

# Unconditional Scene Graph Generation

Sarthak Garg<sup>1,\*</sup>      Helisa Dhama<sup>1,\*</sup>      Azade Farshad<sup>1</sup>  
 Sabrina Musatian<sup>1</sup>      Nassir Navab<sup>1</sup>      Federico Tombari<sup>1,2</sup>  
<sup>1</sup> Technische Universität München    <sup>2</sup> Google

## Abstract

Despite recent advancements in single-domain or single-object image generation, it is still challenging to generate complex scenes containing diverse, multiple objects and their interactions. Scene graphs, composed of nodes as objects and directed-edges as relationships among objects, offer an alternative representation of a scene that is more semantically grounded than images. We hypothesize that a generative model for scene graphs might be able to learn the underlying semantic structure of real-world scenes more effectively than images, and hence, generate realistic novel scenes in the form of scene graphs. In this work, we explore a new task for the unconditional generation of semantic scene graphs. We develop a deep auto-regressive model called SceneGraphGen which can directly learn the probability distribution over labelled and directed graphs using a hierarchical recurrent architecture. The model takes a seed object as input and generates a scene graph in a sequence of steps, each step generating an object node, followed by a sequence of relationship edges connecting to the previous nodes. We show that the scene graphs generated by SceneGraphGen are diverse and follow the semantic patterns of real-world scenes. Additionally, we demonstrate the application of the generated graphs in image synthesis, anomaly detection and scene graph completion.

## 1. Introduction

Scene graphs encompass a representation that describes a scene, where nodes are categorical object instances and the edges describe categorical relationships between them. This representation allows for an extended high-level understanding of the scenes, which goes beyond object-level reasoning. The computer vision community has explored various approaches for scene graph generation from images [38, 27] as well as tasks in which this representation has proven to be suitable, such as image retrieval [18] and

\*The first two authors contributed equally to this work

<https://SceneGraphGen.github.io/>

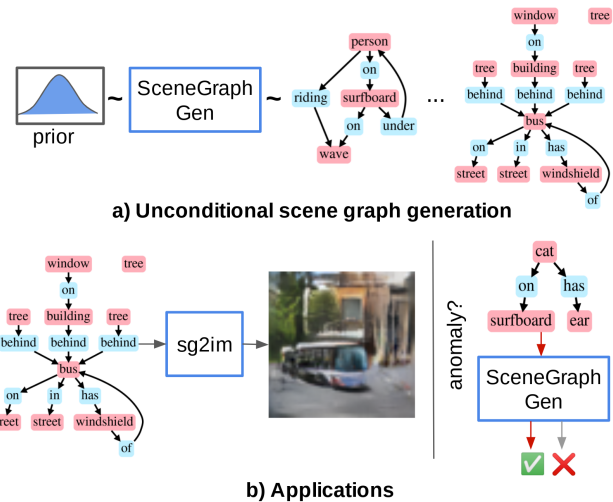


Figure 1. **Method overview.** **a)** SceneGraphGen generates scene graphs unconditionally from a randomly sampled seed object. **b)** Applications in *Left*: translation of a generated scene graph to an image, using an off-the-shelf graph-to-image network *Right*: detection of out-of-distribution samples

VQA [11]. A scene graph allows for modular, specified, and high-level semantic control on the image components, which makes it a good representation for semantic-driven image generation [17] and manipulation [7].

A less explored field has been *unconditional generation of scene graphs*, *i.e.* generating a scene graph without any input image, but instead, from a *random* input. Such scene graph modelling can aid learning patterns from real scenes, such as object co-occurrence, relative placements and interactions. In this paper, scene graph generation is explored under the light of generative models, with the aim to generate novel and realistic instances of scene graphs *unconditionally*. Recent work explored generation of relational graphs [35] or probabilistic grammar [19], designed for a particular domain. To the best of our knowledge, we are the first to investigate the use of a generative model for generating semantic, language-based scene graphs [18, 21].

As scene graphs describe scenes, it is possible to translate the generated graph to another domain, using state-of-the-art models specialized, for instance, in the graph-to-image task [17]. In the context of unconditional image generation, recent works demonstrate impressive results mostly in object-centric images that contain one main subject, or uni-modal distributions, such as datasets of faces or cars. On the other hand, complex and diverse scenes which contain multiple objects are more difficult to capture by these models. We show in our experiments that modelling unconditional image scene generation through scene graphs instead leads to more distinguishable object instances, as it enables understanding of complex and often abstract semantic concepts such as objects, their interactions, and attributes. Additionally, such a generative model can detect out-of-distribution scene graphs and complete partial scene graphs.

Recently, deep generative models have been proposed for graph data [41, 12, 9, 30, 3], which aim to synthesize realistic graphs of a certain domain while capturing graph patterns, such as degree distribution and clustering. However, each model comes with caveats making it unsuitable for certain applications. The size of scene graphs often varies significantly, the object and relationship categories are inherently unbalanced, and the edges are directed. For this purpose, we develop a specialized model called SceneGraphGen (Figure 1). The overall auto-regressive structure is inspired from GraphRNN [41], as it accommodates varying graph sizes, unlike [9, 30, 4]. Specifically, the model is adapted to consume categories for nodes and edges, as well as to support the direction of the edges. In this auto-regressive formulation, the scene graphs are represented as a sequence of sequences. The history of the sequence is carried in a hidden state using a Gated Recurrent Unit (GRU) [5], which is used to generate categorical distribution over the nodes and edges at each step, from which the node and edge categories can be sampled. The nodes are generated using a multi-layer perceptron (MLP) and the edges are generated sequentially using a GRU.

Since unconditional scene graph generation is a new task, metrics to evaluate the quality of the generated graphs have not been proposed yet. Thus, following [41] we leverage a Maximum Mean Discrepancy (MMD) metric, adapted with a random-walk graph kernel and a node kernel, which are appropriate for the scene graph structure. We verify the validity of these kernels using sets of corrupted datasets.

Our contributions can be summarized as follows:

- We introduce SceneGraphGen to tackle the unexplored task of unconditional semantic scene graph generation. We adopt a graph auto-regressive model to enable processing the scene graphs structure.
- We demonstrate the use of the learned scene graph

model in three applications, namely image generation, anomaly detection, and scene graph completion.

- We propose and verify an MMD metric to evaluate the generated scene graphs, which operates on the node and graph level.

We evaluate our model on Visual Genome [21] and show that it can generate semantically plausible scene graphs. We show how these scene graphs can transfer to novel images, using state-of-the-art scene graph to image models [17]. Additionally, we show that the model can be used to detect unusual scene graphs and extend incomplete scene graphs.

## 2. Related Work

**Scene graphs and applications** Scene graphs as proposed in [18] offer a semantic description for an image, in the form of semantic labels for objects and their relationships. Large-scale datasets, *e.g.* Visual Genome [21] annotated with scene graphs enabled deep learning tasks. A line of works explores scene graph prediction conditioned on an image [38, 27, 39, 23, 42] or point clouds [34, 37]. Most works first detect the objects in the scene, and then reason about their relationships. In contrast, our work focuses on unconditional graph generation—*i.e.* does not rely on an input at test time – to model the scene graph distribution.

Scene graphs can be used for a variety of tasks. Johnson *et al.* [17] propose image generation from a scene graph, further explored in an interactive setup for generation [2] and semantic manipulation [7]. Wang *et al.* [35] explore relational graphs for indoor scene planing. Other works exploit scene graphs for image and domain-agnostic retrieval [18, 34]. Often scene graphs are combined with language, such as in Visual Question Answering (VQA) [11, 40] or prediction of object type, given query locations [43].

**Generative models on graphs** Traditional approaches [8, 36, 22, 1] are engineered to capture some specific patterns, often domain specific, and they fall short in generalizing to all graph patterns.

Deep auto-regressive models for graph generation [24, 25, 41, 15] are typically flexible in number of nodes, but need to impose a node ordering. GraphRNN [41] represents the graphs as a sequence of sequences, and uses a hierarchical GRU architecture to model the node and edge dependencies. The model is scalable with  $O(N^2)$  complexity and can output variable size graphs. They utilize a breadth first search (BFS) ordering to significantly reduce the possible orderings. However, the paper only addresses generation of unlabelled graphs.

Another line of work uses Variational Autoencoders (VAE), which embed the graphs into a vector [30], junction tree of node clusters [16] or combine VAE with autoregressive approach [26]. These models usually enable graph

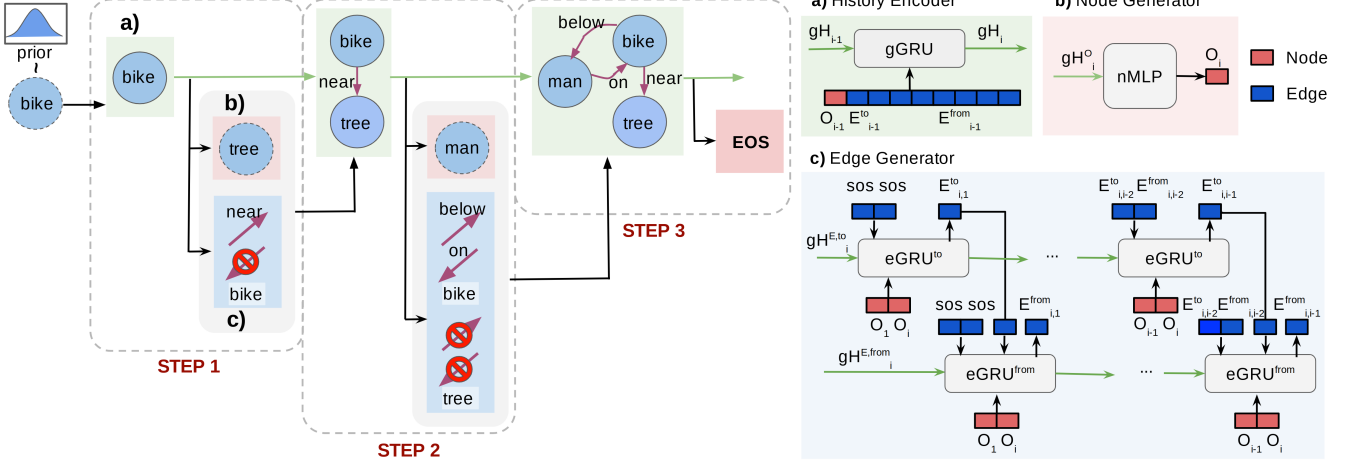


Figure 2. **Illustration of the SceneGraphGen architecture.** *Left:* The overview of the auto-regressive generation process, which takes the current graph sequence in each step and generates a new node and a set of connecting edges. *Right:* Architecture details of each of the modules: **a)** history encoder, **b)** node generator, **c)** edge generator.

modeling with attributes/categories, but the model captures inexact likelihood and do not scale well to larger graphs.

GAN based works are limited to either a single sample [3], or a small and fixed size graph [9, 4], making them unsuitable for variable-sized and diverse scene graphs. Additionally, they struggle with training stability.

In this work, we explore an auto-regressive approach which allows for a flexible number of graph nodes. Unlike GraphRNN, our method supports semantic labels for nodes and edges, models directed edges, and leverages an edge generation GRU which is aware of the node categories.

Some recent works explore generative tasks with similar scene structures (relational graphs [35], probabilistic grammar [19, 6]). Our focus is however on semantic graphs associated with in-the-wild images, which are considerably more diverse in node/edge labels and exhibit less regular graph size and connectivity patterns, compared to the synthetic datasets utilized in other works.

### 3. Methodology

Given a set of  $n$  scene graph samples  $\mathbb{G}_s = \{G_s^1, G_s^2, \dots, G_s^m\}$ , which are assumed to represent the data distribution of scene graphs  $p_{data}(G_s)$ , the goal is to learn a generative model from dataset  $\mathbb{G}_s$ , which can generate new samples of scene graphs. A scene graph sample  $G_s$  corresponding to an image  $I$  is defined as  $G_s = (\mathcal{O}, \mathcal{E})$  by the set of objects (nodes)  $\mathcal{O}$  and relationships (edges)  $\mathcal{E}$ . Each object  $o_i$  from the set  $\mathcal{O} = \{o_1, o_2, \dots, o_m\}$  holds an object category  $o_i \in \mathcal{C}$ , where  $\mathcal{C} = \{1, 2, \dots, C\}$ . The edges are an ordered triplet  $\mathcal{E} \subseteq \{(o_i, r_k, o_j) | o_i, o_j \in \mathcal{O}, o_i \neq o_j\}$  denoting the directed edge from  $o_i$  to  $o_j$ , where  $r_k \in \mathcal{R}$  is the relationship category between the objects,  $\mathcal{R} = \{1, 2, \dots, R\}$  (e.g. man - wearing - shirt). We explore this task via an auto-regressive model, which enables flexibility in num-

ber of nodes and graph connectivity. We give a network overview in section 3.1, then describe each component in more detail in the proceeding sections.

#### 3.1. Model formulation

SceneGraphGen learns a distribution  $p_\phi(G_s)$  over scene graphs which is close to  $p_{data}(G_s)$ . For an auto-regressive formulation, we convert scene graphs to a sequence representation. Under a permutation  $\pi$ , the node set  $\mathcal{O}$  becomes a sequence  $O = (\pi(o_1), \pi(o_2), \dots, \pi(o_m))$ . The edges  $\mathcal{E}$  are represented using two upper triangular sparse matrices  $E^{from}$  and  $E^{to}$ , with relationship label  $r_k$  at index  $(i, j)$  if  $(\pi(o_i), r_k, \pi(o_j)) \in \mathcal{E}$  and  $(\pi(o_j), r_k, \pi(o_i)) \in \mathcal{E}$  respectively. The matrices correspond to the two possible edges (to and from) between any node pair. A scene graph  $G_s$  can, thus, be represented as a sequence  $X = (O, E^{to}, E^{from})$ . Each element of the sequence  $X$  is itself a sequence, given as  $X_i = (O_i, E_i^{to}, E_i^{from})$ , where  $O_i$  is the object node,  $E_i^{to}$  and  $E_i^{from}$  are the sequence of edges between  $O_i$  and previous nodes. We translate the task of learning  $p_\phi(G_s)$  to learning a sequence distribution  $p_\phi(X)$ , that can be modeled auto-regressively. The probability over sequence  $X$  is decomposed into successive conditionals

$$p_\phi(X) = p(X_1) \prod_{i=2}^n p_\phi(X_i | X_{<i}). \quad (1)$$

$X_{<i} = (X_1, \dots, X_{i-1})$  depicts the partial scene graph sequence up to step  $i$ . We further split each conditional  $p_\phi(X_i | X_{<i})$  into three parts for each of the components as

$$p_\phi(X) = p(O_1) \prod_{i=2}^n p_{\phi_1}(O_i | X_{<i}) p_{\phi_2}(E_i^{to} | O_i, X_{<i}) p_{\phi_3}(E_i^{from} | O_i, E_i^{to}, X_{<i}). \quad (2)$$

We model Equation 2 and learn the complete probability distribution over scene graph sequences  $p_\phi(X)$ . We do not make any conditional independence assumptions in our model formulation, which allows SceneGraphGen to potentially capture all the complex object and relationship dependencies present in the data.  $p(O_1)$  is an assumed prior distribution over the first node, *e.g.* the categorical distribution over the object occurrences. The auto-regressive modeling is performed in three stages.

1. History Encoding: the information from the previous steps  $X_{<i}$  is captured in a hidden representation  $h_i$ . This stage models  $h_i = \text{Encoder}(X_{<i})$ .
2. Node generation: the hidden state is used to generate the next object node. This stage models  $p_{\phi_1}(O_i|h_i)$ .
3. Edge generation: the hidden state, along with the explicit node information is used to sequentially generate the edges. This stage models  $p_{\phi_2}(E_i^{to}|O_i, h_i)$  and  $p_{\phi_3}(E_i^{from}|O_i, E_i^{to}, h_i)$ .

The overall architecture and the each of the components are schematically depicted in Figure 2. Now, we describe each of the stages in detail.

### 3.2. History encoding

To carry the past information  $X_{<i}$  at each time step  $i$ , we use a Gated Recurrent Unit (GRU) architecture. We harness three separate GRUs, which are specialized to carry information for each of the three outputs  $O_i$ ,  $E_i^{from}$  and  $E_i^{to}$ . This allows for decoupling of the information for the three outputs. Note that the parameters of each of the three GRUs are different but shared across all time steps. We call these networks  $\text{gGRU}_O$ ,  $\text{gGRU}_{E^{from}}$  and  $\text{gGRU}_{E^{to}}$ . The initial hidden states of the GRUs are taken to be a zero (empty) vector. In further steps, the input to all the three GRUs is the past sequence  $X_{i-1}$ , which is taken as the concatenation of the three outputs from the previous step  $i-1$ , *i.e.*  $O_{i-1}$ ,  $E_{i-1}^{from}$  and  $E_{i-1}^{to}$ . For the first step, the edge inputs are empty (zero), and the node is sampled from a prior distribution. This prior is computed based on the occurrence of first node category for the selected ordering strategy.

### 3.3. Node generation

We use a multi-layer perceptron (MLP) which takes the hidden state from  $\text{gGRU}_O$  as input, and outputs a categorical distribution over the object categories  $\mathcal{C}$ , as prediction scores  $\theta_i^O$ . We call this network nMLP. The parameters of nMLP are shared across all time steps. The node  $O_i$  is then sampled from this distribution. The generation of new nodes (and hence, the scene graph) is halted when an EOS (end of sequence) token is encountered.

### 3.4. Node-aware edge generation

The output at each step  $i$  is itself a *sequence* of edges  $E_i^{from}$ . This forms a sequence of sequences, which in turn needs a sequential model for each step of the sequence. For this hierarchical requirement, we use a GRU which shares the parameters, not only for all time steps  $i$ , but also for each step inside that sequence  $i$ , which we call  $j$ . A similar GRU is used to produce  $E_i^{to}$  at each step  $i$ . We call these networks  $\text{eGRU}_{E^{from}}$  and  $\text{eGRU}_{E^{to}}$ . The initial hidden state for these GRUs is taken as the hidden state resulting from  $\text{gGRU}_{E^{from}}$  and  $\text{gGRU}_{E^{to}}$  networks respectively. The initial input to these GRUs is an SOS (start of sequence) token, which starts the generation process. For any step  $j$ , the input for  $\text{eGRU}_{E^{to}}$  is the concatenation of four elements, the edges from the previous step  $E_{i,j-1}^{from}$  and  $E_{i,j-1}^{to}$ , the node  $O_i$  at current step  $i$ , as well as the node  $O_j$  from the  $j^{\text{th}}$  time step. The additional information about the nodes provides *additional conditioning* for the edge generation decision. This design choice is inspired from the high predictability of co-occurrence of relationships for a given object pair, *i.e.* additional knowledge of node pairs might facilitate relationship predictions. From this input, the GRUs generate categorical distribution parameters over the relation categories  $\mathcal{R}$  (and a no-edge category), given by  $\theta_{i,j}^{E^{to}}$ . The next edge  $E_{i,j}^{to}$  is then sampled from this distribution. The inputs to the  $\text{eGRU}_{E^{from}}$  at step  $j$  are the same inputs as for  $\text{eGRU}_{E^{to}}$ , together with the edge  $E_{i,j}^{to}$  generated by  $\text{eGRU}_{E^{to}}$  at the current step  $j$ . The next edge  $E_{i,j}^{from}$  is generated by  $\text{eGRU}_{E^{from}}$  in a similar fashion. This process is carried out for  $i-1$  steps, which is the number of previous nodes that can connect to the current node  $O_i$ .

### 3.5. Loss objective

We learn the parameters  $\phi$  of the model  $p_\phi(X)$  by maximizing the likelihood of training data  $\mathbb{G}$  on the model. The training is performed by using teacher forcing, *i.e.* at each step  $i$ , we use the ground truth sequences  $O_i, E_i^{to}, E_i^{from}$  instead of sampling from the model’s prediction scores  $\theta_i^O, \theta_i^{E^{from}}, \theta_i^{E^{to}}$ . The loss is computed using cross-entropy (CE) between the predicted scores and the ground truth sequence at each step, and then summing it up. For a sample scene graph sequence  $X$ , the loss is given by

$$\mathcal{L}(X; \phi) = \mathcal{L}_O(O; \phi_1) + \mathcal{L}_E(E_i^{to}, E_i^{from}; \phi_2, \phi_3) \quad (3)$$

$$\mathcal{L}_O = \sum_{i=2}^n CE(\theta_i^O, O_i) \quad (4)$$

$$\mathcal{L}_E = \sum_{i=2}^n \sum_{j=1}^{i-1} CE(\theta_{i,j}^{E^{to}}, E_{i,j}^{to}) + \sum_{i=2}^n \sum_{j=1}^{i-1} CE(\theta_{i,j}^{E^{from}}, E_{i,j}^{from}) \quad (5)$$

### 3.6. Inference

After learning the distribution over scene graph sequences  $p_\phi(X)$ , SceneGraphGen can be used to sample new scene graph instances. The sampling procedure essentially follows the auto-regressive generation formulation described so far. First, we sample the first object node from the prior distribution which makes the first sequence  $X_1$ . The prior distribution is empirically computed from the training set, based on the occurrences of first node when ordering the nodes using the ordering scheme chosen during training. In other steps we use the previous sequence  $X_{i-1}$  as input to compute the hidden states using the three gGRU’s. These hidden states are used to generate the next node  $O_i$  using nMLP by sampling from the categorical output. Similarly, we generate the sequences of edges  $E_i^{to}$  and  $E_i^{from}$  using eGRU $_{E^{to}}$  and eGRU $_{E^{from}}$  respectively from the respective categorical outputs. The node and sequence of edges are combined to form the next sequence  $X_i$ . This process is continued until the node generator outputs an EOS token.

SceneGraphGen can additionally be used to evaluate the likelihood of a given sample. For a given scene graph converted to a sequence  $X$ , each element of the sequence  $X_{i-1}$  is passed to the network to obtain the node categorical outputs  $\theta_i^O$  as well as edge categorical output sequences  $\theta_i^{E^{from}}$  and  $\theta_i^{E^{to}}$ . Then, we pick the probability of the true node/edge category for each output and take the negative of the sum of logarithm over the whole sequence  $X$  to obtain the negative log-likelihood (NLL). This NLL value depicts the ‘unlikeliness’ of occurrence of a given scene graph sample. We use the NLL to detect anomalies in Section 4.4.

### 3.7. Implementation details

SceneGraphGen is trained for 300 epochs, with a batch size of 256, where 256 batches are sampled with replacement per epoch. The initial learning rate is 0.001. We use a decrement rate of 0.95 every 1710 steps. The node and edge inputs are namely embedded to a size of 64 and 8. All gGRUs and eGRUs have 4 GRU layers with a hidden size of 128. nMLP consists of 2 layers with ReLU activation.

## 4. Results

Here we describe the experiments we carried out to measure the performance of our model. First we introduce our evaluation protocol, including evaluation metrics, baselines and datasets. Then we present qualitative and quantitative results in graph generation. Finally, we illustrate the usefulness of the learned model in three applications: image generation, anomaly detection and graph completion.

### 4.1. Evaluation protocol

We evaluate our model on the Visual Genome (VG) dataset [21]. Specifically, we use the scene graphs from

the widely adopted VG split by Xu *et al.* [38] containing 150 object categories and 50 relationship categories, and split into train and test sets with 58k and 26k samples respectively. Further, since the scene graphs in VG dataset have incomplete relationships, we populate the edges using the Unbiased causal TDE model [32] from the SGG benchmark by Kaihua Tang [31] as it claims less bias in PredCls for rarely occurring relationship categories. Since the different instances of objects (bounding boxes) in scene graphs can correspond to the same underlying object, we use a conservative bounding box intersection-over-union (IoU) of 0.5 along with membership test in hand-made group categories to detect duplicate objects and randomly select one of them. Having duplicate objects also leads to a *child* object (e.g. shirt, ear, shoe) having relationships with multiple *parent* objects (e.g. man, child). We remove such relationships to a certain extent using heuristics.

Since there are no existing work for unconditional scene graph generation, we compare our model with GraphRNN, adapted to include node and edge categories as well as edge directions. In contrast, SceneGraphGen includes node information during edge generation and conditioning of  $E^{from}$  edges on  $E^{to}$  edges. We also test our model under various node ordering schemes, namely BFS order, hierarchical order and random order. The BFS node ordering as introduced in GraphRNN [41], visits the graph in a Breadth-First Search order. The hierarchical order sorts the nodes based on their belonging relationships, *i.e.* background (sky), objects/beings (elephant), parts (arm).

Evaluating graph generative models based on sample quality is a difficult task [33], as it requires a comparison of the generated set against the test set. We formulate such comparison between the generated and test dataset using Maximum Mean Discrepancy (MMD), which is defined between two distributions  $p$  and  $q$ , and for a given kernel  $k$  as

$$\begin{aligned} \text{MMD}^2(p, q) = & \mathbb{E}_{x, y \sim p} [k(x, y)] - 2\mathbb{E}_{x \sim p, y \sim q} [k(x, y)] \\ & + \mathbb{E}_{x, y \sim q} [k(x, y)]. \end{aligned} \tag{6}$$

Since MMD has not been used with kernels for directed and labeled graphs, we devise two kernels for comparing any two samples of scene graphs:

**Random-walk graph kernel**, inspired from [10], computes the similarity between two scene graph samples by comparing the *directed* random walks generated from the graphs.

**Object set kernel** computes the similarity between the sets of objects in the scene graphs by comparing the object categories and the number of instances of the objects.

We refer to these metrics as MMD graph and MMD node. MMD graph compares the distributions based on overall graph similarity between samples (triplets, bigger clusters, etc) and MMD node compares the distributions

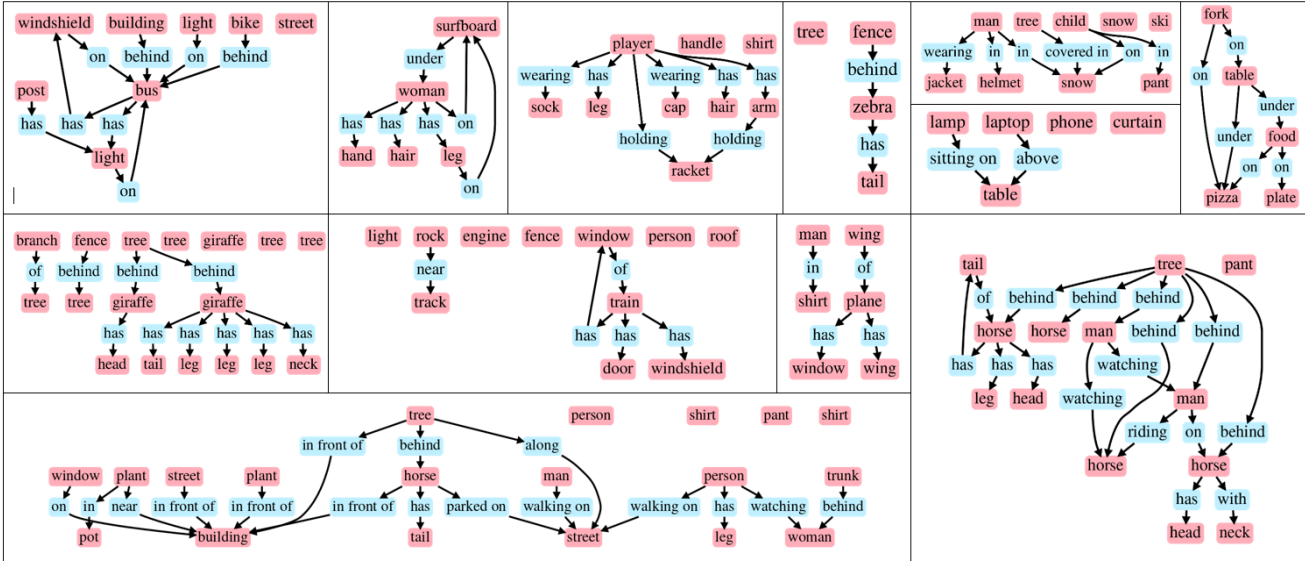


Figure 3. Visualizations of scene graphs generated by SceneGraphGen on Visual Genome. The generated graphs are diverse in size and content and represent reasonable scenes.

Model	Ordering	Graph		FID ↓	IS ↑	Image	
		MMD node ( $\times 10^3$ ) ↓	MMD graph ( $\times 10^3$ ) ↓			Precision ↑	Recall ↑
GraphRNN [41]	BFS	2.3	1.3	75.8	4.88	0.680	0.660
	Random	0.39	1.2	74.5	4.85	0.679	0.664
SceneGraphGen	BFS	2.05	1.82	73.3	5.04	0.679	0.690
	Hierarchical	1.85	0.63	72.2	<b>5.26</b>	0.717	<b>0.714</b>
	Random	<b>0.37</b>	<b>0.11</b>	<b>71.2</b>	4.95	<b>0.727</b>	<b>0.714</b>
Ground Truth		0.018	0.023	73.0	5.22	0.693	0.707

Table 1. Quantitative evaluation of the graph samples (left) and image samples (right)

based on the similarity between the respective set of objects (co-occurrence, number of instances, etc). We verify the metrics by observing their outcome on randomly corrupted datasets and the test dataset (Table 2). A dataset corruption of  $x\%$  is done by randomly choosing  $x\%$  of the nodes and edges in the graph and replacing them with a random label. As expected, MMD values between two disjoint sets of test datasets as well as between two fully-corrupted datasets are very low, as the comparing datasets are very similar. As the level of corruption increases, the MMD metrics between test and corrupted set increases, since the distributions become more dissimilar.

Comparison	MMD node ↓	MMD graph ↓
test Vs test	0.018	0.023
100% corrupt Vs 100% corrupt	0.11	0.0098
test Vs 20% corrupt	6.0	3.7
test Vs 50% corrupt	10	6.3
test Vs 100% corrupt	44	25

Table 2. Sanity checks for the MMD metrics ( $\times 10^3$ ), comparing a ground truth set of graphs (test) against a randomly corrupted set.

We further evaluate SceneGraphGen by assessing the quality of images generated by the generated scene graphs. Specifically, we generate  $64 \times 64$  images using the sg2im [17] model with the scene graphs generated by SceneGraphGen as input. For assessing the sample quality, we use established metrics in image generation quality, like the Frechet Inception Distance (FID) [14], Inception Score (IS) [29], Precision ( $F_{1/8}$ ) and Recall ( $F_8$ ) scores [28]. We compare the generated images against the ground truth Visual Genome images.

## 4.2. Graph generation

SceneGraphGen can generate realistic and diverse samples of scene graphs. Figure 3 shows some visualizations of the generated graphs. We observe that the model can generate graphs corresponding to both indoor (table, food) and outdoor scenes (building, fence, horse), while generating reasonable object co-occurrences and semantically meaningful relationships between the objects. SceneGraphGen outperforms GraphRNN both in terms of MMD graph and MMD node (Table 1, left), showing that node in-

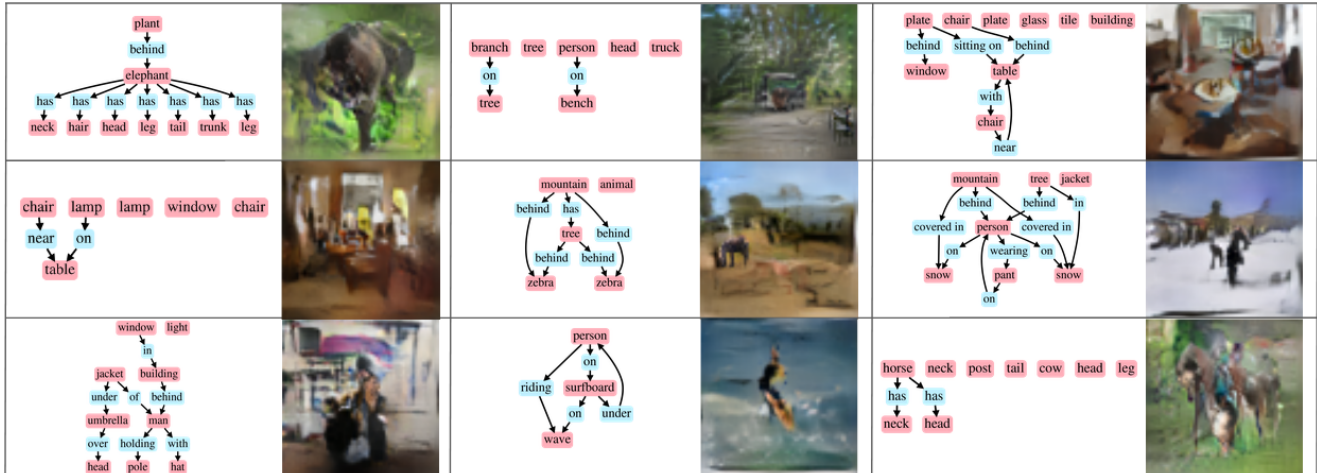


Figure 4. Some examples of 64x64 images synthesized using sg2im on the corresponding scene graphs generated by SceneGraphGen



Figure 5. Some examples of 64x64 resolution images synthesized via a.) Unconditional StyleGAN2 trained on Visual Genome (left) b.) sg2im on scene graphs generated by SceneGraphGen trained on Visual Genome (right).

formation significantly improve edge generation. We also observe that random ordering outperforms other ordering schemes for scene graph data. Note that BFS and hierarchical ordering are biased towards some node orderings due to their dependence on node connectivity and node semantics respectively. Such bias is not useful for a dataset like Visual Genome which varies a lot in structure and little regularity assumptions can be made. We believe that a random ordering performs better as it has no such bias. We refer the readers to the supplement for a comparison of data statistics of generated dataset and test dataset.

### 4.3. Image generation from generated graphs

Here we want to show that the graphs generated from SceneGraphGen can be used to synthesize novel images. We use sg2im [17], an off-the-shelf method for image generation from scene graphs to translate the learned graphs in the image domain. We dub this method sg2im-SGG. Figure 4 shows some examples of novel scenes generated by the corresponding newly generated scene graphs. Quantitative evaluation on the generated images (Table 1, right)

shows that SceneGraphGen outperforms GraphRNN on all metrics, and is fairly similar to ground truth, *i.e.* images generated from the ground truth test set.

We additionally compare these images against a state-of-the-art model in unconditional image generation [20]. For fairness, we train the StyleGAN2 model on Visual Genome on the same training split used to train sg2im. Figure 5 illustrates results for both methods. We observe that, though StyleGAN2 generates images in good quality for simple image constellations (*e.g.* landscapes), it fails to capture semantic properties in more complex scenes (*e.g.* multiple instances of green screens or indistinguishable objects). On the contrary the images produced by the generated graphs have more well grounded compositions, as dictated by the scene graphs. We also compare the object statistics of StyleGAN vs sg2im-SGG by detecting objects using off-the-shelf Faster-RCNN model trained on COCO. The detector was able to detect a total of 50 object categories for sg2im-SGG images as compared to 40 categories for StyleGAN images, illustrating that using scene graphs as an intermediate representation aids in generating more semantically

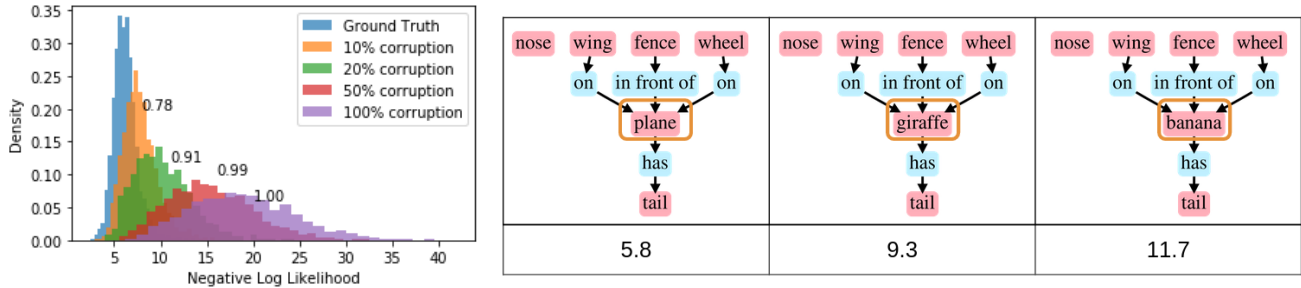


Figure 6. **Anomaly detection using SceneGraphGen** *Left*: The distribution of negative log-likelihood (NLL) of datasets under varied levels of corruption. The value next to each distribution is the area under ROC curve. *Right*: An example with increasing value of NLL.

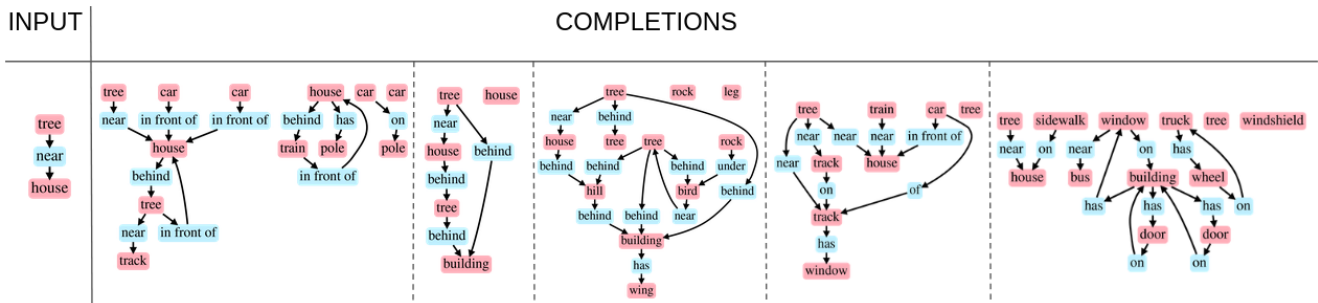


Figure 7. Samples of diverse completed scene graphs from an input partial scene graph.

diverse scenes. Additionally, comparing the object occurrences against the ground truth test set of VG, shows that sg2im-SGG is more in line with the ground truth compared to the StyleGAN2 model, with the ground-truth average error of 1.2 as compared to 1.4. We refer the readers to supplement for a comparison of object occurrence bar plots.

#### 4.4. Anomaly detection

We demonstrate that the learned model can also be used to detect anomalies in scenes. We evaluate the likelihood over the test dataset using SceneGraphGen, and detect anomalies as scene graphs whose likelihoods are outliers to the likelihood distribution. We illustrate the effectiveness of likelihood in anomaly detection by computing the negative log-likelihood (NLL) over datasets with various levels of node and edge corruption. We compute the area under ROC curve (AUROC) by using NLL scores to classify corrupt samples as anomalies, as a measure of effectiveness of NLL in anomaly detection [13]. Figure 6 (left) shows the NLL distribution along with AUROC values. The respective GraphRNN plot can be found in the supplement. Figure 6 (right) illustrates an example of anomaly, where the more surprising samples lead to a higher NLL value.

#### 4.5. Scene graph completion

SceneGraphGen can also generate a scene graph conditioned on a partial graph input, which allows for a more guided generation of scenes. Starting with a partial graph,

we randomly sample a sequence of nodes/edges ( $X_i$ ) conditioned on the previous sequence ( $X_{i-1}$ ) – same as for graph generation – which accumulates stochasticity at each step leading to diverse completions. Figure 7 shows how the model is capable of generating diverse possible scene graph completions starting from the same triplet `tree - near - house`. For more examples, please refer to the supplement.

### 5. Conclusions

We introduced SceneGraphGen, a model for unconditional scene graph generation. We showed that the model is capable of learning semantic patterns from real scenes, and synthesizing reasonable and diverse scene graphs. The model outperforms the baselines on graph-based metrics as well as image-based metrics. We demonstrated the applicability of the learned model and the generated graphs in image generation, detection of out-of-distribution samples and scene graph completion. Future work will be dedicated to exploring conditional variants of the model, such as using text descriptions to constrain the scene category.

### 6. Acknowledgement

We gratefully acknowledge the Deutsche Forschungsgemeinschaft (DFG) for supporting this research work, under the project 381855581. We are also thankful to Munich Center for Machine Learning (MCML) with funding from the Bundesministerium für Bildung und Forschung (BMBF) under the project 01IS18036B.



## References

- [1] Emmanuel Abbe. Community detection and stochastic block models: Recent developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.
- [2] Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4561–4569, 2019.
- [3] Aleksandar Bojchevski, Oleksandr Shchur, Daniel Zügner, and Stephan Günnemann. NetGAN: Generating graphs via random walks. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 610–619, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- [4] Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular graphs. *ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018.
- [5] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches, 2014.
- [6] Jeevan Devaranjan, Amlan Kar, and Sanja Fidler. Meta-sim2: Learning to generate synthetic datasets. In *ECCV*, 2020.
- [7] Helisa Dhamo, Azade Farshad, Iro Laina, Nassir Navab, Gregory D. Hager, Federico Tombari, and Christian Rupprecht. Semantic Image Manipulation Using Scene Graphs. In *Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [8] P. Erdős and A. Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290, 1959.
- [9] Shuangfei Fan and Bert Huang. Labeled graph generative adversarial networks, 2019.
- [10] Matthew Fisher, Manolis Savva, and Pat Hanrahan. Characterizing structural relationships in scenes using graph kernels. In *ACM SIGGRAPH 2011 Papers*, SIGGRAPH '11, New York, NY, USA, 2011. Association for Computing Machinery.
- [11] Shalini Ghosh, Giedrius Burachas, Arijit Ray, and Avi Ziskind. Generating natural language explanations for visual question answering using scene graphs and visual attention, 2019.
- [12] Aditya Grover, Aaron Zweig, and Stefano Ermon. Graphite: Iterative generative modeling of graphs. *arXiv preprint arXiv:1803.10459*, 2018.
- [13] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *ICLR*, 2019.
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [15] John Ingraham, Vikas K Garg, Regina Barzilay, and Tommi Jaakkola. Generative models for graph-based protein design. In *Advances in Neural Information Processing Systems*, 2019.
- [16] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. *ICML*, 2018.
- [17] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image Generation from Scene Graphs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] J. Johnson, R. Krishna, M. Stark, L. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [19] Amlan Kar, Aayush Prakash, Ming-Yu Liu, Eric Cameracci, Justin Yuan, Matt Rusiniak, David Acuna, Antonio Torralba, and Sanja Fidler. Meta-sim: Learning to generate synthetic datasets. In *ICCV*, 2019.
- [20] 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.
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalanidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision (IJCV)*, 2017.
- [22] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: An approach to modeling networks. *J. Mach. Learn. Res.*, 11:985–1042, Mar. 2010.
- [23] Yikang Li, Wanli Ouyang, Bolei Zhou, Jianping Shi, Chao Zhang, and Xiaogang Wang. Factorizable Net: An Efficient Subgraph-Based Framework for Scene Graph Generation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [24] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs, 2018.
- [25] Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, Charlie Nash, William L. Hamilton, David Duvenaud, Raquel Urtasun, and Richard Zemel. Efficient graph generation with graph recurrent attention networks. In *NeurIPS*, 2019.
- [26] Qi Liu, Miltiadis Allamanis, Marc Brockschmidt, and Alexander L. Gaunt. Constrained graph variational autoencoders for molecule design. *The Thirty-second Conference on Neural Information Processing Systems*, 2018.
- [27] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [28] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *NeurIPS*, 2018.
- [29] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016.
- [30] Martin Simonovsky and Nikos Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. *ICANN*, 2018.

- [31] Kaihua Tang. A scene graph generation codebase in pytorch, 2020. <https://github.com/KaihuaTang/Scene-Graph-Benchmark.pytorch>.
- [32] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Conference on Computer Vision and Pattern Recognition*, 2020.
- [33] Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models, 2016.
- [34] Johanna Wald, Helisa Dharmo, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs from 3d indoor reconstructions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [35] Kai Wang, Y. Lin, Ben Weissmann, M. Savva, Angel X. Chang, and D. Ritchie. Planit: planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Trans. Graph.*, 38:132:1–132:15, 2019.
- [36] Stan Wasserman and Philippa Pattison. Logit models and logistic regressions for social networks: I. an introduction to markov graphs and p\*. *Psychometrika*, 61:401–425, 09 1996.
- [37] Shun-Cheng Wu, Johanna Wald, Keisuke Tateno, Nassir Navab, and Federico Tombari. Scenegraphfusion: Incremental 3d scene graph prediction from rgb-d sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7515–7525, June 2021.
- [38] Danfei Xu, Yuke Zhu, Christopher Choy, and Li Fei-Fei. Scene Graph Generation by Iterative Message Passing. In *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [39] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph R-CNN for Scene Graph Generation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [40] Zhuoqian Yang, Zengchang Qin, Jing Yu, and Yue Hu. Scene graph reasoning with prior visual relationship for visual question answering, 2019.
- [41] Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and Jure Leskovec. Graphrnn: Generating realistic graphs with deep auto-regressive models. *ICML*, 2018.
- [42] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural Motifs: Scene Graph Parsing with Global Context. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [43] Yang Zhou, Zachary While, and Evangelos Kalogerakis. Scenegraphnet: Neural message passing for 3d indoor scene augmentation. In *IEEE Conference on Computer Vision (ICCV)*, 2019.