Wu, Lawrence Carin, David Carlson, and Jianfeng Gao. Storygan: A sequential conditional gan for story visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6329–6338, 2019. 2
- [36] Yikang Li, Tao Ma, Yeqi Bai, Nan Duan, Sining Wei, and Xiaogang Wang. Pastegan: A semi-parametric method to generate image from scene graph. *Advances in Neural Information Processing Systems*, 32, 2019. 8
- [37] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, pages 1261–1270, 2017. 3
- [38] Huan Ling, Karsten Kreis, Daiqing Li, Seung Wook Kim, Antonio Torralba, and Sanja Fidler. Editgan: High-precision semantic image editing. *arXiv preprint arXiv:2111.03186*, 2021. 2
- [39] Zelun Luo, Wanze Xie, Siddharth Kapoor, Yiyun Liang, Michael Cooper, Juan Carlos Niebles, Ehsan Adeli, and Fei-Fei Li. Moma: Multi-object multi-actor activity parsing. *Advances in Neural Information Processing Systems*, 34:17939–17955, 2021. 3
- [40] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jianjun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2, 3
- [41] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2
- [42] Gaurav Mittal, Shubham Agrawal, Anuva Agarwal, Sushant Mehta, and Tanya Marwah. Interactive image generation using scene graphs. *arXiv preprint arXiv:1905.03743*, 2019. 3
- [43] Ron Mokady, Amir Hertz, and Amit H Bermano. Clip-cap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021. 3
- [44] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Structure guided image inpainting using edge prediction. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. 2
- [45] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *NeurIPS*, pages 2171–2180, 2017. 3
- [46] 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. 1, 2, 3, 6, 7
- [47] Evangelos Ntavelis, Andrés Romero, Iason Kastanis, Luc Van Gool, and Radu Timofte. SESAME: Semantic Editing of Scenes by Adding, Manipulating or Erasing Objects. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 394–411, Cham, 2020. Springer International Publishing. 2
- [48] Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. *arXiv preprint arXiv:1905.10947*, 2019. 3
- [49] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019. 1, 2, 3
- [50] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [51] Mengshi Qi, Weijian Li, Zhengyuan Yang, Yunhong Wang, and Jiebo Luo. Attentive relational networks for mapping images to scene graphs. *arXiv preprint arXiv:1811.10696*, 2018. 3
- [52] 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 *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 3, 4, 5, 7

- [53] 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. 1
- [54] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 4, 5, 6, 7, 8
- [55] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. (available on arXiv:1505.04597 [cs.CV]). 7
- [56] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. 3
- [57] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [58] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3, 7
- [59] Mohammed Suhail, Abhay Mittal, Behjat Siddiquie, Chris Broaddus, Jayan Eledath, Gerard Medioni, and Leonid Sigal. Energy-based learning for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13936–13945, 2021. 3
- [60] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10531–10540, 2019. 1, 3
- [61] Tristan Sylvain, Pengchuan Zhang, Yoshua Bengio, R Devon Hjelm, and Shikhar Sharma. Object-centric image generation from layouts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2647–2655, 2021. 3
- [62] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3716–3725, 2020. 3
- [63] Zhu Teng, Yani Duan, Baopeng Zhang, and Jianping Fan. Global to local: Clip-1stm-based object detection from remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13, 2021. 3
- [64] 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. 2
- [65] Danfei Xu, Yuke Zhu, Christopher Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, 2017. 3
- [66] Xiaogang Xu and Ning Xu. Hierarchical image generation via transformer-based sequential patch selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. 2
- [67] 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/CVF conference on computer vision and pattern recognition*, pages 3732–3741, 2021. 3
- [68] Yousef Yeganeh, Azade Farshad, and Nassir Navab. Shape-aware masking for inpainting in medical imaging. *arXiv preprint arXiv:2207.05787*, 2022. 2
- [69] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. Semantic image inpainting with deep generative models. In *CVPR*, pages 5485–5493, 2017. 2
- [70] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [71] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, pages 5831–5840, 2018. 3
- [72] Zizhao Zhang, Yuanpu Xie, and Lin Yang. Photographic text-to-image synthesis with a hierarchically-nested adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6199–6208, 2018. 2
- [73] Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. Image generation from layout. In *CVPR*, 2019. 3
- [74] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 2
- [75] Tehseen Zia, Shahan Arif, Shakeeb Murtaza, and Mirza Ahsan Ullah. Text-to-image generation with attention based recurrent neural networks. *arXiv preprint arXiv:2001.06658*, 2020. 1