

Improving Weakly Supervised Visual Grounding by Contrastive Knowledge Distillation

Liwei Wang^{1*} Jing Huang^{2*} Yin Li³ Kun Xu⁴ Zhengyuan Yang⁵ Dong Yu⁴

¹ The Chinese University of Hong Kong ² University of Illinois at Urbana-Champaign

³ University of Wisconsin-Madison ⁴ Tencent AI Lab, Bellevue ⁵ University of Rochester

Abstract

Weakly supervised phrase grounding aims at learning region-phrase correspondences using only image-sentence pairs. A major challenge thus lies in the missing links between image regions and sentence phrases during training. To address this challenge, we leverage a generic object detector at training time, and propose a contrastive learning framework that accounts for both region-phrase and image-sentence matching. Our core innovation is the learning of a region-phrase score function, based on which an image-sentence score function is further constructed. Importantly, our region-phrase score function is learned by distilling from soft matching scores between the detected object names and candidate phrases within an image-sentence pair, while the image-sentence score function is supervised by ground-truth image-sentence pairs. The design of such score functions removes the need of object detection at test time, thereby significantly reducing the inference cost. Without bells and whistles, our approach achieves state-of-the-art results on visual phrase grounding, surpassing previous methods that require expensive object detectors at test time.

1. Introduction

Visual phrase grounding — finding regions associated with phrases in a sentence description of the image, is an important problem at the intersection of computer vision and natural language processing. Most of the existing approaches [15, 39, 46] follow a fully supervised paradigm that requires the labeling of bounding boxes for each phrase. These fine-grained annotations are unfortunately expensive to obtain and thus difficult to scale. Consequently, weakly supervised grounding has recently received considerable attention [42, 51, 50, 56, 6, 53, 54, 14, 45]. In this setting, only images and their sentence descriptions are given at training time. At inference time, given an image sentence pair, a method is asked to link regions to sentence phrases.

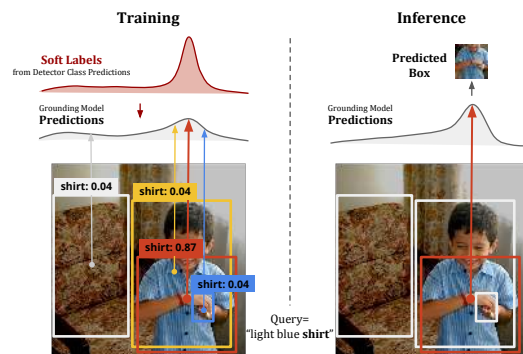


Figure 1: Our method uses object detector predictions to guide the learning of region-phrase matching in training. At the inference time, our method no longer requires object detectors and directly predicts the box with the highest score.

A major challenge of weakly supervised grounding is to distinguish among many “concurrent” visual concepts. For example, the region of a dog and that of its head are likely to co-occur in images associated with the phrase “a running puppy.” Without knowing the ground-truth region-phrase matching, learning to link the region of dog (but not dog head) to its corresponding phrase becomes very challenging. To address this challenge, recent methods [6, 14, 45, 18] leverage generic object detectors for training and/or inference. A detector provides high quality object regions, as well as their category labels that can be further matched to candidate phrases, thereby bringing in external knowledge about region-phrase matching and thus helping to disambiguate those “concurrent” concepts. However, it remains unclear about the best practices of using an external object detector for weakly supervised grounding.

In this paper, we focus on developing a principled approach to distill knowledge from a generic object detector for weakly supervised phrase grounding. To this end, we present a simple method under the framework of contrastive learning. Specifically, our model learns a score function between region-phrase pairs, guided by two levels of similarity constraints encoded using noise-contrastive estimation (NCE) loss [19] during training. The first level of region-

*equal contribution. Contact: lwwang@cse.cuhk.edu.hk

phrase similarity is distilled from object detection outputs. This is done by aligning predicted region-phrase scores to a set of soft targets, computed by matching object names and candidate phrases. The second level of image-sentence similarity is computed from a greedy matching between all region-phrase pairs, and supervised by ground-truth image-sentence pairs. During inference, our method compares each image region to candidate phrases using the learned score function, without the need of object detection. Our training and inference stages are shown in Fig. 1.

To evaluate our method, we conduct extensive experiments on Flickr30K Entities [39] and ReferItGame [26] datasets. We compare our results to the latest methods of weakly supervised phrase grounding. Our experiments show that our method establishes new state-of-the-art results and outperforms all previous methods, including those using strong object detectors at test time [14, 46, 50] or using a similar contrastive loss [18]. For example, on Flickr30K Entities, without additional training data, our method outperforms the best reported results by a large margin. On ReferItGame, our method significantly beats the best reported results. Moreover, we systematically vary the components of our model and demonstrate several best practices for weakly supervised phrase grounding. We hope that our simple yet strong method will shed light on new ideas and practices for weakly supervised image-text grounding.

2. Related Work

We discuss relevant works on weakly supervised phrase grounding, and provide a brief review of recent works on contrastive learning and knowledge distillation — the two main pillars of our method.

Visual Grounding of Phrases. Grounding of textual phrases, also referred to as phrase localization, is an important problem in vision and language. Several datasets, e.g., Flickr30K Entities [39], ReferItGame [26] and Visual Genome [29], have been constructed to capture dense phrase-region correspondences. Building on these datasets, many recent approaches learn a similarity function between regions and phrases by using the ground-truth region-phrase pairs [46, 15, 49, 12, 38, 3, 48]. More recent works considered phrase grounding in videos [57]. This fully supervised paradigm has shown impressive results, yet requires labor-intensive annotations of bounding boxes for all phrases.

Weakly Supervised Phrase Grounding. Weakly supervised phrase grounding has thus received considerable attention recently [42, 50, 47, 6, 56, 14, 45, 17]. In this setting, a method learns from only images and paired sentence descriptions, without explicit region-to-phrase correspondence.

Recent works [6, 50, 51] show that weakly supervised phrase localization can benefit from side information, such as object segmentation or detection. For example, Chen et al. [6] leveraged pre-trained deep models and proposed to

enforce visual and language consistency. Yeh et al. [50] proposed to link words in text and detection classes using co-occurrence statistics from paired captions. Moreover, Xiao et al. [47] investigated the linguistic structure of the sentences. They proposed a structure loss to model the compositionality of the phrases and their attention masks. Zhao et al. [56] presented a model that jointly learns to propose object regions and matches the regions to phrases. Fang et al. [14] explored the weakly supervised grounding by decomposing the problem into several modules and taking additional information, like the color module, to improve the performance. More recently, Wang et al. [45] made use of off-the-shelf models to detect objects, scenes and colors in images, and achieves the goal of grounding via measuring semantic similarity between the categories of detected visual elements and the sentence phrases. Gupta et al. [18] proposed to use contrastive loss for weakly supervised phrase grounding. Datta et al. [9] adopted image-sentence retrieval to guide the phrase localization.

Our method shares the key idea behind these previous works—we seek an explicit alignment between regions and phrases given image-sentence pairs. Our method differs from previous works by explicitly modeling the knowledge distillation from external off-the-shelf object detectors into a unified contrastive learning framework. The most relevant work to us is InfoGround [18]. Both our method and InfoGround [18] consider contrastive learning for weakly supervised grounding. Different from [18], our work moves beyond the contrastive loss and focuses on knowledge distillation from an external object detector under a contrastive learning framework. Our method is also different from WPT [45] though we both need to use object detectors: our approach learns the distillation in the training stage which makes our model free of detectors in the testing stage, while WPT requires using detectors during inference.

Contrastive Learning. There has been a recent trend of exploring contrastive loss for representation learning. For instance, Oord et al. [35] proposed Contrastive Predictive Coding (CPC) that learns representations for sequential data. Hjelm et al. [22] presented Deep InfoMax for unsupervised representation learning by maximizing the mutual information between the input and output of a deep network. More recently, Chen et al. [7] proposed to learn visual representations by maximizing agreement between differently augmented views of the same image via a contrastive loss. He et al. [20] proposed Momentum Contrast (MoCo) for unsupervised visual representation learning. Tian et al. [44] extends the input to more than two views. These methods are all based on a similar contrastive loss related to Noise Contrastive Estimation (NCE) [19]. The NCE loss has also been explored for phrase grounding by Gupta et al. [18] (InfoGround). Our work is relevant to these works since our mathematical framework is also built upon the general

idea of infoNCE [35] and NCE [19]. However, our work leverages a contrastive loss for knowledge distillation in the context of cross-view weakly supervised grounding.

Knowledge Distillation. Knowledge distillation was proposed and popularized by [5, 21, 2, 43, 52]. Several recent works [17, 16, 32, 30, 44] explored knowledge distillation for multi-modal learning. Knowledge distillation has also shown its effectiveness in various vision-language tasks, such as VQA [34, 11], grounded image captioning [59] and video captioning [36, 55]. Different from previous approaches, we consider knowledge distillation for region-phrase grounding by matching the outputs of a region-phrase score function to soft targets computed from object detection results.

3. Approach

Consider $X = [X_1, \dots, X_i, \dots, X_N]$ as the set of images and $Y = [Y_1, \dots, Y_j, \dots, Y_M]$ as the set of sentences. Each image i consists of a set of regions with their features $X_i = [x_i^1, \dots, x_i^l, \dots, x_i^n]$. Similarly, each sentence j includes multiple phrase features $Y_j = [y_j^1, \dots, y_j^k, \dots, y_j^m]$. Thus, i, j index images and sentences, and l, k index regions and phrases. Oftentimes, we have multiple sentences describing the same image and many more image regions than sentence phrases. Moreover, with minor abuse of notations, we denote $p(X_i, Y_j)$ as the probability of a valid image-sentence pair (X_i, Y_j) , i.e., $p(X_i, Y_j) = 1$ if and only if Y_j can be used to describe X_i . Similarly, we use $p(x_i^l, y_j^k)$ as the probability of a valid region-phrase pair (x_i^l, y_j^k) .

Our goal is to learn a score function that measures the similarity between region features x_i^l and phrase features y_j^k . However, we only have access to ground-truth image-sentence pairs $p(X_i, Y_j)$ without knowing the matching between regions and phrases $p(x_i^l, y_j^k)$. To address this challenge of *weakly supervised* grounding, we leverage a generic object detector \mathcal{D} to label candidate image regions, then generate “pseudo” labels of region-phrase correspondence by matching the region object labels to the sentence phrases. Therefore, our key innovation is the design of a contrastive loss that learns to distill from object detection outputs. A main advantage of using knowledge distillation is that our method no longer requires object detection at inference time and thus is very efficient during inference.

Fig. 2 presents an overview of our method. We now present the details of our method by first introducing the design of our score functions for image-text matching, followed by our contrastive learning loss using knowledge distillation.

3.1. Score Functions for Image-Text Matching

Our model builds on a two-branch network [46] for image-text matching at both region-phrase and image-sentence levels. The key idea is learning a score function to match region-phrase pairs. Based on the region-phrase matching scores, we further construct an image-sentence similarity

score. Specifically, our network has two branches f and g that take the inputs of region and phrase features x_i^l and y_j^k , respectively. Each branch is realized by a deep network by stacking multiple fully connected layers with ReLU activation in-between, followed by a L2 normalization at the end. We define the similarity between a region-phrase pair (x_i^l, y_j^k) as the cosine similarity between the transformed features $f(x_i^l)$ and $g(y_j^k)$, given by

$$s(x_i^l, y_j^k) = f(x_i^l)^T g(y_j^k). \quad (1)$$

We further aggregate the region-phrase matching scores $s(x_i^l, y_j^k)$ into a similarity score between a image-sentence pair (X_i, Y_j) , defined as

$$S(X_i, Y_j) = \sum_{k=1}^m \max_{1 \leq l \leq n} s(x_i^l, y_j^k). \quad (2)$$

This image-sentence score $S(X_i, Y_j)$ is computed using greedy matching. Concretely, for each phrase k in the sentence j , we find its best matching region in an image. The scores of best matching regions are further summed across all phrases. Note that phrases and regions are not interchangeable in this score function, i.e., $S(X_i, Y_j) \neq S(Y_j, X_i)$, because each phrase must be matched to at least one region, while some regions, such as background regions, are not matched to any phrase. Similar image-sentence scoring functions were discussed in [25, 58] for image-sentence retrieval.

3.2. Distillation using Contrastive Learning

A major challenge of weakly supervised grounding is the lack of ground-truth region-phrase pairs. Our key idea is to make use of an object detector during training that can provide “pseudo” labels. Our model further distills from these “pseudo” labels for learning region-phrase matching. Once learned, we can directly use the region-phrase score function without object detection. In what follows, we describe the generation of pseudo labels, the contrastive loss used for distillation, and the training and inference of our model.

Pseudo Labels for Region-Phrase Matching. An object detector \mathcal{D} predicts the probability of region x_i^l having an object label z_i^l in the form of nouns (including “background”), i.e., $p(z_i^l | x_i^l) = \mathcal{D}(x_i^l)$. z_i^l can be further matched to the phrase y_j^k , e.g., using similarity scores between object noun and the head noun of the phrase. In particular, our implementation uses the WordNet [33] to define such similarity scores, as discussed in section 4. Let $p(y_j^k, z_i^l)$ be the matching probability between y_j^k and z_i^l . We propose to approximate the unknown region-phrase matching ground-truth $p(x_i^l, y_j^k)$ by soft “pseudo” label $\hat{p}(x_i^l, y_j^k)$, approximated by

$$\hat{p}(x_i^l, y_j^k) \propto \sum_z p(y_j^k, z_i^l) p(z_i^l | x_i^l) p(x_i^l). \quad (3)$$

where we assume $p(z_i^l)$ is a constant value for every z_i^l since object classes are fixed and limited and therefore we omit

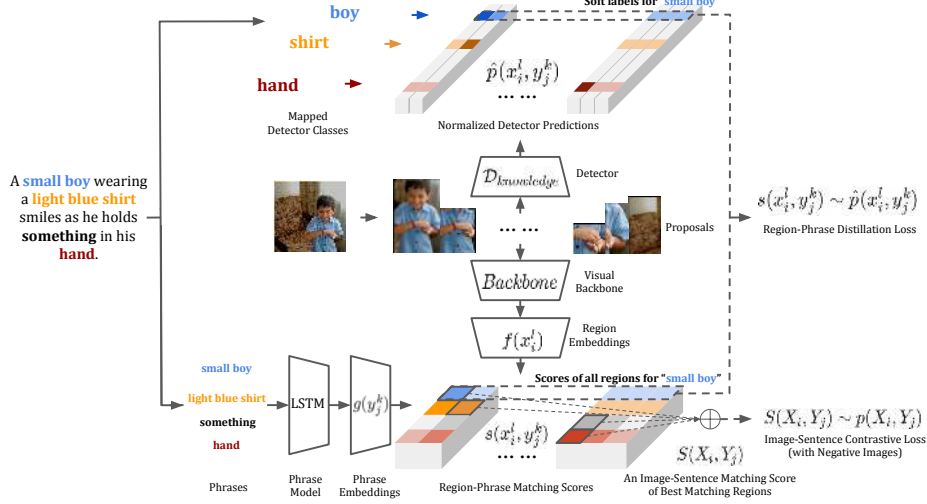


Figure 2: Overview of our method in training. A contrastive learning framework is designed to account for both region-phrase and image-sentence matching. The top part illustrates region-phrase matching learned by distilling from object detection outputs, while the bottom part shows image-sentence matching supervised by ground-truth image-sentence pairs.

the $p(z_i^l)$ in the denominator of Eq.3. Note that this approximation holds when x_i^l and y_j^k is conditionally independent given z_i^l — a rather strong assumption. Nonetheless, we use $\hat{p}(x_i^l, y_j^k)$ as a “pseudo” soft target distribution for training our model. In practice, $\hat{p}(x_i^l, y_j^k)$ is approximated and computed by detecting objects and matching their names to candidate phrases.

Distilling Knowledge from Pseudo Labels. We propose to distill from the pseudo label $\hat{p}(x_i^l, y_j^k)$ by aligning the region-phrase matching scores $s(x_i^l, y_j^k)$ to the soft pseudo label $\hat{p}(x_i^l, y_j^k)$. Specifically, given a matching image-sentence pair (X_i, Y_j) , we propose the following distillation loss function for region-phrase matching

$$\mathcal{L}_{RP}(X_i, Y_j) = - \sum_{y_j^k \in Y_j} \sum_{x_i^l \in R_i^l} \hat{p}(x_i^l, y_j^k) \log \hat{h}(x_i^l, y_j^k), \quad (4)$$

where $\hat{h}(x_i^l, y_j^k)$ is given by

$$\hat{h}(x_i^l, y_j^k) = \frac{\exp(s(x_i^l, y_j^k)/\tau)}{\exp(s(x_i^l, y_j^k)/\tau) + \sum_{x_i^{l'} \in R_i^l \setminus \{x_i^l\}} \exp(s(x_i^{l'}, y_j^k)/\tau)}.$$

Here τ is the temperature scale factor (0.5 in all our experiments). R_i^l controls how we select $x_i^{l'}$. A simple choice for $x_i^{l'}$ is to use all regions in X_i except x_i^l . In this case, our loss can be interpreted as the cross entropy loss, where the normalized output of the score function $s(x_i^l, y_j^k)$ is trained to mimic the pseudo label $\hat{p}(x_i^l, y_j^k)$ given by object detection outputs. This is the same idea as knowledge distillation [21], where the soft target from a teacher detection model is used as a learning objective.

Contrastive Loss for Image Sentence Matching. Moving beyond region-phrase matching, we enforce additional

constraints for image-sentence matching scores $S(X_i, Y_j)$, where the ground truth pairs $p(X_i, Y_j)$ is readily available. To this end, we make use of the noise contrastive estimation loss [19] to contrast samples from data distribution (matched pairs) and noise distribution (non-matched pairs). The NCE loss for image-sentence matching is thus given by

$$\mathcal{L}_{IS}(X_i, Y_j) = -\mathbb{E}_{N(X_i) \in X} [p(X_i, Y_j) \log h(X_i, Y_j)], \quad (5)$$

where $h(X_i, Y_j)$ is defined as

$$h(X_i, Y_j) = \frac{\exp(S(X_i, Y_j)/\tau)}{\exp(S(X_i, Y_j)/\tau) + \sum_{X_i' \in N(X_i)} \exp(S(X_i', Y_j)/\tau)}.$$

Again, τ is the temperature scale factor (0.5). $p(X_i, Y_j)$ is reduced to binary values during training, i.e., $p(X_i, Y_j) = 1$ if and only if (X_i, Y_j) is a ground-truth image-sentence pair. $N(X_i)$ includes a set of negative samples, i.e., those images not matched to the current sentence Y_j , sampled from the set of images X . In practice, we always sample a fixed number of negative pairs from the current mini-batch.

Loss Function. We note that both Eq. 4 and Eq. 5 share a similar form and can be both considered as a variant of contrastive loss. Concretely, the two loss functions seek to align the normalized scores in the form of NCE to a target distribution. The difference is how the target distribution is defined and how the samples are selected for normalization. For region-phrase matching, the target distribution is given by pseudo labels from object detection and local image regions are used for normalization. For image-sentence matching, the target distribution is defined by ground-truth image-sentence pairs and non-matched image-sentence pairs are sampled for normalization.

By combining the distillation loss $\mathcal{L}_{RP}(X_i, Y_j)$ for region-phrase matching and the NCE loss $\mathcal{L}_{IS}(X_i, Y_j)$ for

image-sentence matching, our final loss function is given by

$$\mathcal{L}(X_i, Y_j) = \mathcal{L}_{IS}(X_i, Y_j) + \lambda \mathcal{L}_{RP}(X_i, Y_j), \quad (6)$$

where λ is the coefficient balancing the two loss terms. During training, we gradually increase the coefficient λ , such that our model learns to optimize image-sentence matching during the early stage of training, and to focus on region-phrase matching during the late stage of training.

Inference without Object Detection. During inference, given an input image-sentence pair, we apply the learned region-phrase score function $s(x_i^l, y_j^k)$ between every region-phrase pair. The image region with the highest score to each phrase is then selected as the grounding results, i.e., $\arg \max_{x_i^l} s(x_i^l, y_j^k)$. We must point out that unlike previous methods [45, 18] *the inference of our model does not require running object detection*, therefore our method is very efficient at test time.

4. Experiments and Results

We now present our experiments and results. We first discuss our datasets, experiment setup and implementation details, followed by a comparison of our results to latest methods and an ablation study of our model.

Datasets. Our experiments are conducted on two major visual grounding datasets: Flickr30K Entities [39] and the ReferItGame dataset [26]. Flickr30K Entities [39] includes around 30K images. Each image is associated with five sentences. We follow the same train/val/test splits from [39]. ReferItGame [26] includes 20K images and 120K phrases. For ReferItGame, we follow the standard split of [42].

Experiment Setup. We follow the setting of weakly supervised grounding, and do not use the region-phrase annotations of both datasets during training. For evaluation, we follow the standard protocol used in [6, 42, 45, 18], we report accuracy as the evaluation metric. Accuracy is defined as the fraction of query phrases whose predicted bounding box overlaps ground-truth box with $IoU > 0.5$. For methods that select the predicted bounding box from a set of region proposals, the reported accuracy metric is equivalent to $Recall@1$, as used in [18].

4.1. Implementation Details

We first describe our implementation details, including the features and object detectors, the network architecture and training scheme, and details of object-phrase matching.

Features and Object Detectors. To establish a fair comparison with previous work using region features extracted from different backbones, we benchmark our methods by varying the backbone networks. We follow the same settings in [6, 45] to extract activations from the last layer before the classification head in Faster R-CNN [41] with VGG16 and ResNet-101 backbones pre-trained on PASCAL

VOC (PV) [13] or MS COCO (CC)¹ [31]. To compare with WPT [45] using object detectors trained on Open Images Dataset [28], we also extract classifier logits from Faster R-CNN with Inception-ResNet-V2 (IRV2) backbone pre-trained on the Open Images Dataset (OI)². To compare with the InfoGround from [18], we also experiment with their released features³ on Flickr30K and follow their protocols to use the same VisualGenome (VG) pre-trained Faster R-CNN model⁴ to extract features on ReferItGame. Unlike [6, 42, 56] using a large amount of bounding box proposals, typically 100 per image, InfoGround only extracts 30 proposals per image for the grounding model to select from. We compare the two settings separately in the result tables.

We denote these feature choices as “VGG16”, “Res101”, “IRV2” respectively plus the object data set when reporting our results. For example, “IRV2 OI” means that the backbone is Inception-ResNet-V2 (IRV2) pre-trained on the Open Images (OI) Dataset.

Network Architecture. For visual representation, we normalized the region features to zero-mean and unit-variance using stats from training samples. This normalization helps our model to converge faster. We attached two fully connected layers on top of the region features to get 512-D region embeddings. For phrase representation, we tokenized each query phrase into words and encoded using LSTM [23] with the GloVe embeddings [37]. The embedding vocabulary contains the most frequent 13K tokens from the Flickr30K Entities training split. The same vocabulary is used for ReferItGame. The LSTM has two layers, with both embedding and hidden dimension as 300. Max pooling is applied over the hidden states of all tokens, followed by two fully connected layers to get 512-D phrase embeddings.

Training Details. We trained our model using Adam [27] with a learning rate of 0.0001. We used a mini-batch size of 32 image-sentence pairs (31 negative images per sentence for the contrastive loss). Unlike [14], we did not fine-tune our vision backbone during training for efficiency. Similarly, the GloVe embeddings [37] are also fixed during training. We observed that the model converges quickly within a few epochs on both datasets. For the λ in Eq. 6, we gradually increased the value using a staircase function $\lambda = \min(\lfloor step/a \rfloor, b)$, where a, b are selected based on the validation set. We observed that $a=200, b=1$ for VG features and $a=200, b=3$ for others work the best.

Object-Phrase Matching. We use the WordNet [33] to define the similarity scores between phrases and region object labels. Specifically, we identify the head noun of each phrase using the off-the-shelf POS tagger provided by NLTK [4],

¹<https://github.com/endernewton/tf-faster-rcnn>

²https://github.com/tensorflow/models/blob/master/research/object_detection/g3doc/detection_model_zoo.md

³<https://github.com/BigRedT/info-ground>

⁴<https://github.com/BigRedT/bottom-up-features>

Method	Backbone	Detector _K	Require Detector _K in Inference	ACC (%)
Grounder [42]	VGG16 PV	-	-	28.94
MATN [56]	VGG16 PV	-	-	33.10
UTG [50]	-	VGG16 PV YOLOv2 CC	Yes Yes	35.90 36.93
KAC [6]	VGG16 PV VGG16 PV	VGG16 PV VGG16 CC	Yes Yes	36.14 38.71
MTG [14]	Res101 CC+Res50 CC +Res50 CL	-	-	48.66
WPT [45]	-	IRV2 CC	Yes	37.57
(w2v-max union)	-	IRV2 CC+IRV2 OI	Yes	48.20
	-	IRV2 CC+IRV2 OI +WRN18 PL	Yes	50.49
InfoGround [18] (Trained on Flickr30K)	Res101 VG	-	-	47.88
(Trained on COCO)	Res101 VG	-	-	51.67
NCE+Distillation (Ours)	VGG16 PV	VGG16 CC	No	40.38
	Res101 CC	IRV2 OI	No	50.96
	Res101 VG	IRV2 OI	No	53.10

Table 1: Results on Flickr30K Entities. We report phrase localization accuracy and list the settings of different methods. “Backbone” denotes the visual backbone used to extract region features. Detector_K denotes the detector that provides external knowledge. “-” denotes the method does not use backbone or Detector_K. Dataset notations: PV=PASCAL VOC, CC=COCO, OI=Open Images, CL=Color Name, PL=Place365, and VG=Visual Genome.

which uses the Penn Treebank tag set. If the head noun matches one of the detector class labels, the phrase is mapped to the class. Otherwise, we look up the head noun in the WordNet [33] to find its corresponding synset, the synset’s lemmas, and hypernyms. If any of these exists in the detector classes, the phrase is mapped to the class. For phrases with multiple synsets, the most frequent one is used. The WordNet synset helps to match phrases such as “spectators” to “person” and “sweater” to “clothing”. With the 545 classes in Open Images Dataset [28], the WordNet-based matching algorithm covers 18k out of 70k unique phrases in Flickr30k Entities and 7k out of 27k in ReferItGame training set.

We empirically observe that using WordNet is more reliable than word embeddings for noun matching. Similarity in the word embedding space is not necessarily aligned with the visual similarity of the entities described. A similar observation was made in [14] where the word2vec embeddings are not discriminative for gender related visual concepts.

4.2. Comparison to Other Methods

We further compare our results to the latest methods of weakly supervised phrase grounding on both Flickr30K Entities and ReferItGame datasets.

Baselines We consider a number of baselines. Our main competitors are those methods using object detectors, including KAC [6], UTG [50], and WPT [45]. Among these methods, KAC and UTG used detectors during both training and inference. WPT applied detectors during inference. While these baselines have very different sets of detectors and backbones, we try to match their settings in our experiments. Our baselines also include previous methods that do not use object detectors, such as Grounder [42], MATN [56], MTG [14], and InfoGround [18] for completeness.

Results on Flickr30K Entities. Our results on Flickr30K

Entities are summarized in Table 1. Table 1 compares both the settings of different methods and their phrase localization accuracy. When using detectors pre-trained on COCO (YOLOv2 CC, VGG16 CC, IRV2 CC), our method outperforms previous works (+3.5% / +1.7% / +2.81% for UTG / KAC / WPT, respectively). In comparison to MTG which uses COCO pre-trained backbones, our results are better (+2.6%) with only a single VGG16 backbone in inference (vs. three ResNet backbones in MTG).

When using a stronger backbone (Res101 CC) and a better detector pre-trained on a larger scale dataset (IRV2 OI), our results are further improved by 10.6% on Flickr30K Entities. Our final results thus outperform the latest method of WPT under a similar training setting. Moreover, in contrast to WPT, our method does not require running the three cumbersome object detectors during inference, thus is more applicable for real world deployment.

Lastly, to compare fairly with [18], using the same VisualGenome pre-trained backbone (Res101 VG) and under the same proposals setting (30 per image) with [18], our NCE+Distill model significantly outperforms the latest method of InfoGround by 5.22% when trained on the Flickr30K Entities dataset. Moreover, with the help of distillation, our results also outperform their best results by 1.4%, which is trained on COCO Caption dataset [8] using a strong language model (BERT [10]).

To further illustrate the benefits of our model, we compare the computation complexity at inference time in terms of floating point operations (FLOPs). For each model, the estimation is a combination of the backbone and the object detector, since the rest part of the model, including the language feature extractor, is computationally insignificant (0.3 GFLOPs for our NCE+Distill model). For Faster R-CNN based detectors, we use the numbers reported in [24], under

Method	Backbone	Detector _K	Require Detector _K in Inference	ACC (%)
GroundeR [42]	VGG16 PV	-	-	10.70
MATN [56]	VGG16 PV	-	-	13.61
UTG [50]	-	VGG16 CC+YOLOv2 CC	Yes	20.91
KAC [6]	VGG16 PV	VGG16 PV	Yes	13.38
	VGG16 PV	VGG16 CC	Yes	15.83
WPT [45] (w2v-max union)	-	IRV2 CC	Yes	15.40
	-	IRV2 CC+IRV2 OI +WRN18 PL+CL	Yes	26.48
NCE+Distillation (Ours)	VGG16 PV	VGG16 CC	No	24.52
	Res101 CC	IRV2 OI	No	27.59
	Res101 VG	IRV2 OI	No	38.39

Table 2: Results on ReferItGame. We report phrase localization accuracy and settings of different methods. “Backbone” denotes the visual backbone used to extract region features. Detector_K denotes the detector that provides external knowledge. “-” denotes the method does not use backbone or Detector_K. Dataset notations: PV=PASCAL VOC, CC=COCO, OI=Open Images, CL=Color Name, and PL=Place365.

Method	Detector _K GFLOPs	Backbone GFLOPs	Total GFLOPs
Detector _K =YOLOv2 CC UTG	60	-	60
Detector _K =Faster R-CNN VGG16 CC KAC	180	180 (VGG16 PV)	360
NCE+Distill	0 (Not needed)	180 (VGG16 PV)	180
Detector _K =Faster R-CNN IRV2 OI WPT	1600+ (+IRV2 CC,WRN18 PL)	-	1600+
NCE+Distill	0 (Not needed)	500 (Res101 CC)	500

Table 3: Estimated number of FLOPs in inference.

the high-resolution input (600×600) and 300 proposals setting. For YOLOv2, we use the number reported in [40]. The comparison focuses on methods that incorporate external knowledge, namely UTG, KAC, WPT, and ours. As shown in Table 3, our proposed NCE+Distill method reduces the computational complexity by 50% and 70%, while being a bit more expensive than UTG due to the detector meta-architecture difference (Faster R-CNN vs. YOLOv2).

Results on ReferItGame We summarize the results on ReferItGame under different settings in Table 2. When using COCO pre-trained detectors, our method significantly outperforms UTG, KAC, and WPT by 3.6%, 8.7%, and 9.12% respectively. When using a stronger backbone (Res101 CC) and a better detector pre-trained on a larger scale dataset (IRV2 OI), our results are improved by 3.1%, outperforming the best results from WPT using four cumbersome knowledge detectors in inference by 1.1%.

Finally, using the VisualGenome pre-trained backbone (Res101 VG), our results can be further improved by 10.8%. This gain is significantly larger than the one on Flickr30K Entities. We conjecture that the VisualGenome pre-trained backbone provides more discriminative features in certain categories that appear more frequently in the ReferItGame than Flickr30K Entities. To verify this hypothesis, we compare the most frequent phrases from Flickr30K Entities and ReferItGame. In Flickr30K Entities, top phrases are mostly people related: *man, woman, boy, girl, person, people, dog, two men, street, young boy, child*, whereas top phrases in ReferItGame are mostly scene related: *sky, water, people, ground, person, trees, building, face, road, grass, clouds*.

Many scene related objects are not in the COCO label set, but are available in the VisualGenome categories.

4.3. Ablation Study

To fully understand our model, we conduct ablation studies on both Flickr30K Entities and ReferItGame datasets. Specifically, we consider four different variants of our model: (1) our model with only image-sentence score function (Eq. 2) supervised by a max margin loss following [25, 56], denoted as “Max Margin”, i.e. modeling positive and negative image-sentence pairs distances via max margin loss. The full loss function is defined in the Eq.7 below. (2) our model with only image-sentence score function (Eq. 2) supervised by the NCE loss (Eq. 5), denoted as “NCE”; (3) our model with only region-phrase score function (Eq. 1) supervised by the distillation loss (Eq. 4), denoted as “Distill”; and (4) our full model with both region-phrase and image-sentence score functions supervised by our joint loss (Eq. 6), denoted as “NCE+Distill”.

We present our ablation results on the four model variations in Table 4.

Contrastive vs. Ranking loss. We first define the max margin loss following notations in the image-sentence level contrastive loss (Eq. 5) as

$$\mathcal{L}_{IS}(X_i, Y_j) = \mathbb{E}_{\mathcal{N}(X_i) \in X} [h(X_i, Y_j)] \quad (7)$$

$$h(X_i, Y_j) = \sum_{i' \in \mathcal{N}(X_i)} \max\{0, m - S(X_i, Y_j) + S(X_{i'}, Y_j)\},$$

where m is the margin. In our experiment, we fix $m = 0.05$.

We observe that NCE loss substantially outperforms the standard max margin loss by +6.2%/+3.7% on Flickr30K Entities and ReferItGame, respectively. These results suggest the effectiveness of contrastive learning, as also demonstrated in the concurrent work [18].

Effects of Knowledge Distillation. Our full model combining both region-phrase and image-sentence matching brings further improvements over both the NCE only model and the Distill only model. We conjecture that NCE and Distill

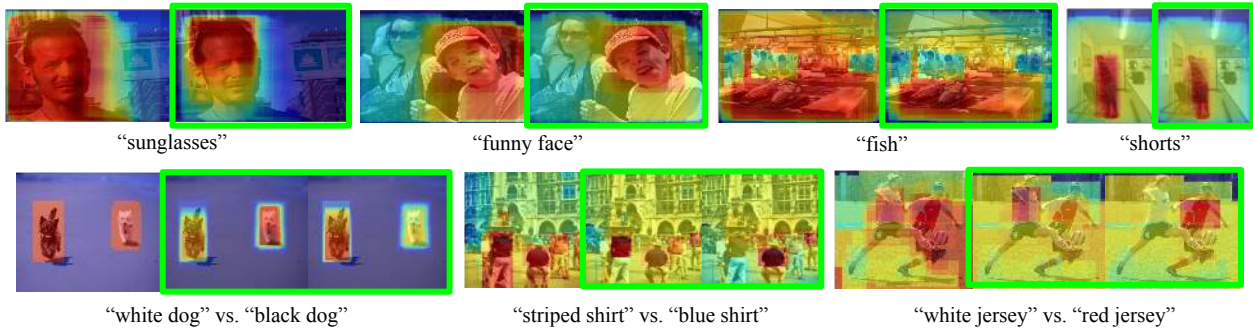


Figure 3: Visualization of region-phrase matching. First row: results of the NCE only (left) and the NCE+Distill (green) on phrases mapped to Open Images Dataset classes. Second row: results of the Distill only (left) and NCE+Distill (green) on phrases mapped to the same class but with different attributes. For each pixel, we compute a matching score by averaging scores from all proposals covering the pixel. Red colors indicate high matching scores. Our knowledge distillation helps to better identify the extent of objects, while contrastive learning helps to distinguish finer attributes.

Method	Flickr30K ACC (%)	ReferItGame ACC (%)
Max Margin	42.11	22.94
NCE	48.35	26.63
Distill	45.05	17.25
NCE+Distill	50.96	27.59

Table 4: Ablation our proposed methods on Flickr30K and ReferItGame. All models use Res101 CC as backbone and models with distillation use IRV2 OI as Detector_K.

provide complementary information for phrase grounding. Specifically, Distill helps learn the extent of objects better, distinguishing parts frequently co-occurred in the same image, such as “person” and “face”. NCE helps to learn finer-grained attributes, such as “striped shirt” and “blue shirt”, as well as concepts not covered by detector classes.

To understand where the accuracy improvement of the distillation method comes from, we compute per phrase accuracy on frequent phrases that are mapped to one of the classes in the Open Images Dataset. The results can be found in the supplementary materials. We also perform qualitative analysis with side-by-side heatmaps of region-phrase matching scores on such mapped phrases, shown in Figure 3. Our full model (NCE+Distill) can better localize objects corresponding to the query phrase.

Generalization to Different Backbones. We vary the object detectors used by our model and present the results in Table 5. As the knowledge gap between visual backbone and external detector becomes smaller, e.g., the visual backbones are trained on tasks that involve finer-grained object and attribute annotations, the effects of distillation become less predominant. Nevertheless, our method can consistently improve performance for all detectors.

5. Conclusion, Limitation and Future Work

In this paper, we presented a novel contrastive learning framework for weakly supervised visual phrase grounding. The key idea of our method is to learn a score function

Backbone	Detector _k	NCE ACC (%)	NCE+Distill ACC (%)
Flickr30k			
VGG16 PV	VGG16 CC	35.96	40.38 (+4.42)
Res101 CC	IRV2 OI	48.35	50.96 (+2.61)
Res101 VG	IRV2 OI	51.31	53.10 (+1.79)
ReferItGame			
VGG16 PV	VGG16 CC	23.56	24.52 (+0.96)
Res101 CC	IRV2 OI	26.63	27.59 (+0.96)
Res101 VG	IRV2 OI	36.65	38.39 (+1.74)

Table 5: Results of our proposed distillation methods on Flickr30K and ReferItGame using various backbones and detectors providing external knowledge. “Backbone” denotes the visual backbone used to extract region features. Detector_K denotes the detector that provides external knowledge. Dataset notations: PV=PASCAL VOC, CC=COCO, OI=Open Images, and VG=Visual Genome.

measuring the similarity between region-phrase pairs, distilled from object detection outputs and further supervised by image-sentence pairs. Once learned, this score function can be used for visual grounding without the need of object detectors at test time. While conceptually simple, our method demonstrated strong results on major benchmarks, surpassing state-of-the-art methods that use expensive object detectors. Our work offers a principled approach to leverage object information, as well as an efficient method for weakly supervised grounding.

Limitation and Future Work. A main limitation of our work is the need of a generic object detector to cover most of the object classes during training. While object detection is currently more mature than visual grounding, learning open-set object detectors from large scale web data and further bridging the gap between object detection and visual grounding is a promising direction. Moving forward, we also plan to extend our method to solve other multi-modal grounding tasks, including video grounding [1].

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017. [8](#)
- [2] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *NeurIPS*, pages 2654–2662, 2014. [3](#)
- [3] Mohit Bajaj, Lanjun Wang, and Leonid Sigal. G3raphGround: Graph-based language grounding. In *ICCV*, October 2019. [2](#)
- [4] Steven Bird and Edward Loper. NLTK: The natural language toolkit. In *ACL Interactive Poster and Demonstration Sessions*, pages 214–217, 2004. [5](#)
- [5] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *SIGKDD*, 2006. [3](#)
- [6] Kan Chen, Jiyang Gao, and Ram Nevatia. Knowledge aided consistency for weakly supervised phrase grounding. In *CVPR*, 2018. [1](#), [2](#), [5](#), [6](#), [7](#)
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. [2](#)
- [8] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. [6](#)
- [9] S. Datta, K. Sikka, A. Roy, K. Ahuja, D. Parikh, and A. Divakaran. Align2Ground: Weakly supervised phrase grounding guided by image-caption alignment. In *ICCV*, pages 2601–2610, 2019. [2](#)
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *ACL*, 2019. [6](#)
- [11] Tuong Do, Thanh-Toan Do, Huy Tran, Erman Tjiputra, and Quang D Tran. Compact trilinear interaction for visual question answering. In *ICCV*, pages 392–401, 2019. [3](#)
- [12] Pelin Dogan, Leonid Sigal, and Markus Gross. Neural sequential phrase grounding (seqground). In *CVPR*, 2019. [2](#)
- [13] Mark Everingham, LV Gool, CKI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes VOC dataset and challenge. *IJCV*, 2010. [5](#)
- [14] Zhiyuan Fang, Shu Kong, Charless Fowlkes, and Yezhou Yang. Modularized textual grounding for counterfactual resilience. In *CVPR*, pages 6378–6388, 2019. [1](#), [2](#), [5](#), [6](#)
- [15] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016. [1](#), [2](#)
- [16] Nuno Garcia, Pietro Morerio, and Vittorio Murino. Modality distillation with multiple stream networks for action recognition. In *ECCV*, 2018. [3](#)
- [17] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *CVPR*, pages 2827–2836, 2016. [2](#), [3](#)
- [18] Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding. In *ECCV*, 2020. [1](#), [2](#), [5](#), [6](#), [7](#)
- [19] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, pages 297–304, 2010. [1](#), [2](#), [3](#), [4](#)
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. [2](#)
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [3](#), [4](#)
- [22] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. 2019. [2](#)
- [23] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. [5](#)
- [24] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy. Speed/accuracy trade-offs for modern convolutional object detectors. In *CVPR*, pages 3296–3297, 2017. [6](#)
- [25] Andrej Karpathy, Armand Joulin, and Li F Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. In *NeurIPS*, 2014. [3](#), [7](#)
- [26] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014. [2](#), [5](#)
- [27] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [5](#)
- [28] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, et al.

- OpenImages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2(3):18, 2017. 5, 6
- [29] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yanis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 2
- [30] Jinyu Li, Michael L Seltzer, Xi Wang, Rui Zhao, and Yifan Gong. Large-scale domain adaptation via teacher-student learning. *arXiv preprint arXiv:1708.05466*, 2017. 3
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. 5
- [32] Zelun Luo, Jun-Ting Hsieh, Lu Jiang, Juan Carlos Niebles, and Li Fei-Fei. Graph distillation for action detection with privileged modalities. In *ECCV*, 2018. 3
- [33] George A Miller. Wordnet: A lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. 3, 5, 6
- [34] Jonghwan Mun, Kimin Lee, Jinwoo Shin, and Bohyung Han. Learning to specialize with knowledge distillation for visual question answering. In *NeurIPS*, pages 8081–8091, 2018. 3
- [35] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 3
- [36] Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, and Juan Carlos Niebles. Spatio-temporal graph for video captioning with knowledge distillation. In *CVPR*, 2020. 3
- [37] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global vectors for word representation. In *EMNLP*, 2014. 5
- [38] Bryan A Plummer, Paige Kordas, M Hadi Kiapour, Shuai Zheng, Robinson Piramuthu, and Svetlana Lazebnik. Conditional image-text embedding networks. In *ECCV*, 2018. 2
- [39] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30K entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015. 1, 2, 5
- [40] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *CVPR*, pages 7263–7271, 2017. 7
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 5
- [42] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *ECCV*, 2016. 1, 2, 5, 6, 7
- [43] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 3
- [44] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, 2020. 2, 3
- [45] Josiah Wang and Lucia Specia. Phrase localization without paired training examples. In *CVPR*, pages 4663–4672, 2019. 1, 2, 5, 6, 7
- [46] Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. Learning two-branch neural networks for image-text matching tasks. *IEEE TPAMI*, 2018. 1, 2, 3
- [47] Fanyi Xiao, Leonid Sigal, and Yong Jae Lee. Weakly-supervised visual grounding of phrases with linguistic structures. In *CVPR*, 2017. 2
- [48] Zhengyuan Yang, Tianlang Chen, Liwei Wang, and Jiebo Luo. Improving one-stage visual grounding by recursive sub-query construction. *ECCV*, 2020. 2
- [49] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *ICCV*, 2019. 2
- [50] Raymond Yeh, Minh N Do, and Alexander G Schwing. Unsupervised textual grounding: Linking words to image concepts. In *CVPR*, 2018. 1, 2, 6, 7
- [51] Raymond Yeh, Jinjun Xiong, Wen-Mei Hwu, Minh Do, and Alexander Schwing. Interpretable and globally optimal prediction for textual grounding using image concepts. In *NeurIPS*, 2017. 1, 2
- [52] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017. 3
- [53] Hanwang Zhang, Yulei Niu, and Shih-Fu Chang. Grounding referring expressions in images by variational context. In *CVPR*, 2018. 1
- [54] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. *IJCV*, 2018. 1
- [55] Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, and Zhengjun Zha. Object relational graph with teacher-recommended learning for video captioning. In *CVPR*, 2020. 3

- [56] Fang Zhao, Jianshu Li, Jian Zhao, and Jiashi Feng. Weakly supervised phrase localization with multi-scale anchored transformer network. In *CVPR*, 2018. 1, 2, 5, 6, 7
- [57] Luowei Zhou, Yannis Kalantidis, Xinlei Chen, Jason J Corso, and Marcus Rohrbach. Grounded video description. In *CVPR*, 2019. 2
- [58] Luowei Zhou, Nathan Louis, and Jason J Corso. Weakly-supervised video object grounding from text by loss weighting and object interaction. In *BMVC*, 2018. 3
- [59] Yuanen Zhou, Meng Wang, Daqing Liu, Zhenzhen Hu, and Hanwang Zhang. More grounded image captioning by distilling image-text matching model. In *CVPR*, 2020. 3