

From Coarse to Fine-grained Concept based Discrimination for Phrase Detection

Maan Qraitem, Bryan A. Plummer
Boston University
{mqraitem, bplum}@bu.edu

Abstract

Phrase detection requires methods to identify if a phrase is relevant to an image and localize it, if applicable. A key challenge for training more discriminative detection models is sampling negatives. Sampling techniques from prior work focus primarily on hard, often noisy, negatives disregarding the broader distribution of negative samples. Our proposed CFCD-Net addresses this through two novel methods. First, we generate groups of semantically similar words we call concepts (e.g., {dog, cat, horse} and {car, truck, SUV}), and then train our CFCD-Net to discriminate between a region of interest and its unrelated concepts. Second, for phrases containing fine-grained mutually-exclusive words (e.g., colors), we force the model to select only one applicable phrase for each region using our novel fine-grained module (FGM). We evaluate our approach on Flickr30K Entities and RefCOCO+, where we improve mAP over the state-of-the-art by 1.5-2 points. When considering only the phrases affected by our FGM module, we improve by 3-4 points on both datasets.

1. Introduction

The goal of phrase detection is to localize all instances of a phrase in a database of images. A key challenge differentiating phrase detection from other phrase grounding tasks is that many (and oftentimes most) of the images do not contain any instances of a phrase. This requires a phrase detector to discriminate between image regions that may contain entities that are semantically similar to a phrase. For example, being able to tell apart “a semi-truck” and “a pickup truck,” or identifying that a vehicle is “a gray car” and not “a blue car.” While recent methods have been very successful in the localization-only task (e.g. [14, 18]), phrase detection remains a significant challenge. This is due, in part, to the long-tailed distribution of phrases, *i.e.*, there may be tens of thousands of annotated queries in a dataset, but many occur only a few times. Prior work used hard-negative sampling to boost performance (e.g., [11]). However, this generates many false negatives during training, and, as illustrated in

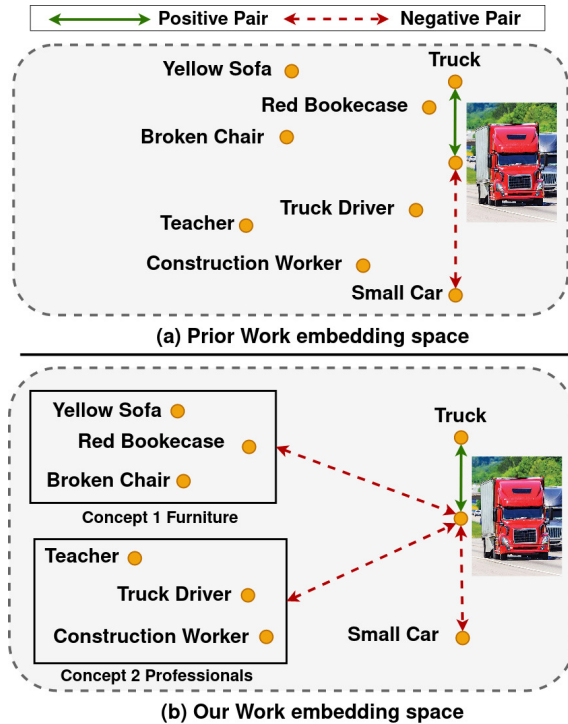


Figure 1. **Comparing embedding spaces.** (a) prior work used hard-negative mining to make their model more discriminative (e.g., [11]), which misses a large distribution of negatives (b), our CFCD-Net creating a set of semantically coherent words (concepts) (e.g. {Sofa, table, Chair}) and then learns to discriminate a region from its unrelated concepts. The resulting embedding space better separates the region from its unrelated phrases.

Figure 1(a), the larger pool of semantically similar negative phrases may still embed near the ground truth phrase, causing many false positives during inference.

To address this issue, we introduce a Coarse-to-Fine-grained Concept-based Discrimination Network (CFCD-Net), which uses semantically coherent groupings of words (concepts) to perform both coarse and fine grained discrimination. A straightforward approach to distinguish between a large set of phrases would be simply using all the phrases in a batch during training. However, this approach has two major drawbacks. First, phrase grounding datasets are very

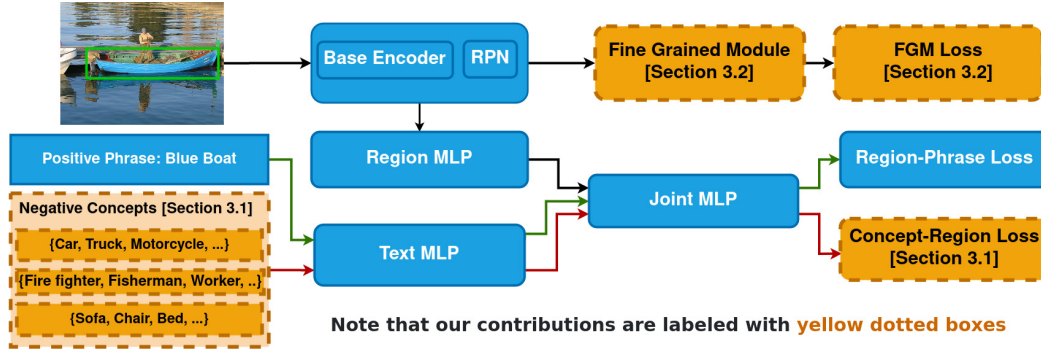


Figure 2. **CFCD-Net overview.** We begin by passing the input image to a base encoder (ResNet-101 [10]) and a region proposal network [29] to generate a set of candidate boxes. To learn a more discerning region-phrase model, we introduce an approach that trains a model to discriminate between a set of automatically constructed visually-related coarse concepts (Section 3.1). We also introduce a fine-grained reasoning module (Section 3.2) that forces the CFCD-Net to delineate between mutually-exclusive fine-grained tokens.

sparingly labeled, so many unannotated phrases can simply be false negatives (see supplementary for discussion). Second, using all phrases in a single batch would not fit into GPU memory as phrase grounding datasets can have tens or hundreds of thousands of unique training phrases [26].

Prior work avoids memory issues by augmenting a batch with a limited number of positive phrases [26] or negative phrases [11] during training. However, as mentioned earlier, this may result in semantically similar phrases being embedded nearby the ground truth phrases, reducing performance. Instead, CFCD-Net separates the words belonging to the training phrases into groups with the same part-of-speech. Then we automatically construct a set of concepts, *i.e.*, semantically coherent bags of words, for each group. For the concepts related to nouns (*e.g.*, (dog, cat, horse) and (car, truck, bike)), we perform a coarse discrimination. First, for each phrase we obtain a set of unrelated concepts (*i.e.*, concepts that share no semantically similar nouns with the phrase), which we refer to as negative coarse concepts (NCC). Then, given a region/phrase pair, we use the phrase’s NCCs as negatives. As illustrated in Figure 1(b), this encourages separation between the region and a wide array of unrelated phrases.

This approach enables CFCD-Net to minimize false negatives during training, as most unannotated positive phrases belong to the same concept (*e.g.*, the phrase “a cow” would not use the concept “domestic animals” as a negative). In addition, since the set of concepts is small (< 70 in our experiments), we can include all of them in every batch. Our approach is similar in spirit to the distributional sampling approach of Wu *et al.* [33], which demonstrated that balanced batches representing the entire dataset boosts performance over hard-negative mining. However, we create concept groups only once in a short preprocessing step (only a few seconds on a CPU). In contrast, Wu *et al.* periodically computes pairwise distances between all instances in the training set, which would take hours on a single GPU

every time it was recomputed (details in Section 4.1).

While NCCs provides a more robust representations of nouns (which typically refer to objects), another challenge in phrase detection the need to recognize attributes, which are commonly captured by adjectives in phrases. We find that adjective concepts work in the opposite way as nouns, where within an adjective concept the words refer to difficult fine-grained differences, while phrases can have words in multiple concept bags. For example, let us assume we have concepts for colors (*e.g.*, (red, green)) and some texture patterns (*e.g.*, (striped, plaid)). A shirt could be both red and striped, but a striped shirt should not also be a plaid shirt. We take advantage of these adjective concepts containing mutually exclusive words by adding a fine-grained module (FGM) to our model. Instead of discriminating between the concepts themselves as outlined by our method NCC above, FGM differentiates between the fine-grained words within the concept and then augments the main model score with its prediction. While there are other potentially useful parts-of-speech that may belong to a phrase (*e.g.*, verbs), we find that existing datasets contain too few of them to make a significant impact on performance. Figure 2 provides an overview of our approach.

Our contributions can be summarized as:

- A novel model, CFCD-Net, that improves over SOTA by 1.5-2 average mAP on phrase detection by mining semantically coherent concepts and then learning to discriminate between a region of interest and its unrelated negative coarse concepts (NCC).
- A fine grained reasoning module (FGM) that boosts performance by 1-4 average mAP over affected phrases by learning to distinguish visually similar instances.
- A novel method for automatically mining semantically coherent concepts that are visually similar and with minimum outliers that improves the distribution of our mini-batches and, thus, better represents the training data.




Figure 3. Comparison across methods [14, 18, 26] on phrase grounding (left) and Binary Image Selection (BISON) [12] (right). Methods along the x-axis are ordered by their localization-only performance. We see that methods designed for localization are not correlated with detection performance (left). However, we see comparing the left and right plots that improvements on phrase detection are correlated with improved performance on downstream tasks like BISON. See Section 2 for discussion

2. Related work

In this paper, we refer to phrase grounding as the general task of relating phrases to image regions. The most popular phrase grounding tasks are phrase localization [27] and referring expression comprehension [15, 21, 37]. In these tasks, methods given a ground truth image-phrase pair and must identify the relevant image region (e.g., [2, 4, 13, 15, 24, 25, 35]). This localization-only evaluation does not align well with how these methods might be used for downstream tasks. For example, image captioning requires reasoning about whether an image contains “a yacht” vs “a small boat.” However, these methods are typically only evaluated on if they could localize the boat, but not if they could determine which phrase is relevant to the image.

In contrast, phrase detection methods are ranked precisely on their ability to determine relevance as well as localize the phrase. Plummer *et al.* [26] demonstrated that using hard-negative mining to boost phrase detection performance was a challenge due to the high number of false negatives due to the sparse annotations in phrase grounding datasets. Instead, they proposed an approach that alleviated some label sparsity issues by augmenting the datasets with likely unannotated positive phrases. In contrast, our CFCD-Net reduces the number of false negatives suffered from hard-negative mining by grouping together semantically similar words (e.g., “dog, cat, horse”) so that any likely false negatives would be within the same concept. Thus, any concepts not related to a phrase could safely be used as negatives during training, boosting performance.

Generalizing across phrase grounding tasks. Plummer *et al.* [26] demonstrated that many methods designed for the localization-only are prone to overfitting to that setting, *i.e.*, they improve performance on localization, but reduce phrase detection performance. This is partly because dis-

tinguishing between similar phrases in the localization-only task is unnecessary since most images only contain a reference to a single object of the same type [26]. Thus, a localization model that is given the phrase *a young teenager* as input could look for the object category “person” to identify the right object most of the time, whereas a detector also has to determine if it exists at all. As shown in Figure 3(left), this means that improved localization performance does not guarantee improvements on detection.

There is also some work that falls between phrase detection and localization-only tasks, where a ground truth image-phrase pair is not provided [7, 11, 38], but they severely limit the number of negative phrases seen by each image at test time. Thus, methods from these tasks often do not generalize to phrase detection [26].

Effect on downstream applications. We argue the disconnect between localization-only and phrase detection performance partially explains why improving localization performance leads to small improvements on downstream tasks [5, 20, 27]. Once any grounding method is integrated, even doubling localization performance leads to negligible differences downstream performance [27]. Thus, most vision-language work uses object and attribute detectors to represent images, such as the commonly used bottom-up features [1], rather than models trained for localization. In contrast, since phrase detection can also identify what phrases are relevant to an image, it is directly correlated with improved downstream task performance. To illustrate this, we compared several grounding methods on the Binary Image Selection (BISON) [12] benchmark. In this task a model is given a sentence and is asked to choose between two semantically similar images. In our experiment, we extracted the set of noun phrases from the given sentence, and then ranked images by averaging the grounding

scores of the phrases. In Figure 3(right), we see methods that performed well on phrase detection also did well on BISON. See the supplementary for additional details.

Comparison to open-vocabulary detection. Some recent work (e.g., [14, 18]) also evaluated on long-tailed object detection benchmarks like LVIS [8]. However, phrase detection is considerably more challenging, as phrases include information about attributes (e.g., “a red shirt”) and spatial relationships (e.g., “a cup on top of a white table”), in addition to a long tailed set of object categories. Thus, as our experiments show, methods that work well on LVIS (e.g., [18]) often do not generalize to phrase detection.

3. Coarse-to-Fine-grained Concept-based Discrimination Network (CFCD-Net)

The goal of phrase detection is to detect and localize all instances of a phrase within a dataset of images. More formally, assume we have a dataset of images X where for each $x \in X$, regions of interests R_x are annotated with boxes B_x and phrases P_x . Denote the space of possible phrases as $P = \bigcup_{x \in X} P_x$. Thus, for each $p \in P$ and $x \in X$, the task of phrase detection involves determining whether p is relevant to x , and if so localize it with a bounding box. Note that this task involves zero-shot evaluation by definition, since while there are phrases shared between P_{train} and P_{test} , $P_{train} \neq P_{test}$, as well as evaluating few-shot and common phrases. Moreover, as we pointed out earlier, detection is a generalization of localization. More concretely, in localization, for each image x , only $P_x \subset P$ are evaluated, whereas, in detection, the entirety of P is used. This difference is key for the disconnect between improvements on localization and improvements for downstream tasks as we discussed in Section 2.

We improve the discriminative power of detection models by using Concepts. We define a concept c as a bag of semantically coherent words that share the same part-of-speech (e.g., nouns). We mine for a set of concepts C where each $c \in C$ represents a unique semantic concept (e.g., Vehicles: {car, bike, truck}). Based on these concepts, we introduce two novel methods of discrimination: a coarse method that uses noun based concepts C^n outlined in Section 3.1, and a fine grained methods that uses adjective based Concepts C^a outlined in Section 3.2.

3.1. Expanding Batch Coverage through Coarse Negative Concepts (NCC)

The goal of phrase detection is ensuring a model can discriminate between a region of interest r and all unrelated phrases in P . Thus, a natural solution to this problem is to expand the batch coverage to include all possible phrases P as potential negatives. However, there can be tens or hundreds of thousands of unique phrases, making it

impossible to fit all of them in a single due to GPU memory constraints. Hinami and Satoh [11] sampled phrases using hard-negative mining to increase their model’s discriminative power. However, as we discussed in the introduction, these methods miss a large pool of negative phrases that are not “hard” but may still end up embedded reasonably close to the ground truth region embedding. In addition, phrase grounding datasets are sparsely labeled. Thus, many selected hard negatives may be unannotated positive phrases, e.g., a region labeled as *skier* but not *woman*.

To mitigate this problem, we propose grouping the nouns in the training dataset into semantically coherent groups, i.e. noun based concepts C^n . Then, we pair every phrase in the dataset with its set of unrelated concepts. Consequently, given a positive phrase-region pair (p, r) , most of the unannotated positive phrases will be contained in one concept (e.g. *domestic animals*). Thus, we are more confident that all unrelated concepts should be true negatives. Furthermore, since our concepts group only nouns, we are able to limit the concepts to at most 70 per dataset. This makes it feasible to fit all the concepts in one batch. Moreover, since our concepts span the majority of the language space in the training set, our batch now has a balance of easy to hard negatives, which has been shown to improve performance in prior work [33]. We outline our full pipeline of concept generation and phrase-concept assignment below.

Concept Generation and Assignment: To obtain the noun based concept set C^n , we first extract nouns from our dataset and use a language embedding to represent them. Then, we use the embeddings to cluster the nouns into our semantically similar set of concepts. Given the resulting set of concepts C^n , we assign each phrase to its set of relevant concepts C_p^n . A simple solution is to pair a phrase with a concept with which it shares a noun. This approach, however, disregards phrase-concept pairs that don’t share nouns but are nevertheless viable pairs. Thus, we make these assignments using semantic similarity as detailed below.

First, given a phrase p , and phrase nouns p_n , we collect concepts related to phrase p from simple noun-phrase matching: $C_{p1}^n = \{c : c \cap p_n \neq \emptyset, c \in C^n\}$. We also consider concepts C_{p2}^n which share some semantic similarity (e.g. *hoody* with concept: *shirt and sweatshirt*). Specifically, we look up the language representation \tilde{n} for each noun $n \in p_n$. We also compute a concept c representation by averaging its nouns’ textual feature vectors. We will denote this as \tilde{c} . Now assume $\text{sim}(\tilde{n}, \tilde{c}) = \frac{\tilde{n} \cdot \tilde{c}}{\|\tilde{n}\| \|\tilde{c}\|}$, we compute an “association” score as follows:

$$\text{Assoc}(n, c) = \frac{\exp(\text{sim}(\tilde{n}, \tilde{c})/\tau)}{\sum_{j=0}^M \exp(\text{sim}(\tilde{n}, \tilde{c}_j)/\tau)}. \quad (1)$$

Thus, the concepts associated with phrase p are:

$$C_{p2}^n = \{c : \text{Assoc}(n, c) > \gamma, n \in p_n, c \in C^n\}. \quad (2)$$

This way, a phrase is assigned to concepts it shares the greatest semantic similarity. Thus, the final set of phrase concepts is: $C_p^n = C_{p1}^n \cup C_{p2}^n$. Note that phrases near multiple concepts will be assigned to all of them. This ensures we are confident in all negative concepts for a phrase.

A traditional choice of clustering-algorithm/textual embedding is GLoVe [23]/K-Means. However, this combination would result in many noisy concepts (discussed further in Section 4.2). Alternatively, we note that textual embeddings that are visually grounded and designed to make fine-grained distinctions, like ViCO [9]. In addition, density based clustering, like DBSCAN [6], only clusters items that are very close together rather than forcing every item to belong to a cluster like K-means, so we can be more certain they belong to the same concept. As we will show, these choices enabled us to produce a set of concepts with minimal noise. This, in turn, has improved our concept-phrase assignment process in Equation 2. In addition to ViCO, we experimented with transformer based embeddings such as BERT [30]. However, we noticed that there was no semantic similarity in the generated concepts. This is because individual nouns do not have enough context which is critical for transformers to work well. Therefore, we do not report the performance of these embeddings in our experiments.

After assigning each phrase p to its set of related concepts C_p , we use the phrase’s unrelated concepts $\overline{C_p} = C \setminus C_p$ as negative samples. More concretely, we pair $\overline{C_p}$ with the region r associated with phrase p to obtain negative concept-region pairs. These pairs are concatenated with the ground truth phrase-region positive pairs during training.

3.2. Concept based Fine-Grained Discrimination

The mined negative concepts obtained in Section 3.1 which improve our CFCD-Net’s discriminative power against a wide spectrum of nouns, but phrases may contain other important cues, such as attributes found in adjectives. However, noun concepts refer to semantically similar objects (*i.e.*, potential unannotated positives), whereas adjective concepts group together words that refer to fine grained differences (*e.g.*, color, texture patterns). Thus, we can not discriminate between concepts like we did for nouns since a phrase can have multiple adjectives from multiple concepts (*e.g.*, *red striped shirt*). However, we can confidently differentiate words within a single concept since its members are often mutually exclusive (*e.g.*, red vs. green). We obtain the set of adjective concepts C^a by replacing nouns with adjectives in the procedure outlined in Section 3.1.

Having obtained the set of adjective concepts, we introduce a novel module, Fine Grained Module (FGM), to discriminate between each concept members. Formally, the FGM module encodes image regions using a set of convolutional layers. It then performs multi-label classification on the members of each adjective based concept (*e.g.*, colors)

$c \in C^a$, then uses them to augment the main model region-phrase scores. Formally, given concept $c \in C^a$, $a \in c$, let R be the number of regions/fine-grained adjective pairs, s^a be the region-adjective score, and l^a be its 0/1 label indicating whether it is a positive/negative region-adjective pair, then:

$$L_{FGM} = \sum_i^R l_i^a \log s_i^a + (1 - l_i^a) \log(1 - s_i^a). \quad (3)$$

Module Inference with FGM: At test time, the FGM module’s scores are augmented with the base model’s phrase-region scores. Given a phrase p , the phrase adjectives p_a and $a \in p_a$, a phrase-region pair score s^p , and adjective-region score s^a , then the final score s^f :

$$s^f = (1 - \lambda_c) s^p + \lambda_c s^a, \quad (4)$$

where λ_c is a scalar that applies for every $a \in c$. With this, the final loss for CFCD-net is:

$$L_{final} = L_{base} + L_{FGM}, \quad (5)$$

where L_{base} is the loss used to train the phrase detection backbone which our model is agnostic to.

3.3. Hard Negative Mining

Using our proposed Negative Coarse Concepts (NCC) we ensure we learn to discriminate between representatives of the full distribution of possible phrases. However, discriminating between closely related phrases (*e.g.* car vs. truck) is important for good performance. Despite the noise hard-negative mining methods have shown to produce in phrase grounding [26], the benefits have been shown to overcome these issues [11]. Thus, we adapt Negative Phrase Augmentation (NPA) to help our model distinguish between phrases that are often confused with each other. For each phrase p in the validation set, we record the non-ground truth regions that the model is likely to associate p with. Then, we register the phrases associated with these regions as hard-negative candidates for p . We store these candidates in a "Confusion Table" and update it every 3 epochs. Following [11], we reduce some noise from known mutually non-exclusive phrases using WordNet [22]. Specifically, we remove any potential hard-negatives that share words with a parent-child relationship in the WordNet hierarchy. In other words, the phrase *vehicle* would remove *car* as a potential hard-negative since they hold parent-child relationship. After this process, for each positive phrase-region pair (p, r) in the batch, we sample a candidate hard negative phrase \bar{p} and concatenate the negative pair (\bar{p}, r) to the batch.

4. Experiments

Datasets: We evaluate CFCD-net on two common phrase grounding datasets. First, we use Flickr30K Entities [27]

#Train Samples	Flickr30K Entities [27]				RefCOCO+ [37]			
	zero-shot	few-shot	common	mean	zero-shot	few-shot	common	mean
	0	1-100	> 100		0	1-100	> 100	
(a) QA R-CNN [11]	3.9	4.3	8.9	5.7	0.9	1.5	9.3	3.9
Subquery [34]	–	–	–	–	0.7	1.3	9.2	3.7
FAOG [31]	3.2	3.5	7.6	4.8	0.7	1.1	8.9	3.6
MDETR [14]	1.5	2.1	4.8	2.8	1.4	2.6	10.1	4.7
R-CLIP [3]	8.8	6.2	3.9	6.3	4.2	3.5	2.8	3.5
GLIP [18]	4.3	6.9	11.3	7.5	2.0	4.5	13.0	6.5
SimNet [32]	4.7	4.4	8.6	5.9	2.0	3.3	13.1	6.1
CCA [27]	8.6	10.5	17.2	12.1	5.7	8.4	20.3	11.5
(b) SimNet w/CCA [26]	9.7	11.2	17.3	12.7	6.0	10.2	20.1	12.1
+ NCC	10.2	11.9	18.1	13.5	6.1	10.2	21.7	12.7
+ NPA	10.1	12.0	18.9	13.7	5.9	10.5	22.9	13.1
+ NPA + NCC	10.6	12.6	19.6	14.1	6.3	10.5	23.4	13.4
CFCD-Net (ours, NPA + NCC + FGM)	10.9	12.9	19.8	14.6	6.6	10.6	23.7	13.7
# Unique Phrases	1783	2764	472	5019	5653	2293	48	7994
# Instances	1860	4373	8248	14481	5758	3686	1171	10615

Table 1. mAP Split by frequency of training instances. **(a)** contains results reported in prior work or produced using their code. **(b)** contains ablations of our model that compares the performance of our its three components (NPA, NCC, and FGM). See Section 4.1 for discussion

Model	Training Data	Inference Time
MDETR [14]	COCO, VG, F30K Entities(200k)	7 Days
GLIP [18]	FourODs,GoldG+, COCO(24M)	1 Day
ours	F30K Entities(30k)	< 1 hour

Table 2. Comparison between CFCD-Net and transformer based phrase localization models evaluated on Flickr30K Entities. See Section 4.1 for discussion

that consists of 276K bounding boxes in 32K images for the noun phrases associated with each image’s descriptive captions (5 per image) from the Flickr30K dataset [36]. We use the official splits [27] that consist of 30K/1K/1K train/test/validation images. Second, we evaluate on RefCOCO+ [37], which consists of 19,992 images from the COCO dataset [19] that have been labeled with 141,564 region descriptions. We use the official split [37], which splits the train/val and testing sets 16K/1.5K/1.5K. Both datasets are licensed under creative commons.

Metrics: We follow the evaluation protocols of Plummer *et al.* [26]. For every image we obtain the most likely region and confidence score for every phrase in our test split. For any ground truth phrases in an image, we consider them successfully localized if the predicted bounding box has at least 0.5 intersection-over-union with its ground truth bounding box. Then, we compute average precision (AP) for each phrase and then split them into zero-shot, few-shot, and common sets, based on if they didn’t occur in our training split, if they had between 1-100 occurrences, or if they

occurred more than 100 times, respectively. We then report an overall mAP for each set of phrases, as well as the average of them for an overall performance score. This procedure ensures that the zero-shot and few-shot phrases are not over represented compared to the common phrases, since the zero- and few-shot sets have more unique phrases, but represent a smaller portion of overall instances.

Implementation details. We use the same framework as the state-of-the-art [26] for a fair comparison. This combines a ResNet-101 [10] based Faster R-CNN model [29] to obtain a set of image regions. Phrases are representing using the average of HGLMM Fisher Vectors [17]. The model is optimized using ADAM [16] and we keep all hyperparameters introduced by prior work fixed. All hyperparameters introduced by our paper are set via grid search on the validation set (see Section 4.2 for a sensitivity analysis). This resulted in $\tau = 0.01$, $\gamma = 0.2$ from Eq. (1) and Eq. (2), respectively, for both datasets. We set DBSCAN’s sensitivity to 0.43 for Flickr30K and 0.53 for RefCOCO+. Our model was trained using a single NVIDIA RTX 8000 GPU.

4.1. Results

Table 1 compares our model (CFCD-Net) (averaged over 3 runs) with the state-of-the-art in phrase grounding. Comparing the last line of Table 1(b) to the results from prior work in Table 1(a) and the first line of Table 1(b) we get 1.5-2 point gain in mAP over the state of the art. This gain came mostly from common phrases, where we achieved a 2-3 point gain over prior work. Furthermore, note the re-

#Train Samples	Flickr30K Entities [27]				RefCOCO+ [37]			
	zero-shot	few-shot	common		zero-shot	few-shot	common	
	0	1-100	> 100	mean	0	1-100	> 100	mean
SimNet w/CCA [26]	13.0	14.7	12.0	13.2	7.0	10.3	11.0	9.4
+ NPA + NCC	12.9	13.9	17.1	14.7	7.1	10.6	17.1	11.6
CFCD-Net (NPA + NCC + FGM)	15.0	17.2	18.5	16.8	7.8	11.7	17.9	12.5

Table 3. Performance of the phrases impacted by our FGM module compared with the previous SOTA. See Section 4.1 for discussion

sults of transformer-based models MDETR [14], GLIP [18], and an adaptation of CLIP [28] to phrase grounding, R-CLIP [3]. These methods performed well on phrase localization and object detection benchmarks like LVIS [8]. This is further evidence that simply optimizing for the localization objective as in prior work does not yield discriminative models even as data is scaled up. Furthermore, as reported in Table 2, our model inference time is far shorter. This is because both MDETR and GLIP use cross modal transformers to fuse language/vision data adding significant computational complexity. In contrast, our computationally efficient fusion based on element-wise product followed by lightweight fully connected layers results in an inference time less than 1.1% that of either model.

In Table 1(b) we also report the contribution of each component of our model. We observe that NCC makes a significant improvements over both datasets. Although the improvements are quite similar to those made by NPA [11], they come at a fraction of the cost. Specifically, NPA must update a confusion table to select good negatives for every phrase in the training pool during training, which took 3 hours on Flickr and 4 hours on RefCOCO+ using 4 NVIDIA RTX8000 GPUs every time the confusion table was updated (every 3 epochs in our experiments). In contrast, the groups used by NCC are computed only once and takes a few seconds on a CPU using precomputed language features. Nevertheless, as shown in Table 1(b), we find we get best performance when both methods are combined (NPA + NCC).

In addition to the gains from using NCC and NPA, Table 1(b) reports that the FGM module (described in Section 3.2) further boosts performance for both datasets over all phrases. Note though that the FGM module only affects phrases with adjectives. Thus, in Table 3 we report performance over affected phrases, where we report a 3-3.5 gain on Flickr30K Entities and RefCOCO+. See supplementary for qualitative examples of our detection results.

4.2. NCC Hyperparameter Analysis

Clustering-Algorithm/Textual-Embedding. Table 4 reports the effect of different choices of clustering algorithm and language embedding used to create our NCC concepts (outlined in Section 3.1) has on phrase detection performance. A traditional choice of clustering-

algorithm/textual-embedding is K-Means clustering over GloVe features [23], but we show a small but mostly consistent gain with DBSCAN [6] + ViCo [9] instead.

To better understand the causes in these performance differences, we begin by computing a "visual coherence noise" metric over concepts. We manually inspected each concept and counted the number of times at least %50 of the words were not visually similar. Since there are relatively few concepts (less than 100), this only takes a few minutes for a single annotator. Figure 4 reports the performance of the different embedding and clustering combinations, where ViCO + DBSCAN also obtains best performance.

Figure 5 provides examples of the the different concepts. Note that the left concept related to K-means + Glove in Figure 5(a) contains largely unrelated words. This is because K-means does not impose a constraint over each cluster density, *i.e.*, outliers in the embedding space are clustered with other concepts even though there is little evidence they belong together. Furthermore, even when the concept is less noisy (concept on the right), it includes words that are semantically similar but not visually similar (*e.g.* Baker vs Dough). This likely harms our concept-phrase assignment process. For example, given a phrase *tall baker*, a textual embedding that is not visually grounded would likely assign a concept containing ("bread, baker, dough") as related to the phrase. However, a visually grounded embedding would only assign concepts that exclusively contain humans as related. ViCO [9] helps to address these shortcomings because it was explicitly trained to consider visual similarity, which should ensure that concepts rule out words that are semantically similar but not visually similar. Furthermore, DBSCAN [6], whose clusters are formed from words that fall within a density threshold ϵ . Thus, the method is more effective at ruling out outliers, as illustrated in Figure 5(d), which are more visually similar and include fewer outliers.

Number of Concepts. We also investigate the effect of changing the number of concepts in Figure 6. We progressively increase DBSCAN's density threshold for each dataset until we can not generate more concepts. In other words, as we increase the threshold, we create more concepts by force them to be more tightly coupled. This breaks apart less semantically related words, thus improving the accuracy of our phrase-concept assignment in Section 3.1.

#Train Samples	Flickr30K Entities [27]				RefCOCO+ [37]			
	zero-shot	few-shot	common	mean	zero-shot	few-shot	common	mean
	0	1-100	> 100		0	1-100	> 100	
Baseline [SimNet w/CCA]	9.7	11.2	17.3	12.7	6.0	10.2	20.1	12.1
NCC w/ K-means + GLoVe	10.0	11.8	18.5	13.4	6.0	9.9	20.4	12.1
NCC w/ K-means + ViCo	10.1	11.9	18.3	13.4	6.0	9.8	20.7	12.2
NCC w/ DBSCAN + GLoVe	9.9	11.9	18.4	13.4	6.1	10.4	20.7	12.4
NCC w/ DBSCAN + ViCo	10.2	11.9	18.1	13.5	6.1	10.2	21.7	12.7

Table 4. Comparison of the effect of clustering algorithm and language embedding when creating the concepts used by our NCC sampling strategy. See Section 4.2 for discussion.

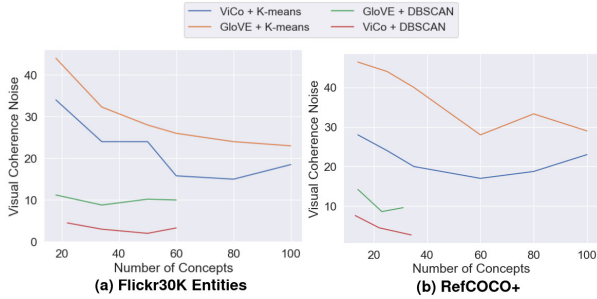


Figure 4. Quantitative comparison between different choices of clustering-algorithm/textual-embedding for concept generation. Refer to Section 3.1 for discussion



Figure 5. Qualitative comparison between different choices of clustering-algorithm/textual-embedding for concept generation. Refer to Section 3.1 for discussion

Concept-Phrase Assignment. Figure 7 reports the effect of changing γ in Eq. (2), which controls whether a given concept is assigned to a phrase (*i.e.* they are related). We vary γ between 0 (a phrase is assigned to every given concept) and 1 (a phrase is only assigned to concepts that it share nouns with). We report best performance for γ values between (0.2, 0.6). Performance drops for $\gamma < 0.2$ as we incorrectly assign concepts to many phrases. Performance also drops when $\gamma > 0.8$, demonstrating the importance of our similarity-based matching approach. As discussed in Section 3.1, this is due to many concepts being relevant to phrases despite not sharing words.




Figure 6. Effect of Concept Size $|C|$




Figure 7. Effect of γ in Equation 2

5. Conclusion

In this work, we introduced a new phrase detection mode (CFCD-NET) that improves performance by 1.5-2 points on two phrase detection datasets. It does so by incorporating visually coherent clusters (concepts) to sample negative concept-region that effectively improve the model discriminative abilities when compared to prior work. Our model further improves performance by incorporating a novel fine grained module that learns to discriminate between adjective fine grained tokens. Notably, our approach even outperforms recent transformer-based methods like MDETR [14] and GLIP [18] on phrase detection performance while requiring less training data and significantly faster inference speeds. In addition, although our experiments used the Faster R-CNN framework to fairly compare to prior work, our contributions are modular and can be adapted to any underlying detection framework.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [3](#)
- [2] Mohit Bajaj, Lanjun Wang, and Leonid Sigal. G3raphground: Graph-based language grounding. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. [3](#)
- [3] Zhaowei Cai, Gukyeong Kwon, Avinash Ravichandran, Erhan Bas, Zhuowen Tu, Rahul Bhotika, and Stefano Soatto. X-detr: A versatile architecture for instance-wise vision-language tasks, 2022. [6](#), [7](#)
- [4] Kan Chen, Rama Kovvuri, and Ram Nevatia. Query-guided regression network with context policy for phrase grounding. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 824–832, 2017. [3](#)
- [5] Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. Align2ground: Weakly supervised phrase grounding guided by image-caption alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [3](#)
- [6] Dingsheng Deng. Dbscan clustering algorithm based on density. In *2020 7th International Forum on Electrical Engineering and Automation (IFEAA)*, pages 949–953, 2020. [5](#), [7](#)
- [7] Zhiyuan Fang, Shu Kong, Charless Fowlkes, and Yezhou Yang. Modularized textual grounding for counterfactual resilience. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [3](#)
- [8] Agrim Gupta, Piotr Dollár, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [4](#), [7](#)
- [9] Tanmay Gupta, Alexander Schwing, and Derek Hoiem. Vico: Word embeddings from visual co-occurrences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [5](#), [7](#)
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. [2](#), [6](#)
- [11] Ryota Hinami and Shin’ichi Satoh. Discriminative learning of open-vocabulary object retrieval and localization by negative phrase augmentation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2605–2615, Brussels, Belgium, Oct.-Nov. 2018. Association for Computational Linguistics. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [12] Hexiang Hu, Ishan Misra, and Laurens van der Maaten. Binary image selection (bison): Interpretable evaluation of visual grounding. *ArXiv*, abs/1901.06595, 2019. [3](#)
- [13] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [3](#)
- [14] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr - modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1780–1790, October 2021. [1](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [15] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 787–798, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. [3](#)
- [16] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. [6](#)
- [17] Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf. Associating neural word embeddings with deep image representations using fisher vectors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4437–4446, 06 2015. [6](#)
- [18] Liumian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [19] 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 David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing. [6](#)
- [20] Chenxi Liu, Junhua Mao, Fei Sha, and Alan Yuille. Attention correctness in neural image captioning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, page 4176–4182, 2017. [3](#)
- [21] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11–20, 2016. [3](#)
- [22] George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, Nov. 1995. [5](#)
- [23] Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. [5](#), [7](#)
- [24] Bryan A. Plummer, Paige Kordas, M. Hadi Kiapour, Shuai Zheng, Robinson Piramuthu, and Svetlana Lazebnik. Conditional image-text embedding networks. In *The European Conference on Computer Vision (ECCV)*, 2018. [3](#)
- [25] Bryan A. Plummer, Arun Mallya, Christopher M. Cervantes, Julia Hockenmaier, and Svetlana Lazebnik. Phrase localization and visual relationship detection with comprehensive

- image-language cues. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 3
- [26] Bryan Allen Plummer, Kevin Shih, Yichen Li, Ke Xu, Svetlana Lazebnik, Stan Sclaroff, and Kate Saenko. Revisiting image-language networks for open-ended phrase detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, page 1–1, 2020. 2, 3, 5, 6, 7
- [27] Bryan A. Plummer, Liwei Wang, Christopher M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *International Journal of Computer Vision*, 123(1):74–93, 2017. 3, 5, 6, 7, 8
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021. 7
- [29] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 2, 6
- [30] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv:1908.08530*, 2019. 5
- [31] Josiah Wang and Lucia Specia. Phrase localization without paired training examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 6
- [32] Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. Learning two-branch neural networks for image-text matching tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 41(2):394–407, 2018. 6
- [33] Chao-Yuan Wu, R. Manmatha, Alexander J. Smola, and Philipp Krähenbühl. Sampling matters in deep embedding learning. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 4
- [34] Z. Yang and et al. Improving one-stage visual grounding by recursive sub-query construction. In *The European Conference on Computer Vision (ECCV)*, 2020. 6
- [35] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [36] Peter Young, Alice Lai, Micah Hodosh, and J. Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014. 6
- [37] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. Modeling context in referring expressions. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 69–85, Cham, 2016. Springer International Publishing. 3, 6, 7, 8
- [38] Yuting Zhang, Luyao Yuan, Yijie Guo, Zhiyuan He, I-An Huang, and Honglak Lee. Discriminative bimodal networks

for visual localization and detection with natural language queries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 3