

The Devil is in the Labels: Noisy Label Correction for Robust Scene Graph Generation

Lin Li¹, Long Chen^{2†}, Yifeng Huang¹, Zhimeng Zhang¹, Songyang Zhang³, Jun Xiao¹
¹Zhejiang University, ²Columbia University, ³University of Rochester

{mukti,yfhuang,zhimeng,junx}@zju.edu.cn zjuchenlong@gmail.com szhang83@ur.rochester.edu

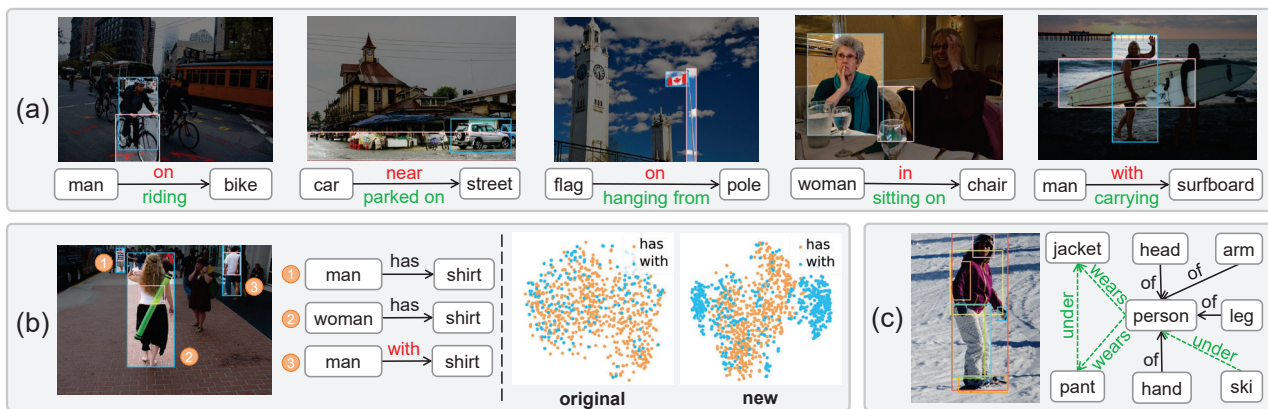


Figure 1: Illustration of three types of noisy annotations in SGG datasets (take VG as an example). (a) **Common-prone:** For some triplets, the annotators tend to select less informative coarse-grained predicates (**red**) instead of fine-grained ones (**green**). The subject and object for each triplet are denoted by **blue** and **pink** boxes, respectively. (b) **Synonym-random:** For some triplets, annotators usually randomly choose one predicate from the several synonyms (e.g., *has* and *with* are synonyms for $\langle \text{man}/\text{woman}/\text{shirt} \rangle$). *Original:* The t-SNE visualization of original triplets $\langle \text{man}/\text{has}/\text{with}/\text{shirt} \rangle$ features. For brevity, we randomly sample part of triplets for each type. *New:* The t-SNE visualization of same triplets after NICE. (c) **Negative:** Some negative triplets may not be background (**green** dash arrows).

Abstract

Unbiased SGG has achieved significant progress over recent years. However, almost all existing SGG models have overlooked the ground-truth annotation qualities of prevailing SGG datasets, i.e., they always assume: 1) all the manually annotated positive samples are equally correct; 2) all the un-annotated negative samples are absolutely background. In this paper, we argue that both assumptions are inapplicable to SGG: there are numerous “noisy” ground-truth predicate labels that break these two assumptions, and these noisy samples actually harm the training of unbiased SGG models. To this end, we propose a novel model-agnostic *Noisy* label *Corr*ection strategy for SGG: *NICE*. *NICE* can not only detect noisy samples but also reassign more high-quality predicate labels to them. After the *NICE* training, we can obtain a cleaner version of SGG dataset for

model training. Specifically, *NICE* consists of three components: negative Noisy Sample Detection (*Neg-NSD*), positive *NSD* (*Pos-NSD*), and Noisy Sample Correction (*NSC*). Firstly, in *Neg-NSD*, we formulate this task as an out-of-distribution detection problem, and assign pseudo labels to all detected noisy negative samples. Then, in *Pos-NSD*, we use a clustering-based algorithm to divide all positive samples into multiple sets, and treat the samples in the noisiest set as noisy positive samples. Lastly, in *NSC*, we use a simple but effective weighted *KNN* to reassign new predicate labels to noisy positive samples. Extensive results on different backbones and tasks have attested to the effectiveness and generalization abilities of each component of *NICE*.

1. Introduction

Scene Graph Generation (SGG), i.e., detecting all object instances and their pairwise visual relations, is a crucial step

[†] Corresponding author. This work started when Long Chen at ZJU. Codes available at: <https://github.com/muktilin/NICE>.

towards comprehensive visual scene understanding. In general, each scene graph is a visually-grounded graph, where each node and edge refer to an object and visual relation, respectively. Recently, with the release of several large-scale SGG benchmarks (e.g., Visual Genome (VG) [15]) and advanced object detectors [28, 1, 35], SGG has received unprecedented attention [7]. However, due to the compositional nature of pairwise visual relations, the number distributions of different triplets in SGG datasets are much more imbalanced (i.e., long-tailed) than other recognition tasks. Accordingly, the performance of many state-of-the-art SGG models [42, 2, 32, 23] degrades significantly on the *tail* categories¹ compared to the *head* categories counterparts.

Currently, the mainstream solutions to mitigate the long-tailed problem in SGG can be coarsely categorized into two types: 1) *Re-balancing strategy*: It utilizes class-aware sample re-sampling or loss re-weighting to balance the proportions of different predicate categories in the network training. The former attempts to balance the number of training samples in instance-level² or image-level [17], and the latter leverages prior commonsense knowledge (e.g., frequency of predicates [22], predicate correlations [39], or rule-based predicate priority [23, 14]) to re-weight the contributions of different categories in loss calculations. 2) *Biased-model-based strategy*: It infers debiased predictions from pre-trained biased SGG models. For instance, using counterfactual causality to disentangle frequency biases [31], deriving more balanced loss weights for different predicates [41], or adjusting the probabilities of predicate predictions [4].

Although these methods have dominated performance on debiasing metrics (e.g., mean Recall@K), it is worth noting that almost all existing models have taken two plausible assumptions about the ground-truth annotations for granted:

Assumption 1: *All the manually annotated positive samples are equally correct.*

Assumption 2: *All the un-annotated negative samples are absolutely background.*

For the first assumption, by “equally”, we mean that the confidence (or quality) of annotated ground-truth predicate label for each positive sample² is the same as others, i.e., all positive predicate labels are of high quality. Unfortunately, unlike other close-set classification tasks where each sample has only a unique ground-truth label, a subject-object pair in SGG sometimes has multiple reasonable predicates. This phenomenon has led to two inevitable annotation characteristics in SGG datasets: 1) *Common-prone*: When these reasonable relations are in different semantic granularities,

¹For brevity, we directly use “tail”, “body”, and “head” categories to represent the predicate categories in the tail, body, and head parts of the number distributions of different predicates in SGG datasets, respectively.

²We use “instance” to denote an instance of visual relation triplet, and we also use “sample” to represent the triplet instance interchangeably.

the annotators tend to select the most common predicate (or coarse-grained) as ground-truth. As shown in Figure 1(a), both `riding` and `on` are “reasonable” for `man` and `bike`, but the annotated ground-truth predicate is less informative `on` instead of more convincing `riding`. And this characteristic is very common in SGG datasets (more examples in Figure 1(a)). 2) *Synonym-random*: When these reasonable relations are synonymous for the subject-object pair, the annotators usually randomly choose one predicate as ground-truth, i.e., the annotations for some similar visual patterns are inconsistent. For example, in Figure 1(b), both `has` and `with` denote “be dressed in” for `man/woman` and `shirt`, but the ground-truth annotations are inconsistent even in the same image. We further visualize thousands of sampled instances of `(man-has / with-shirt)` in VG, and these instances are all randomly distributed in the feature space (cf. Figure 1(b)). Thus, we argue that all the positive samples are NOT equally correct, i.e., a part of positive samples are not high-quality — their labels can be more fine-grained (cf. *common-prone*) or more consistent (cf. *synonym-random*).

For the second assumption, although all SGG works have agreed that visual relations in existing datasets are always sparsely identified and annotated [25] (Figure 1(c)), almost all of them still train their models by regarding all the un-annotated pairs as `background`, i.e., there is no visual relation between the subject and object. In contrast, we argue that all negative samples are NOT absolutely background, i.e., a part of negative samples are not high-quality — they are actually foreground with missing annotations.

In this paper, we try to get rid of these two questionable assumptions, and reformulate SGG as a noisy label learning problem. To the best of our knowledge, we are the first work to take a deep dive into the ground-truth annotation qualities of both positive and negative samples in SGG. Specifically, we propose a novel model-agnostic NoIsy label CorrEction strategy, dubbed as NICE. NICE can not only detect numerous noisy samples, but also reassign more high-quality predicate labels to them. By “noisy”, we mean that these samples break these two assumptions. After the NICE training, we can obtain a cleaner version of dataset for SGG training. Particularly, we can: 1) increase the number of fine-grained predicates (*common-prone*); 2) decrease annotation inconsistency among similar visual patterns (*synonym-random*); 3) increase the number of positive samples (*assumption 2*).

NICE consists of three components: negative noisy sample detection (Neg-NSD), positive noisy sample detection (Pos-NSD), and noisy sample correction (NSC). Firstly, in Neg-NSD, we reformulate the negative NSD as an out-of-distribution (OOD) detection problem, i.e., regarding all the positive samples as in-distribution (ID) training data, and all the un-annotated negative samples as OOD test data. In this way, we can detect the missing annotated (ID) samples with pseudo labels. Then, in Pos-NSD, we use a clustering-based

algorithm to divide all positive samples (including the outputs of Neg-NSD) into multiple sets, and regard samples in the noisiest set as noisy positive samples. The clustering results are based on the local density of each sample. Lastly, in NSC, we use a simple but effective weighted KNN to reassign new predicate labels to all noisy positive samples.

We evaluate NICE on the most prevalent SGG benchmark: VG [15]. Since NICE only focuses on refining noisy annotations of the dataset, it can be seamlessly incorporated into any SGG architecture to boost their performance. Extensive ablations have attested to the effectiveness and generalization abilities of each component of NICE.

In summary, we make three contributions in this paper:

1. We are the first to reformulate SGG as a noisy label learning problem, and point out the two plausible assumptions are not applicable for SGG, *i.e.*, the devil is in the labels.
2. We propose a novel model-agnostic strategy NICE. Extensive ablations on several baselines, tasks, and metrics have demonstrated its excellent generalization abilities.
3. Each part of NICE can serve as an independent plug-and-play module to improve SGG annotation qualities³.

2. Related Work

Scene Graph Generation. SGG aims to transform visual data into semantic graph structures. Early methods [25, 43] always ignore the visual context, *i.e.*, they regard each object as an individual and predict pairwise relations directly. Subsequent SGG works start to utilize the overlooked visual context by resorting to different advanced techniques, *e.g.*, message passing [37, 18, 42, 3, 2], or tree/graph structure modelling [40, 32]. Recently, unbiased SGG has drawn unprecedented attention, *i.e.*, they focus on the performance gaps in different predicate categories. As above mentioned, existing unbiased SGG models can be categorized into: rebalancing strategy [17, 22, 39, 23, 14] and biased-model-based strategy [31, 41, 4, 10]. Different from existing SGG works, we are the first to explicitly refine the original noisy ground-truth annotations on SGG datasets. Although some previous works also have discussed the issue of sparse annotations [34, 4] or semantic imbalance [10], they still heavily rely on these original noisy annotations in model training.

Learning with Noisy Labels. Existing noisy label learning methods can be roughly divided into two categories: 1) Utilizing an explicit or implicit noise model to estimate the distribution of noisy and clean labels, and then deleting or correcting the noise samples. These models can be: neural networks [8, 13, 16, 27], conditional random field [33] or knowledge graphs [20]. However, they always need abundant clean samples for training, which is always inapplicable for many noisy label learning datasets. 2) Constructing a

³For example, the Pos-NSD can help models gain good results with much fewer training samples, and the Neg-NSD can generate plentiful unseen reasonable visual triplets. More details are left in Sec. 4 and appendix.

more balanced loss function to reduce the influence of noisy samples [26, 44, 36, 38]. In this paper, we are the first to formulate SGG as a noisy label learning problem, and propose a novel noisy sample detection and correction strategy.

3. Approach

Given an image dataset \mathcal{I} , SGG task aims to convert each image $\mathcal{I}_i \in \mathcal{I}$ into a graph $\mathcal{G}_i = \{\mathcal{N}_i, \mathcal{E}_i\}$, where \mathcal{N}_i and \mathcal{E}_i denote the node set (*i.e.*, objects) and edge set (*i.e.*, visual relations) of image \mathcal{I}_i , respectively. In general, each graph \mathcal{G}_i can also be viewed as a set of visual relation triplets (*i.e.*, $\langle \text{subject-predicate-object} \rangle$), denoted as \mathcal{T}_i . For each triplet set \mathcal{T}_i , we can further divide it into two subsets: \mathcal{T}_i^+ and \mathcal{T}_i^- , where \mathcal{T}_i^+ denotes all the annotated positive triplets (or samples) in image \mathcal{I}_i , and \mathcal{T}_i^- denotes all the un-annotated negative triplets in image \mathcal{I}_i . Analogously, we use $\mathcal{T}^+ = \{\mathcal{T}_i^+\}$ and $\mathcal{T}^- = \{\mathcal{T}_i^-\}$ to represent all positive and negative triplets in the whole dataset \mathcal{I} .

The whole pipeline of NICE is illustrated in Figure 2⁴. In this section, we sequentially introduce each part of NICE, including negative noisy sample detection (Neg-NSD), positive NSD (Pos-NSD), and noisy sample correction (NSC). Specifically, given an image and its corresponding ground-truth triplet annotations (*i.e.*, $\mathcal{T}^+ \cup \mathcal{T}^-$)⁵, we first use the Neg-NSD to detect all possible noisy negative samples, *i.e.*, missing annotated foreground triplets. The \mathcal{T}^- can be divided into $\mathcal{T}_{\text{clean}}^-$ and $\mathcal{T}_{\text{noisy}}^-$. Meanwhile, Neg-NSD will assign pseudo positive predicate labels for all samples in $\mathcal{T}_{\text{noisy}}^-$ (*e.g.*, painted on for $\langle \text{letter-window} \rangle$ in Figure 2). The $\mathcal{T}_{\text{noisy}}^-$ with pseudo positive labels and original \mathcal{T}^+ compose a new positive set $\tilde{\mathcal{T}}^+$. Then, we use the Pos-NSD to detect all possible noisy positive samples in $\tilde{\mathcal{T}}^+$, *i.e.*, the positive samples which suffer from either common-prone or synonym-random characteristics (*e.g.*, at for $\langle \text{women-laptop} \rangle$ in Figure 2). Similarly, $\tilde{\mathcal{T}}^+$ can be divided into $\tilde{\mathcal{T}}_{\text{clean}}^+$ and $\tilde{\mathcal{T}}_{\text{noisy}}^+$. Next, we use NSC to reassign more high-quality predicate labels to all samples in $\tilde{\mathcal{T}}_{\text{noisy}}^+$, denote as $\tilde{\mathcal{T}}_{\text{noisy} \rightarrow \text{clean}}^+$. Lastly, after processing the ground-truth triplet annotations of all images, we can obtain a cleaner version of dataset ($\tilde{\mathcal{T}}_{\text{clean}}^+ \cup \tilde{\mathcal{T}}_{\text{noisy} \rightarrow \text{clean}}^+ \cup \mathcal{T}_{\text{clean}}^-$) for SGG training.

3.1. Negative Noisy Sample Detection (Neg-NSD)

In this module, we aim to discover all possible *noisy* negative samples, *i.e.*, missing annotated visual relation triplets. Meanwhile, Neg-NSD also assigns a pseudo positive predicate label for each detected noisy negative sample. Due to the nature of missing annotations in existing negative samples, it is difficult to directly train and evaluate a binary clas-

⁴In Figure 2, we use a single image as input for a clear illustration. In real experiments, we directly process the whole dataset in each module.

⁵For brevity, we omit the subscripts i for image \mathcal{I}_i .

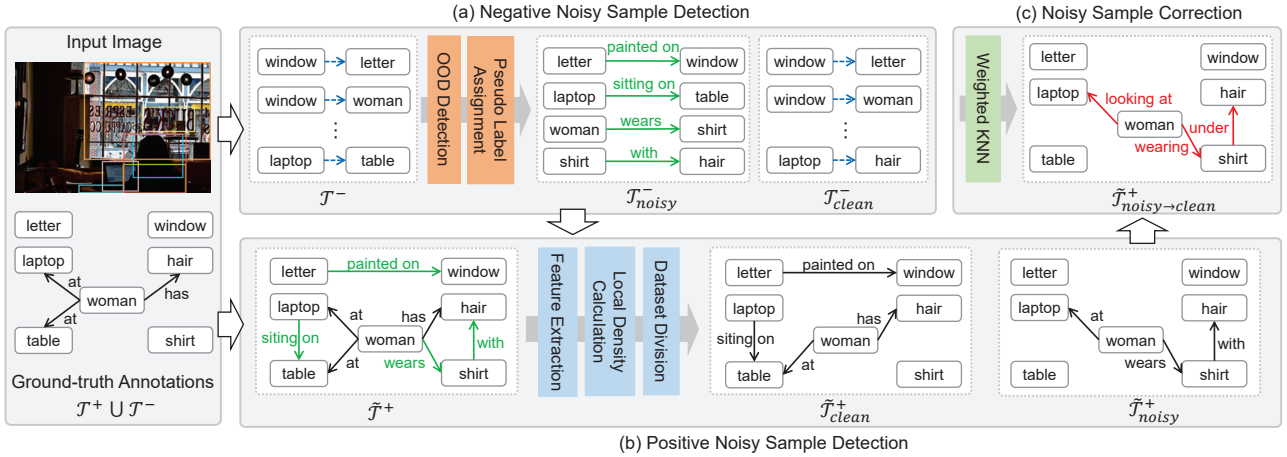


Figure 2: The pipeline of NICE (take an image from VG as an example). (a) **Neg-NSD:** Given all negative triplets (blue dash arrows), the OOD detection model detects missing annotated triplets ($\mathcal{T}_{\text{noisy}}^-$) and assigns pseudo labels to them (green predicates). (b) **Pos-NSD:** Given the new composed positive triplet set ($\tilde{\mathcal{T}}^+$), Pos-NSD detects all noisy positive samples ($\tilde{\mathcal{T}}_{\text{noisy}}^+$). (c) **NSC:** NSC reassigns more high-quality predicate labels to all noisy positive samples (red predicates). Finally, we obtain a new cleaner version of ground-truth annotations.

sifier based on these noisy samples. To this end, we propose to formulate the negative noisy sample detection as an out-of-distribution (OOD) detection problem [12]. Specifically, we regard all annotated positive samples as in-distribution (ID) training data, and all un-annotated negative samples as OOD test data. The Neg-NSD is built on top of a plain SGG model (denoted as F_{sgg}^n), but it is trained with only the annotated positive samples \mathcal{T}^+ . In the inference stage, Neg-NSD will predict a score of being foreground and a pseudo positive predicate category for each triplet $t_i^- \in \mathcal{T}^-$.

Following existing OOD detection methods [6], we also utilize a confidence-based model, *i.e.*, Neg-NSD consists of two network output branches: 1) a classification branch to predict a probability distribution \mathbf{p} over all positive predicate categories, and 2) a confidence branch to predict a confidence score $c \in [0, 1]$, which indicates the confidence of being an ID category (foreground). In the inference stage, for each sample t_i^- , if its confidence score c_i is larger than a threshold θ , we then regard this negative sample as a noisy negative sample, *i.e.*, the detection function $g(\cdot)$ is as:

$$g(t_i^-) = \begin{cases} 1, & c_i \geq \theta \\ 0, & c_i < \theta. \end{cases} \quad (1)$$

When $g(t_i^-) = 1$, the pseudo label of t_i^- is directly derived from the classification branch, *i.e.*, $\arg \max(\mathbf{p}_i)$. Since the predicted average confidence scores vary considerably for different predicate categories, we set different thresholds for head, body, and tail categories. (More details are in Sec. 4.)

Training of Neg-NSD. To train the classification branch and confidence branch, we combine predicted probabilities \mathbf{p}_i and corresponding target probability distribution \mathbf{y}_i , *i.e.*,

$$\mathbf{p}'_i = c_i \cdot \mathbf{p}_i + (1 - c_i) \cdot \mathbf{y}_i, \quad (2)$$




Figure 3: Left: Multidimensional scaling visualization of features of randomly sampled triplets with predicate *in*. Right: Detected detected green samples and noisy samples by Pos-NSD.

where \mathbf{p}'_i is the adjusted probabilities by confidence c_i . The motivation of Eq. (2) is that if the model is given a chance to ask for a hint of the ground-truth probability with some penalty, the model will definitely choose to ask for the hint if it is not confident about its output (*i.e.*, c_i is small). And the training objective for Neg-NSD consists of a weighted cross-entropy loss and a regularization penalty loss:

$$\mathcal{L} = -\sum_{j=1} w_j \log(p'_{ij})y_{ij} - \lambda \log(c_i), \quad (3)$$

where p'_{ij} and y_{ij} are the j -th element of \mathbf{p}'_i and \mathbf{y}_i , respectively. w_j is the reciprocal of the frequency of the j -th predicate category, which mitigates the impact of the long-tail issues on confidence. The penalty loss is utilized to prevent the network from always choosing $c = 0$ and using ground-truth probability distribution to minimize task loss.

3.2. Positive Noisy Sample Detection (Pos-NSD)

As shown in Figure 2⁴, the original positive set \mathcal{T}^+ and outputs of the Neg-NSD (*i.e.*, $\mathcal{T}_{\text{noisy}}^-$) compose a new positive sample set $\tilde{\mathcal{T}}^+$. The Pos-NSD module aims to detect all




Figure 4: Above: The multidimensional scaling visualization of features of randomly sampled triplets with predicate `in` with cut-off distance ranked at 50% (left) and 1% (right). Below: The triplet categories of the randomly sampled visual relation triplets from the corresponding red circle and green circle.

noisy samples in $\tilde{\mathcal{T}}^+$. In general, we use a clustering-based solution to divide all these positive samples into multiple subsets with different degrees of noise, and treat all samples in the noisiest subset as noisy positive samples. Intuitively, if a predicate label is consistent with other visually-similar samples of the same predicate category (*i.e.*, visual features of these samples are close to each other), this predicate is more likely to be a clean sample, because these annotations are consistent with each others. Otherwise, it is likely to be a noisy sample. As shown in Figure 3, the two clean triplets (`window-in-room`) have more visually-similar neighbors than the noisy triplets (*e.g.*, `plant-in-window`).

Based on these observations, we propose a local density based solution for positive noisy sample detection. Specifically, we utilize an off-the-shelf pretrained SGG model (denoted as F_{sgg}^p) to extract all visual relation triplet features, and use \mathbf{h}_i^k to represent the visual feature of i -th sample of predicate category k (this sample is denoted as t_i^k). Then, we utilize a distance matrix $\mathbf{D}^k = (d_{ij}^k)_{N \times N} \in \mathbb{R}^{N \times N}$ to measure the similarity between all positive samples of the same predicate k , and d_{ij}^k is calculated by:

$$d_{ij}^k = \|\mathbf{h}_i^k - \mathbf{h}_j^k\|^2, \quad (4)$$

where $\|\cdot\|$ is the Euclidean distance. Thus, a smaller distance d_{ij}^k means a relatively higher similarity between sample t_i^k and sample t_j^k . Then following [29], we define the local density ρ_i^k of each sample t_i^k as the number of samples (within the same predicate category) whose similarity distance to sample t_i^k are closer than a threshold d_c^k , *i.e.*,

$$\rho_i^k = \sum_j \mathbf{1}((d_c^k - d_{ij}^k) > 0), \quad (5)$$

where $\mathbf{1}(\cdot)$ is the indicator function and d_c^k is the cutoff distance for predicate k , which is ranked at $\alpha\%$ of sorted $N \times N$ distances in \mathbf{D}^k from small to large. Thus, a sample with higher local density ρ means that this sample is more

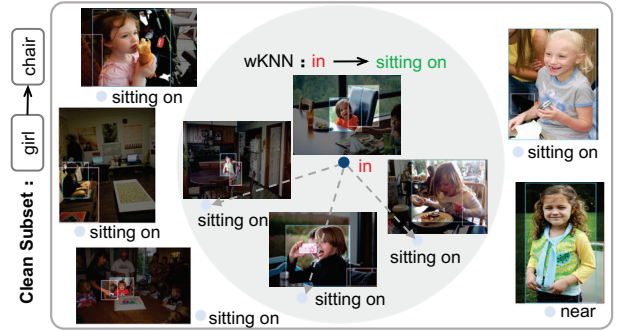


Figure 5: The illustration of NSC. Dashed lines indicate the distances between the noisy sample and other samples in clean subset with `girl-chair`. wKNN replaces noisy `in` to `sitting on`.

similar to the samples of the same predicate category. Analogously, samples with low local density ρ are considered as noisy samples. Finally, we use an unsupervised K-means algorithm [11] to divide all data samples into multiple subsets with respect to different ρ values, *i.e.*, different degrees of noise [9]. And all samples in the subset with the lowest ρ are regarded as noisy positive samples (*i.e.*, $\tilde{\mathcal{T}}_{\text{noisy}}^+$), and fed into the following NSC module for label correction.

Influence of the Cutoff Distance d_c^k . From Eq. (5), we can observe that the distribution of local density ρ is directly decided by the selection of the cutoff distance d_c (or the hyperparameter $\alpha\%$). As shown in Figure 4, when the cutoff distance ranked at 50% and 1%, local densities of samples diffuse outwards from large to small with one and two centers, respectively, *i.e.*, a smaller cutoff distance (*e.g.*, 1% for $\alpha\%$) may divide the whole feature space into more clusters. Meanwhile, different predicate categories may contain variable types of semantic meanings. For example, in Figure 4, predicate `in` of samples inside the red circle represents “inside” (*e.g.*, `plant-in-vase`), while predicate `in` of the samples inside the green circle represents “wearing” (*e.g.*, `man-in-shirt`). Thus, we set different cutoff distances to different categories. More details are in Sec. 4.

In addition, more detailed discussions about the influence of d_c^k on the clustering results are left in the appendix.

3.3. Noisy Sample Correction (NSC)

Given all the detected noisy positive samples from PosNSD, the NSC module aims to correct these noisy positive predicate labels. The motivation of our NSC is that the predicate label of a sample should be consistent with other visually-similar samples, especially for those samples with the same subject and object categories. For example, in Figure 5, for the noisy sample `girl-in-chair`, we can retrieval all other samples with the same `girl-chair`, and find most of the visually-similar samples are annotated as `girl-sitting on-chair`. Thus, we

use the simple but effective weighted K-Nearest Neighbor (wKNN) algorithm to derive the most possible labels for noisy positive samples. The wKNN assigns larger weights to the closest samples and smaller weights to the ones that are far away. Specifically, let $N(i)$ be the set of K neighbors of sample t_i , then the new assigned label for t_i is:

$$r'_i = \arg \max_v \sum_{t_j \in N(i)} w_{ij} \cdot \mathbf{1}(v = r_j), \quad (6)$$

where v is a predicate category, r_j is predicate label of t_j , and $\mathbf{1}(\cdot)$ is an indicator function. The weight w_{ij} is assigned to each neighbor, defined as $a \cdot \exp(-\frac{(d_{ij}-b)^2}{2c^2})$. The d_{ij} is the Euclidean distance between \mathbf{h}_i and \mathbf{h}_j (cf. Eq. (4)), and a, b, c are hyperparameters. It is worth noting that since we only consider the samples with the same subject-object categories, we can solve NSC with fast inference speed. If the new label is the same as the old, no label will be assigned.

4. Experiments

4.1. Experimental Settings and Details

Dataset. We conducted all experiments on the Challenging VG dataset [15]. It contains 108,073 images in total. In this paper, we followed widely-used splits [37], which keep the 150 most frequent object categories and the 50 most frequent predicate categories. Specifically, 70% of images are training set and 30% of images are test set. Following [42], we sampled 5,000 images from the training set as the val set. Besides, we followed [24] to divide all predicate categories into three parts based on the number of samples in training set: head ($>10k$), body ($0.5k \sim 10k$), and tail ($<0.5k$).

Tasks. We evaluated NICE on three SGG tasks [37]: 1) *Predicate Classification (PredCls)*: Given the ground-truth objects with labels, we need to only predict pairwise predicate categories. 2) *Scene Graph Classification (SGCls)*: Given the ground-truth object bounding boxes, we need to predict both the object categories and predicate categories. 3) *Scene Graph Generation (SGGen)*: Given an image, we need to detect all object bounding boxes, and predict both the object categories and predicate categories.

Metrics. We evaluated all results on three metrics: 1) *Recall@K (R@K)*: It calculates the proportion of top-K confident predicted relation triplets that are in the ground-truth. Following prior works, we used $K = \{50, 100\}$. 2) *mean Recall@K (mR@K)*: It calculates the recall for each predicate category separately, and then averages R@K over all predicates, *i.e.*, it puts relatively more emphasis on the tail categories. 3) **Mean**: It is the mean of all mR@K and R@K scores. R@K favors head predicates, while mR@K favors tail ones. Thus, it is a comprehensive metric that can better reflect model performance on different predicates.

Implementation Details. It is described in the appendix.

4.2. Comparisons with State-of-the-Arts

Settings. Since NICE is a model-agnostic strategy, it can be seamlessly incorporated into any advanced SGG model. In this section, we equipped NICE into two baselines: **Motifs** [42] and **VCTree** [32], and compared them with the state-of-the-art SGG methods. According to the generalization of these methods, we group them into two categories: 1) **TDE** [31], **PCPL** [39], **CogTree** [41], **DLFE** [4], and **BPL-SA** [10]. These methods are all model-agnostic SGG debiasing strategies. For fair comparisons, we also reported their performance on the Motifs and VCTree baselines. 2) **KERN** [3], **G-RCNN** [40], **MSDN** [19], **BGNN** [17], and **DT2-ACBS** [5]. These methods are specifically designed SGG models. All results are reported in Table 1.

Results. From the results in Table 1, we have the following observations: 1) Compared to the two strong baselines (*i.e.*, Motifs and VCTree), our NICE can consistently improve model performance on metric mR@K over all three tasks (*e.g.*, 6.7% \sim 14.5% and 4.7% \sim 14.6% absolute gains on metric mR@100 over Motifs and VCTree, respectively). 2) Compared to other state-of-the-art model-agnostic debiasing strategies, NICE can not only always achieve top performance on mR@K metrics, but also keep relatively high performance on R@K metrics, *i.e.*, NICE can improve the tail categories performance significantly, and maintain good performance on head categories. Thus, NICE can realize a better trade-off between accuracy among different predicate categories, and always achieve the best mean scores.

4.3. Ablation Studies

4.3.1 Ablation Studies on Neg-NSD

Hyperparameter Settings of Neg-NSD. The hyperparameter in Neg-NSD is the threshold θ for the confidence score (cf. Eq. (1)). Particularly, when the threshold θ for one category is set to 100%, which means we never assign this category as pseudo labels. Without loss of generality, we choose three representative hyperparameter settings, *i.e.*, we mine missing annotated triplets on 1) all predicate categories, or 2) only body and tail categories, or 3) only tail categories. The thresholds θ for corresponding head, body, and tail categories were set as 95%, 90%, and 60%, respectively. To disentangle the influence of the other two modules (*i.e.*, Pos-NSD and NSC), we directly use the outputs of Neg-NSD and pristine positive samples for SGG training.

Results. From the results in Table 3(a), we can observe: 1) Different threshold settings have a slight influence on the mR@K metrics, but relatively more influence on the R@K metrics. 2) When only mining the missing tail predicates in Neg-NSD, the model gains the best performance.

Effectiveness of Neg-NSD. We evaluated the effectiveness of the Neg-NSD by using the same refined samples by Neg-NSD and pristine positive samples for SGG training.

B	Models	PredCls			SGCls			SGGen		
		mR@50/100	R@50/100	Mean	mR@50/100	R@50/100	Mean	mR@50/100	R@50/100	Mean
VGG-16	Motif [42] _{CVPR'18}	14.0 / 15.3	65.2 / 67.1	40.4	7.7 / 8.2	35.8 / 36.5	22.1	5.7 / 6.6	27.2 / 30.3	17.5
	VCtree [32] _{CVPR'19}	17.9 / 19.4	66.4 / 68.1	43.0	10.1 / 10.8	38.1 / 38.8	24.5	6.9 / 8.0	27.9 / 31.3	18.5
	KERN [3] _{CVPR'19}	17.7 / 19.2	65.8 / 67.6	42.6	9.4 / 10.0	36.7 / 37.4	23.4	6.4 / 7.3	29.8 / 27.1	17.7
	PCPL [39] _{MM'20}	35.2 / 37.8	50.8 / 52.6	44.1	18.6 / 19.6	27.6 / 28.4	23.6	9.5 / 11.7	14.6 / 18.6	13.6
	MSDN [19] _{ICCV'17}	15.9 / 17.5	64.6 / 66.6	41.2	9.3 / 9.7	38.4 / 39.8	24.3	6.1 / 7.2	31.9 / 36.6	20.5
X-101-FPN	G-RCNN [40] _{ECCV'18}	16.4 / 17.2	64.8 / 66.7	41.3	9.0 / 9.5	38.5 / 37.0	23.5	5.8 / 6.6	29.7 / 32.8	18.7
	BGNN [17] _{CVPR'21}	30.4 / 32.9	59.2 / 61.3	45.9	14.3 / 16.5	37.4 / 38.5	26.7	10.7 / 12.6	31.0 / 35.8	22.5
	DT2-ACBS [5] _{ICCV'21}	35.9 / 39.7	23.3 / 25.6	31.1	24.8 / 27.5	16.2 / 17.6	21.5	22.0 / 24.4	15.0 / 16.3	19.4
	Motifs [42] _{CVPR'18}	16.5 / 17.8	65.5 / 67.2	41.8	8.7 / 9.3	39.0 / 39.7	24.2	5.5 / 6.8	32.1 / 36.9	20.3
	+TDE [31] _{CVPR'20}	24.2 / 27.9	45.0 / 50.6	36.9	13.1 / 14.9	27.1 / 29.5	21.2	9.2 / 11.1	17.3 / 20.8	14.6
	+PCPL [39] _{MM'20}	24.3 / 26.1	54.7 / 56.5	40.4	12.0 / 12.7	35.3 / 36.1	24.0	10.7 / 12.6	27.8 / 31.7	20.7
	+CogTree [41] _{IJCAI'21}	26.4 / 29.0	35.6 / 36.8	32.0	14.9 / 16.1	21.6 / 22.2	18.7	10.4 / 11.8	20.0 / 22.1	16.1
	+DLFE [4] _{MM'21}	26.9 / 28.8	52.5 / 54.2	40.6	15.2 / 15.9	32.3 / 33.1	24.1	11.7 / 13.8	25.4 / 29.4	20.1
	+BPL-SA [10] _{ICCV'21}	29.7 / 31.7	50.7 / 52.5	41.2	16.5 / 17.5	30.1 / 31.0	23.8	13.5 / 15.6	23.0 / 26.9	19.8
	+NICE (ours)	29.9 / 32.3	55.1 / 57.2	43.6	16.6 / 17.9	33.1 / 34.0	25.4	12.2 / 14.4	27.8 / 31.8	21.6
	VCtree [32] _{CVPR'19}	17.1 / 18.4	65.9 / 67.5	42.2	10.8 / 11.5	45.6 / 46.5	28.6	7.2 / 8.4	32.0 / 36.2	20.9
	+TDE [31] _{CVPR'20}	26.2 / 29.6	44.8 / 49.2	37.5	15.2 / 17.5	28.8 / 32.0	23.4	9.5 / 11.4	17.3 / 20.9	14.8
	+PCPL [39] _{MM'20}	22.8 / 24.5	56.9 / 58.7	40.7	15.2 / 16.1	40.6 / 41.7	28.4	10.8 / 12.6	26.6 / 30.3	20.1
	+CogTree [41] _{IJCAI'21}	27.6 / 29.7	44.0 / 45.4	36.7	18.8 / 19.9	30.9 / 31.7	25.3	10.4 / 12.1	18.2 / 20.4	15.3
	+DLFE [4] _{MM'21}	25.3 / 27.1	51.8 / 53.5	39.4	18.9 / 20.0	33.5 / 34.6	26.8	11.8 / 13.8	22.7 / 26.3	18.7
+BPL-SA [10] _{ICCV'21}	30.6 / 32.6	50.0 / 51.8	41.3	20.1 / 21.2	34.0 / 35.0	27.6	13.5 / 15.7	21.7 / 25.5	19.1	
+NICE (ours)	30.7 / 33.0	55.0 / 56.9	43.9	19.9 / 21.3	37.8 / 39.0	29.5	11.9 / 14.1	27.0 / 30.8	21.0	

Table 1: Performance (%) of state-of-the-art SGG models on three SGG tasks. “B” denotes the backbone of object detector (Faster R-CNN [28]) in each SGG model: *i.e.*, VGG-16 [30] and ResNeXt-101-FPN [21]. “Mean” is the average of mR@50/100 and R@50/100. The **best** and **second best** methods under each setting are marked according to formats.

#	Components			PredCls		
	N-NSD	P-NSD	NSC	mR@50/100	R@50/100	Mean
1	✗	✗	✗	16.5 / 17.8	65.5 / 67.2	41.8
2	✓	✗	✗	23.3 / 25.2	62.3 / 64.5	43.8
3	✗	✓	✗	20.3 / 22.0	57.6 / 59.2	39.8
4	✓	✓	✗	20.4 / 23.4	56.7 / 61.4	40.5
5	✗	✓	✓	23.3 / 25.2	59.6 / 61.3	42.4
6	✓	✓	✓	29.9 / 32.3	55.1 / 57.2	43.6

Table 2: Ablation studies on each component of NICE. “#” is the line number. The baseline model (# 1) is Motifs [42].

Results. All results are reported in Table 2. Compared to the baseline model (# 1), the Neg-PSD module (# 2) can significantly improve model performance on mR@K metrics (*e.g.*, 25.2% vs. 17.8% in mR@100), which proves that these harvested “positive” samples (noisy negative samples with pseudo labels) are indeed beneficial for SGG training.

4.3.2 Ablation Studies on Pos-NSD

Hyperparameter Settings of Pos-NSD. The hyperparameter in Pos-NSD is the cutoff distance d_c ranked at $\alpha\%$ for different categories (cf. Eq. (5)). As mentioned in Sec. 3.2, different d_c directly influence the clustering results of each

predicate category, and smaller d_c is more suitable for predicates with multiple semantic meanings. Thus, without loss of generality, we choose five typical settings for different categories, including $(\mathcal{L}, \mathcal{L}, \mathcal{L})$, $(\mathcal{M}, \mathcal{M}, \mathcal{M})$, and so on. The \mathcal{L} , \mathcal{M} , \mathcal{S} denote the d_c are large, medium, and small. In our experiments, their corresponding $\alpha\%$ were set to 50.0%, 25.0% and 12.5%. Similarly, we disentangle the influence of Neg-NSD, and use pristine negative samples and refined positive samples (outputs of NSC) for SGG training.

Results. When the cutoff distance d_c for all predicate categories are set to large or small or from large to small, the model achieves relatively worse results. These results are also consistent with our expectation, *i.e.*, for the predicate categories with multiple semantic meanings (head categories), small d_c is better for noisy sample detection. Instead, for predicate categories with unique semantic meaning (tail categories), larger d_c is better. Thus, we utilize the $(\mathcal{S}, \mathcal{M}, \mathcal{L})$ setting for all experiments.

Effectiveness of Pos-NSD. We evaluated the effectiveness of Pos-NSD by just using clean positive samples detected from \mathcal{T}^+ with pristine negative samples or from the new positive samples set $\tilde{\mathcal{T}}^+$ for SGG training.

Results. As shown in Table 2 (# 3), the single Pos-NSD component can still improve mR@K metrics by using much fewer positive training samples. Besides, baseline can also

Categories			PredCls		
head	body	tail	mR@50/100	R@50/100	Mean
✓	✓	✓	22.2 / 25.8	58.5 / 62.5	42.3
✗	✓	✓	23.7 / 26.2	59.5 / 62.3	42.9
✗	✗	✓	23.3 / 25.2	62.3 / 64.5	43.8

(a) The ablation studies on mining missing annotated triplets in head, body and tail categories in Neg-NSD, respectively.

Size of d_c			PredCls		
head	body	tail	mR@50/100	R@50/100	Mean
\mathcal{L}	\mathcal{L}	\mathcal{L}	19.9 / 21.5	64.0 / 65.7	42.8
\mathcal{M}	\mathcal{M}	\mathcal{M}	21.0 / 22.8	62.2 / 64.0	42.5
\mathcal{S}	\mathcal{S}	\mathcal{S}	21.6 / 23.4	61.1 / 62.8	42.2
\mathcal{S}	\mathcal{M}	\mathcal{L}	23.3 / 25.2	59.6 / 61.3	42.4
\mathcal{L}	\mathcal{M}	\mathcal{S}	18.6 / 20.1	64.4 / 66.1	42.3

(b) The ablation studies on different cutoff distance d_c in Pos-NSD for head, body and tail predicates respectively.

K	PredCls		
	mR@50/100	R@50/100	Mean
$K = 1$	23.3 / 25.0	59.3 / 60.9	42.1
$K = 3$	23.3 / 25.2	59.6 / 61.3	42.4
$K = 5$	22.9 / 24.7	59.8 / 61.5	42.2

(c) The ablation studies on different K in wKNN.

Table 3: Ablation studies on the influence of different hyperparameters of each component of NICE. Motifs [42] is the baseline model which are used in all experiments.

be exceeded on mR@K (# 4) after Neg-NSD and Pos-NSD alone with fewer training samples. It proves that the clean subset divided by Pos-NSD is better for SGG training, and numerous noisy positive samples actually hurt performance.

4.3.3 Ablation Studies on NSC

Hyperparameter of NSC. The hyperparameter of NSC is the K in wKNN. We investigated $K = \{1, 3, 5\}$. All results are reported in Table 3(c). From the results, we can observe that the SGG performance is robust to different K . To better trade-off different metrics, we set K to 3.

Effectiveness of NSC. Based on Pos-NSD, NSC replaces the original labels of the noisy positive samples with cleaner labels. As shown in Table 2, NSC markedly improves SGG performance on both mR@K and R@K (# 5 vs. # 3), *i.e.*, NSC always reassigns better labels to noisy samples.

T-SNE Visualizations. We visualized the t-SNE distributions of features of $\langle \text{man-on/riding-skateboard} \rangle$ as well as $\langle \text{man-has/with-arm} \rangle$ on the feature space before and after the NSC in Figure 6. And the features are from the previous layer of classification, like P-NSD. The former is a pair of predicates of different granularity (common-prone), and the latter is a pair of predicates of the same granularity (synonym-random). As shown in Figure 6, NSC can help to alleviate the inconsistency of ground-truths, *i.e.*, similar visual patterns always have more consistent ground-truth predicate annotations, which is beneficial for SGG training.




Figure 6: The t-SNE visualization of randomly sampled instances of different triplet categories on the feature space before and after NSC.

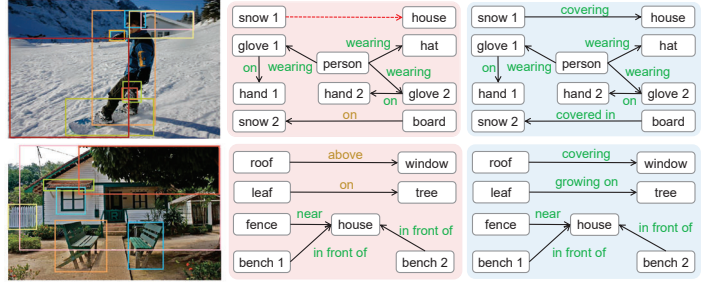


Figure 7: Scene graphs generated by Motifs (left) and Motifs+NICE (right) on PredCls. Green predicates are correct (*i.e.*, GT), and brown predicates are more coarse-grained (not in GT). Red dashed arrows are the relations that are not detected. Only detected boxes overlapped with GT are shown.

4.3.4 Qualitative Results

Figure 7 demonstrates some qualitative results generated by Motifs and Motifs+NICE. From Figure 7, we can observe that NICE can not only help to detect more missing false positive predicates (*e.g.*, covering), but also more fine-grained and informative predicates (*e.g.*, growing on vs. on, and covering vs. above).

5. Conclusions and Limitations

In this paper, we argued that two plausible assumptions about the ground-truth annotations are inapplicable for existing SGG datasets. To this end, we reformulated SGG as a noisy label learning problem and proposed a novel model-agnostic noisy label correction strategy: NICE. NICE can not only detect noisy samples, but also reassign better predicate labels to them. We have validated the effectiveness of each part of NICE through extensive experiments.

Limitations. Although NICE can mine some potential unseen triplets in Neg-NSD, there is no guarantee that the harvested triplets must be reasonable. Meanwhile, some hyperparameters have different impacts on each predicate category, which makes it difficult to achieve the best trade-off between different metrics (cf. mR@K & R@K in Table 3).

Acknowledgement This work was supported by National Key Research & Development Project of China (2021ZD0110700), National Natural Science Foundation of China (U19B2043, 61976185), Zhejiang Natural Science Foundation (LR19F020002), Zhejiang Innovation Foundation (2019R52002), and Fundamental Research Funds for Central Universities.

References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020. [2](#)
- [2] Long Chen, Hanwang Zhang, Jun Xiao, Xiangnan He, Shiliang Pu, and Shih-Fu Chang. Counterfactual critic multi-agent training for scene graph generation. In *ICCV*, pages 4613–4623, 2019. [2, 3](#)
- [3] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *CVPR*, pages 6163–6171, 2019. [3, 6, 7](#)
- [4] Meng-Jiun Chiou, Henghui Ding, Hanshu Yan, Changhu Wang, Roger Zimmermann, and Jiashi Feng. Recovering the unbiased scene graphs from the biased ones. In *ACM MM*, 2021. [2, 3, 6, 7](#)
- [5] Alakh Desai, Tz-Ying Wu, Subarna Tripathi, and Nuno Vasconcelos. Learning of visual relations: The devil is in the tails. In *ICCV*, pages 15404–15413, 2021. [6, 7](#)
- [6] Terrance DeVries and Graham W Taylor. Learning confidence for out-of-distribution detection in neural networks. In *arXiv*, 2018. [4](#)
- [7] Kaifeng Gao, Long Chen, Yulei Niu, Jian Shao, and Xiao Jun. Classification-then-grounding: Reformulating video scene graphs as temporal bipartite graphs. In *CVPR*, 2022. [2](#)
- [8] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *ICLR*, 2017. [3](#)
- [9] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. [5](#)
- [10] Yuyu Guo, Lianli Gao, Xuanhan Wang, Yuxuan Hu, Xing Xu, Xu Lu, Heng Tao Shen, and Jingkuan Song. From general to specific: Informative scene graph generation via balance adjustment. In *ICCV*, pages 16383–16392, 2021. [3, 6, 7](#)
- [11] Klaus Hechenbichler and Klaus Schliep. Weighted k-nearest-neighbor techniques and ordinal classification. 2004. [5](#)
- [12] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv*, 2016. [4](#)
- [13] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313, 2018. [3](#)
- [14] Boris Knyazev, Harm de Vries, Cătălina Cangea, Graham W Taylor, Aaron Courville, and Eugene Belilovsky. Graph density-aware losses for novel compositions in scene graph generation. In *arXiv*, 2020. [2, 3](#)
- [15] 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. *IJCV*, 2017. [2, 3, 6](#)
- [16] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, pages 5447–5456, 2018. [3](#)
- [17] Rongjie Li, Songyang Zhang, Bo Wan, and Xuming He. Bipartite graph network with adaptive message passing for unbiased scene graph generation. In *CVPR*, pages 11109–11119, 2021. [2, 3, 6, 7](#)
- [18] Yikang Li, Wanli Ouyang, Xiaogang Wang, and Xiao’ou Tang. Vip-cnn: Visual phrase guided convolutional neural network. In *CVPR*, pages 1347–1356, 2017. [3](#)
- [19] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, pages 1261–1270, 2017. [6, 7](#)
- [20] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *ICCV*, pages 1910–1918, 2017. [3](#)
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. [7](#)
- [22] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. [2, 3](#)
- [23] Xin Lin, Changxing Ding, Jinquan Zeng, and Dacheng Tao. Gps-net: Graph property sensing network for scene graph generation. In *CVPR*, pages 3746–3753, 2020. [2, 3](#)
- [24] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, pages 2537–2546, 2019. [6](#)
- [25] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, pages 852–869, 2016. [2, 3](#)
- [26] Xingjun Ma, Yisen Wang, Michael E Houle, Shuo Zhou, Sarah Erfani, Shutao Xia, Sudanthi Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *ICML*, pages 3355–3364, 2018. [3](#)
- [27] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pages 4334–4343, 2018. [3](#)
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015. [2, 7](#)
- [29] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. *Science*, pages 1492–1496, 2014. [5](#)
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv*, 2014. [7](#)
- [31] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *CVPR*, pages 3716–3725, 2020. [2, 3, 6, 7](#)
- [32] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *CVPR*, pages 6619–6628, 2019. [2, 3, 6, 7](#)

- [33] Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *NeurIPS*, 2017. 3
- [34] Tzu-Jui Julius Wang, Selen Pehlivan, and Jorma Laaksonen. Tackling the unannotated: Scene graph generation with bias-reduced models. In *BMVC*, 2020. 3
- [35] Wenxiao Wang, Lu Yao, Long Chen, Binbin Lin, Deng Cai, Xiaofei He, and Wei Liu. Crossformer: A versatile vision transformer hinging on cross-scale attention. In *ICLR*, 2022. 2
- [36] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *ICCV*, pages 322–330, 2019. 3
- [37] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, pages 5410–5419, 2017. 3, 6
- [38] Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L_{dmi}: An information-theoretic noise-robust loss function. In *arXiv*, 2019. 3
- [39] Shaotian Yan, Chen Shen, Zhongming Jin, Jianqiang Huang, Rongxin Jiang, Yaowu Chen, and Xian-Sheng Hua. Pcpl: Predicate-correlation perception learning for unbiased scene graph generation. In *ACM MM*, pages 265–273, 2020. 2, 3, 6, 7
- [40] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *ECCV*, pages 670–685, 2018. 3, 6, 7
- [41] Jing Yu, Yuan Chai, Yue Hu, and Qi Wu. Cogtree: Cognition tree loss for unbiased scene graph generation. In *IJCAI*, 2021. 2, 3, 6, 7
- [42] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, pages 5831–5840, 2018. 2, 3, 6, 7, 8
- [43] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *CVPR*, pages 5532–5540, 2017. 3
- [44] Zhilu Zhang and Mert R Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*, 2018. 3