

Semantic Image Manipulation Using Scene Graphs

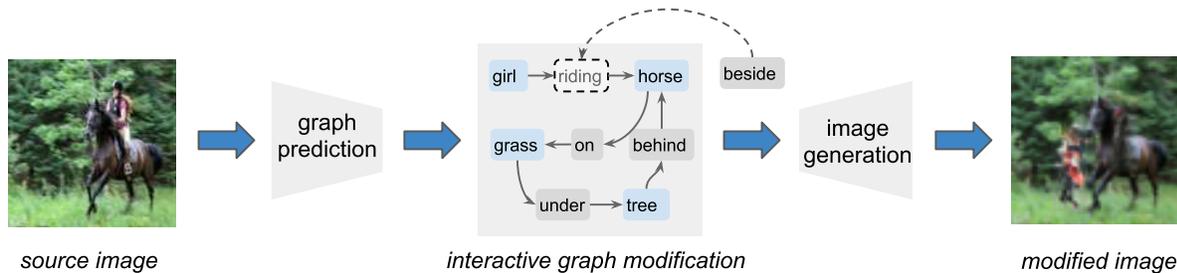
Helisa Dhamo^{1, *}Azade Farshad^{1, *}Iro Laina^{1,2}Nassir Navab^{1,3}Gregory D. Hager³Federico Tombari^{1,4}Christian Rupprecht²¹ Technische Universität München ² University of Oxford ³ Johns Hopkins University ⁴ Google

Figure 1: **Semantic Image Manipulation.** Given an image, we predict a semantic scene graph. The user interacts with the graph by making changes on the nodes and edges. Then, we generate a modified version of the source image, which respects the constellations in the modified graph.

Abstract

Image manipulation can be considered a special case of image generation where the image to be produced is a modification of an existing image. Image generation and manipulation have been, for the most part, tasks that operate on raw pixels. However, the remarkable progress in learning rich image and object representations has opened the way for tasks such as text-to-image or layout-to-image generation that are mainly driven by semantics. In our work, we address the novel problem of image manipulation from scene graphs, in which a user can edit images by merely applying changes in the nodes or edges of a semantic graph that is generated from the image. Our goal is to encode image information in a given constellation and from there on generate new constellations, such as replacing objects or even changing relationships between objects, while respecting the semantics and style from the original image. We introduce a spatio-semantic scene graph network that does not require direct supervision for constellation changes or image edits. This makes it possible to train the system from existing real-world datasets with no additional annotation effort.

1. Introduction

The goal of image understanding is to extract rich and meaningful information from an image. Recent techniques based on deep representations are continuously pushing the boundaries of performance in recognizing objects [37] and their relationships [27] or producing image descriptions [18]. Understanding is also necessary for image synthesis, e.g. to generate natural looking images from an abstract semantic canvas [4, 44, 57] or even from language descriptions [11, 24, 36, 53, 55]. High-level image manipulation, however, has received less attention. Image manipulation is still typically done at pixel level via photo editing software and low-level tools such as in-painting. Instances of higher-level manipulation are usually object-centric, such as facial modifications or reenactment. A more abstract way of manipulating an image from its semantics, which includes objects, their relationships and attributes, could make image editing easier with less manual effort from the user.

In this work, we present a method to perform semantic editing of an image by modifying a scene graph, which is a representation of the objects, attributes and interactions in the image (Figure 1). As we show later, this formulation allows the user to choose among different editing functions. For example, instead of manually segmenting, deleting and in-painting unwanted tourists in a holiday photo, the user

*The first two authors contributed equally to this work
Project page: <https://he-dhamo.github.io/SIMSG/>

can directly manipulate the scene graph and delete selected `<person>` nodes. Similarly, graph nodes can be easily replaced with different semantic categories, for example replacing `<clouds>` with `<sky>`. It is also possible to re-arrange the spatial composition of the image by swapping people or object nodes on the image canvas. To the best of our knowledge, this is the first approach to image editing that also enables semantic relationship changes, for example changing “*a person walking in front of the sunset*” to “*a person jogging in front of the sunset*” to create a more scenic image. The capability to reason and manipulate a scene graph is not only useful for photo editing. The field of robotics can also benefit from this kind of task, *e.g.* a robot tasked to tidy up a room can — prior to acting — manipulate the scene graph of the perceived scene by moving objects to their designated spaces, changing their relationships and attributes: “*clothes lying on the floor*” to “*folded clothes on a shelf*”, to obtain a realistic future view of the room.

Much previous work has focused either on generating a scene graph from an image [25, 29] or an image from a graph [1, 16]. Here we face challenges unique to the combined problem. For example, if the user changes a relationship attribute — *e.g.* `<boy, sitting on, grass>` to `<boy, standing on, grass>`, the system needs to generate an image that contains the *same* boy, thus preserving the identity as well as the content of the rest of the scene. Collecting a fully supervised data set, *i.e.* a data set of “before” and “after” pairs together with the associated scene graph, poses major challenges. As we discuss below, this is not necessary. It is in fact possible to learn how to modify images using only training pairs of images and scenes graphs, which is data already available.

In summary, we present a novel task; given an image, we manipulate it using the respective scene graph. Our contribution is a method to address this problem that does not require full supervision, *i.e.* image pairs that contain scene changes. Our approach can be seen as semi-automatic, since the user does not need to manually edit the image but indirectly interacts with it through the nodes and edges of the graph. In this way, it is possible to make modifications with respect to visual entities in the image and the way they interact with each other, both spatially and semantically. Most prominently, we achieve various types of edits with a single model, including semantic relationship changes between objects. The resulting image preserves the original content, but allows the user to flexibly change and/or integrate new or modified content as desired.

2. Related Work

Conditional image generation The success of deep generative models [8, 20, 35, 42, 43] has significantly contributed to advances in (un)conditional image synthesis. Conditional image generation methods model the conditional

distribution of images given some prior information. For example, several practical tasks such as denoising or inpainting can be seen as generation from noisy or partial input. Conditional models have been studied in literature for a variety of use cases, conditioning the generation process on image labels [28, 30], attributes [46], lower resolution images [23], semantic segmentation maps [4, 44], natural language descriptions [24, 36, 53, 55] or generally translating from one image domain to another using paired [13] or unpaired data [60]. Most relevant to our approach are methods that generate natural scenes from layout [11, 57] or scene graphs [16].

Image manipulation Unconditional image synthesis is still an open challenge when it comes to complex scenes. Image manipulation, on the other hand, focuses on image parts in a more constrained way that allows to generate better quality samples. Image manipulation based on semantics has been mostly restricted to object-centric scenarios; for example, editing faces automatically using attributes [5, 22, 56] or via manual edits with a paintbrush and scribbles [3, 59]. Also related is image composition which also makes use of individual objects [2] and faces the challenge of decoupling appearance and geometry [52].

On the level of scenes, the most common examples based on generative models are inpainting [33], in particular conditioned on semantics [49] or user-specified contents [15, 58], as well as object removal [7, 40]. Image generation from semantics also supports interactive editing by applying changes to the semantic map [44]. Differently, we follow a semi-automatic approach to address all these scenarios using a single general-purpose model and incorporating edits by means of a scene graph. On another line, Hu *et al.* [12] propose a hand-crafted image editing approach, which uses graphs to carry out library-driven replacement of image patches. While [12] focuses on copy-paste tasks, our framework allows for high-level semantic edits and deals with object deformations. Our method is trained by reconstructing the input image so it does not require paired data. A similar idea is explored by Yao *et al.* [48] for 3D-aware modification of a scene (*i.e.* 3D object pose) by disentangling semantics and geometry. However, this approach is limited to a specific type of scenes (streets) and target objects (cars) and requires CAD models. Instead, our approach addresses semantic changes of objects and their relationships in natural scenes, which is made possible using scene graphs.

Images and scene graphs Scene graphs provide abstract, structured representations of image content. Johnson *et al.* [19] first defined a scene graph as a directed graph representation that contains objects and their attributes and relationships, *i.e.* how they interact with each other. Following this graph representation paradigm, different meth-

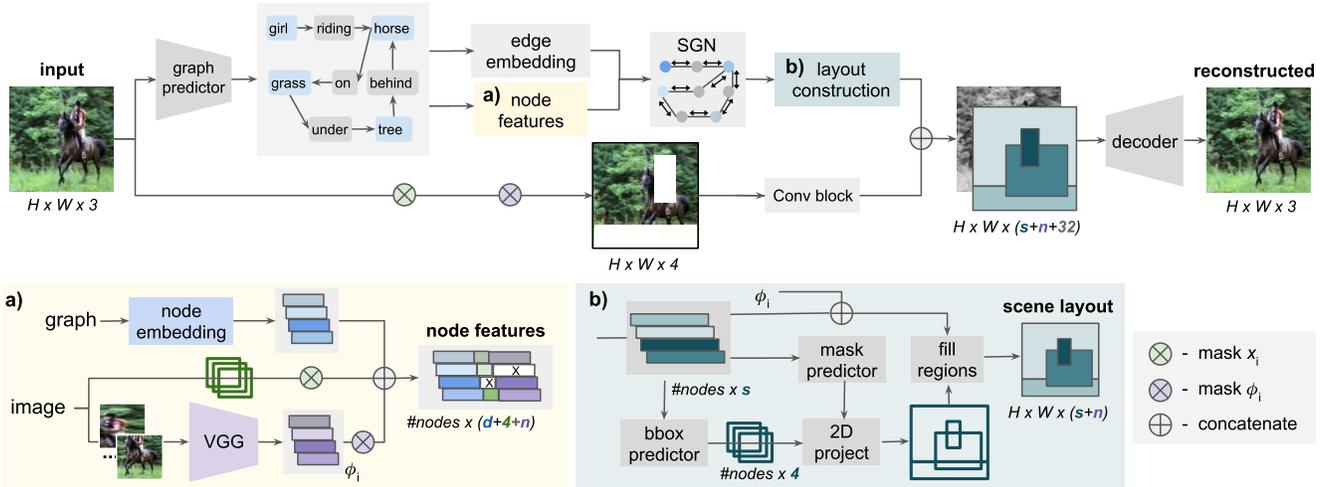


Figure 2: **Overview of the training strategy.** *Top:* Given an image, we predict its scene graph and reconstruct the input from a masked representation. a) The graph nodes o_i (blue) are enriched with bounding boxes x_i (green) and visual features ϕ_i (violet) from cropped objects. We randomly mask boxes x_i , object visual features ϕ_i and the source image; the model then reconstructs the same graph and image utilizing the remaining information. b) The per-node feature vectors are projected to 2D space, using the bounding box predictions from SGN.

ods have been proposed to generate scene graphs from images [9, 25, 26, 29, 34, 45, 47, 51]. By definition, scene graph generation mainly relies on successfully detecting visual entities in the image (object detection) [37] and recognizing how these entities interact with each other (visual relationship detection) [6, 14, 27, 38, 50].

The reverse and under-constrained problem is to generate an image from its scene graph, which has been recently addressed by Johnson *et al.* using a graph convolution network (GCN) to decode the graph into a layout and consecutively translate it into image [16]. We build on this architecture and propose additional mechanisms for information transfer from an image that act as conditioning for the system, when the goal is image editing and not free-form generation. Also related is image generation directly from layouts [57]. Very recent related work focuses on interactive image generation from scene graphs [1] or layout [41]. These methods differ from ours in two aspects. First, while [1, 41] process a graph/layout to generate multiple variants of an image, our method manipulates an *existing* image. Second, we present complex semantic relationship editing, while they use graphs with simplified spatial relations — *e.g.* relative object positions such as `left of` or `above` in [1] — or without relations at all, as is the case for the layout-only approach in [41].

3. Method

The focus of this work is to perform semantic manipulation of images without direct supervision for image edits, *i.e.* without paired data of original and modified content.

Starting from an input image I , we generate its scene graph \mathcal{G} that serves as the means of interaction with a user. We then generate a new image I' from the user-modified graph representation $\hat{\mathcal{G}}$ and the original content of I . An overview of the method is shown in Figure 1. Our method can be split into three interconnected parts. The first step is scene graph generation, where we encode the image contents in a spatio-semantic scene graph, designed so that it can easily be manipulated by a user. Second, during inference, the user manipulates the scene graph by modifying object categories, locations or relations by directly acting on the nodes and edges of the graph. Finally, the output image is generated from the modified graph. Figure 2 shows the three components and how they are connected.

A particular challenge in this problem is the difficulty in obtaining training data, *i.e.* matching pairs of source and target images together with their corresponding scene graphs. To overcome these limitations, we demonstrate a method that learns the task by image reconstruction in an unsupervised way. Due to readily available training data, graph prediction instead is learned with full supervision.

3.1. Graph Generation

Generating a scene graph from an image is a well-researched problem [25, 29, 45, 51] and amounts to describing the image with a directed graph $\mathcal{G} = (\mathcal{O}, \mathcal{R})$ of objects \mathcal{O} (nodes) and their relations \mathcal{R} (edges). We use a state-of-the-art method for scene graph prediction (F-Net) [25] and build on its output. Since the output of the system is a generated image, our goal is to encode as much image

information in the scene graph as possible — additional to semantic relationships. We thus define objects as triplets $o_i = (c_i, \phi_i, x_i) \in \mathcal{O}$, where $c_i \in \mathbb{R}^d$ is a d -dimensional, learned embedding of the i -th object category and $x_i \in \mathbb{R}^4$ represents the four values defining the object’s bounding box. $\phi_i \in \mathbb{R}^n$ is a visual feature encoding of the object which can be obtained from a convolutional neural network (CNN) pre-trained for image classification. Analogously, for a given relationship between two objects i and j , we learn an embedding ρ_{ij} of the relation class $r_{ij} \in \mathcal{R}$.

One can also see this graph representation as an augmentation of a simple graph—that only contains object and predicate categories—with image features and spatial locations. Our graph contains sufficient information to preserve the identity and appearance of objects even when the corresponding locations and/or relationships are modified.

3.2. Spatio-semantic Scene Graph Network

At the heart of our method lies the spatio-semantic scene graph network (SGN) that operates on the (user-) modified graph. The network learns a graph transformation that allows information to flow between objects, along their relationships. The task of the SGN is to learn robust object representations that will be then used to reconstruct the image. This is done by a series of convolutional operations on the graph structure.

The graph convolutions are implemented by an operation τ_e on edges of the graph

$$(\alpha_{ij}^{(t+1)}, \rho_{ij}^{(t+1)}, \beta_{ij}^{(t+1)}) = \tau_e \left(\nu_i^{(t)}, \rho_{ij}^{(t)}, \nu_j^{(t)} \right), \quad (1)$$

with $\nu_i^{(0)} = o_i$, where t represents the layer of the SGN and τ_e is implemented as a multi-layer perceptron (MLP). Since nodes can appear in several edges, the new node feature $\nu_i^{(t+1)}$ is computed by averaging the results from the edge-wise transformation, followed by another projection τ_n

$$\nu_i^{(t+1)} = \tau_n \left(\frac{1}{N_i} \left(\sum_{j|(i,j) \in \mathcal{R}} \alpha_{ij}^{(t+1)} + \sum_{k|(k,i) \in \mathcal{R}} \beta_{ki}^{(t+1)} \right) \right) \quad (2)$$

where N_i represents the number of edges that start or end in node i . After T graph convolutional layers, the last layer predicts one latent representation per node, *i.e.* per object. This output object representation consists of predicted bounding box coordinates $\hat{x}_i \in \mathbb{R}^4$, a spatial binary mask $\hat{m}_i \in \mathbb{R}^{M \times M}$ and a node feature vector $\psi_i \in \mathbb{R}^s$. Predicting coordinates for each object is a form of reconstruction, since object locations are known and are already encoded in the input o_i . As we show later, this is needed when modifying the graph, for example for a new node to be added. The predicted object representation will be then reassembled into the spatial configuration of an image, as the scene layout.

3.3. Scene Layout

The next component is responsible for transforming the graph-structured representations predicted by the SGN back into a 2D spatial arrangement of features, which can then be decoded into an image. To this end, we use the predicted bounding box coordinates \hat{x}_i to project the masks \hat{m}_i in the proper region of a 2D representation of the same resolution as the input image. We concatenate the original visual feature ϕ_i with the node features ψ_i to obtain a final node feature. The projected mask region is then filled with the respective features, while the remaining area is padded with zeros. This process is repeated for all objects, resulting in $|\mathcal{O}|$ tensors of dimensions $(n+s) \times H \times W$, which are aggregated through summation into a single layout for the image. The output of this component is an intermediate representation of the scene, which is rich enough to reconstruct an image.

3.4. Image Synthesis

The last part of the pipeline is the task of synthesizing a target image from the information in the source image I and the layout prediction. For this task, we employ two different decoder architectures, cascaded refinement networks (CRN) [4] (similar to [16]), as well as SPADE [32], originally proposed for image synthesis from a semantic segmentation map. We condition the image synthesis on the source image by concatenating the predicted layout with extracted low-level features from the source image. In practice, prior to feature extraction, regions of I are occluded using a mechanism explained in Section 3.5. We fill these regions with Gaussian noise to introduce stochasticity for the generator.

3.5. Training

Training the model with full supervision would require annotations in the form of quadruplets $(I, \mathcal{G}, \mathcal{G}', I')$ where an image I is annotated with a scene graph \mathcal{G} , a modified graph \mathcal{G}' and the resulting modified image I' . Since acquiring ground truth (I', \mathcal{G}') is difficult, our goal is to train a model supervised only by (I, \mathcal{G}) through reconstruction. Thus, we generate annotation quadruplets $(\tilde{I}, \tilde{\mathcal{G}}, \mathcal{G}, I)$ using the available data (I, \mathcal{G}) as the *target* supervision and simulate $(\tilde{I}, \tilde{\mathcal{G}})$ via a random masking procedure that operates on object instances. During training, an object’s visual features ϕ_i are masked with probability p_ϕ . Independently, we mask the bounding box x_i with probability p_x . When “hiding” input information, image regions corresponding to the hidden nodes are also occluded prior to feature extraction.

Effectively, this masking mechanism transforms the editing task into a reconstruction problem. At run time, a real user can directly edit the nodes or edges of the scene graph. Given the edit, the image regions subject to modification are occluded, and the network, having learned to reconstruct the image from the scene graph, will create a plausible modified image. Consider the example of a person riding a horse

(Figure 1). The user wishes to apply a change in the way the two entities interact, modifying the predicate from `riding` to `beside`. Since we expect the spatial arrangement to change, we also discard the localization x_i of these entities in the original image; their new positions \hat{x}_i will be estimated given the layout of the rest of the scene (e.g. grass, trees). To encourage this change, the system should automatically mask the original image regions related to the target objects. However, to ensure that the *visual identities* of horse and rider are preserved through the change, their visual feature encodings ϕ_i must remain unchanged.

We use a combination of loss terms to train the model. The bounding box prediction is trained by minimizing the L_1 -norm: $\mathcal{L}_b = \|x_i - \hat{x}_i\|_1$, with weighting term λ_b . The image generation task is learned by adversarial training with two discriminators. A local discriminator D_{obj} operates on each reconstructed region to ensure that the generated patches look realistic. We also apply an auxiliary classifier loss [31] to ensure that D_{obj} is able to classify the generated objects into their real labels. A global discriminator D_{global} encourages consistency over the entire image. Finally, we apply a photometric loss term $\mathcal{L}_r = \|I - I'\|_1$ to enforce the image content to stay the same in regions that are not subject to change. The total synthesis loss is then

$$\begin{aligned} \mathcal{L}_{\text{synthesis}} = & \mathcal{L}_r + \lambda_g \min_G \max_D \mathcal{L}_{\text{GAN,global}} \\ & + \lambda_o \min_G \max_D \mathcal{L}_{\text{GAN,obj}} + \lambda_a \mathcal{L}_{\text{aux,obj}}, \end{aligned} \quad (3)$$

where $\lambda_g, \lambda_o, \lambda_a$ are weighting factors and

$$\mathcal{L}_{\text{GAN}} = \mathbb{E}_{q \sim p_{\text{real}}} \log D(q) + \mathbb{E}_{q \sim p_{\text{fake}}} \log(1 - D(q)), \quad (4)$$

where p_{real} corresponds to the ground truth distribution (of each object or the whole image) and p_{fake} is the distribution of generated (edited) images or objects, while q is the input to the discriminator which is sampled from the real or fake distributions. When using SPADE, we additionally employ a perceptual loss term $\lambda_p \mathcal{L}_p$ and a GAN feature loss term $\lambda_f \mathcal{L}_f$ following the original implementation [32]. Moreover, D_{global} becomes a multi-scale discriminator.

Full implementation details regarding the architectures, hyper-parameters and training can be found in the Appendix.

4. Experiments

We evaluate our method quantitatively and qualitatively on two datasets, CLEVR [17] and Visual Genome [21], with two different motivations. As CLEVR is a synthetic dataset, obtaining ground truth pairs for image editing is possible, which allows quantitative evaluation of our method. On the other hand, experiments on Visual Genome (VG) show the performance of our method in a real, much less constrained, scenario. In absence of source-target image pairs in VG, we

Method	All pixels				RoI only	
	MAE ↓	SSIM ↑	LPIPS ↓	FID ↓	MAE ↓	SSIM ↑
Full-sup	6.75	97.07	0.035	3.35	9.34	93.49
Ours (CRN)	7.83	96.16	0.036	6.32	10.09	93.54
Ours (SPADE)	5.47	96.51	0.035	4.73	7.22	94.98

Table 1: **Image manipulation on CLEVR.** We compare our method with a fully-supervised baseline. Detailed results for all modification types are reported in the Appendix.

evaluate an image in-painting proxy task and compare to a baseline based on `sg2im` [16]. We report results for standard image reconstruction metrics: the structural similarity index (SSIM), mean absolute error (MAE) and perceptual error (LPIPS) [54]. To assess the image generation quality and diversity, we report the commonly used inception score (IS) [39] and the FID [10] metric.

Conditional `sg2im` baseline (Cond-`sg2im`). We modify the model of [16] to serve as a baseline. Since their method generates images directly from scene graphs without a source image, we condition their image synthesis network on the input image by concatenating it with the layout component (instead of noise in the original work). To be comparable to our approach, we mask image regions corresponding to the target objects prior to concatenation.

Modification types. Since image editing using scene graphs is a novel task, we define several modification modes, depending on how the user interacts with the graph. **Object removal:** A node is removed entirely from the graph together with all the edges that connect this object with others. The source image region corresponding to the object is occluded. **Object replacement:** A node is assigned to a different semantic category. We do not remove the full node; however, the visual encoding ϕ_i of the original object is set to zero, as it does not describe the novel object. The location of the original entity is used to keep the new object in place, while size comes from the bounding box estimated from the SGN, to fit the new category. **Relationship change:** This operation usually involves re-positioning of entities. The goal is to keep the subject and object but change their interaction, e.g. `<sitting>` to `<standing>`. Both the original and novel appearance image regions are occluded, to enable background in-painting and target object generation. The visual encodings ϕ_i are used to condition the SGN and maintain the visual identities of objects on re-appearance.

4.1. Synthetic Data

We use the CLEVR framework [17] to generate a dataset (for details please see the Appendix) of image and scene

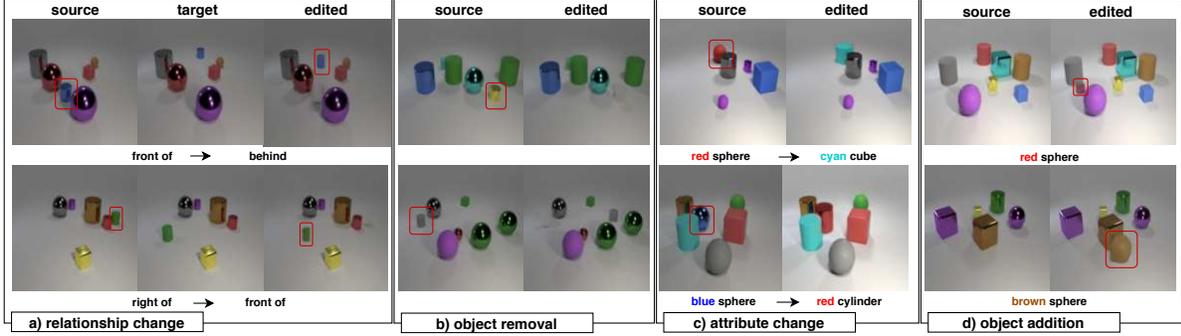


Figure 3: **Image manipulation on CLEVR** We compare different changes in the scene including changing the relationship between two objects, node removal and changing a node (corresponding to attribute changing).

graph editing pairs $(I, \mathcal{G}, \mathcal{G}', I')$, to evaluate our method with exact ground truth.

We train our model *without making use of image pairs* and compare our approach to a fully-supervised setting. When training with full supervision the complete source image and target graph are given to the model and the model is trained by minimizing the \mathcal{L}_1 loss to the ground truth target image instead of the proposed masking scheme.

Table 1 reports the mean SSIM, MAE, LPIPS and FID on CLEVR for the manipulation task (replacement, removal, relationship change and addition). Our method performs better or on par with the fully-supervised setting, on the reconstruction metrics, which shows the capability of synthesizing meaningful changes. The FID results suggest that additional supervision for pairs, if available, would lead to improvement in the visual quality. Figure 3 shows qualitative results of our model on CLEVR. At test time, we apply changes to the scene graph in four different modes: changing relationships, removing an object, adding an object or changing its identity. We highlight the modification with a bounding box drawn around the selected object.

4.2. Real Images

We evaluate our method on Visual Genome [21] to show its performance on natural images. Since there is no ground truth for modifications, we formulate the quantitative evaluation as image reconstruction. In this case, objects are occluded from the original image and we measure the quality of the reconstruction. The qualitative results better illustrate the full potential of our method.

Feature encoding. First, we quantify the role of the visual feature ϕ_i in encoding visual appearance. For a given image and its graph, we use all the associated object locations x_i and visual features (w/ ϕ_i) to condition the SGN. However, the region of the conditioning image corresponding to a candidate node is masked. The task can be interpreted as conditional in-painting. We test our approach in two scen-



Figure 4: **Visual feature encoding.** Comparison between the baseline (top) and our method (center). The scene graph remains unchanged; an object in the image is occluded, while ϕ_i and x_i are active. Our latent features ϕ_i preserve appearance when the objects are masked from the image.

arios; using ground truth graphs (GT) and graphs predicted from the input images (P). We evaluate over all objects in the test set and report the results in Table 2, measuring the reconstruction error a) over all pixels and b) in the target area only (RoI). We compare to the same model without using visual features (w/o ϕ_i) but only the object category to condition the SGN. Naturally, in all cases, including the missing region’s visual features improves the reconstruction metrics (MAE, SSIM, LPIPS). In contrast, inception score and FID remain similar, as these metrics do not consider similarity between direct corresponding pairs of generated and ground truth images. From Table 2 one can observe that while both decoders perform similarly in reconstruction metrics (CRN is slightly better), SPADE dominates for the FID and inception score, indicating higher visual quality.

To evaluate our method in a fully generative setting, we mask the whole image and only use the encoded features ϕ_i for each object. We compare against the state of the art in interactive scene generation (ISG) [1], evaluated in the same setting. Since our main focus is on semantically rich relations, we trained [1] on Visual Genome, utilizing their publicly available code. Table 2 shows comparable recon-

Method	Decoder	All pixels					RoI only	
		MAE ↓	SSIM ↑	LPIPS ↓	FID ↓	IS ↑	MAE ↓	SSIM ↑
ISG [1] (Generative, GT)	Pix2pixHD	46.44	28.10	0.32	58.73	6.64±0.07	-	-
Ours (Generative, GT)	CRN	41.57	33.9	0.34	89.55	6.03±0.17	-	-
Ours (Generative, GT)	SPADE	41.88	34.89	0.27	44.27	7.86±0.49	-	-
<hr/>								
Cond-sg2im [16] (GT)	CRN	14.25	84.42	0.081	13.40	11.14±0.80	29.05	52.51
Ours (GT) w/o ϕ_i	CRN	9.83	86.52	0.073	10.62	11.45±0.61	27.16	52.01
Ours (GT) w/ ϕ_i	CRN	7.43	88.29	0.058	11.03	11.22±0.52	20.37	60.03
Ours (GT) w/o ϕ_i	SPADE	10.36	86.67	0.069	8.09	12.05±0.80	27.10	54.38
Ours (GT) w/ ϕ_i	SPADE	8.53	87.57	0.051	7.54	12.07±0.97	21.56	58.60
<hr/>								
Ours (P) w/o ϕ_i	CRN	9.24	87.01	0.075	18.09	10.67±0.43	29.08	48.62
Ours (P) w/ ϕ_i	CRN	7.62	88.31	0.063	19.49	10.18±0.27	22.89	55.07
Ours (P) w/o ϕ_i	SPADE	13.16	84.61	0.083	16.12	10.45±0.15	32.24	47.25
Ours (P) w/ ϕ_i	SPADE	13.82	83.98	0.077	16.69	10.61±0.37	28.82	49.34

Table 2: **Image reconstruction on Visual Genome.** We report the results using ground truth scene graphs (GT) and predicted scene graphs (P). (Generative) indicates experiments in full generative setting, *i.e.* the whole input image is masked out.

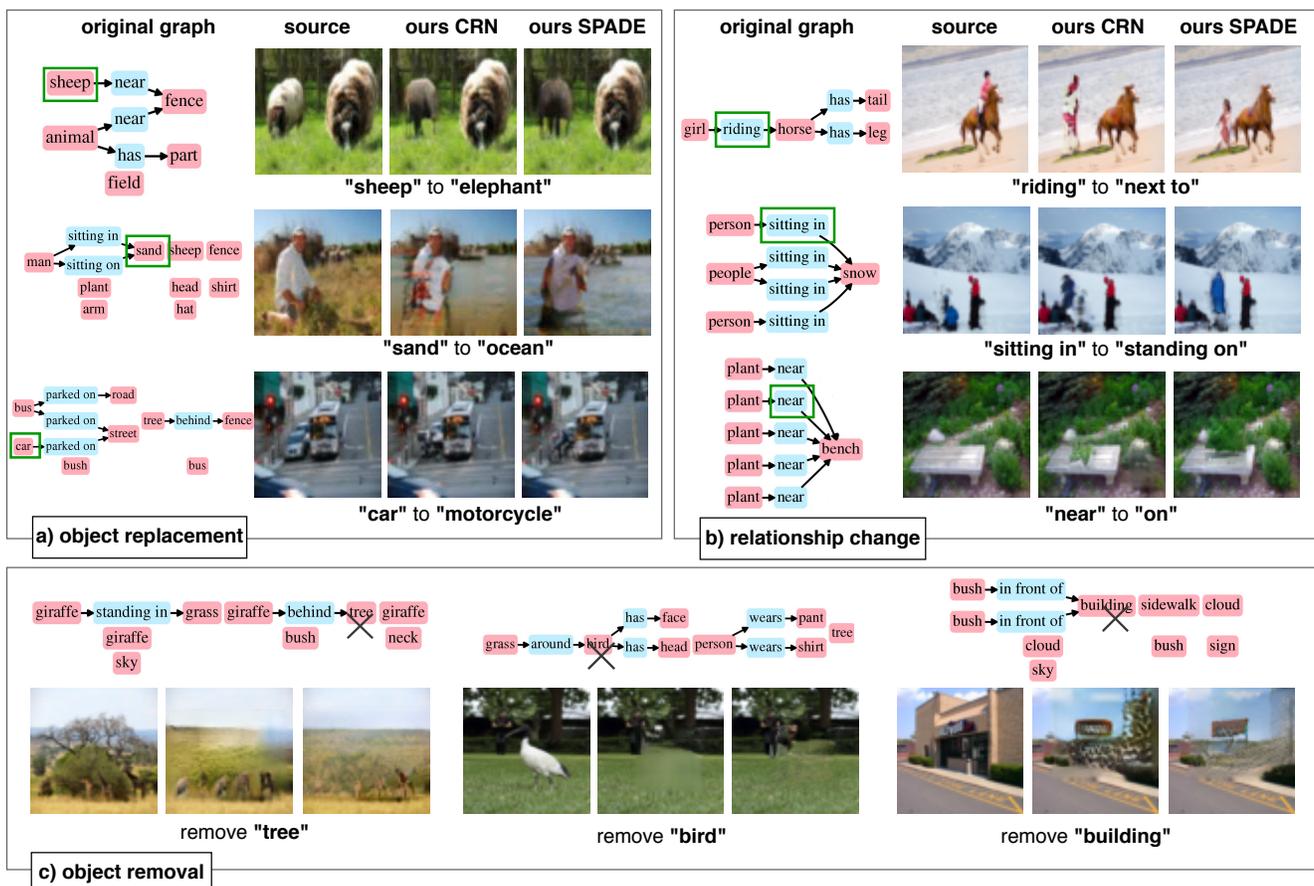


Figure 5: **Image manipulation** Given the source image and the GT scene graph, we semantically edit the image by changing the graph. **a)** object replacement, **b)** relationship changes, **c)** object removal. Green box indicates the changed node or edge.

struction errors for the generative task, while we clearly outperform [1] when a source image is given. This motivates

our choice of directly manipulating an existing image, rather than fusing different node features, as parts of the image

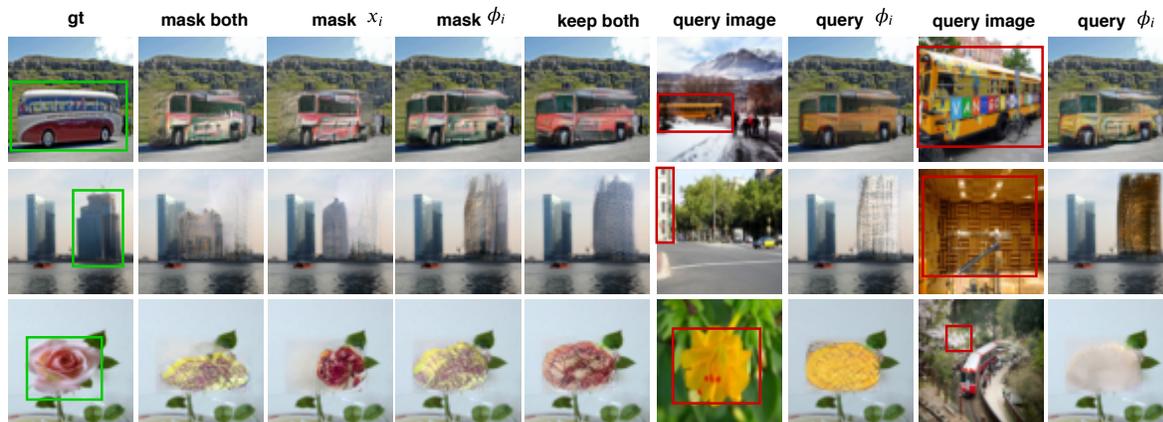


Figure 6: **Ablation of the method components** We present all the different combinations in which the method operates - *i.e.* masked vs. active bounding boxes x_i and/or visual features ϕ_i . When using a query image, we extract visual features of the object annotated with a red bounding box and update the node of an object of the same category in the original image.

need to be preserved. Inception score and FID mostly depend on the decoder architecture, where SPADE outperforms Pix2pixHD and CRN.

Figure 4 illustrates qualitative examples. It can be seen that both our method and the cond-sg2im baseline, generate plausible object categories and shapes. However, with our approach, visual features from the original image can be successfully transferred to the output. In practice, this property is particularly useful when we want to re-position objects in the image without changing their identity.

Main task: image editing. We illustrate visual results in three different settings in Figure 5 — object removal, replacement and relationship changes. All image modifications are made by the user at test time, by changing nodes or edges in the graph. We show diverse replacements (a), from small objects to background components. The novel entity adapts to the image context, *e.g.* the ocean (second row) does not occlude the person, which we would expect in standard image inpainting. A more challenging scenario is to change the way two objects interact, which typically involves re-positioning. Figure 5 (b) shows that the model can differentiate between semantic concepts, such as *sitting vs. standing* and *riding vs. next to*. The objects are rearranged meaningfully according to the change in relationship type. In the case of object removal (c), the method performs well for backgrounds with uniform texture, but can also handle more complex structures, such as the background in the first example. Interestingly, when the building on the rightmost example is removed, the remaining sign is improvised standing in the bush. More results are shown in the Appendix.

Component ablation. In Figure 6 we qualitatively ablate the components of our method. For a certain image, we mask

out a certain object instance which we aim to reconstruct. We test the method under all the possible combinations of masking bounding boxes x_i and/or visual features ϕ_i from the augmented graph representation. Since it might be of interest to in-paint the region with a different object (changing either the category or style), we also experiment with an additional setting, in which external visual features ϕ are extracted from an image of the query object. Intuitively, masking the box properties leads to a small shift in the location and size of the reconstructed object, while masking the object features can result in an object with a different identity than that in the original image.

5. Conclusion

We have presented a novel task — semantic image manipulation using scene graphs — and have shown a novel approach to tackle the learning problem in a way that does not require training pairs of original and modified image content. The resulting system provides a way to change both the content and relationships among scene entities by directly interacting with the nodes and edges of the scene graph. We have shown that the resulting system is competitive with baselines built from existing image synthesis methods, and qualitatively provides compelling evidence for its ability to support modification of real-world images. Future work will be devoted to further enhancing these results, and applying them to both interactive editing and robotics applications.

Acknowledgements We gratefully acknowledge the Deutsche Forschungsgemeinschaft (DFG) for supporting this research work, under the project #381855581. Christian Rupprecht is supported by ERC IDIU-638009.

References

- [1] Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4561–4569, 2019. [2](#), [3](#), [6](#), [7](#)
- [2] Samaneh Azadi, Deepak Pathak, Sayna Ebrahimi, and Trevor Darrell. Compositional gan: Learning conditional image composition. *arXiv preprint arXiv:1807.07560*, 2018. [2](#)
- [3] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016. [2](#)
- [4] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1511–1520, 2017. [1](#), [2](#), [4](#)
- [5] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018. [2](#)
- [6] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3076–3086, 2017. [3](#)
- [7] Helisa Dhano, Nassir Navab, and Federico Tombari. Object-driven multi-layer scene decomposition from a single image. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, 2019. [2](#)
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, 2014. [2](#)
- [9] Roei Herzig, Moshiko Raboh, Gal Chechik, Jonathan Berant, and Amir Globerson. Mapping images to scene graphs with permutation-invariant structured prediction. In *Advances in Neural Information Processing Systems*, 2018. [3](#)
- [10] 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 *NIPS*, 2017. [5](#)
- [11] 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*, 2018. [1](#), [2](#)
- [12] 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 Trans. Graph.*, 32, 2013. [2](#)
- [13] 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](#)
- [14] Seong Jae Hwang, Sathya N Ravi, Zirui Tao, Hyunwoo J Kim, Maxwell D Collins, and Vikas Singh. Tensorize, factorize and regularize: Robust visual relationship learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1014–1023, 2018. [3](#)
- [15] Youngjoo Jo and Jongyoul Park. Sc-fegan: Face editing generative adversarial network with user’s sketch and color. *arXiv preprint arXiv:1902.06838*, 2019. [2](#)
- [16] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *CVPR*, 2018. [2](#), [3](#), [4](#), [5](#), [7](#)
- [17] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [5](#)
- [18] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4565–4574, 2016. [1](#)
- [19] J. Johnson, R. Krishna, M. Stark, L. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. [2](#)
- [20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [2](#)
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017. [5](#), [6](#)
- [22] Guillaume Lample, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic Denoyer, et al. Fader networks: Manipulating images by sliding attributes. In *Advances in Neural Information Processing Systems*, pages 5967–5976, 2017. [2](#)
- [23] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [2](#)
- [24] Yitong Li, Zhe Gan, Yelong Shen, Jingjing Liu, Yu Cheng, Yuxin Wu, Lawrence Carin, David Carlson, and Jianfeng Gao. Storygan: A sequential conditional gan for story visualization. *arXiv preprint arXiv:1812.02784*, 2018. [1](#), [2](#)
- [25] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 335–351, 2018. [2](#), [3](#)
- [26] 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. [3](#)
- [27] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *European Conference on Computer Vision*, 2016. [1](#), [3](#)
- [28] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. [2](#)
- [29] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *Advances in neural information processing systems*, pages 2171–2180, 2017. [2](#), [3](#)
- [30] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine*

- Learning-Volume 70*, pages 2642–2651. JMLR. org, 2017. 2
- [31] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *ICML*, 2017. 5
- [32] 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*, 2019. 4, 5
- [33] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [34] 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
- [35] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 2
- [36] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*, 2016. 1, 2
- [37] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1, 3
- [38] Mohammad Amin Sadeghi and Ali Farhadi. Recognition using visual phrases. In *CVPR*, 2011. 3
- [39] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016. 5
- [40] Rakshith R Shetty, Mario Fritz, and Bernt Schiele. Adversarial scene editing: Automatic object removal from weak supervision. In *Advances in Neural Information Processing Systems*, pages 7717–7727, 2018. 2
- [41] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 3
- [42] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixellcn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016. 2
- [43] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756, 2016. 2
- [44] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2018. 1, 2
- [45] 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. 3
- [46] Xinchen Yan, Jimei Yang, Kihyuk Sohn, and Honglak Lee. Attribute2image: Conditional image generation from visual attributes. In *European Conference on Computer Vision*, pages 776–791. Springer, 2016. 2
- [47] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 670–685, 2018. 3
- [48] Shunyu Yao, Tzu Ming Hsu, Jun-Yan Zhu, Jiajun Wu, Antonio Torralba, Bill Freeman, and Josh Tenenbaum. 3d-aware scene manipulation via inverse graphics. In *Advances in Neural Information Processing Systems*, 2018. 2
- [49] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5485–5493, 2017. 2
- [50] Guojun Yin, Lu Sheng, Bin Liu, Nenghai Yu, Xiaogang Wang, Jing Shao, and Chen Change Loy. Zoom-net: Mining deep feature interactions for visual relationship recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 322–338, 2018. 3
- [51] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018. 3
- [52] Fangneng Zhan, Hongyuan Zhu, and Shijian Lu. Spatial fusion gan for image synthesis. *arXiv preprint arXiv:1812.05840*, 2018. 2
- [53] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiao lei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5907–5915, 2017. 1, 2
- [54] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5
- [55] 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*, 2018. 1, 2
- [56] Bo Zhao, Bo Chang, Zequn Jie, and Leonid Sigal. Modular generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [57] Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. Image generation from layout. 2019. 1, 2, 3
- [58] Yanan Zhao, Brian L. Price, Scott Cohen, and Danna Gurari. Guided image inpainting: Replacing an image region by pulling content from another image. *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1514–1523, 2018. 2
- [59] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision*, pages 597–613. Springer, 2016. 2
- [60] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2017. 2