

End-to-End Optimization of Scene Layout

Andrew Luo¹ Zhoutong Zhang² Jiajun Wu³ Joshua B. Tenenbaum²

¹Carnegie Mellon University

²Massachusetts Institute of Technology

³Stanford University

Abstract

We propose an end-to-end variational generative model for scene layout synthesis conditioned on scene graphs. Unlike unconditional scene layout generation, we use scene graphs as an abstract but general representation to guide the synthesis of diverse scene layouts that satisfy relationships included in the scene graph. This gives rise to more flexible control over the synthesis process, allowing various forms of inputs such as scene layouts extracted from sentences or inferred from a single color image. Using our conditional layout synthesizer, we can generate various layouts that share the same structure of the input example. In addition to this conditional generation design, we also integrate a differentiable rendering module that enables layout refinement using only 2D projections of the scene. Given a depth and a semantics map, the differentiable rendering module enables optimizing over the synthesized layout to fit the given input in an analysis-by-synthesis fashion. Experiments suggest that our model achieves higher accuracy and diversity in conditional scene synthesis and allows exemplar-based scene generation from various input forms.

1. Introduction

Interior scene layout generation is primarily concerned with the positioning of objects that are commonly encountered indoors, such as furniture and appliances. It is of great interest due to its important role in simulated navigation, home automation, and interior design. The predominant approach is to perform unconditioned layout generation using implicit likelihood models [34]. These unconditional models can produce diverse possible layouts, but often lack the fine-grained control that allows a user to specify additional requirements or modify the scene. In contrast to the unconditional models, conditional layout generation uses various types of inputs, such as activity traces, partial layouts, or text-based descriptions, enabling more flexible synthesis.

In this work, we use the 3D scene graph representation as a high-level abstraction, with the graph not only encoding object attributes and identities, but also 3D spatial relationships. Not only does this design enable more control over

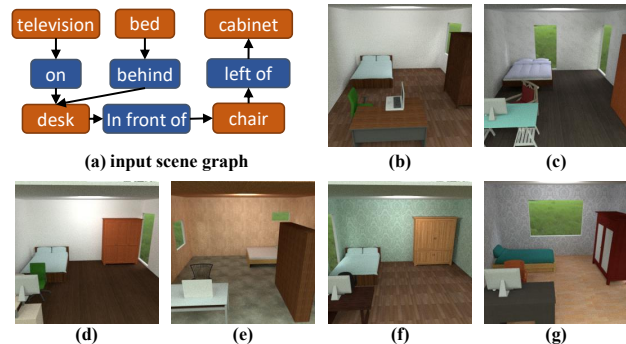


Figure 1: Conditional scene synthesis. (a) The input is a scene graph describing object relationships. (b)–(g) Diverse layouts synthesized conforming to the input scene graph.

the generated content, as users can directly manipulate the input scene graph, it also serves as a general intermediate representation between various modalities of scene descriptions, such as text-based descriptions and exemplar images.

Many previous methods have used scene graphs as an intermediate representation for downstream vision tasks such as image synthesis, where they mostly formulate the problem in a deterministic manner. In contrast, our model respects the stochastic nature of the actual scene layout conditioned on the abstract description of a scene graph. The model we introduce, named 3D Scene Layout Network (3D-SLN), is a general framework for scene layout synthesis from scene graphs. 3D-SLN combines a variational autoencoder with a graph convolutional network to generate diverse and plausible layouts that are described by the relationships given in the 3D scene graph, as shown in Figure 1.

We further demonstrate how a differentiable renderer can be used to refine the generated layout using a single 2.5D sketch (depth/surface normal) and the semantic map of a 3D scene. In addition, our framework can be applied to perform exemplar-based layout generation, where we synthesize different scene layouts that share the same scene graph extracted from text or inferred from a reference image.

In summary, our contributions are threefold. First, we introduce 3D-SLN, a conditional variational autoencoder-based network that generates diverse and realistic scene layouts conditioned on a scene graph. Second, we demonstrate our model can be fine-tuned to generate 3D scene layouts

that match the given depth and semantic information. Finally, we showed that our model can be useful for several applications, such as exemplar-based layout synthesis and scene graph-based image synthesis.

2. Related Work

Our method is related to the multiple areas in computer vision and graphics, including scene graph representations, scene synthesis, and differentiable rendering.

Scene graphs. Scenes can be represented as scene graphs—directed graphs whose nodes are objects and edges are relationships between objects. Scene graphs have found wide applications such as image retrieval [12] and image captioning [1]. There have also been attempts to generate scene graphs from text [25], images [32, 21, 17], and partially completed graphs [30]. In this paper, we use scene graphs to guide our synthesis of 3D indoor scene layout.

Scene synthesis. In computer graphics, there has been extensive research on indoor scene synthesis. Much of this work is associated with producing plausible layouts constrained by a statistical prior learned from data. Typical techniques used include probabilistic models [4, 5, 34], stochastic grammar [23], and recently, convolutional networks [31, 24].

Many of these approaches build upon the recent advancement of large-scale scene repositories [28, 20, 36] and indoor scene rendering methods [35, 16, 10]. These methods typically focus on modeling the possible distribution of objects given a particular room type (e.g., bedrooms). Some recent papers [2] have studied 3D scene layout generation directly from text. Some concurrent work also uses relational graphs for modelling scenes, however our approach is capable of single-pass scene synthesis in a fully differentiable manner, where as [30] tries to generate the scene graph in an autoregressive fashion, and has a non-differentiable sampling step.

Differentiable rendering. Traditional graphics engines do not produce usable gradients for optimization purposes. A variety of renderers that allow for end-to-end differentiation have been proposed [19, 15, 13, 18]. These differentiable renderers have been used for texture optimization [13], face inference [29], single image mesh reconstruction [6], and scene parsing [9]. Additionally non-differentiable rendering has been used with approximated gradients [14] for instance level 3D construction. We utilize the neural mesh renderer [13], which allows us to manipulate the layout and rotation of individual objects given depth and semantic maps as reference.

3. Methods

We propose **3D Scene Layout Networks** (3D-SLN), a conditional variational autoencoder network tailored to op-

erate on scene graphs. We first use a graph convolution network to model the posterior distribution of the layout conditioned on the given scene graph. Then, we generate diverse scene layouts, which includes each object’s size, location and rotation, by sampling from the prior distribution.

3.1. Scene Layout Generator

While previous methods generate 2D bounding boxes from a scene graph [11] or text descriptions [7], our model generates 3D scene layouts, consisting of the 3D bounding box and rotation along the vertical axis for each object. In addition, we augment traditional 2D scene graphs to *3D scene graphs*, encoding object relationships in 3D space.

Specifically, we define the X and Y axis to span the plane consisting of the room’s floor, and an up-direction Z for objects above the floor. Under such definition, the relationship ‘left of’ constrains the X and Y coordinates between pairs of objects, while the relationship ‘on’ constrains the Z coordinate between them. Each node in the scene graph will not only define what *type* of object it is, it may optionally define object’s attributes regarding the object height (*tall*, *short*) but also volume (*large*, *small*). The scene graph y is represented by a set of relationship triplets, where each triplet is in the form of (o_i, p, o_j) . Here p denotes the spatial relationship and o_i denotes the i -th object’s type and attributes.

In order to operate on the input graph and to generate multiple scenes from the same input, we propose a novel framework, named 3D Scene Layout Network (3D-SLN), combining a graph convolution network (GCN) [11] with a conditional variational autoencoder (cVAE) [26]. The architecture is shown in Figure 2. During training, the encoder is tasked to generate the posterior distribution of a given scene layout conditioned on the corresponding scene graph. The encoder therefore takes a scene graph and an exemplar layout as input, and outputs the posterior distribution for each object, represented by the mean and log-variance of a diagonal Gaussian distribution. A latent vector is then sampled from the Gaussian for each object. The decoder then takes the sampled latent vectors and the scene graph as input and generates a scene layout, represented by the 3D bounding box and its rotation for each object.

We define x to be the input scene graph, y to be an exemplar layout, \hat{y} to be the generated layout, and θ_e, θ_d to be the weights of the encoder P_{θ_e} and decoder Q_{θ_d} of 3D-SLN, respectively. Each element in y_i in layout y is defined by a 7-tuple, representing the bounding box and the rotation of each object i :

$$y_i = (\min_{X_i}, \min_{Y_i}, \min_{Z_i}, \max_{X_i}, \max_{Y_i}, \max_{Z_i}, \omega_i), \quad (1)$$

where $\min_{X_i}, \min_{Y_i}, \min_{Z_i}, \max_{X_i}, \max_{Y_i}, \max_{Z_i}$ denotes the 3D bounding box coordinates, and ω_i denotes the rotation around the Z axis.

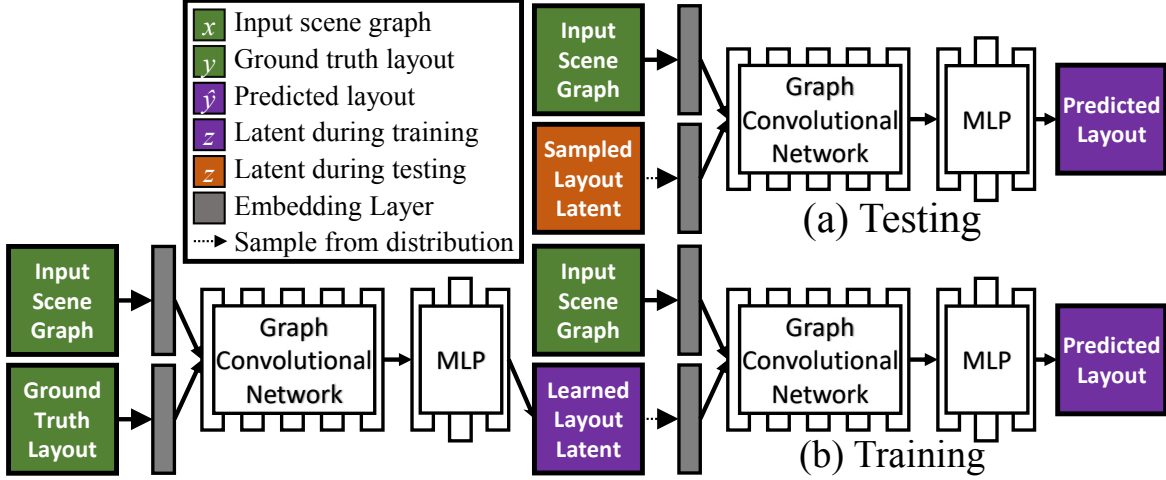


Figure 2: Network architecture of the scene layout generator. (a) At test time, a latent code is sampled from a learned distribution and is sent to a decoder with the scene graph to generate scene layout. (b) During training, an encoder converts the ground truth scene layout and the scene graph into a distribution, from which the latent code is sampled and decoded.

To train the graph-based conditional variational autoencoder described above, we optimize

$$\mathcal{L}(x, y; \theta) = \lambda D_{KL}(P_{\theta_e}(z|x, y)|p(z|x)) + L_{\text{layout}}(Q_{\theta_d}(x, z), y), \quad (2)$$

where λ is the weight of the Kullback-Liebler divergence, $p(z|x)$ is the prior distribution of the latent vectors, which is modeled as diagonal Gaussian distribution, and $\mathcal{L}_{\text{layout}}$ is the loss function defined over layouts. $\mathcal{L}_{\text{layout}}$ consists two parts: $\mathcal{L}_{\text{position}}$ and $\mathcal{L}_{\text{rotation}}$. $\mathcal{L}_{\text{position}}$ is defined as the $L1$ loss over each object’s bounding box parameters. For the rotation, we first discretize the range of the rotation angles to 24 bins, and define $\mathcal{L}_{\text{rotation}}$ as the negative log-likelihood loss between the discretized angles for all the objects. We apply learned embedding layers to process object type, rotation, attribute, and relations; and a linear layer to process the bounding box. The rotation and box embeddings are used for the encoder only. The object type, bounding box, rotation, attribute, and relational embeddings have dimensions [48, 48, 16, 16, 128]. Embeddings are computed separately for the encoder and decoder. The intermediate latent representation is a 64 dimensional vector for each object. Both the encoder and decoder contain five graph convolution layers with average pooling and batch normalization.

At test time, we use the decoder to sample scene layouts from scene graphs. We first sample latent vectors from the prior distribution, modeled as a Gaussian distribution. Given the sampled latent vectors and the 3D scene graph, the decoder then generates multiple possible layouts.

3.2. Gradient-Based Layout Refinement

Here we consider the case where we would like to generate a layout that fits a target layout, represented as a depth image D with corresponding semantics S . Using our scene graph and an inferred layout, we first retrieve object

meshes from the SUNCG dataset to construct a complete scene model. Specifically, for each object i in the generated layout, we retrieve its 3D model M_i from the SUNCG dataset by finding the models with the most similar bounding box parameters within its class.

After instantiating a full 3D scene model, we then utilize a differentiable renderer [13] R to render the corresponding semantic image \tilde{S} and the depth image \tilde{D} from the scene. The rendered images are then used to compare with the target semantics and depth image. This provides the gradients to update both the sampled latent vectors and the weights of the decoder, making the generated 3D layout to be consistent with the input semantics and depth.

Specifically, we note the entire generate process as

$$\tilde{S} = R_S(\hat{y}_1, M_1, \hat{y}_2, M_2, \dots, \hat{y}_N, M_N), \quad (3)$$

$$\tilde{D} = R_D(\hat{y}_1, M_1, \hat{y}_2, M_2, \dots, \hat{y}_N, M_N), \quad (4)$$

$$\hat{y} = Q_{\theta_d}(x, z), \quad (5)$$

where R_S denotes the rendered semantic map, R_D denotes the rendered depth map, \hat{y} denotes the generated layout, and N denotes the number of objects in the scene graph x .

To optimize the decoder and the latent vectors, we aim to calculate the gradient of \tilde{D} and \tilde{S} with respect to z and θ_d . Note that since the output of ω_i in y_i is discretized into 24 bins, i.e. $\omega_i = \text{bin}_k$, where $k = \text{argmax}_k(\{\omega_{ik} | k \in [1, 2, \dots, 24]\})$, the entire rendering process is not differentiable due to the argmax operator. To overcome this, we use a softmax-based approximation to compute the angle for the i -th object, defined as (assuming a k -way prediction)

$$x_k = \sum_{n=1}^k 1, \quad \omega_{ik} = \frac{e^{\omega_{ik}}}{\sum_k e^{\omega_{ik}}},$$

$$\omega_i = \sum (\omega_{ik} \times x_k) - 1.0.$$




Figure 3: We can fine-tune object positions, sizes, and rotations by computing the difference in estimated and ground-truth 2.5D sketches and back-propagating the gradients.

By doing so, we can take gradients of the loss function between the rendered images \tilde{D}, \tilde{S} and the target images D, S . A simple loss function can be defined as $\mathcal{L}_{\text{total}} = \mathcal{L}_2(D, \tilde{D}) + \mathcal{L}_{\text{cross-entropy}}(S, \tilde{S})$. This however, leads to highly unstable gradients in practice. To stabilize the layout refinement process, we calculate the loss between the depth images in a per-class manner. More specifically, for each class c , we calculate its class-conditioned depth map D_c as

$$D_c[S == c] = D \otimes S[S == c], \quad (6)$$

$$D_c[S \neq c] = \text{mean}(D_c[S == c]). \quad (7)$$

That is, we keep the depth values that lies within a particular semantic class c , and fill the rest of the values with the mean depth of this class. Therefore, we rewrite the depth loss as $\mathcal{L}_{\text{depth}} = \sum_c \mathcal{L}_2(\tilde{D}_c, D_c)$.

This can be understood as a class-wise isolation of the depth gradient, and can prevent spurious optima during the layout refinement process. We also impose a soft constraint on the change in object sizes with an additional \mathcal{L}_1 penalty on the size of each object during optimization at a given time step \tilde{s}_t compared to the original size \tilde{s}_0 . To facilitate the optimization process when target and proposed layouts may not have an exact match in shape, we apply multi-scale average pooling on both the candidate and target. The total refinement loss with respect to a target depth and semantic is $\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{size}} + \beta \mathcal{L}_{\text{depth}} + \gamma \mathcal{L}_{\text{sem}}$, where α, β , and γ are hyper-parameters, and $\mathcal{L}_{\text{size}} = |\tilde{s}_t, \tilde{s}_0|_1$, $\mathcal{L}_{\text{depth}} = \sum_c |\tilde{D}_c, D_c|_2^2$, and $\mathcal{L}_{\text{sem}} = \mathcal{L}_{\text{cross-entropy}}(\tilde{S}, S)$.

This allows us to obtain meaningful gradients from an exemplar 2D projection of a scene to optimize \hat{y} by calculating a gradient with respect to the sampled latent vector z and the decoder. The framework for our gradient based layout refinement is shown in Figure 3.

4. Experiments

In this section, we compare our approach with state-of-the-art scene layout synthesis algorithms to demonstrate the quality and diversity of our synthesized scenes. Additional ablation studies show that each component in our model

contributes to its performance. Finally, we demonstrate our algorithm also enables exemplar based layout synthesis and refinement.

4.1. Setup

For layout generation, we learn from bedroom layouts in SUNCG [28]. The training dataset consists of 53,860 bedroom scenes with 13.15 objects in each scene on average. During training, we use synthetic scene graphs sampled from the ground truth scene layout, which can avoid human labeling and also serve as data augmentation. At test time, we can either use human-created scene graphs or sample scene graphs from the validation set as model input.

The cVAE 3D graph network is trained on a total of 600k batches, which takes around 64 hours with a single Titan Xp. For each batch we sample 128 scene graphs. A learning rate of 10^{-4} is used with the Adam optimizer. We use three losses with the following weights: $\lambda_{\text{pos}} = 1, \lambda_{\text{rot}} = 1$, and $\lambda_{\text{KL}} = 0.1$.

4.2. Scene Layout Synthesis

We first evaluate our 3D-SLN on scene layout synthesis from a scene graph. We sample 10 layouts per scene, and calculate the average standard deviation for object size, position, and rotation. Layout synthesis alone during testing is highly efficient, taking about 70ms on a GPU for a batch of 128 graphs.

Baselines. We compare our model with the state-of-the-art scene layout synthesis algorithm, DeepSynth [31]. Following Qi et al. [23], we also include two additional baselines: Random, where every object is distributed randomly in a room; and Perturbed, where we perturb object positions against their ground truth positions with a variance of 0.1 on their spatial location (all locations are normalized to $[0, 1]$) and with a standard deviation of 3 bins on their rotation (approximately 0.785 radians).

Metrics. We analyze both the accuracy and diversity of the results through three metrics:

- Scene graph accuracy measures the percentage of scene graph relationships a given layout respects, and is a metric that measures input-output alignment.
- \mathcal{L}_1 loss of the proposed and ground truth bounding boxes. It should be noted that since the goal is the generate multiple plausible layouts, \mathcal{L}_1 is not necessarily a meaningful metric and is provided for reference only.
- The standard deviation of the size, position, and rotation of objects in predicted scene layouts. Because DeepSynth produces layouts in an autoregressive fashion, a particular object of interest (e.g., a bed) might appear at various steps across multiple trials. Due to the lack of correspondence, we can only compute the standard deviation for all objects within each semantic

Model	Scene Graph Acc. (%)	\mathcal{L}_1 box loss	STD (size)	STD (position)	STD (rotation)
Random Layout	57.1	0.317	0.000	0.244	6.48
Perturbed Layout	82.6	0.080	0.000	0.100	3.00
DeepSynth	N/A	N/A	N/A	0.129	2.27
GCN	86.3	0.111	0.000	0.000	0.00
GCN+noise	86.9	0.109	0.001	0.002	0.18
3D-SLN (Ours)	94.3	0.148	0.026	0.078	4.77

Table 1: Quantitative results on scene layout generation. We use scene graph accuracy and \mathcal{L}_1 bounding box loss to evaluate the accuracy of generated scene layouts. Standard deviation of boxes and angles are used to measure the diversity of scene layouts. In the above evaluation, bounding boxes are normalized in the range $[0, 1]$, while angles are represented as integers ranging from 0 to 23. DeepSynth [31] is used as a baseline.

category, and average the standard deviations across all categories. For our model and the random/perturbed layout models, we calculate standard deviations for each object of interest and compute their mean.

Results. Table 1 shows that our model has the highest scene graph accuracy and diversity. This indicates that our model has successfully learned to position objects according to a distribution rather than approximating a fixed location. While DeepSynth has a higher standard deviation in object positions, it has a lower standard deviation in rotations. It also does not allow fine-grained control of the synthesis process. Although ‘Perturbed Layout’ has the lowest \mathcal{L}_1 loss, it has a significantly lower scene graph accuracy.

4.3. Ablation Study

We perform ablation studies on our scene layout generation network. By utilizing a graph convolution network combined with a VAE, it is able to generate multiple plausible layouts from a given scene graph.

Baselines. Following Johnson et al. [11], we run an ablated version of our network (denoted as GCN) that consists of a single graph convolution network followed by an MLP to predict the layout conditioned on a scene graph. This baseline is deterministic. We also propose a different method, GCN+noise, which samples noises from $\mathcal{N}(0, 1)$ to perturb the layout of the GCN baseline.

Results. Table 1 shows that our full model (3D-SLN) achieves the highest scene graph accuracy, indicating that most of the synthesized scene layouts indeed follow the input scene graph. Our full model also achieves the highest diversity, as measured in the standard deviation in the size, position, and rotation of objects in the synthesized scene.

4.4. End-to-End Layout Refinement

We now demonstrate that our model can be guided by 2.5D sketches and semantic maps when synthesizing 3D

Metrics	Pre-Finetune	Post-Finetune	Improve (%)
3D IoU	0.2353	0.3035	28.9
Depth MSE	0.0525	0.0480	8.64
Semantic CE	2.9471	2.8504	3.28

Table 2: Quantitative results on finetuning with 2.5D sketches of a target layout. We measure the Intersection-over-Union (IoU) of the 3D bounding boxes, class-specific mean squared error (MSE) of the depth maps, as well as the cross-entropy loss (CE) on semantic maps.

scene layout from a scene graph. We perform optimization over 150 randomly selected scene graphs. Here, for sampling different scene layouts from our stochastic model, we use the latent sampled from the ground truth bounding boxes of a given scene. To prevent cases when wall or object occlusion negatively impact optimization performance, we take six attempts at a given scene graph, and select the best. The analysis-by-synthesis process requires a forward (rendering) pass to produce depth and semantic maps, then a backwards pass to produce gradients. We optimize for 60 steps, taking three minutes for each scene on average.

Metrics. We use three metrics for this problem: The first is done in 3D, and captures the Intersection-over-Union (3D IoU) of objects and their target after all transformations (rotations and translations) are applied. The latter two are performed in the 2D projection: we calculate a mean-squared error (MSE) on the predicted and ground-truth depth maps; we also calculate the cross-entropy (CE) loss of our current proposed layout against the target layout.

Results. The results can be seen in Table 2. We also perform qualitative evaluations on layouts in Figure 4. As expected, the initial proposed layout shares the same scene graph with the target layout, the action location of the objects can be different from the target, because our layout synthesis is conditioned on scene graph only. After optimization, we are able to fit the target layout reasonably well. Multiple views of synthesized scenes are shown in Figure 5.

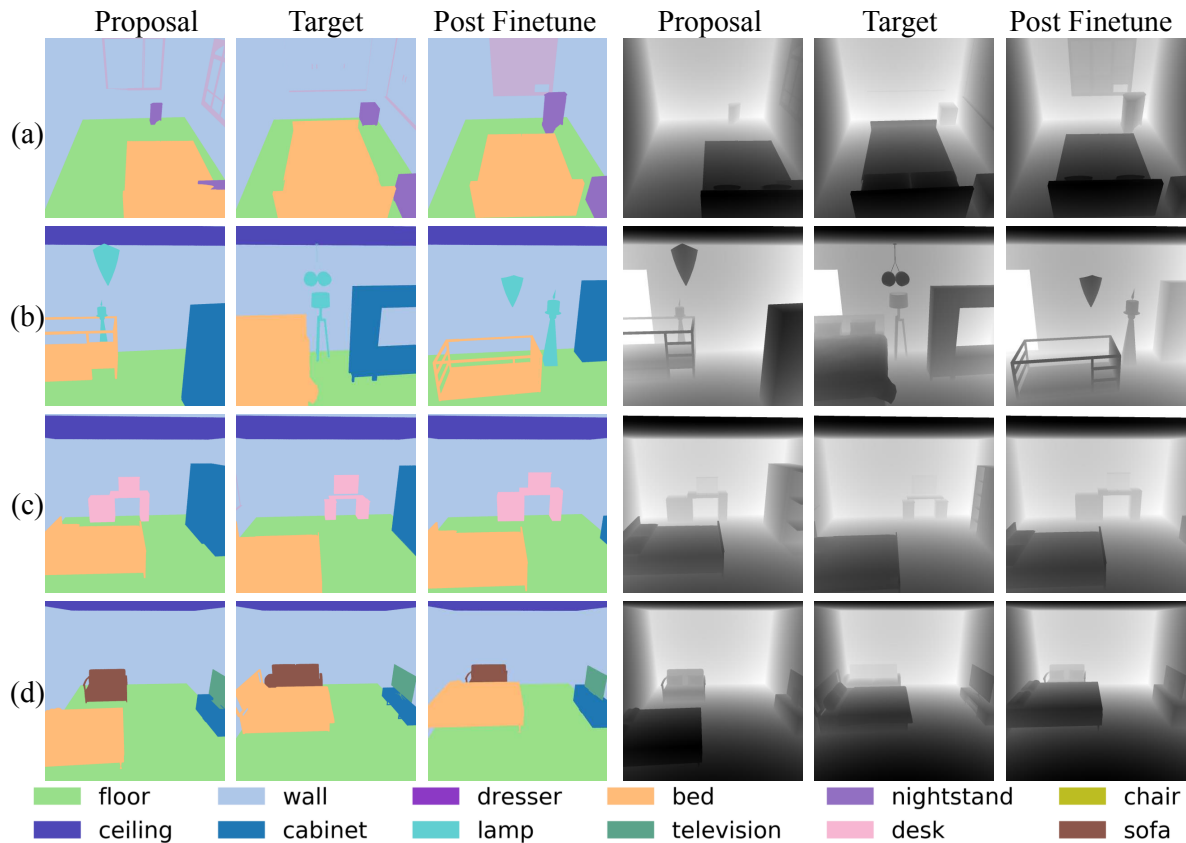


Figure 4: Each row shows a test case for exemplar-based layout fine-tuning. The left three columns represent the 2D semantic map of the initially proposed layout, the ground truth target, and the semantic map after layout optimization. The right three columns represent the 2.5D depth map similarly. **(a)** After fine-tuning, the bed has moved to the center as in the target, and the two night stands become more prominent; **(b)** The lamp (in light blue) has moved downwards; **(c)** The desk has moved to the right after optimization and the bed has become closer; **(d)** The bed has moved closer to the sofa.



Figure 5: Multiple views of synthesized scenes.

Analyzing the latent space. We examine the latent representation of the scene layouts. In Figure 6, we demonstrate that the layout of objects can smoothly change as we interpolate two random latent vectors. In Figure 8, we demonstrate the effect of manipulating dimension 11 and 62 of the latent vector for the bed object.

User study. We randomly sample 300 scene graphs and, for each, generate five layouts using the GCN+noise model




Figure 6: Top down visualization of a room as we linearly interpolate between the latent vector representing layouts for the same scene graph.

and five using our 3D-SLN model, respectively. We then present the layouts in a top down view along with the scene graph in sentences to subjects on Amazon Mechanical Turk, asking them which set of layouts is more diverse. Each subject is shown twelve scene graphs. 78.9% of responses suggested layouts generated by 3D-SLN be more diverse.




Figure 7: Qualitative results for conditional image synthesis. Top: input scene graph; Middle: images generated by our model; Bottom: images generated by Johnson et al. [11], which does not incorporate 3D information. Our model generates better results via its understanding of 3D scene layout.




Figure 8: Top down visualization of a room as we manipulate individual dimensions of latent vector for bed object (orange). For dimension 11, note bed elongation.




Figure 9: Top down visualization of failure cases (left of dotted line) compared to a good layout. All layouts are synthesized from the same scene graph.

Failure cases. As shown in Figure 9, scene synthesis performance decreases when a graph contains too many objects that might overlap. As part of future work, this could be improved during training by adding an adversarial loss, or during inference by rejecting implausible layouts with the use of physical simulation as in [3], or by performing simple collision detection on generated layouts.

5. Applications

Our scene graph-based layout synthesis algorithm enables many downstream applications. In this section,

we show results on scene graph-based image synthesis, sentence-based scene layout synthesis, and exemplar-based scene layout synthesis.

5.1. Scene Graph-Based Image Synthesis


As our model produces not only 3D scene layouts, but also 2.5D sketches and semantic maps, we train an image translation network, SPADE [22], that takes in depth and semantic maps and synthesizes an RGB image. Training data for the SPADE model, including RGB, depth, and semantic maps, are all taken from the Structured3D dataset [36], and are randomly cropped to 256×256 after we resize the longest edge to 480 pixels. The training dataset consists of 82,838 images total. We compare our model with the state-of-the-art, scene graph-to-image model [11].

Results are shown in Figure 7. The images generated by our model are sharp and photo-realistic, with complex lighting. Meanwhile, the baseline [11] can only generate blurry images, where sometimes the objects are hardly recognizable and fail to preserve the 3D structure.

5.2. Sentence-Based Scene Layout Synthesis

Conventional text-to-image synthesis methods use a text encoder to convert an input sentence into a latent code, which is then fed into a conditional GAN to generate an image. However existing methods only work when the input sentence has only one or a few objects. The task becomes more challenging when input text consists of multiple objects and contains complex relationships. We compare our approach against AttnGAN [33], the state-of-the-art image synthesis algorithm that takes in sentences as input.

Qualitative results are shown in Figure 10. As AttnGAN suffers from deterioration when there are too many objects,



a tall television is on a short wooden cabinet, a short fabric bed is in front of a small wooden desk, the bed is on the right of the television, the television is in front of a tall wooden chair, the chair is in front of the desk

a tall wooden bed is behind touching a short fabric chair, the chair is front touching the bed

a short fabric bed is in front of a large fabric desk, the desk is behind a large wooden cabinet, the cabinet is left touching the bed

the second tall fabric bed is on the left of a small television, a large wooden desk is behind the first tall fabric bed, a tall wooden cabinet is on the right of the second bed, the cabinet is on the right of the first bed, the cabinet is in front of the desk

Figure 10: Comparison with AttnGAN [33]. The images on the top row are generated by our model, while the images on the bottom row are generated by AttnGAN [33]. In comparison, our model generates higher-quality scenes.

we have constricted each individual description to have at most five sentences. Our 3D-SLN generates more realistic images compared with AttnGAN.

5.3. Exemplar-Based Scene Layout Synthesis

Our model can also be used to reconstruct and create new layouts based on an example image. We use *Cooperative Scene Parsing* [8] to predict object classes and 3D bounding boxes from an image. For our purposes, we test on bedroom images sampled from the SUN RGB-D dataset [27]. After extracting 3D bounding boxes for each object, we infer a 3D scene graph with the same object classes and relationships that our model is trained on. This scene graph is sent to our model to generate layouts that observe the relational constraints present in the scene graph.

We present some qualitative results in Figure 11. Our model is not only capable of recovering the original layout in the example image, but it can also create new layouts according to the scene graph (notice the different locations and rotations of the bed and nightstand).

6. Conclusion

In this paper, we have introduced a novel, stochastic scene layout synthesis algorithm conditioned on scene graphs. Using scene graphs as input allows flexible and controllable scene generation. Experiments demonstrate that our model generates more accurate and diverse 3D scene

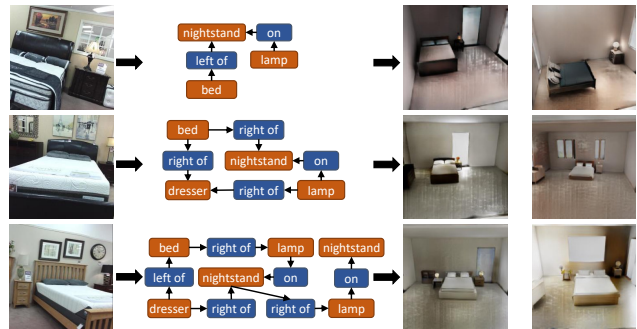


Figure 11: On the left most column are images of the class ‘bedroom’ from the SUN RGB-D dataset [27]. 3D bounding boxes are calculated per object, and are fed to a rule-based parser, which generates the relationships and creates a scene graph. The scene graph is then fed to our 3D-SLN to generate diverse layouts. Final images are rendered with SPADE [22].

layouts compared with baselines. Our model can also be integrated with a differentiable renderer to refine 3D layout conditioned on a single example. Our model finds wide applications in downstream scene layout and image synthesis tasks. We hope our work will inspire future work in conditional scene generation.

Acknowledgements. This work is supported by NSF #1447476, ONR MURI N00014-16-1-2007, and NIH T90-DA022762.

References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, pages 382–398. Springer, 2016. 2
- [2] Angel Chang, Will Monroe, Manolis Savva, Christopher Potts, and Christopher D Manning. Text to 3d scene generation with rich lexical grounding. *arXiv preprint arXiv:1505.06289*, 2015. 2
- [3] Yilun Du, Zhijian Liu, Hector Basevi, Ales Leonardis, Bill Freeman, Josh Tenenbaum, and Jiajun Wu. Learning to exploit stability for 3d scene parsing. In *Advances in Neural Information Processing Systems*, pages 1726–1736, 2018. 7
- [4] Matthew Fisher, Daniel Ritchie, Manolis Savva, Thomas Funkhouser, and Pat Hanrahan. Example-based synthesis of 3d object arrangements. *ACM Transactions on Graphics (TOG)*, 31(6):135, 2012. 2
- [5] Matthew Fisher, Manolis Savva, Yangyan Li, Pat Hanrahan, and Matthias Nießner. Activity-centric scene synthesis for functional 3d scene modeling. *ACM Transactions on Graphics (TOG)*, 34(6):179, 2015. 2
- [6] Paul Henderson and Vittorio Ferrari. Learning to generate and reconstruct 3d meshes with only 2d supervision. *arXiv preprint arXiv:1807.09259*, 2018. 2
- [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. 2
- [8] Siyuan Huang, Siyuan Qi, Yinxue Xiao, Yixin Zhu, Ying Nian Wu, and Song-Chun Zhu. Cooperative holistic scene understanding: Unifying 3d object, layout, and camera pose estimation. In *Advances in Neural Information Processing Systems*, pages 207–218, 2018. 8
- [9] Siyuan Huang, Siyuan Qi, Yixin Zhu, Yinxue Xiao, Yuanlu Xu, and Song-Chun Zhu. Holistic 3d scene parsing and reconstruction from a single rgb image. In *ECCV*, 2018. 2
- [10] Chenfanfu Jiang, Siyuan Qi, Yixin Zhu, Siyuan Huang, Jenny Lin, Lap-Fai Yu, Demetri Terzopoulos, and Song-Chun Zhu. Configurable 3d scene synthesis and 2d image rendering with per-pixel ground truth using stochastic grammars. *International Journal of Computer Vision*, 126(9):920–941, 2018. 2
- [11] 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. 2, 5, 7
- [12] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015. 2
- [13] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3907–3916, 2018. 2, 3
- [14] Abhijit Kundu, Yin Li, and James M Rehg. 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3559–3568, 2018. 2
- [15] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. In *SIGGRAPH Asia 2018 Technical Papers*, page 222. ACM, 2018. 2
- [16] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang, and Stefan Leutenegger. Interiornet: Mega-scale multi-sensor photo-realistic indoor scenes dataset. *arXiv preprint arXiv:1809.00716*, 2018. 2
- [17] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1261–1270, 2017. 2
- [18] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. *arXiv preprint arXiv:1904.01786*, 2019. 2
- [19] Matthew M Loper and Michael J Black. Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision*, pages 154–169. Springer, 2014. 2
- [20] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J Davison. Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2678–2687, 2017. 2
- [21] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *Advances in neural information processing systems*, pages 2171–2180, 2017. 2
- [22] 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. 7, 8
- [23] Siyuan Qi, Yixin Zhu, Siyuan Huang, Chenfanfu Jiang, and Song-Chun Zhu. Human-centric indoor scene synthesis using stochastic grammar. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5899–5908, 2018. 2, 4
- [24] Daniel Ritchie, Kai Wang, and Yu-an Lin. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6182–6190, 2019. 2
- [25] Sebastian Schuster, Ranjay Krishna, Angel Chang, Li Fei-Fei, and Christopher D Manning. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *Proceedings of the fourth workshop on vision and language*, pages 70–80, 2015. 2
- [26] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pages 3483–3491, 2015. 2

- [27] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015. 8
- [28] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017. 2, 4
- [29] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2549–2559, 2018. 2
- [30] Kai Wang, Yu-An Lin, Ben Weissmann, Manolis Savva, Angel X Chang, and Daniel Ritchie. Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)*, 38(4):132, 2019. 2
- [31] Kai Wang, Manolis Savva, Angel X Chang, and Daniel Ritchie. Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)*, 37(4):70, 2018. 2, 4, 5
- [32] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5410–5419, 2017. 2
- [33] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1316–1324, 2018. 7, 8
- [34] Lap-Fai Yu, Sai Kit Yeung, Chi-Keung Tang, Demetri Terzopoulos, Tony F Chan, and Stanley Osher. Make it home: automatic optimization of furniture arrangement. *ACM Trans. Graph.*, 30(4):86, 2011. 1, 2
- [35] Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, and Thomas Funkhouser. Physically-based rendering for indoor scene understanding using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5287–5295, 2017. 2
- [36] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. *arXiv preprint arXiv:1908.00222*, 2019. 2, 7