This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Scene Graph Expansion for Semantics-Guided Image Outpainting

Chiao-An Yang¹, Cheng-Yo Tan¹, Wan-Cyuan Fan¹, Cheng-Fu Yang¹ Meng-Lin Wu², Yu-Chiang Frank Wang¹ ¹ National Taiwan University, ² Qualcomm Technologies, Inc.

joeyang@ntu.edu.tw, {cy.ugo.tan, christine5200312, joeyy5588}@gmail.com

menglinw@qti.qualcomm.com, ycwang@ntu.edu.tw

Abstract

In this paper, we address the task of semantics-guided image outpainting, which is to complete an image by generating semantically practical content. Different from most existing image outpainting works, we approach the above task by understanding and completing image semantics at the scene graph level. In particular, we propose a novel network of Scene Graph Transformer (SGT), which is designed to take node and edge features as inputs for modeling the associated structural information. To better understand and process graph-based inputs, our SGT uniquely performs feature attention at both node and edge levels. While the former views edges as relationship regularization, the latter observes the co-occurrence of nodes for guiding the attention process. We demonstrate that, given a partial input image with its layout and scene graph, our SGT can be applied for scene graph expansion and its conversion to a complete layout. Following state-of-the-art layout-toimage conversions works, the task of image outpainting can be completed with sufficient and practical semantics introduced. Extensive experiments are conducted on the datasets of MS-COCO and Visual Genome, which quantitatively and qualitatively confirm the effectiveness of our proposed SGT and outpainting frameworks.

1. Introduction

Given an incomplete image or a partial image input, humans generally are able to picture the context of the corresponding complete version. Such reasoning skill is largely based on our prior experience and knowledge observed from diverse images and their semantics. In the scope of machine learning, the objective is typically applied for the task of image completion, aiming to generate or predict reasonable missing image regions based on the observed input. In the areas of computer vision and image processing, several content creation applications such as object removal editing [21], image panorama creation [30], texture cre-



Figure 1. **Illustration of semantics-guided image outpainting.** Our work can be divided into stages of (a) Scene Graph Expansion (SGE), (b) Scene Graph to Layout (G2L), and (c) Layout to Image (L2I) conversions. The blue node and red edges in the scene graph indicate the generated objects and relationship, respectively.

ation [24], and view expansion [29] are closely related to the aforementioned task.

Depending on where the missing parts are to be recovered, the task of image completion is typically divided into two categories, image inpainting (also known as image hole-filling) and outpainting (also known as image extrapolation). Compared to image inpainting, image outpainting needs to synthesize unknown regions in single-sided fashions and thus is considered to be more challenging. Based on image inpainting works [8, 16, 18, 27, 31], researchers advance local and global GAN [8], Partial Convolution [16], Gated Convolution [31] and edge information [18] for outpainting tasks [11, 17, 20, 23, 26, 29]. However, despite impressive performances, most existing approaches are not designed to predict novel semantic regions in the output images. That is, they mainly focus on extending the surrounding texture or completing the fractional objects, resulting in extrapolated image regions with repeating structures or patterns. It is not clear how to introduce novel semantics with reasonable relationships with the existing ones during outpainting. As a result, we choose to approach this challenging semantics-oriented image outpainting problem by modeling and manipulating images at the semantic level.

In order to tackle the above task, a scene graph would be a desirable representation due to their ability in describing the presence of semantic objects and their relationships in an image. Thus, based on recent works such as [9], [6] and [19], one can describe and categorize a given image into three levels. The first level is the image level, containing pixel-level information. The second one is the layout level, which describes the locations/sizes of the objects of interest, including their corresponding category labels. The final level is the scene graph level, which describes semantic objects and their relationships (e.g., *right of, throw*) in an image. The higher the level is, the more abstract and semantic information it would contain.

In this paper, we choose to decompose the semanticsguided image outpainting task into three stages, as depicted in Figure 1. Given the scene graph extracted from the partial image and its layout, the first stage of scene graph expansion (SGE) utilizes the proposed Scene Graph Transformer (SGT), which uniquely performs node and edge-level attention, for expanding the input scene graph. The following stage of G2L further transforms such an expanded scene graph into a complete layout. Finally, layout-to-image (L2I) models can be applied for producing the final image output. We note that both SGE and G2L stages utilize our proposed SGT module, taking scene graph data as inputs with unique objectives introduced to enforce the desirable object/relationship properties, as later discussed in Sect. 3.

The contributions of our work are highlighted as follows:

- We approach the task of semantics-guided image outpainting, which is able to synthesize novel yet semantically practical objects with associated relationships for completing an image output.
- We propose a Scene Graph Transformer (SGT), which takes node and edge features with unique node-level and edge-level attention mechanisms for modeling the associated structural information.
- Expecting the sparsity of the object relationships in a scene graph, our SGT is designed to exploit the converse relationships between objects, so that semantically practical nodes and their corresponding edges can be properly recovered or expanded.

2. Related Work

2.1. Image Outpainting

Adversarial learning [20] has been applied for image outpainting, generating image regions toward horizontal directions. By adopting a recurrent neural network, [29] extends the output image in a single direction with varying lengths. As for [17], it fills in the intermediate gap between left and right partial image inputs for outpainting purposes. Although the method of [26] allows outpainting in all four directions, they require extra information (i.e., the image margins) during both training and testings. While such requirements are later alleviated by [22], most existing works are only capable of extending background textural regions or mending fractional objects. It is not clear whether novel yet semantically practical image regions can be added to the output image. Recently, [11] proposes to outpaint images based on the extrapolated segmentation map, serving as guidance for generating novel object instances.

2.2. From Scene Graphs to Images

As noted in [10], scene graph is a data structure with each node encoding an object in the image, and each edge describing the associated relationship. Scene graph generation can be viewed as a task of image-to-text conversion. However, generating an image from a scene graph is a more challenging task, and is first tackled by [9] in an end-to-end learning fashion. Taking the image layout as an intermediate representation, one typically converts a scene graph to an image layout, followed by a layout-to-image conversion task. For scene graph to layout, [6] leverages the converse and transition property of relationships. [19] proposes Spade, an architecture for describing image semantic layouts. [6] extends Spade for manipulating the attributes of the generated objects.

With the recent advances of the Transformer [25], recent approaches like [2, 28, 32] utilize Transformer based architectures for handling scene graph data, either for scene graph generation or scene graph to layout generation. However, these methods cannot be easily applied for scene graph expansion, which is critical in our focus on semanticsguided image outpainting. Nevertheless, since the Transformer deals with sequential data, one needs to convert the input scene graph to a sequence of triplets, each consisting of a subject, a predicate, and an object. Moreover, since "no relation" would be also viewed as a predicate, describing a scene graph from an image would inevitably result in a large number of triplets. This would result in long triplet sequences, making the learning of the Transformer inefficient. Another potential problem for triplet representations is that if one object node has multiple relationship edges, the object node will appear in multiple triplets, which might result in redundant representations with inconsistent semantic outputs. In this paper, we propose an alternative yet novel architecture, Scene Graph Transformer (SGT). As detailed in the following section, our SGT would alleviate the aforementioned problems and can be applied for both scene graph expansion and scene graph to layout generation.



Figure 2. Scene Graph Transformer. (a) Node-level attention: attention across nodes under guidance of the associated edges. (b) Edge-level attention: attention across edges conditioned on the sharing nodes. Note that $W_v^n, W_s^n, W_v^e, W_s^e$ denote the MLP transformation layers for the corresponding feature modalities.

3. Methodology

3.1. Notations and Algorithmic Overview

Image outpainting. Given an incomplete image I^{in} of $h_1 \times w_1$ pixels, image outpainting is to generate an extrapolated image I^{op} of $h_2 \times w_2$ pixels with $h_2 > h_1$ and $w_2 > w_1$. During training, we have I^{in} partially cropped out from a complete image I^{gt} (of $h_2 \times w_2$ pixels), aiming at producing I^{op} to recover I^{gt} .

Scene graph and layout. To describe semantic information in an image, a scene graph S = (O, R) consists of a list of N objects (nodes) $O = \{o_i\}_{i=1:N}$ and the associated relationship (edge) matrix $R = (r_{ij}) \in \mathbb{R}^{N \times N}$, where o_i is the object label, and r_{ij} indicates the edge label between objects o_i and o_j . Note that r_{ij} belongs to $\{y_1^R, y_2^R, \cdots, y_M^R\} \cup \{0\}$, where each y_i^R denotes the relation label (e.g., *riding, wear, on,* etc.), and M is its label number. And, $r_{ij} = 0$ indicates no relationship between the corresponding object pair. On the other hand, a layout is a list of bounding boxes of each object in an image, i.e., $\mathbf{B} = \{b_i\}_{i=1:N}$ with each $b_i = (b_i^x, b_i^y, b_i^w, b_i^h)$ describing the center coordinates and the size of the bounding box. We also compute the bounding box disparity for

each relationship $\mathbf{D} = (\mathbf{d}_{ij}) \in \mathbb{R}^{N \times N \times 4}$, where each $\mathbf{d}_{ij} = \{b_i^x - b_j^x, b_i^y - b_j^y, \log(b_i^w/b_j^w), \log(b_i^h)/\log(b_j^h)\}$ describes the spatial displacement between the bounding boxes of each subject-object pair.

Algorithmic overview. To perform semantic-guided image outpainting, our model would introduce novel object instances with realistic relationships with semantic practicality, which can be decomposed into the following three stages: The scene graph expansion (SGE) stage deploys the Scene Graph Transformer based on incomplete images I^{in} with their layouts $L^{in} = (B^{in}, D^{in})$ and scene graphs $\mathcal{S}^{in} = (\mathbf{O}^{in}, \mathbf{R}^{in})$, so that the model T_{SGE} would expand S^{in} into $S^{op} = (O^{op}, R^{op})$. In the stage of scene graph to layout (G2L), we learn a second SGT-based model T_{G2L} which converts the expanded scene graph into layout L^{op} under the guidance of I^{in} . Finally, for the layout to image (L2I) stage, we produce the final outpainted image I^{op} via the model G_{L2I} . While not being the main focus of this work, our model G_{L2I} is based on SPADE [19] resnet blocks and consists of an image encoder and a generator.

3.2. Scene Graph Transformer

In this paper, we propose a novel architecture of Scene Graph Transformer (SGT), which is particularly designed to handle graph-structured data. With the ability to describe the nodes and their relationships in an image scene graph, our SGT performs separate yet mutually related selfattention between node levels and edge levels. That is, SGT views edges in scene graphs as regularization during the self-attention between different nodes, while the cooccurrence of nodes would guide the self-attention across different edges. Since both stages of SGE and G2L in our outpainting task take scene graph data as the inputs, our SGT will be utilized in both stages with objectives properly introduced and enforced.

For the sake of completeness, we briefly review the standard Transformer and explain how it can be applied to handle graph-structured data with N nodes. As a sequenceto-sequence model, the Transformer consists of multiple transformation layers mapping an input sequence H = $\{h_i\}_{i=1:3N^2}$ to the output $\hat{H} = \{\hat{h}_i\}_{i=1:3N^2}$. Note that, with N nodes and $N^2 = N \times N$ edges in the input graph, the Transformer in [2, 28, 32] needs to convert such input data into a sequence, whose length is at least $3N^2$ due to the triplet representation "subject-predicate-object". For each transformer layer, the input vector h is first converted into the query vector q, the key vector k, and the value vector v through an MLP layer. The output vector h is computed as the weighting sum of the value vector v_i , i.e., $h_i =$ $\sum_{j} s_{ij} \boldsymbol{v}_{j}$, with the weight $s_{ij} = softmax(\boldsymbol{q}_{i} \cdot \boldsymbol{k}_{j}/\sqrt{d_{\mathbf{k}}})$, where $d_{\mathbf{k}}$ is the dimension of k, and \cdot stands for the inner product operation.

Instead of viewing the scene graph as a single sequence

of triplets, the transformation layers in our SGT consider node (object) and edge (relationship) features as distinct yet mutually related data modalities. Thus, we have input and output of node feature sequences denoted as $H^n =$ $\{h_i^n\}_{i=1:N}$ and $\hat{H}^n = \{\hat{h}_i^n\}_{i=1:N}$, respectively. As for those of the edge feature matrices, they are denoted as $H^e = \{h_{ij}^e\}_{i,j=1:N}$ and $\hat{H}^e = \{\hat{h}_{ij}^e\}_{i,j=1:N}$. For each modality, we deploy unique attention mechanisms based on the scene graph structure, as we present below.

3.2.1 Node-level attention

The first type of attention in our SGT is performed at the node level, while the cross-attention between nodes is enforced by the edge relationships observed. Recall that, for the standard Transformer, it simply "flattens" the scene graph as a sequence of $(node_i - edge_{ij} - node_j)$ triplets with its attention mechanism not distinguishing between data modalities, nor considering the intrinsic graph structure.

With the scene graph nodes as the inputs, our SGT calculates the similarity between node features h_i^n and h_j^n under the guidance of the associated edge feature h_{ij}^e , with the output for that node \hat{h}_i^n as the weighted summation of the *values* across each node j. Thus, we have

$$\hat{\boldsymbol{h}}_{i}^{n} = \sum_{j} \boldsymbol{s}_{ij}^{n} \odot \boldsymbol{v}_{j}^{n}, \qquad (1)$$

where v_j^n is the value features for node j, s_{ij}^n indicates the *attention weight* derived from each triplet with node i (i.e., node_i, edge_{ij}, and node_j), and \odot denotes an element-wise multiplication.

As depicted in Figure 2(a), the above calculation allows the edges associated with the node of interest to be incorporated into the attention process, which effectively regularizes the attention across nodes based on their corresponding relationships. To provide additional details, we calculate the value vector v_i^n for each node h_i^n through a single MLP W_v^n . Instead of utilizing query or key vectors for calculating the attention weight s_{ij}^n , we take the triplet features of node_i-edge_{ij}-node_j by concatenating their representations as $t_{ij}^n = h_i^n \oplus h_j^n \oplus h_{ij}^e$. With another MLP W_s^n taking t_{ij}^n as the input, the output weight vector s_{ij}^n thus attends across h_i^n and h_j^n , i.e. $s_{ij}^n = W_s^n(t_{ij}^n)$.

As a final remark, we do not follow the standard Transformer for using inner-product followed by *softmax* to produce the attention weight. This is because our edge-regularized attention mechanism provides guidance of structural information, and the use of inner-product operation would dilute such information. Thus, we have the output vector \hat{h}_i^n as the summation of element-wise multiplication between s_{ij}^n and v_j^n , as shown in Equation (1).

3.2.2 Edge-level attention

For the edge features H^e in a scene graph, only h^e_{ij} contributes to the computation of $t^n_{ij} = h^n_i \oplus h^n_j \oplus h^e_{ij}$. If one simply performs cross-attention on t^n_{ij} , similar nodes would imply and result in the same \hat{h}^e_{ij} , which is viewed as the *edge collapse* problem, i.e., resulting in repeating or redundant edges related to the same node *i*. For example, it is possible that, for node *j* and node *k* both linking to node *i*, the same \hat{h}^e_{ij} and \hat{h}^e_{ik} are produced (e.g., two *men* hold the same *tennis racket*).

To tackle the above problem, we propose edge-level attention in our SGT, while the cross-attention between edges is regularized by the nodes sharing the edge of interest, as illustrated in Figure 2(b). To exploit inter-edge information and to take the shared nodes into consideration, we have an input edge feature h_{ij}^e with the node pair i, j, and we consider edges linking to either node i or j for attention. Thus, we have the features h_{kl}^e for such edges expressed as $\{h_{kl}^e|k=i \lor l=j\}$. And, the triplet feature for edge-level attention is computed as follows:

$$\boldsymbol{t}_{ij,kl}^{e} = \boldsymbol{h}_{ij}^{e} \oplus \boldsymbol{h}_{kl}^{e} \oplus \begin{cases} \boldsymbol{h}_{i}^{n}, \text{ if } k = i \\ \boldsymbol{h}_{j}^{n}, \text{ if } l = j. \end{cases}$$
(2)

Instead of having the resulting edge-level attention matrix as $N^2 \times N^2 = N^4$, only N^2 (edge number) $\times 2N$ (*N subjects* +*N objects*)= $2N^3$ edge pairs need to be considered. This greatly reduces the computation load when comparing to the use of the standard Transformer to perform attention across all edges in a graph.

The remaining attention mechanism follows that of node-level attention discussed earlier. As depicted in Figure 2(b), the above calculation allows the nodes associated with the edge of interest to be incorporated into the attention process, which effectively regularizes attention across edges based on their shared nodes. To provide further details, we calculate the value vector v_{kl}^e for each edge h_{kl}^e through a single MLP W_v^e . The "edge triplet feature" of $edge_{ij}$ -shared node- $edge_{kl}$ is obtained by concatenating their representations as shown by Equation 2. With an MLP W_s^e taking $t_{ij,kl}^e$ as the input, the output weight vectors $s_{ij,kl}^e$ thus attends across h_{ij}^e and h_{kl}^e , i.e. $s_{ij,kl}^e = W_s^e(t_{ij,kl}^e)$.

3.3. Semantic-Guided Image Outpainting

3.3.1 Scene Graph Expansion

Aiming at expanding the scene graph extracted from the input image, our SGT-based SGE model T_{SGE} learns to append novel object nodes with the associated and necessary relationship edges introduced. Inspired by masked language model [4], we train this model by observing a complete scene graph $S^{gt} = (O^{gt}, R^{gt})$, with a number of objects in O^{gt} being masked with a special token [MASK] assigned.



Figure 3. Flowcharts for (a) SGE and (b) G2L. Note that the colors for the feature *cubes* indicate their data modalities, and those in gray and dark-color denote the *masked* and *generated* ones. For SGE, the object and relationship classes are used as node and edge inputs, respectively. For G2L, we have the concatenated features of object class, bounding box, and the associated visual features as the node input, while the concatenated relationship label and its bounding box disparity as the edge inputs.

Subsequently, the relationships in \mathbf{R} linking to the masked node (either a *subject* or a *object*) will also be *masked*. This results in the partial input scene graph $S^{in} = (\mathbf{O}^{in}, \mathbf{R}^{in})$.

In order to perform node and edge-level attention, our SGE model T_{SGE} contains object and relationship embedding encoders E_O and E_R to extract features from nodes and edges, object and relationship classifiers C_O and C_R for recognizing the derived output features, as shown in Fig. 3(a). That is, T_{SGE} takes the object category word embeddings $f_i^O = E_O(o_i^{in})$ as the node inputs h_i^n , and the relationship category word embedding $f_{ij}^R = E_R(r_{ij}^{in})$ as the relationship matrix input h_{ij}^e . C_O learns to predict the object class $o_i^{op} = C_O(\hat{h}_i^n)$, and the relationship classifier C_R predicts the relationship label $r_{ij}^{op} = C_R(\hat{h}_{ij}^e)$. From the above process, the objective of training T_{SGE} is to recover the complete scene graph $S^{op} = (O^{op}, R^{op})$ from S^{in} . Thus, the objective can be summarized below:

$$\mathcal{L}_{SGE} = \sum_{i} \mathcal{L}_{CE}(o_i^{op}, o_i^{gt}) + \sum_{i,j} \mathcal{L}_{CE}(r_{ij}^{op}, r_{ij}^{gt}), \quad (3)$$

where \mathcal{L}_{CE} indicates the cross-entropy classification loss.

Exploitation of converse relationships. As noted in Sect. 3.1, $\mathbf{R} = (r_{ij}) \in \mathbb{R}^{N \times N}$ denotes the relationships between each object pair in the scene graph. However, this matrix is *not* necessarily expected to be a symmetric matrix, since $r_{ij} = y^R$ and $r_{ji} = \tilde{y}^R$ are viewed as relational antonyms and thus with *converse relationships*, even both edges are connected to the same node pair. Given an input

scene graph, typically only one of such relationship pairs would be observed. Thus, the above converse relationship can be implicitly inferred when either r_{ij} or r_{ji} is presented, resulting in the relationship matrix towards skew-symmetric (i.e., $r_{ji} = \tilde{r}_{ij}$). In practice, only a limited number of relationships would be specified in a scene graph, one thus observes a sparse ground truth relationship matrix \mathbf{R}^{gt} , lacking converse relationship pairs. Also, one of the attention mechanisms introduced in our SGT is node-level attention, which is specifically guided by the relationship between the nodes of interest. Without properly generating and observing the aforementioned converse relationship pairs, the attention would be partially biased and result in undesirable outputs. The above challenges make the learning of the SGE model T_{SGE} very difficult.

To tackle the above problem, we choose to process \mathbf{R}^{gt} as follows. For each non-*empty* $r_{ij}^{gt} = y^R$, we manually assign the converse label \tilde{y}^R to the associated empty r_{ji}^{gt} . (e.g., *converse-riding* vs. *riding*, and *converse-on* vs. *on*). It is worth noting that, the above label processing is for training T_{SGE} only, not for later G2L and L2I training purposes.

Furthermore, to enforce the one-to-one mapping between a relationship and its converse version, we deploy an additional feature converter E_C which takes the input relationship $E_R(y^R)$ and produces its converse version. This allows the classifier C_R to predict its lalel \tilde{y}^R . Thus, E_C is trained with the classification loss: $\mathcal{L}_{conv} = \sum_i \mathcal{L}_{CE}(C_R \circ E_C \circ E_R(y^R_i), \tilde{y}^R_i)$.

With converse relationship enforced, the skew-symmetry property of the SGE model can be expected, which can be



Figure 4. Illustration of our L2I Model G_{L2I} . Based on AttSpade [6], the decoder of G_{L2I} (i.e., G_{SPADE}) takes semantic maps as guidance for producing image outputs.

calculated by the following loss function:

$$\mathcal{L}_{sym} = \sum_{ij} \mathcal{L}_{CE}(C_R \circ E_C \circ E_R(r_{ij}^{op}), r_{ji}^{gt}).$$
(4)

Note that $r_{ij}^{op} = C_R(h_{ij}^e)$ denotes the relationship label derived from the output relationship h_{ij}^e . Finally, our T_{SGE} is trained with the combination of Equations (3) and (4).

3.3.2 Scene Graph to Layout (G2L)

Given a partial input image I^{in} with the corresponding layout L^{in} , together with the expanded scene graph S^{op} , the second stage of our work is to learn a G2L model T_{G2L} for generating the plausible layout L^{op} . Based on the architecture of SGT but different from T_{SGE} , our T_{G2L} considers a bounding box encoder E_B with a regressor R_B , and a disparity encoder E_D with a regressor R_D . Moreover, as shown in Figure 3(b), an image encoder E_I is deployed to distinguish whether a non-masked object is with missing parts. For example, it is possible the I_{in} consists of a horse with its legs cropped out of the image, thus with a smaller incomplete L_{in} . By feeding T_{G2L} with the visual feature of I_{in} , it is expected to attend the partial horse and thus expand its incomplete bounding boxes accordingly.

Utilizing SGT, the input layout for T_{G2L} is also described as a graph $(\boldsymbol{H}^n, \boldsymbol{H}^e)$. Each node \boldsymbol{h}_i^n is obtained by concatenating the object category embedding $\boldsymbol{f}_i^O = E_O(o_i)$, bounding box feature $\boldsymbol{f}_i^B = E_B(\boldsymbol{b}_i)$, and the visual feature \boldsymbol{f}_i^I , i.e. $\boldsymbol{h}_i^n = \boldsymbol{f}_i^O \oplus \boldsymbol{f}_i^B \oplus \boldsymbol{f}_i^I$. Note that \boldsymbol{f}_i^I can be directly obtained by cropping out the associated region from the input feature map $\boldsymbol{f}^I = E_I(\boldsymbol{I})$. As for the edges in \boldsymbol{H}^e , each edge input \boldsymbol{h}_{ij}^e is obtained by concatenating the relationship category embedding $\boldsymbol{f}_{ij}^R = E_R(r_{ij})$ and the disparity feature $\boldsymbol{f}_{ij}^D = E_D(\boldsymbol{d}_{ij})$, i.e. $\boldsymbol{h}_{ij}^e = \boldsymbol{f}_{ij}^R \oplus \boldsymbol{f}_{ij}^D$. We note that the regressor R_B in T_{G2L} predicts bound-

we note that the regressor R_B in T_{G2L} predicts bounding box information. If node *i* denotes a novel/masked object, the regressor is trained to predict the bounding boxes b_i^{op} , under the supervision of ground truth b_i^{gt} . Otherwise, it would predict the boundary offset $b_i^{off} = b_i^{gt} - b_i^{in}$ (i.e., top, bottom, left, and right) between the input object and the ground truth one. As for the regressor R_D , it is deployed to maintain the consistency between the node outputs and edge outputs, and is trained to predict the bounding box disparities d_{ij}^{op} under the supervision of ground truth d_{ij}^{gt} . With the above definitions, we train T_{G2L} with the following loss:

$$\mathcal{L}_{G2L} = \sum_{i, \text{ if } o_i^{in} \neq mask} \mathcal{L}_{cIoU}(\boldsymbol{b}_i^{off} + \boldsymbol{b}_i^{in}, \boldsymbol{b}_i^{gt}) \\ + \sum_{i, \text{ if } o_i^{in} = mask} \mathcal{L}_{cIoU}(\boldsymbol{b}_i^{op}, \boldsymbol{b}_i^{gt}) + \sum_{i,j} |\boldsymbol{d}_{ij}^{op} - \boldsymbol{d}_{ij}^{gt}|,$$
(5)

where $\mathcal{L}_{cIoU}(\cdot)$ is the complete-IoU loss utilized in [34].

3.3.3 Layout to Image (L2I)

With the expanded scene graph and layout, our final stage is to perform layout to image conversion. Adapted from AttSpade [6], our L2I model G_{L2I} learns to outpaint the partial input image into I^{op} , conditioned on S^{op} and L^{op} . To enforce visual consistency, we choose to concatenate the image feature $f^{I} = E_{I}(I^{in})$ with the layout feature map f^{L} to form the semantic information map f^{S} . This allows our model to generate a realistic output image with the guidance of f^{S} through layers of SPADE blocks. Since the ground truth images are available during training, in addition to the adversarial loss, we are able to train G^{L2I} with the reconstruction loss between I^{op} and I^{gt} .

It is worth repeating that, since we focus on the design of SGT (and its use for SGE and G2L), producing highquality image outputs is *not* within the main scope of this work. Thus, AttSpade-based designs can be replaced with state-of-the-art image conversion models if desirable.

4. Experiments

4.1. Datasets

We evaluate our proposed methods on scene-level image datasets with bounding box annotation, namely COCOstuff [1, 15], VG-MSDN [13, 14] and CityScapes [11]. Please see the supplementary materials for more details.

4.2. Evaluation and Analysis

Scene graph expansion. To compare the output expanded scene graph \mathbb{S}^{op} to the ground truth \mathbb{S}^{gt} , we report the metrics of the averaged rank of correct prediction (rAVG) and the top-k accuracy (Hits@k) for both object and relationship predictions, respectively. Note that We ignore the "empty" relationship in \mathcal{S}^{gt} for accuracy calculation due to the sparsity expected for scene graphs.

To assess that, compared to the training of masked language models (MLM), whether our proposed SGT learning

The authors from NTU downloaded, evaluated, and completed the experiments on the datasets.

		VG-MSDN				COCO-stuff			
	Masking strategy	Object		Relation		Object		Relation	
		$rAVG\downarrow$	Hit@ 1/5↑	$rAVG\downarrow$	Hit@ 1 / 5 ↑	$\mathrm{rAVG}\downarrow$	Hit@ 1/5↑	$rAVG\downarrow$	Hit@ 1 / 3 ↑
Transformer	М	28.96	9.55 / 29.7	5.32	37.4 / 70.1	31.14	11.1 / 29.7	2.41	28.7 / 78.2
LTNet	М	24.72	9.86 / 36.4	4.62	42.7 / 74.8	30.76	12.0/30.3	2.40	27.9 / 77.8
GTwE	М	10.93	28.3 / 58.8	5.26	34.9 / 73.1	12.25	24.6 / 54.1	2.92	20.1 / 61.9
SGT	М	9.40	34.7 / 64.5	3.92	48.7 / 81.7	11.32	26.0 / 58.5	2.20	36.4 / 81.5
Transformer	Е	33.77	10.6 / 28.9	5.30	35.3 / 65.8	22.35	14.7 / 37.8	2.37	29.4 / 78.5
LTNet	Е	24.45	13.9 / 34.8	4.70	34.8 / 74.6	17.22	20.1 / 45.8	2.36	29.1 / 78.4
GTwE	Е	11.91	27.0 / 57.2	5.36	35.8 / 72.5	11.81	28.4 / 57.2	2.89	20.4 / 63.3
SGT	E	8.38	39.7 / 68.9	3.43	55.3 / 84.3	11.03	29.6 / 59.0	2.19	45.5 / 82.2

Table 1. Quantitative evaluation on scene graph expansion. Note that masking strategies of M and E denote uses of standard MLM and our expansion-based learning schemes, as described in Sect. 4.2.

Table 2. Quantitative Results on scene graph to layout. The (\checkmark / \varkappa) indicates whether the method is trained and test with or without input image and layout guidance.

		VG-MSDN	COCO-stuff
	L^{in}	mIoU	mIoU
Sg2Im	X	16.8	29.0
Canonical	X	18.0	41.9
LTNet	X	18.3	49.0
SGT	×	25.3	54.6
Transformer	1	5.1 / 71.2 / 51.9	10.4 / 75.7 / 61.2
GCN	1	11.4 / 70.6 / 50.0	21.1 / 72.3 / 60.8
GTwE	1	12.3 / 79.9 / 62.1	21.3 / 73.2 / 64.8
SGT	1	14.5 / 81.1 / 62.4	28.2 / 85.1 / 74.9

strategy would favor the task of SGE, we consider/compare the following two training schemes. First (and as our proposed one), one object randomly is removed during training (including its corresponding edges). For the second case, we follow an existing MLM work [32] to randomly mask $\alpha = 30\%$ of the total objects and relationships for training the SGE model.

In addition to the standard transformer, we apply LT-Net [28] and GTwE (Graph Transformer with Edge features [5]) for comparison purposes, and the results are listed in Table 1. From this table, it can be seen that our T_{SGE} performed against baseline and state-of-the-art models on both VG-MSDN and COCO with clear margins, for both object generation and relationship prediction. More specifically, between the two training schemes, we found that the setting with a single object (and its edges) intentionally and randomly removed resulted in more effective performance. Since this setting is consistent with the expansion task for scene graphs, the use of SGT-based T_{SGE} can be verified. Selected visualization examples of SGE are shown and compared in Fig. 5(a). Scene graph to layout. To evaluate the performance for this stage, we measure the mIoU between the output layout L^{op} and the ground truth L^{gt} . We note that depending on the object is an introduced or an existing one, we show their mIoU separately, with the total mIoU as their weighted average. We compare our SGT-based T_{G2L} with two other models: Transformer [25], GCN [6, 12] and GTwE [5]. Table 2 lists the performances of the above methods, and we see that our model consistently performed against GCN and Transformer across different settings. From the visual examples shown in Figure 5, it can be observed that our model better comprehends the expanded scene graph and the given partial input, so that the predicted layout would be more semantically practical.

Furthermore, we demonstrate the robustness of our model which is trained without any input image or layout guidance. This is thus consistent with the setting used in Sg2Im [9], Canonical [6], and LTNet [28]. With no input guidance except for our expanded scene graphs, all objects would be considered *masked*(new) and thus only one mIoU score is reported. From the results shown in Table 2, one can see that our model still achieved the highest mIoU and thus would be preferable in such scenarios.

Semantic-guided image outpainting. Finally, we evaluate the performance for semantics-guided image outpainting. From the visual examples shown in Figure 5(c) and (d). We see that, under the guidance of the expanded scene graph and completed layout, our model better generates novel object instances at the pixel level, e.g., a *grassland* (second row of Fig. 5(c)), *sky* (first row of Fig. 5(d)), or expands an existing object to a reasonable size, e.g., the jaw of the bear (first row of Fig. 5(c)). Additional visualization results are provided in the supplementary materials.

We note that, most image outpainting works consider images in restrained scenes (e.g., Cityscapes [11] and ADE20K [11]) or those with single category object (e.g., CUB [26], CelebA [26, 33] or DeepFashion [26]). And,



Figure 5. Visualization examples of SGE, G2L, and image outpainting. From left to right: (a) input scene graph S^{in} , output scene graphs S^{op} from LTNet, GTwE, and ours. Nodes and edges in green denote correct predictions, while those in blue are semantically practical but differ from the ground truth ones. Finally, those in red denote incorrect predictions. (b) input layout L^{in} , input scene graph S^{gt} , output layouts L^{op} from GCN and Ours, and ground truth L^{gt} . Bounding boxes from novel (generated) objects are denoted in blue, while the existing ones are shown in green. (c) and (d): input image I^{in} , output images I^{op} from Boundless, AttSpade, and ours. Note that we also highlight selected nodes and their bounding boxes following the protocol in (b).

to the best of our knowledge, we are the first outpaint image data in the wild with rich interaction between various categories of objects (e.g., VG-MSDN and COCO-stuff). Therefore, only limited quantitative comparisons can be conducted. Specifically, we consider Cityscapes [3] and take the Fréchet inception distance (FID) [7] as the metric. Our model reported FID of 60.99, which surpassed Outpainting-SRN [26] at 66.89, Boundless [23] at 77.86, and a modified AttSpade [6] at 68.91 (equivalent to our G_{L2I} only). While a recent work of SemIE [11] reported an improved FID score of 47.67, it is designed for restrained street-view (Cityscapes) or indoor scenes (ADE20K [35]), and cannot be easily applied to outpaint image data in the wild as ours does. Another requirement of SemIE is the use of segmentation masks as learning guidance, while we only require guidance at scene graph levels. Thus, the effectiveness and practicality of our proposed model can be verified.

Ablation studies. To assess the design of our SGT, we consider VG-MSDN and report the performance on SGE. For Hits@1, the baseline SGT with only node-level attention reported 35.7/46.1 on object/relationship prediction while adding edge-level attention and regularization of skew-symmetry result in 38.7/48.6 and 38.2/52.0, respectively. Finally, our SGT with full objectives achieved 39.7/55.3, which confirms its design and learning schemes. More details can be found in the supplementary materials.

5. Conclusions

We address the task of semantics-guided image outpainting by proposing a novel Scene Graph Transformer (SGT). By decomposing the task into the stages SGE, G2L, and L2I, our proposed model leverages information observed from the nodes and edges in the partial input scene graph, inferring plausible object co-occurrences, and thus producing the final image output. Our SGT uniquely performs attention at both node and edge levels for modeling input structural information. In addition, for completing a semantically practical image, our SGT exploits converse relationships between edges for scene graph expansion. Our experiments confirmed that our proposed SGT performs favorably against state-of-the-art transformer-based models on both SGE and G2L. With novel objects and their relationships introduced, satisfactory image outputs can be achieved.

Acknowledgement This work is supported in part by the Ministry of Science and Technology of Taiwan under grant MOST 110-2634-F-002-036 and in part by Qualcomm Technologies, Inc. through a Taiwan University Research Collaboration Project. We also thank to National Center for High-performance Computing (NCHC) for providing computational and storage resources.

References

- Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocostuff: Thing and stuff classes in context. In *CVPR*, 2018.
 6
- [2] Yuren Cong, Wentong Liao, Hanno Ackermann, Michael Ying Yang, and Bodo Rosenhahn. Spatialtemporal transformer for dynamic scene graph generation. *CoRR*, 2021. 2, 3
- [3] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In CVPR, 2016. 8
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In NAACL-HLT, 2019.
 4
- [5] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *CoRR*, 2020. 7
- [6] Roei Herzig, Amir Bar, Huijuan Xu, Gal Chechik, Trevor Darrell, and Amir Globerson. Learning canonical representations for scene graph to image generation. In *ECCV*, 2020. 2, 6, 7, 8
- [7] 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. 8
- [8] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. 2017. 1
- [9] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In CVPR, 2018. 2, 7
- [10] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *CVPR*, 2015. 2
- [11] Bholeshwar Khurana, Soumya Ranjan Dash, Abhishek Bhatia, Aniruddha Mahapatra, Hrituraj Singh, and Kuldeep Kulkarni. Semie: Semantically-aware image extrapolation. In *ICCV*, 2021. 1, 2, 6, 7, 8
- [12] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
 7
- [13] 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. 2017. 6
- [14] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. 6
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In ECCV, 2014. 6
- [16] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In ECCV, 2018. 1

- [17] Chia-Ni Lu, Ya-Chu Chang, and Wei-Chen Chiu. Bridging the visual gap: Wide-range image blending. In *CVPR*, 2021.
 1, 2
- [18] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Structure guided image inpainting using edge prediction. In *ICCV Workshops*, 2019.
- [19] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 2, 3
- [20] Mark Sabini and Gili Rusak. Painting outside the box: Image outpainting with gans. *CoRR*, 2018. 1, 2
- [21] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In WACV, 2022. 1
- [22] Cheng-Yo Tan, Chiao-An Yang, Shang-Fu Chen, Meng-Lin Wu, and Yu-Chiang Frank Wang. Robust image outpainting with learnable image margins. In *ICIP*, 2021. 2
- [23] Piotr Teterwak, Aaron Sarna, Dilip Krishnan, Aaron Maschinot, David Belanger, Ce Liu, and William T Freeman. Boundless: Generative adversarial networks for image extension. In *ICCV*, 2019. 1, 8
- [24] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *ICML*, 2016. 1
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 2, 7
- [26] Yi Wang, Xin Tao, Xiaoyong Shen, and Jiaya Jia. Widecontext semantic image extrapolation. In *CVPR*, 2019. 1, 2, 7, 8
- [27] Wei Xiong, Jiahui Yu, Zhe Lin, Jimei Yang, Xin Lu, Connelly Barnes, and Jiebo Luo. Foreground-aware image inpainting. In CVPR, 2019. 1
- [28] Cheng-Fu Yang, Wan-Cyuan Fan, Fu-En Yang, and Yu-Chiang Frank Wang. Layouttransformer: Scene layout generation with conceptual and spatial diversity. In *CVPR*, 2021. 2, 3, 7
- [29] Zongxin Yang, Jian Dong, Ping Liu, Yi Yang, and Shuicheng Yan. Very long natural scenery image prediction by outpainting. In *ICCV*, 2019. 1, 2
- [30] Zhenqiang Ying and Alan C. Bovik. 180-degree outpainting from a single image. *CoRR*, 2020. 1
- [31] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *ICCV*, 2019. 1
- [32] Alireza Zareian, Zhecan Wang, Haoxuan You, and Shih-Fu Chang. Learning visual commonsense for robust scene graph generation. In *ECCV*, 2020. 2, 3, 7
- [33] Lingzhi Zhang, Jiancong Wang, Yinshuang Xu, Jie Min, Tarmily Wen, James C Gee, and Jianbo Shi. Nested scaleediting for conditional image synthesis. In CVPR, 2020. 7
- [34] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In AAAI, 2020. 6

[35] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In CVPR, 2017. 8