# Panoptic Scene Graph Generation

Jingkang Yang<sup>1</sup>, Yi Zhe Ang<sup>1</sup>, Zujin Guo<sup>1</sup>, Kaiyang Zhou<sup>1</sup>, Wayne Zhang<sup>2</sup>, and Ziwei Liu<sup>1</sup>  $\boxtimes$ 

<sup>1</sup> S-Lab, Nanyang Technological University, Singapore {jingkang001,yizhe.ang,gu00008,kaiyang.zhou,ziwei.liu}@ntu.edu.sg <sup>2</sup> SenseTime Research, Shenzhen, China wayne.zhang@sensetime.com

Abstract. Existing research addresses scene graph generation (SGG) a critical technology for scene understanding in images-from a detection perspective, *i.e.*, objects are detected using bounding boxes followed by prediction of their pairwise relationships. We argue that such a paradigm causes several problems that impede the progress of the field. For instance, bounding box-based labels in current datasets usually contain redundant classes like hairs, and leave out background information that is crucial to the understanding of context. In this work, we introduce panoptic scene graph generation (PSG), a new problem task that requires the model to generate a more comprehensive scene graph representation based on panoptic segmentations rather than rigid bounding boxes. A high-quality PSG dataset, which contains 49k wellannotated overlapping images from COCO and Visual Genome, is created for the community to keep track of its progress. For benchmarking, we build four two-stage baselines, which are modified from classic methods in SGG, and two one-stage baselines called PSGTR and PS-GFormer, which are based on the efficient Transformer-based detector, *i.e.*, DETR. While PSGTR uses a set of queries to directly learn triplets, PSGFormer separately models the objects and relations in the form of queries from two Transformer decoders, followed by a prompting-like relation-object matching mechanism. In the end, we share insights on open challenges and future directions. We invite users to explore the PSG dataset on our project page https://psgdataset.org/, and try our codebase https://github.com/Jingkang50/OpenPSG.

### 1 Introduction

The goal of scene graph generation (SGG) task is to generate a graph-structured representation from a given image to abstract out objects—grounded by bounding boxes—and their pairwise relationships [4,64]. Scene graphs aim to facilitate the understanding of complex scenes in images and has potential for a wide range of downstream applications, such as image retrieval [26,51,49], visual reasoning [1,52], visual question answering (VQA) [20], image captioning [15,7], structured image generation and outpainting [25,12,65], and robotics [13,2].

Since the introduction of SGG [26], this problem has been addressed from a detection perspective, *i.e.*, using bounding boxes to detect objects followed by



Fig. 1: Scene graph generation (a. SGG task) vs. panoptic scene graph generation (b. PSG task). The existing SGG task in (a) uses bounding box-based labels, which are often inaccurate—pixels covered by a bounding box do not necessarily belong to the annotated class—and cannot fully capture the background information. In contrast, the proposed PSG task in (b) presents a more comprehensive and clean scene graph representation, with more accurate localization of objects and including relationships with the background (known as stuff), *i.e.*, the trees and pavement.

the prediction of their pairwise relationships [34,63,41]. We argue that such a bounding box-based paradigm is not ideal for solving the problem, and would instead cause a number of issues that impede the progress of the field. Firstly, bounding boxes—as labeled in current datasets [34]—only provide a coarse localization of objects and often contain noisy/ambiguous pixels belonging to different objects or categories (see the bounding boxes of the two persons in Fig. 1–a). Secondly, bounding boxes typically cannot cover the full scene of an image. For instance, the pavement region in Fig. 1–a is crucial for understanding the context but is completely ignored. Thirdly, current SGG datasets often include redundant classes and information like woman-has-hair in Fig. 1–a, which is mostly deemed trivial [41]. Furthermore, inconsistent and redundant labels are also observed in current datasets, *e.g.*, the trees and benches in Fig. 1–a are labeled multiple times, and some extra annotations do not contribute to the graph (see isolated nodes). Using such labels for learning might confuse the model.

Ideally, the grounding of objects should be clear and precise, and a scene graph should not only focus on salient regions and relationships in an image but also be comprehensive enough for scene understanding. We argue that as compared to bounding boxes, panoptic segmentation [32] labels would be a better choice for constructing scene graphs. To this end, we introduce a new problem, *panoptic scene graph generation*, or PSG, with a goal of generating scene graph representations based on panoptic segmentations rather than rigid bounding boxes.

To help the community keep track of the research progress, we create a new PSG dataset based on COCO [42] and Visual Genome (VG) [34], which contains 49k well-annotated images in total. We follow COCO's object annotation schema of 133 classes; comprising 80 thing classes and 53 stuff (background) classes. To construct predicates, we conduct a thorough investigation into existing VG-based datasets, *e.g.*, VG-150 [63], VrR-VG [41] and GQA [23], and summarize 56 predicate classes with minimum overlap and sufficient coverage of semantics. See Fig. 1–b for an example of our dataset. From Fig. 1, it is clear that the panoptic scene graph representation—including both panoptic segmentations and the scene graph—is much more informative and coherent than the previous scene graph representation.

For benchmarking, we build four two-stage models by integrating four classic SGG methods [63,70,57,53] into a classic panoptic segmentation framework [31]. We also turn DETR [3], an efficient Transformer-based detector, into a one-stage PSG model dubbed as PSGTR, which has proved effective for the PSG task. We further provide another one-stage baseline called PSGFormer, which extends PSGTR with two improvements: 1) modeling objects and relations separately in the form of queries within two Transformer decoders, and 2) a prompting-like interaction mechanism. A comprehensive comparison of one-stage models and two-stage models is discussed in our experiments.

In summary, we make the following contributions to the SGG community:

- A New Problem and Dataset: We discuss several issues with current SGG research, especially those associated with existing datasets. To address them, we introduce a new problem that combines SGG with panoptic segmentation, and create a large PSG dataset with high-quality annotations.
- A Comprehensive Benchmark: We build strong two-stage and one-stage PSG baselines, and evaluate them comprehensively on our new dataset, so that the PSG task is solidified in its inception. We find that one-stage models, despite having a simplified training paradigm, have great potential for PSG as it achieves competitive results on the dataset.

# 2 Related Work

Scene Graph Generation Existing scene graph generation (SGG) methods have been dominated by the two-stage pipeline that consists of object detection and pairwise predicate estimation. Given bounding boxes, early work predicts predicates using conditional random fields [26,10] or casts predicate prediction into a classification problem [74,33,48]. Inspired by knowledge graph embeddings, VTransE [73] and UVTransE [24] are proposed for explicit predicate modeling. Follow-up works have investigated various variants based on, *e.g.*, RNN and graph-based modeling [63,70,57,66,43,8,37], energy-based models [53], external knowledge [57,17,68,69,45], and more recently language supervision [76,67]. Recent research has shifted the attention to problems associated with the SGG datasets, such as the long-tailed distribution of predicates [56,11],

Table 1: Comparison between classic SGG datasets and PSG dataset. #PPI counts predicates per image. DupFree checks whether duplicated object groundings are cleaned up. Spvn indicates whether the objects are grounded by bounding boxes or segmentations.

Dataset	#Image	# ObjCls	$\# \mathrm{RelCls}$	$\# \mathrm{PPI}$	DupFree	$\operatorname{Spvn}$	Source
VG [34]	108K	34K	40K	21.4	×	BBox	COCO & Flickr
VG-150 [63]	88K	150	50	5.7	×	$\operatorname{BBox}$	Clean VG
VrR-VG [41]	59K	1,600	117	3.4	×	$\operatorname{BBox}$	Clean VG
GQA [23]	85K	1,703	310	50.6	1	$\operatorname{BBox}$	Re-annotate ${\rm VG}$
PSG	49K	133	56	5.7	1	Seg	Annotate COCO



Fig. 2: Word Cloud for PSG Predicates.

excessive visually-irrelevant predicates [41], and inaccurate localization of bounding boxes [28]. In particular, a very recent study [28] shows that training an SGG model to simultaneously generate scene graphs and predict semantic segmentation masks can bring about improvements, which inspires our research. In our work, we study panoptic segmentation-based scene graph generation in a more systematic way by formulating a new problem and building a new benchmark. We also notice that a closely-related topic human-object interaction (HOI) [18] shares a similar goal with SGG, *i.e.*, to detect prominent relations from the image. However, the HOI task restricts the model to only detect human-related relations while ignoring the valuable information between objects that is often critical to comprehensive scene understanding. Nevertheless, many HOI methods are applicable to SGG tasks [16,27,5,59,39,77,60,21,38,14,29,44,54,22,71,58,72], and some of them have inspired our PSG baselines [30,78].

While early SGG works have constructed several **Scene Graph Datasets** smaller-size datasets such as RW-SGD [26], VRD [45], and UnRel-D [47], the large-scale Visual Genome (VG) [34] quickly became the standard SGG dataset as soon as it was released in 2017, prompting subsequent work to research in a more realistic setting. However, several critical drawbacks of VG were raised by the community, and therefore, some VG variants were gradually introduced to address some problems. Notice that VG contains an impractically large number of 33,877 object classes and 40,480 predicate classes, leading VG-150 [63] to only keep the most frequent 150 object classes and 50 predicate classes for a more realistic setting. Later, VrR-VG [41] argued that many predicates in VG-150 can be easily estimated by statistical priors, hence deciding to re-filter the original VG categories to only keep visually relevant predicate classes. However, by scrutinizing VrR-VG, we find many predicates are redundant (e.q., alongside, beside,besides) and ambiguous (e.g., beyond, down). Similar drawbacks appeared in another large-scale dataset with scene graph annotations called GQA [23]. In summary, although relations play a pivotal role in SGG tasks, a systematic definition of relations is unfortunately overlooked across all the existing SGG datasets. Therefore, in our proposed PSG dataset, we consider both comprehensive coverage and non-overlap between words, and carefully define a predicate dictionary with 56 classes to better formulate the scene graph problem. Fig. 2 shows the word cloud of the predicate classes, where font size indicates frequency.

Apart from the problem of predicate definition, another critical issue of SGG datasets is that they all adopt bounding box-based object grounding, which inevitably causes a number of issues such as coarse localization (bounding boxes cannot reach pixel-level accuracy), inability to ground comprehensively (bounding boxes cannot ground backgrounds), tendency to provide trivial information (current datasets usually capture objects like head to form the trivial relation of person-has-head), and duplicate groundings (the same object could be grounded by multiple separate bounding boxes). These issues together have caused the low-quality of current SGG datasets, which impede the progress of the field. Therefore, the proposed PSG dataset tries to address all the above problems by grounding the images using accurate and comprehensive panoptic segmentations with COCO's appropriate granularity of object categories. Table 1 compares the statistics of the PSG dataset with classic SGG datasets.

**Panoptic Segmentation** The panoptic segmentation task unifies semantic segmentation and instance segmentation [32] for comprehensive scene understanding, and the first approach is a simple combination of a semantic segmentation model and an instance segmentation model to produce stuff masks and thing masks respectively [32]. Follow-up work, such as Panoptic FPN [31] and UPSNet [62], aim to unify the two tasks in a single model through multitask learning to achieve gains in compute efficiency and segmentation performance. Recent approaches (e.g., MaskFormer [9], Panoptic Segformer [40] and K-Net [75]) have turned to more efficient architectures based on Transformers like DETR [3], which simplifies the detection pipeline by casting the detection task as a set prediction problem.

### 3 Problem and Dataset

**Recap: Scene Graph Generation** We first briefly review the goal of the classic scene graph generation (SGG) task, which aims to model the distribution:

$$\Pr\left(\mathbf{G} \mid \mathbf{I}\right) = \Pr\left(\mathbf{B}, \mathbf{O}, \mathbf{R} \mid \mathbf{I}\right),\tag{1}$$

where  $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$  is the input image, and  $\mathbf{G}$  is the desired scene graph which comprises the bounding boxes  $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$  and labels  $\mathbf{O} = \{o_1, \dots, o_n\}$ that correspond to each of the *n* objects in the image, and their relations in the set  $\mathbf{R} = \{r_1, \dots, r_l\}$ . More specifically,  $\mathbf{b}_i \in \mathbb{R}^4$  represents the box coordinates,  $o_i \in \mathbb{C}^{\mathbf{O}}$  and  $r_i \in \mathbb{C}^{\mathbf{R}}$  belong to the set of all object and relation classes.

**Panoptic Scene Graph Generation** Instead of localizing each object by its bounding box coordinates, the new task of panoptic scene graph generation (PSG task) grounds each object with the more fine-grained panoptic segmentation. For conciseness, we refer to both objects and background as objects.

Formally, with panoptic segmentation, an image is segmented into a set of masks  $\mathbf{M} = {\mathbf{m}_1, \ldots, \mathbf{m}_n}$ , where  $\mathbf{m}_i \in {\{0,1\}}^{H \times W}$ . Each mask is associated

with an object with class label  $o_i \in C^O$ . A set of relations **R** between objects are also predicted. The masks do not overlap, *i.e.*,  $\sum_{i=1}^{n} \mathbf{m}_i \leq \mathbf{1}^{H \times W}$ . Hence, PSG task models the following distribution:

$$\Pr\left(\mathbf{G} \mid \mathbf{I}\right) = \Pr\left(\mathbf{M}, \mathbf{O}, \mathbf{R} \mid \mathbf{I}\right).$$
(2)

**PSG Dataset** To address the PSG task, we build our PSG dataset following these three major steps. Readers can find more details in the Appendix.

Step 1: A Coarse COCO & VG Fusion: To create a dataset with both panoptic segmentation and relation annotations, we use the 48,749 images in the intersection of the COCO and VG datasets with an automatic but coarse dataset fusion process. Specifically, we use an object category matching procedure to match COCO's segmentations with VG's bounding boxes, so that part of VG's predicates are applied to COCO's segmentation pairs. Due to the inherent mismatch between the label systems and localization annotations of VG and COCO, the auto-generated dataset is very noisy and requires further cleaning.

**Step 2: A Concise Predicate Dictionary:** Inspired by the appropriate granularity of COCO categories [42], we carefully identify 56 salient relations by taking reference from common predicates in the initial noisy PSG dataset and all VG-based datasets including VG-150 [63], VrR-VG [41] and GQA [23]. The selected 56 predicates are maximally independent (*e.g.*, we only keep "over/on" and do not have "under") and cover most common cases in the dataset.

**Step 3: A Rigorous Annotation Process:** Building upon the noisy PSG dataset, we require the annotators to 1) *filter* out incorrect triplets, and 2) *supplement* more relations between not only object-object, but also object-background and background-background pairs, using the predefined 56 predicates. To prevent ambiguity between predicates, we ask the annotators strictly not to annotate using general relations like on, in when a more precise predicate like **parked on** is applicable. With this rule, the PSG dataset allows the model to understand the scene more precisely and saliently.

**Quality Control:** The PSG dataset goes through a professional dataset construction process. The main authors first annotated 1000 images to construct a detailed documentation (available in project page), and then employed a professional annotation company (sponsored by SenseTime) to annotate the training set within a month (US\$11K spent). Each image is annotated by two workers and examined by one head worker. All the test images are annotated by the authors.

**Summary:** Several merits worth highlighting by virtue of the novel and effective procedure of PSG dataset creation: 1) Good grounding annotation from the pixel-wise panoptic segmentation from COCO dataset [42], 2) Clear category system with 133 objects (*i.e.*, things plus stuff) and 56 predicates with appropriate granularity and minimal overlaps, and 3) Accurate and comprehensive relation annotations from a rigorous annotation process that pays special attention to salient relations between object-object, object-background and background-background. These merits address the notorious shortcomings [36] of classic scene graph datasets discussed in Sec. 2.

**Evaluation and Metrics** This section introduces the evaluation protocol for the PSG task. Following the settings of the classic SGG task [64,4], our PSG task comprises two sub-tasks: *predicate classification* (when applicable) and *scene graph generation* (main task) to evaluate the PSG models. *Predicate classification* (*PredCls*) aims to generate a scene graph given the ground-truth object labels and localization. The goal is to study the relation prediction performance without the interference of the segmentation performance. Notice that this metric is only applicable to two-stage PSG models in Sec. 4.1, since onestage models cannot leverage the given segmentations to predict scene graph. *Scene graph generation* (*SGDet*) aims to generate scene graphs from scratch, which is the main result for the PSG task.

We also notice that classic SGG tasks contain another sub-task of scene graph classification (SGCls), which provide the ground-truth object groundings to simplify the scene graph generation process. We find SGCls is not applicable for PSG baselines. Unlike SGG tasks where object detectors such as Faster-RCNN [50] can utilize ground-truth object bounding boxes to replace predictions from the Region Proposal Network (RPN), panoptic segmentation models are unable to directly use the ground-truth segmentations for classification, so the SGCls task is inapplicable even for two-stage PSG methods.

The classic metrics of R@K and mR@K are used to evaluate the previous two sub-tasks, which calculates the triplet recall and mean recall for every predicate category, given the top K triplets from the PSG model. Notice that PSG grounds objects with segmentation, a successful recall requires both subject and object to have mask-based IOU larger than 0.5 compared to their ground-truth counterparts, with the correct classification on every position in the S-V-O triplet.

While the triplet recall rates that mentioned above are the main metric for PSG task, since objects are required to be grounded by segmentation masks, panoptic segmentation metrics [32] such as PQ [31] can be used for model diagnosis. However, it is not considered as the core evaluation metric of PSG task.

# 4 PSG Baselines

To build a comprehensive PSG benchmark, we refer to frameworks employed in the classic SGG task [4,64] and prepare two-stage and one-stage baselines.

### 4.1 Two-Stage PSG Baselines

Most prior SGG approaches tackle the problem in two stages: first performing object detection using off-the-shelf detectors like Faster-RCNN [50], then pairwise relationship prediction between these predicted objects. As shown in Figure 3, we follow a similar approach in establishing two-stage baselines for the PSG task: 1) using pretrained panoptic segmentation models of classic Panoptic FPN [31] to extract initial object features, masks and class predictions, and then 2) processing them using a relation prediction module from classic scene graph generation methods like IMP [63], MOTIFS [70], VCTree [57], and GPSNet [43]



Fig. 3: Two-stage PSG baselines using Panoptic FPN. a) In stage one, for each thing/stuff object, Panoptic FPN [31] produces a segmentation mask with its tightest bounding box to crop out the object feature. The union of relevant objects can produce relation features. b) In the second stage, the extracted object and relation features are fed into by any existing SGG relation model to predict the relation triplets.

to obtain the final scene graph predictions. In this way, classic SGG methods can be adapted to solve the PSG task with minimal modification. Formally, the two-stage PSGG baselines decompose formulation from Eq. 1 to Eq. 3.

$$\Pr\left(\mathbf{G} \mid \mathbf{I}\right) = \Pr\left(\mathbf{M} \mid \mathbf{I}\right) \cdot \Pr\left(\mathbf{O} \mid \mathbf{M}, \mathbf{I}\right) \cdot \Pr\left(\mathbf{R} \mid \mathbf{O}, \mathbf{M}, \mathbf{I}\right).$$
(3)

### 4.2 A One-Stage PSG Baseline - PSGTR

Unlike classic dense prediction models (e.g., Faster-RCNN [50]) with sophisticated design, the transformer-based architectures support flexible input and output specifications. Based on the end-to-end DETR [3] and its extension to the HOI task [78], we naturally design a one-stage PSG method named PSGTR to predict triples and localizations simultaneously, which can be directly modeled as Eq. 2 without decomposition.

**Triplet Query Learning Block** As shown in Fig. 4, PSGTR first extracts image features from a CNN backbone and then feeds the features along with queries and position encoding into a transformer encoder-decoder. Here we expect the queries to learn the representation of scene graph triplets, so that for each triplet query, the subject/predicate/object predictions can be extracted by three individual Feed Forward Networks (FFNs), and the segmentation task can be completed by two panoptic heads for subject and object, respectively.

**PSG Prediction Block** To train the model, we extend the DETR's Hungarian matcher [35] into a triplet Hungarian matcher. To match the triplet query  $\mathcal{T}_i \in \mathbb{Q}^{\mathrm{T}}$  with ground truth triplets  $\mathcal{G}$ , all contents in the triplet (*i.e.*, all outputs that are predicted from  $\mathcal{T}_i$ ) are used, including the class of subject  $\mathcal{T}_i^{\mathrm{S}}$ , relation  $\mathcal{T}_i^{\mathrm{R}}$ , and object  $\mathcal{T}_i^{\mathrm{O}}$ , and localization of subjects  $\mathcal{T}_i^{\mathrm{S}}$  and objects  $\mathcal{T}_i^{\mathrm{O}}$ . Therefore, the triplet matching (tm) cost  $\mathbf{C}_{\mathrm{tm}}$  is designed with the combination of class matching  $\mathbf{C}_{\mathrm{cls}}$  and segments matching  $\mathbf{C}_{\mathrm{seg}}$ :

$$\mathbf{C}_{\mathrm{tm}}\left(\mathcal{T}_{i},\mathcal{G}_{\sigma(i)}\right) = \sum_{k \in \{\mathrm{S},\mathrm{O}\}} \mathbf{C}_{\mathrm{seg}}\left(\tilde{\mathcal{T}}_{i}^{k},\tilde{\mathcal{G}}_{\sigma(i)}^{k}\right) + \sum_{k \in \{\mathrm{S},\mathrm{R},\mathrm{O}\}} \mathbf{C}_{\mathrm{cls}}\left(\ddot{\mathcal{T}}_{i}^{k},\ddot{\mathcal{G}}_{\sigma(i)}^{k}\right), \quad (4)$$



Fig. 4: PSGTR: One-stage PSG baseline. The one-stage model takes in a) features extracted by CNNs with positional encoding, and a set of queries aiming to represent triplets. b) Query learning block processes image features with Transformer encoder-decoder and use queries to represent triplet information. Then, c) the PSG prediction head concretes the triplet predictions by producing subject/object/predicate classes using simple FFNs, and uses panoptic heads for panoptic segmentation.

where  $\sigma$  is the mapping function to correspond each triplet query  $\mathcal{T}_i \in \mathbb{Q}^T$  to the closest ground truth triplet. The triplet query set  $\mathbb{Q}^T$  collects all the  $|\mathbb{Q}^T|$ triplet queries. The optimization objective is thus:

$$\hat{\sigma} = \arg\max_{\sigma} \sum_{i=1}^{|\mathbb{Q}^{\mathrm{T}}|} \mathbf{C}_{\mathrm{tm}} \left( \mathcal{T}_{i}, \mathcal{G}_{\sigma(i)} \right).$$
(5)

Once the matching is done, the total loss  $\mathcal{L}_{total}$  can be calculated by applying cross-entropy loss  $\mathcal{L}_{cls}$  for labels and DICE/F-1 loss [46] for segmentation  $\mathcal{L}_{cls}$ :

$$\mathcal{L}_{\text{total}} = \sum_{i=1}^{|\mathbb{Q}^{\mathrm{T}}|} \left( \sum_{k \in \{\mathrm{S},\mathrm{O}\}} \mathcal{L}_{\text{seg}} \left( \tilde{\mathcal{T}}_{i}^{k}, \tilde{\mathcal{G}}_{\hat{\sigma}(i)}^{k} \right) + \sum_{k \in \{\mathrm{S},\mathrm{R},\mathrm{O}\}} \mathcal{L}_{\text{cls}} \left( \ddot{\mathcal{T}}_{i}^{k}, \ddot{\mathcal{G}}_{\hat{\sigma}(i)}^{k} \right) \right).$$
(6)

### 4.3 Alternative One-Stage PSG Baseline - PSGFormer

Based on the PSGTR baseline that explained in Section 4.2, we extend another end-to-end HOI method [30] and further propose the alternative one-stage PSG baseline named PSGFormer, featured by an explicit relation modeling with a prompting-like matching mechanism. The model diagram is illustrated in Figure 5 and will be elaborated as follows.

**Object & Relation Query Learning Block** Compared to the classic objectoriented tasks such as object detection and segmentation, the most significant uniqueness of PSG task as well as SGG task is their extra requirements on the predictions of relations. Notice that the relation modeling in our two-stage baselines depends on features from object-pairs, while PSGTR implicitly models the objects and relations altogether within the triplets, the important relation modeling has not been given a serious treatment. Therefore, in the exploration of



Fig. 5: PSGFormer: The proposed one-stage PSG method. a) Two types of queries, *i.e.*, object queries and relation queries, are fed into transformer block with CNN features and positional encoding. b) Query Learning Block processes image features with one encoder and output object or relation queries with the corresponding decoder. c) Object queries and relation queries interact with each other in the prompting-like query matching block, so that the triplets are formed and proceed to d) PSG prediction block to output results or compute loss as PSGTR behaves.

PSGFormer, we explicitly model the relation query  $\mathcal{R}_i \in \mathbb{Q}^{\mathbb{R}}$ , as well as object query  $\mathcal{O}_i \in \mathbb{Q}^{\mathbb{O}}$  separately, in hope that object queries to specially pay attentions to objects (*e.g.*, **person** and **phones**), and relation queries to focus on the area where the relationship takes place in the picture (*e.g.*, **person looking-at phone**). Similar to PSGTR in Figure 4, both object and relation queries with CNN features and position encoding are fed into a transformer encoder, but being decoded with their corresponding decoder, *i.e.*, object or relation decoder, so that the queries can learn the corresponding representations.

**Object & Relation Query Matching Block** In PSGFormer, each object query yields an object prediction with FFN and a mask prediction with a panoptic head, and each relation query yields a relation prediction. However, due to the parallel process of object queries and relation queries, the missing interdependence between different query types makes the triplet still not formed. To connect object and relation queries for compositing triplets, we are inspired by the design in HOTR [30] and implement a prompting-like query matching block.

Query matching block models the triplet composition task as a fill-in-theblank question with prompts, *i.e.*, by prompting a relation, we expect a pair of suitable objects provided by their corresponding object queries can be selected, so that a complete **subject-predicate-object** triplet, can be generated. Therefore, two selectors, *i.e.*, subject selector and object selector, are required.

Given a relation query  $\mathcal{R}_i \in \mathbb{Q}^R$  as prompt, both subject selector and object selector should return the most suitable candidate to form a complete triplet. We use the most standard cosine similarity between object queries and the provided relation query and pick the highest similarity to determine the subject and object candidates. It should also be noticed that subject and object selectors should rely on the level of association between objects and relation queries rather than the semantic similarity. Besides, object queries are regarded as different roles (*i.e.*, subject or object) in different selectors. Therefore, the object queries are expected to pass another two FFNs to extract some specific information for subject (with FFN denoted as  $f_{\rm S}$ ) and object (with FFN denoted as  $f_{\rm O}$ ), so that the distinguishable subject and object representations are obtained from the object queries. With the idea above, a set of subjects S are generated in Eq. 7, with the *i*-th subject corresponding to the *i*-th relation query  $\mathcal{R}_i$ . With a similar process, the object set  $\mathbb{O}$  is also generated.

$$\mathbb{S} = \left\{ \mathcal{S}_i \mid \mathcal{S}_i = \arg \max_{\mathcal{O}} \left( f_{\mathrm{S}}(\mathcal{O}_j) \cdot \mathcal{R}_i \right), \ \mathcal{O}_j \in \mathbb{Q}^{\mathrm{O}}, \ \mathcal{R}_i \in \mathbb{Q}^{\mathrm{R}} \right\}.$$
(7)

Till now, the subject set S and the object set  $\mathbb{O}$  are well-prepared by subject and object selectors, with the *i*-th subject query  $S_i$  and the *i*-th object  $\mathcal{O}_i$  corresponding to the *i*-th relation query  $\mathcal{R}_i$ . Finally, it is straightforward to obtain all the matched triplet  $\mathbb{T}$ , which is shown in Eq. 8.

$$\mathbb{T} = \{ (\mathcal{S}_i, \mathcal{R}_i, \mathcal{O}_i) \mid \mathcal{S}_i \in \mathbb{S}, \ \mathcal{R}_i \in \mathbb{R}, \ \mathcal{O}_i \in \mathbb{O} \} .$$
(8)

Apart from interpreting the query matching as a prompt-like process, it can also be considered as a cross-attention mechanism. For a relation query (Q), the goal is to find the high-attention relations among all subject / object representations, which are considered as keys (K). The subject / object labels predicted by the corresponding representations are regarded as values (V), so that the QKV attention model outputs the labels of selected keys. Fig. 5–c is generally depicted following this interpretation.

**PSG Prediction Block** Similar to PSGTR, with the predicted triplets prepared, the prediction block can finally train the model using  $\mathcal{L}_{\text{total}}$  from Eq. 6. In addition, with object labels and masks predicted by object queries, a standard training loss introduced in panoptic segmentation DETR [3] is used to enhance the object decoder and avoid duplicate object groundings.

### 5 Experiments

In this section, we first report the results of all PSG methods introduced in the paper. Implementation details are available in the appendix, and all codes are integrated in the OpenPSG codebase, which is developed based on MMDetection [6]. Most of the two-stage SGG implementations refer to MMScene-Graph [61] and Scene-Graph-Benchmark.pytorch [55].

### 5.1 Main Results

Table 2 reports the scene graph generation performance of all the methods mentioned in Sec. 4.1, Sec. 4.2, and Sec. 4.3 under the PSG dataset. Fig. 6 reports the panoptic segmentation result using PQ and visualizes the segmentation results of two examples as well as the predicted scene graph in the form of triplet lists.

Table 2: Comparison between all baselines and PSGFormer. Recall (R) and mean recall (mR) are reported. IMP [63] (CVPR'17), MOTIFS [70] (CVPR'18), VC-Tree [57] (CVPR'19), and GPSNet [43] (CVPR'20) all originate from the SGG task and are adapted for the PSG task. Different backbones of ResNet-50 [19] and ResNet-101 [19] are used. Notice that predicate classification task is not applicable to one-stage PSG models, so the corresponding results are marked as '-'. Models are trained using 12 epochs by default. <sup>†</sup> denotes that the model is trained using 60 epochs.

Backbone	Method	Pred	licate Classifi	cation	Scene Graph Generation			
		R/mR@20	R/mR@50	R/mR@100	R/mR@20	R/mR@50	R/mR@100	
ResNet-50	IMP MOTIFS VCTree GPSNet	31.9 / 9.55 44.9 / 20.2 <b>45.3</b> / <b>20.5</b> 31.5 / 13.2	36.8 / 10.9 50.4 / 22.1 <b>50.8</b> / <b>22.6</b> 39.9 / 16.4	38.9 / 11.6 52.4 / 22.9 <b>52.7</b> / <b>23.3</b> 44.7 / 18.3	16.5 / 6.52 20.0 / 9.10 <b>20.6</b> / <b>9.70</b> 17.8 / 7.03	18.2 / 7.05 21.7 / 9.57 <b>22.1</b> / <b>10.2</b> 19.6 / 7.49	18.6 / 7.23 22.0 / 9.69 <b>22.5</b> / <b>10.2</b> 20.1 / 7.67	
	PSGTR PSGFormer	-	-	-	3.82 / 1.29 16.8 / 14.5	4.16 / 1.54 19.2 / 17.4	4.27 / 1.57 20.2 / 18.7	
	$PSGTR^{\dagger}$ $PSGFormer^{\dagger}$	-	-	-	<b>28.4</b> / <b>16.6</b> 18.0 / 14.8	<b>34.4</b> / <b>20.8</b> 19.6 / 17.0	<b>36.3</b> / <b>22.1</b> 20.1 / 17.6	
ResNet-101	IMP MOTIFS VCTree GPSNet	30.5 / 8.97 45.1 / 19.9 <b>45.9</b> / <b>21.4</b> 38.8 / 17.1	35.9 / 10.5 50.5 / 21.5 <b>51.2</b> / <b>23.1</b> 46.6 / 20.2	38.3 / 11.3 52.5 / 22.2 <b>53.1</b> / <b>23.8</b> 50.0 / 21.3	17.9 / 7.35 20.9 / 9.60 <b>21.7</b> / <b>9.68</b> 18.4 / 6.52	19.5 / 7.88 22.5 / 10.1 <b>23.3</b> / <b>10.2</b> 20.0 / 6.97	20.1 / 8.02 23.1 / 10.3 23.7 / 10.3 20.6 / 7.17	
	PSGTR PSGFormer	-	-	-	3.47 / 1.18 18.0 / 14.2	3.88 / 1.56 20.1 / 18.3	4.00 / 1.64 <b>21.0</b> / <b>19.8</b>	
	$\begin{array}{c} \mathrm{PSGTR}^{\dagger} \\ \mathrm{PSGFormer}^{\dagger} \end{array}$	-	-	-	<b>28.2</b> / <b>15.4</b> 18.6 / 16.7	<b>32.1</b> / <b>20.3</b> 20.4 / 19.3	<b>35.3</b> / <b>21.5</b> 20.7 / 19.7	

**Two-Stage Baselines Rely on First-Stage Performance** For predicate classification task (PredCls) that is only applicable to two-stage models, the provided ground-truth segmentation can significantly improve the triplet prediction performance. For example, even the most classic method IMP can reach over 30% R@20, which already exceeds all the available R@20 under the scene graph generation (SGDet) task (cf. 28.4% by PSGTR). This phenomenon indicates that a good panoptic segmentation performance could naturally benefit the PSG task. Further evidence where the performance of IMP on the SGDet task is almost halved (from 32% to 17% on R@20) strengthens the above conjecture.

**Some SGG Techniques for VG are not Effective for PSG** Table 2 shows that the results of some two-stage baselines (*i.e.*, IMP, MOTIFS, and VCTree) are generally proportional to their performance on SGG tasks, indicating that the advantages of the two-stage models (*i.e.*, MOTIFS and VCTree) are transferable to PSG task. However, we notice that another two-stage baseline, GPSNet, does not seem to exceed its SGG baselines of MOTIFS and VCTree in the PSG task. The key advantage of GPSNet over MOTIFS and VCTree is that it explicitly models the direction of predicates. While the design can be effective in the VG dataset where many predicates are trivial with obvious direction of predicates (*e.g.*, of in hair-of-man, has in man-has-head), PSG dataset gets rid of these predicates, so the model may not be effective as expected.



Fig. 6: Visualization of segmentations and the top 20 predicted triplets of 5 PSG methods. The panoptic segmentation metric PQ is also reported. The colors of the subject and object in the triplet corresponds to the mask in the segmentation result. Reasonable triplets are marked by ticks. Triplets that match the ground-truth are marked by green ticks. One-stage models can provide more reasonable and diverse triplet predictions, but they are unable to achieve a good panoptic segmentation result.

**PSGFormer is an Unbiased PSG Model** When the training schedule is limited to 12 epochs, the end-to-end baseline PSGFormer outperforms the best two-stage model VCTree by significant 4.8% on mR@20 and 8.5% on mR@100. Although PSGFormer still cannot exceed two-stage methods on the metrics of R@20/50/100, its huge advantage in mean recall indicates that the model is unbiased in predicting relations. As Fig. 6 shows, PSGFormer can predict unusual but accurate relations such as person-going down-snow in the upper example, and the imperceptible relation person-driving-bus in the lower example. Also, in the upper example, PSGFormer predicts an interesting and exclusive triplet person-wearing-skis. This unique prediction should come from the design of

the separate object / relation decoders, so that relation queries can independently capture the meaning of the predicates.

**PSGTR Obtains SOTA Results with Long Training Time** In PSGTR, every triplet is expected to be captured by a query, which needs to predict everything in the triplet simultaneously, so the model is required to better focus on the connections between objects. Besides, the cross-attention mechanism in the transformer encoder and triplet decoder enable each triplet query access to the information of the entire image. Therefore, PSGTR is considered as the most straightforward and simplest one-stage PSG baseline with minimal constraints or prior knowledge. As a result, although PSGTR only achieves one-digit recall scores in 12 epochs, it surprisingly achieves SOTA results with a prolonged training time of 60 epochs.

# 6 Challenges and Outlook

**Challenges** While some *prior knowledge* introduced by some two-stage SGG methods might not be effective in the PSG task, we expect that more creative knowledge-aided models can be developed for the PSG task in the era of multi-modality, so that more interesting triplets can be extracted with extra priors. However, it should be noted that although priors can be useful to enhance performance, the PSG prediction should heavily rely on *visual clues*. For example, in Fig. 6, person-walking on-pavement should be identified if the model can perceive and understand the subtle visual differences between walking and standing. Also, PSG models are expected to *predict more meaningful and diverse relations*, such as rare relations like feeding and kissing, rather than only being content with statistically common or positional relations.

**Relation between PSG and Panoptic Segmentation** Fig. 6 visualizes the panoptic segmentation results of PSG methods, where PSGTR only obtains a miserable PQ result even with good PSG performance. The reason is that triplet queries in PSGTR produce object groundings independently, so that one object might be referred and segmented by several triplets, and the deduplication or the re-identification (Re-ID) process is non-trivial. Although the performance of Re-ID does not affect PSG metrics, it might still be critical to form an accurate and logical scene understanding for real-world applications.

**Outlook** Apart from attracting more research on the learning of relations (either closed-set or open-set) and pushing the development of scene understanding, we also expect the PSG models to empower more exciting downstream tasks such as visual reasoning and segmentation-based scene graph-to-image generation.

### Acknowledgements

This work is supported by NTU NAP, MOE AcRF Tier 2 (T2EP20221-0033), and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

## References

- Aditya, S., Yang, Y., Baral, C., Aloimonos, Y., Fermüller, C.: Image understanding using vision and reasoning through scene description graph. Computer Vision and Image Understanding (2018)
- 2. Amiri, S., Chandan, K., Zhang, S.: Reasoning with scene graphs for robot planning under partial observability. IEEE Robotics and Automation Letters (2022)
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. arXiv preprint arXiv:2005.12872 (2020)
- Chang, X., Ren, P., Xu, P., Li, Z., Chen, X., Hauptmann, A.: Scene graphs: A survey of generations and applications. arXiv preprint arXiv:2104.01111 (2021)
- Chao, Y.W., Liu, Y., Liu, X., Zeng, H., Deng, J.: Learning to detect human-object interactions. In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV) (2018)
- Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
- Chen, S., Jin, Q., Wang, P., Wu, Q.: Say as you wish: Fine-grained control of image caption generation with abstract scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Chen, T., Yu, W., Chen, R., Lin, L.: Knowledge-embedded routing network for scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Cheng, B., Schwing, A.G., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. vol. abs/2107.06278 (2021)
- Dai, B., Zhang, Y., Lin, D.: Detecting visual relationships with deep relational networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Desai, A., Wu, T.Y., Tripathi, S., Vasconcelos, N.: Learning of visual relations: The devil is in the tails. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
- Dhamo, H., Farshad, A., Laina, I., Navab, N., Hager, G.D., Tombari, F., Rupprecht, C.: Semantic image manipulation using scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Gadre, S.Y., Ehsani, K., Song, S., Mottaghi, R.: Continuous scene representations for embodied ai. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Gao, C., Xu, J., Zou, Y., Huang, J.B.: Drg: Dual relation graph for human-object interaction detection. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Gao, L., Wang, B., Wang, W.: Image captioning with scene-graph based semantic concepts. In: Proceedings of the International Conference on Machine Learning and Computing (2018)
- Gkioxari, G., Girshick, R., Dollár, P., He, K.: Detecting and recognizing humanobject interactions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- 17. Gu, J., Zhao, H., Lin, Z., Li, S., Cai, J., Ling, M.: Scene graph generation with external knowledge and image reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)

- 16 J. Yang et al.
- Gupta, S., Malik, J.: Visual semantic role labeling. arXiv preprint arXiv:1505.04474 (2015)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- Hildebrandt, M., Li, H., Koner, R., Tresp, V., Günnemann, S.: Scene graph reasoning for visual question answering. ICML Workshop Graph Representation Learning and Beyond (GRL+) (2020)
- Hou, Z., Peng, X., Qiao, Y., Tao, D.: Visual compositional learning for humanobject interaction detection. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Hou, Z., Yu, B., Qiao, Y., Peng, X., Tao, D.: Affordance transfer learning for human-object interaction detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
- Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- 24. Hung, Z.S., Mallya, A., Lazebnik, S.: Contextual translation embedding for visual relationship detection and scene graph generation. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2020)
- Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Johnson, J., Krishna, R., Stark, M., Li, L.J., Shamma, D., Bernstein, M., Fei-Fei, L.: Image retrieval using scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
- Kato, K., Li, Y., Gupta, A.: Compositional learning for human object interaction. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
- Khandelwal, S., Suhail, M., Sigal, L.: Segmentation-grounded scene graph generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
- Kim, B., Choi, T., Kang, J., Kim, H.J.: Uniondet: Union-level detector towards real-time human-object interaction detection. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Kim, B., Lee, J., Kang, J., Kim, E.S., Kim, H.J.: Hotr: End-to-end human-object interaction detection with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
- Kirillov, A., Girshick, R., He, K., Dollár, P.: Panoptic feature pyramid networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- 33. Kolesnikov, A., Kuznetsova, A., Lampert, C., Ferrari, V.: Detecting visual relationships using box attention. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (CVPR-W) (2019)
- 34. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., Bernstein, M., Fei-Fei, L.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. International Journal of Computer Vision (IJCV) (2017)

- 35. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly (1955)
- Li, L., Chen, L., Huang, Y., Zhang, Z., Zhang, S., Xiao, J.: The devil is in the labels: Noisy label correction for robust scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- 37. Li, Y., Ouyang, W., Zhou, B., Shi, J., Zhang, C., Wang, X.: Factorizable net: an efficient subgraph-based framework for scene graph generation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
- Li, Y.L., Liu, X., Lu, H., Wang, S., Liu, J., Li, J., Lu, C.: Detailed 2d-3d joint representation for human-object interaction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Li, Y.L., Zhou, S., Huang, X., Xu, L., Ma, Z., Fang, H.S., Wang, Y., Lu, C.: Transferable interactiveness knowledge for human-object interaction detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- 40. Li, Z., Wang, W., Xie, E., Yu, Z., Anandkumar, A., Alvarez, J.M., Lu, T., Luo, P.: Panoptic segformer. arXiv preprint arXiv:2109.03814 (2021)
- Liang, Y., Bai, Y., Zhang, W., Qian, X., Zhu, L., Mei, T.: Vrr-vg: Refocusing visually-relevant relationships. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Proceedings of the European Conference on Computer Vision (ECCV) (2014)
- Lin, X., Ding, C., Zeng, J., Tao, D.: Gps-net: Graph property sensing network for scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Liu, Y., Chen, Q., Zisserman, A.: Amplifying key cues for human-object-interaction detection. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 852–869 (2016)
- Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: International Conference on 3D Vision (3DV) (2016)
- Peyre, J., Sivic, J., Laptev, I., Schmid, C.: Weakly-supervised learning of visual relations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2017)
- Qi, M., Li, W., Yang, Z., Wang, Y., Luo, J.: Attentive relational networks for mapping images to scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Qi, M., Wang, Y., Li, A.: Online cross-modal scene retrieval by binary representation and semantic graph. In: Proceedings of the ACM International Conference on Multimedia (ACM MM) (2017)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. Advances in neural information processing systems (2015)
- Schuster, S., Krishna, R., Chang, A., Fei-Fei, L., Manning, C.D.: Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In: Proceedings of the fourth workshop on vision and language (2015)

- 18 J. Yang et al.
- 52. Shi, J., Zhang, H., Li, J.: Explainable and explicit visual reasoning over scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Suhail, M., Mittal, A., Siddiquie, B., Broaddus, C., Eledath, J., Medioni, G., Sigal, L.: Energy-based learning for scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
- Tamura, M., Ohashi, H., Yoshinaga, T.: Qpic: Query-based pairwise human-object interaction detection with image-wide contextual information. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
- 55. Tang, K.: A scene graph generation codebase in pytorch (2020), https://github.com/KaihuaTang/Scene-Graph-Benchmark.pytorch
- Tang, K., Niu, Y., Huang, J., Shi, J., Zhang, H.: Unbiased scene graph generation from biased training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- 57. Tang, K., Zhang, H., Wu, B., Luo, W., Liu, W.: Learning to compose dynamic tree structures for visual contexts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Wang, S., Duan, Y., Ding, H., Tan, Y.P., Yap, K.H., Yuan, J.: Learning transferable human-object interaction detector with natural language supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- 59. Wang, T., Anwer, R.M., Khan, M.H., Khan, F.S., Pang, Y., Shao, L., Laaksonen, J.: Deep contextual attention for human-object interaction detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
- 60. Wang, T., Yang, T., Danelljan, M., Khan, F.S., Zhang, X., Sun, J.: Learning human-object interaction detection using interaction points. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- 61. Wang, W.: Mmscenegraph (2021), https://github.com/Kenneth-Wong/ MMSceneGraph
- 62. Xiong, Y., Liao, R., Zhao, H., Hu, R., Bai, M., Yumer, E., Urtasun, R.: Upsnet: A unified panoptic segmentation network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- 63. Xu, D., Zhu, Y., Choy, C.B., Fei-Fei, L.: Scene graph generation by iterative message passing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Xu, P., Chang, X., Guo, L., Huang, P.Y., Chen, X., Hauptmann, A.G.: A survey of scene graph: Generation and application. IEEE Transactions on Neural Networks and Learning Systems (TNNLS) (2020)
- Yang, C.A., Tan, C.Y., Fan, W.C., Yang, C.F., Wu, M.L., Wang, Y.C.F.: Scene graph expansion for semantics-guided image outpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Yang, J., Lu, J., Lee, S., Batra, D., Parikh, D.: Graph r-cnn for scene graph generation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
- 67. Ye, K., Kovashka, A.: Linguistic structures as weak supervision for visual scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)

19

- Zareian, A., Karaman, S., Chang, S.F.: Bridging knowledge graphs to generate scene graphs. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Zareian, A., Wang, Z., You, H., Chang, S.: Learning visual commonsense for robust scene graph generation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Zellers, R., Yatskar, M., Thomson, S., Choi, Y.: Neural motifs: Scene graph parsing with global context. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Zhang, A., Liao, Y., Liu, S., Lu, M., Wang, Y., Gao, C., Li, X.: Mining the benefits of two-stage and one-stage hoi detection. Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2021)
- Zhang, F.Z., Campbell, D., Gould, S.: Efficient two-stage detection of humanobject interactions with a novel unary-pairwise transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Zhang, H., Kyaw, Z., Chang, S.F., Chua, T.S.: Visual translation embedding network for visual relation detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Zhang, J., Elhoseiny, M., Cohen, S., Chang, W., Elgammal, A.: Relationship proposal networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Zhang, W., Pang, J., Chen, K., Loy, C.C.: K-Net: Towards unified image segmentation. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2021)
- Zhong, Y., Shi, J., Yang, J., Xu, C., Li, Y.: Learning to generate scene graph from natural language supervision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
- 77. Zhou, T., Wang, W., Qi, S., Ling, H., Shen, J.: Cascaded human-object interaction recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Zou, C., Wang, B., Hu, Y., Liu, J., Wu, Q., Zhao, Y., Li, B., Zhang, C., Zhang, C., Wei, Y., et al.: End-to-end human object interaction detection with hoi transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)