

Zero-Shot Grounding of Objects from Natural Language Queries

Arka Sadhu¹ Kan Chen^{2†} Ram Nevatia¹
¹University of Southern California ²Facebook Inc.
 {asadhu|nevatia}@usc.edu kanchen18@fb.com

Abstract

A phrase grounding system localizes a particular object in an image referred to by a natural language query. In previous work, the phrases were restricted to have nouns that were encountered in training, we extend the task to **Zero-Shot Grounding (ZSG)** which can include novel, “unseen” nouns. Current phrase grounding systems use an explicit object detection network in a 2-stage framework where one stage generates sparse proposals and the other stage evaluates them. In the ZSG setting, generating appropriate proposals itself becomes an obstacle as the proposal generator is trained on the entities common in the detection and grounding datasets. We propose a new single-stage model called **ZSGNet** which combines the detector network and the grounding system and predicts classification scores and regression parameters. Evaluation of ZSG system brings additional subtleties due to the influence of the relationship between the query and learned categories; we define four distinct conditions that incorporate different levels of difficulty. We also introduce new datasets, sub-sampled from Flickr30k Entities and Visual Genome, that enable evaluations for the four conditions. Our experiments show that ZSGNet achieves state-of-the-art performance on Flickr30k and ReferIt under the usual “seen” settings and performs significantly better than baseline in the zero-shot setting.

1. Introduction

Detecting objects in an image is a fundamental objective of computer vision. A variation of this task is *phrase grounding* (also called *visual grounding* and *referring expressions*) where the objective is to detect objects referenced by noun phrases in a text query [7, 18, 39, 44]. It can be directly applied to other tasks such as visual question answering [1, 50] and image retrieval [5] and has thus garnered wide interest.

While existing phrase grounding systems accept novel query phrases as inputs, they are limited to the nouns en-

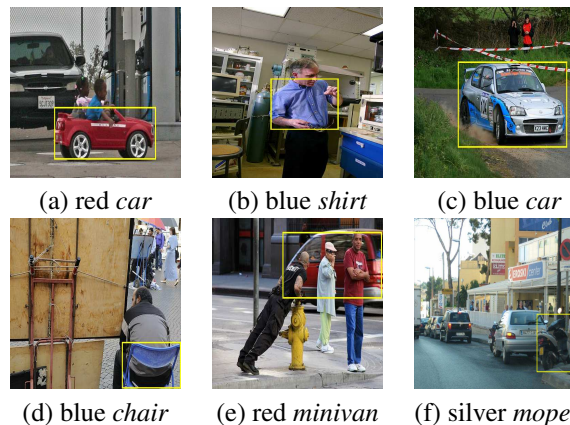


Figure 1. Illustration of the key difference between current scope of phrase grounding and the proposed zero-shot grounding. The query word is italicized in all cases. (a)-(f) denote the image-query pairs input to the system. (a) and (b) are examples of training images. A test image query pair for phrase grounding could be (c). Zero-shot grounding additionally can be tested on (d), (e) and (f) in which “chair”, “minivan” and “moped” are object categories not annotated in the training data. (f) additionally contains a “car” object which is a trained category, indicating that both novel and related trained category objects may be present in a test image.

countered in the training data (*i.e.* the referred object types need to have been “seen” in training images before). As an important extension, we define zero-shot grounding (ZSG) to allow the use of phrases with nouns that the grounding system has not encountered in training set before. Fig 1 illustrates this concept with examples.

To enable grounding of novel object categories, we need to relate the appearance of referred objects to their linguistic descriptions. Current, state-of-art phrase grounding systems [7, 9, 33, 44, 48] rely on an explicit object detector to obtain proposed object bounding boxes and their ROI-pooled features as a pre-processing step. This essentially limits these systems to a fixed set of object categories that the detector was trained on. In ZSG, we need to have a reasonable proposal box for the novel object, classify it as a foreground and regress the box to be a more accurate spatial fit. In traditional phrase grounding, a key challenge is to dis-

[†]This work was done while the author was at USC.

ambiguate between similar objects using the query phrase, but ZSG requires us to also first find likely image regions that may contain the referenced objects.

To address the above issues, we replace the traditional two-stage approach, where the first stage generates proposal bounding boxes and the second stage does the classification, by a single-stage network with dense proposals; we call this network **ZSGNet**. It takes combined language query features and visual features from the image proposals and predicts classification scores and regression parameters. The system is trained directly on the grounding training data, in an end-to-end manner, and does not utilize any externally trained object detector. We show that, besides enabling grounding of novel categories, it does not degrade performance on learned categories even though our method does not utilize external training data. Moreover, our design is computationally efficient especially during inference owing to its single-stage architecture akin to SSD [28].

Evaluating the performance of a ZSG method is complex due to the influence of the relationship of the new query category to the learned categories. To make the evaluations and distinctions clearer, we define four specific cases for different conditions: (i) when the query word is novel (Fig 1 d-f) (ii) when the referred object belongs to a novel category (Fig 1-d) (iii) when the referred object is “similar” to objects seen during training but none of the latter are present (Fig 1-e) (iv) when at least one similar object also exists in the test image (Fig 1-f)(more details in Section 3.1).

To support evaluation of zero-shot grounding for the four cases, we introduce new datasets which are sub-sampled from the existing Visual Genome [22] and Flickr30k Entities [34]. We create examples of the four cases outlined above (dataset creation details are in Section 4.1, experiments on these datasets are in Section 4.5).

Our contributions can be summarized as follows: (i) we introduce the problem of Zero-shot grounding, (ii) we propose a simple yet effective architecture ZSGNet to address limitations of current phrase grounding systems for this task, (iii) we create new datasets suitable for evaluating zero-shot grounding and (iv) we evaluate performance on these datasets and show the effectiveness of our approach. Our code and datasets are publicly released¹.

2. Related Work

Phrase grounding: Extensive work in creating grounding datasets like Flickr30k, ReferIt, RefCoCo, RefCoCo+, RefCoCog, Visual Genome, GuessWhat [8, 19, 22, 30, 34, 43, 45] have been crucial to the success of phrase grounding. Early works use reconstruction based approach [39] or integrate global context with the spatial configurations [18]. Recent approaches [7, 33, 48] learn directly in the multi-

modal feature space and use attention mechanisms [9, 44] which have also been extended to phrase grounding in dialogue systems [8, 51]. Few approaches also look at unsupervised learning using variational context [49] and semi-supervised learning via gating mechanisms [6].

Above techniques use an object detector like FasterRCNN [38] or MaskRCNN [15] as a pre-processing step to get the bounding boxes and ROI-pooled features which effectively limits them to the object categories of the detector. We combine the detection and the grounding networks and learn directly from the grounding dataset and thus no pre-processing step is involved.

Multi-modal feature representation has many flavors like linear transformation, concatenation, hadamard product [20], bilinear pooling [27] and have shown success in vision-language tasks like VQA [3, 4, 13, 46, 47], Scene Graph Generations [24, 42] and Image Captioning [29]. We stick to feature concatenation for simplicity and a fair comparison with previous works in phrase grounding.

Single stage networks used in object detection are popular for their real-time inference speed. Prominent works include SSD [28], YOLO [35–37] and more recently FPN [25] and RetinaNet [26]. In this work, we combine a single-stage detection network directly into the grounding framework; besides enabling zero-shot grounding, it also results in highly efficient inference.

Zero-shot grounding is unexplored but there are a few similar works. [14] aims at open-vocabulary object retrieval though it still assumes the entity of the referred object is seen at train time. Recently [2] proposed zero-shot detection (ZSD) where they consider a set of unseen classes for which no bounding box information at train time. At test time, all the objects including the unseen classes must be detected. However, the set of background classes is needed prior to training, but this is not needed in ZSG.

3. Design Considerations for ZSG

We first discuss the zero-shot grounding cases and then describe the limitations in extending current phrase grounding systems to ZSG. Finally, we present a new architecture to address the limitations.

3.1. ZSG Cases

We now describe the four cases for zero-shot grounding in detail. For brevity, we use the notations in Table 1. Each case defines the scope for what is classified as a zero-shot grounding example. Further, we assume that Q (the word which refers to the object in the image) is not an OOV (out of vocabulary word) which is reasonable if we use word embeddings which are trained on a large language corpus.

Case 0: $Q \notin W$. The query noun, Q , is not included in any training example before. We only look at the lemmatized word so synonyms are considered to be different

¹<https://github.com/TheShadow29/zsgnet-pytorch>

Notation	Meaning	Example
T	Test Image	Fig 1-(f)
P	Test query phrase	silver moped
A	Referred object at test time	Moped
Q	Word in P referring to A	moped
B	Set of objects close to A and seen during training	{Car}
C	Set of categories seen during training	{Vehicles, Clothing}
W	Set of words seen during training	{red, blue, car, shirt}

Table 1. Notations used to describe ZSG with examples (Fig 1). By close objects we mean their word embeddings are similar.

(novel) words. Fig 1(d)-(f) are examples of this case. Fig 1-c with the phrase “blue automobile” would also be considered zero-shot since we haven’t seen the word automobile before even though it is a synonym of “car”.

Case 1: $A \notin C$. Here, we assume that objects seen at train time belong to a set of pre-defined categories and the referred object A doesn’t belong to these categories. In Fig 1-d, “chair” is considered zero-shot as this category was not seen at train time.

Case 2: $\exists B$ but $\forall b \in B$ we have $b \notin T$. Here, objects that are semantically close (similar) to the referred object A are present in the training set but not in the test image. Fig 1-e is an example as “minivan” (novel object) is semantically close to “car” (seen in train set) but there is no other similar object like “car” in the test image.

Case 3: $\exists B$ and $\exists b \in B$ such that $b \in T$. Same as Case 2 but at least one of the objects semantically close (similar) to A is also present in the test image. For example, Fig 1-f containing “moped” (a novel object) and “car” (seen in the training set) which are semantically close.

For Case 2 and Case 3, there can be multiple interpretations for being “semantically close”. Here, we assume two objects are “close” if their word embeddings are similar. In our implementation, we cluster the word embeddings of the objects and objects belonging to the same cluster are considered semantically close (more details in Section 4.1).

3.2. Limitations in Phrase Grounding Systems

Prior works view phrase grounding either as an entity selection problem [9, 44] or that of regressing sparse proposals to a tighter bounding box [7, 33, 48]. In either case, given an image I , we have N candidate boxes and their ROI-pooled features $\{o_i\}_{i=1}^N$. Given a query phrase P , the problem reduces to finding a good candidate box with a possible additional regression step.

Grounding systems using this framework don’t have a mechanism to generalize to object categories not in the detection dataset. Consider a novel category X whose in-

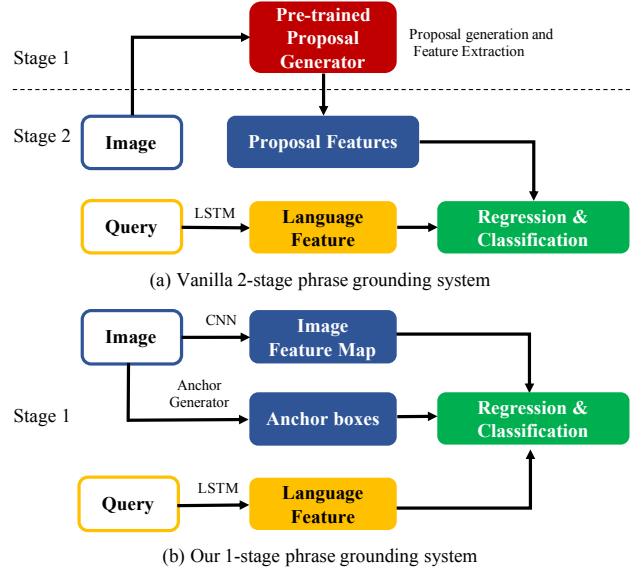


Figure 2. Previous phrase grounding systems (a) produce a small subset of proposals without considering the query restricting it to the entities of the detection network. Our system (b) considers dense proposals, looks at the query to disambiguate and learns directly from grounding dataset

stances may be present in the training images but not annotated. The object detector learns to classify X as background and this error is propagated to the grounding system. [7, 33] suggest fine-tuning the detector on the grounding categories (entities) but the grounding datasets are not densely annotated, *i.e.* not all object instances of X in every image are annotated. Additionally, some grounding datasets like ReferIt [19] don’t have entity information so fine-tuning is not feasible.

Object detectors also favor features invariant to intra-class changes but grounding systems need to capture intra-class variations as well.

3.3. Model Design

We propose the following new formulation: Given an image I with fixed candidate boxes (also called anchor boxes or dense proposals) $DP = [dp_1, \dots, dp_N]$ and a query phrase P , the task is to choose the best candidate box dp_i and regress it to a tight bounding box b_i . Since our candidate boxes depend only on the size of the image, we can use any image encoder to compute the visual features at runtime and remove the need for a pre-trained object detector as illustrated in Fig 2. This design is similar to single-shot architecture used in object detection [26, 28].

Framework: Our model consists of a language module to encode the query phrase, a visual module to obtain image feature maps followed by fully convolutional networks to output a $5d$ vector (for each candidate box) one for score

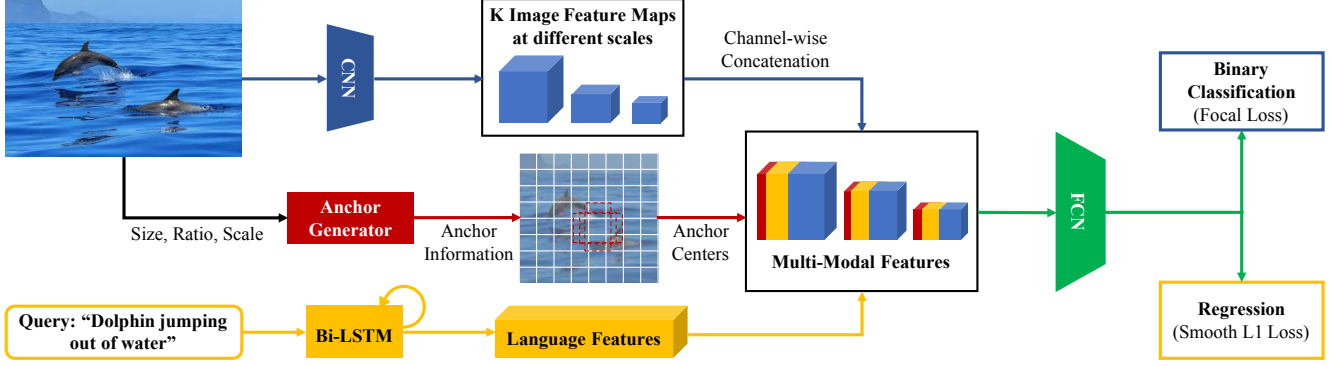


Figure 3. A schematic of the ZSGNet Architecture. Input to the system is an image-query pair. A deep network is used to produce K image feature maps taken at different resolutions. The anchor generator uses the image size to produce anchors at different scales and resolution. We append the anchor centers at each cell of each feature map. The query phrase is encoded using a bidirectional LSTM (Bi-LSTM) and the language feature obtained is appended at every cell location of every feature map along the channel dimension. The resulting multi-modal feature maps are input to a Fully Convolutional Network (FCN) block to output a prediction score and regression parameters which are trained using focal-loss (L_{pred}) and SmoothL1-loss (L_{reg}) respectively.

and the rest for regressing to a tighter bounding box. Fig 3 provides an overview of our proposed architecture.

ZSGNet directly learns about the entities in a grounding dataset in an end-to-end fashion. Moreover, since a query phrase refers to a particular object with possibly different attributes, the visual features are no longer invariant to intra-class changes. This way, we address the limitations posed by previous systems. Finally, owing to its single-stage approach, the network is computationally efficient.

Language module consists of an embedding layer followed by a Bi-LSTM [17, 41] to encode the input query phrase. Given a query phrase $P = \{w_i\}_{i=1}^n$ we use GloVe vectors [32] to encode each word in P as word embedding vectors $\{w_i\}_i^n \in \mathbb{R}^{d_q}$, where d_q is the dimension of the embedding vector. These are fed into a Bi-LSTM [17, 41]. We use the normalized last hidden state vector $\{\hat{h}\} \in \mathbb{R}^{2d_l}$ of Bi-LSTM as the query feature, where d_l is the dimension of the hidden layer of a single LSTM.

Visual Module consists of an image encoder to produce K feature maps $\{v_i\}_{i=1}^K$ at different resolutions. We use ResNet-50 [16] with FPN [25] as our default image encoder. We first normalize the visual feature maps along the channel dimension. Then we expand the language feature to the same dimensions of the visual feature maps and concatenate it along the channel dimension for each normalized visual feature map \hat{v}_i . Finally, we append the normalized locations of the feature maps ($C_x, C_y = [c_x/W, c_y/H]$) along the channel dimension to aid in location based grounding (phrases which contain location information) and obtain the multi-modal feature maps m_i . At a particular index of the i^{th} feature map (indexed by x, y) we have

$$m_i[x, y] = [\hat{v}_i[x, y]; \hat{h}; C_x; C_y] \quad (1)$$

where ; denotes the concatenation operation.

Anchor Matching Following [26] we match 9 candidate boxes to every index of a feature map. We use a fully convolutional network (FCN) to process the multi-modal features to output a 5 dimensional output (score and regression parameters) for each box. In the anchor matching step, we use an IoU threshold of 0.5 (found ideal via experimentation).

Loss Function For the binary classification into foreground and background we use the focal loss as described in [26]. For regressing to a tighter box we use the same encoding scheme as [38] with smooth-L1 loss.

Let dp_j denote the j^{th} anchor and gt denote the ground truth bounding box. Let

$$g_{dp_j} = \mathbb{1}_{IoU(dp_j, gt) \geq 0.5} \quad (2)$$

$$G = \{g_{dp_j} | g_{dp_j} = 1\} \quad (3)$$

Here $\mathbb{1}$ denotes the indicator random variable. Thus, $g_{dp_j} = 1$ means the candidate box dp_j matches with the ground-truth box and G is the set of all such candidate boxes. Now denoting focal loss [26] with default parameters ($\alpha = 0.25, \gamma = 2$) by L_F and the predicted score for the box dp_j as p_{dp_j} we get

$$L_{pred} = \frac{1}{|G|} \sum_{j=1}^{|DP|} L_F(p_{dp_j}, g_{dp_j}) \quad (4)$$

Similarly, denoting SmoothL1-loss by L_S , the predicted regression parameters by r_{dp_j} and the ground-truth regression parameters by gt_{dp_j} we get

$$L_{reg} = \frac{1}{|G|} \sum_{j=1}^{|DP|} g_{dp_j} L_S(r_{dp_j}, gt_{dp_j}) \quad (5)$$

The final loss is calculated as $L = L_{pred} + \lambda L_{reg}$. Here λ is a hyper-parameter (we set $\lambda = 1$).

Training: We match the candidate boxes (anchors) to each feature map generated by the feature encoder. We classify each candidate box as a foreground or a background using a prediction loss (L_{pred}) and regress it to get a tighter box (L_{reg}). Foreground means that the candidate box shares $IoU \geq 0.5$ with the ground truth box. For the regression loss we only consider the foreground candidate boxes.

Testing At test time, we choose the candidate box with the highest score and use its regression parameters to obtain the required bounding box.

4. Experiments

This section describes the dataset construction methods, followed by experiments and visualization.

4.1. Dataset Construction

We sub-sample Flickr30k [34] and Visual Genome [22] to create datasets for the cases described in Section 3.1.

Flickr30k Entities contains 5 sentences per image with every sentence containing 3.6 queries and has bounding box information of the referred object and its category (e.g. “people”, “animal”).

Visual Genome (VG) has a scene graph for every image. The objects in the scene-graph are annotated with bounding boxes, region description and a synset (obtained from [31]).

We briefly describe the steps taken to create the ZSG datasets (more details can be found in the supplementary material). We follow the notation described in Table 1.

Case 0 is sampled from Flickr30k Entities [34]. We need to ensure that $Q \notin W$. For this, we first obtain lemmatized representation of the query words. As the query phrases are noun-phrases of a complete annotated sentence, the query word Q referring to the noun is almost always the last word of the query phrase P , we take it be so. We do a 70:30 split of the extracted words to obtain “included” and “excluded” word lists respectively. We then create a training set from the included list and validation and tests from the excluded list, removing images that have overlap between train, validation or test lists. We call the resulting split **Flickr-Split-0**.

Case 1 is also sampled from Flickr30k Entities [34] but this time we also use predefined entity information. We need the referred object A to belong to a category which is not in C . Flickr30k has 7 common object categories (e.g. “people”, “animals”) and one category called “other” for objects which do not belong to the seven categories. We take images with “other” objects and split them evenly to create validation and test sets, The remaining images comprise the train set; we remove any box annotations that belong to the “other” category to create **Flickr-Split-1**.

Case 2 and Case 3 are sampled from Visual Genome [22]. In addition to region phrases, visual genome also provides entity names mapped to synsets in wordnet [31]. We count all the objects present in the dataset, choose

$topI (= 1000)$ objects and get their word embeddings, skipping without an embedding (we use GloVe [32] trained on common crawl corpus). We apply K-Means clustering ($K = 20$) to cluster similar words. We sort the words in each cluster k by their frequency in the dataset and take the top half to be in the seen objects set (S_k) and the bottom half to be in the unseen objects set (U_k). If an image contains an object o_i such that $o_i \in U_k$ and another object $o_j \in S_k$ then it is an example of Case 3. If no such o_j exists then it is Case 2. Finally, we take the union of images of the two cases to constitute the test set. We call the resulting splits **VG-Split-2** and **VG-Split-3**. This design ensures that for both **Cases 2 and 3** the referred object A (in the test set) has a set of objects B (in the training set) which are in the same semantic cluster.

We consider the remaining images to be candidates for the training set and include them if they contain at least one object o_i in S (where $S = \cup_k S_k$) and remove the annotations for any object $o_j \in U$ (where $U = \cup_k U_k$). This ensures that the training set contains objects in S and does not contain any objects in U . However, such a training set turns out to be extremely imbalanced with respect to the clusters as clusters containing common entities such as “person” are much more prevalent than clusters containing “cakes”. We balance the training set following a simple threshold based sampling strategy (details in supplementary material) which results in most clusters (except 2) to have similar number of query phrases. We follow the same strategy to create balanced test splits of **VG-Split-2** and **VG-Split-3**.

Dataset Caveats: (i) We note that polysemy is not taken care of *i.e.* the same word can have different meanings. (ii) Neither Visual Genome nor Flickr30k is a true referring expressions dataset *i.e.* the query phrase may not always uniquely identify an object.

4.2. Datasets Used

Flickr30k Entities [34] contains 30k images each with 5 sentences and each sentence has multiple query phrases. We use the same splits used in [7, 18, 39].

ReferIt(RefClef) [19] is a subset of Imageclef [11] containing 20k images with 85k query phrases. We use the same split as [7, 18, 39].

Flickr-Split-0 We create an unseen split of Flickr30k based on the method outlined in Section 4.1. It contains 19K train images with 11K queries, 6K validation images with 9K queries and 6K test images with 9K queries.

Flickr-Split-1 This split of Flickr30k has “other” category only in the validation and test images. It contains 19k training images with 87k query phrases and 6k images with 26k query phrases for validation and test each.

VG-Split We use a balanced training set (as described in Section 4.1) containing 40K images and 264K query phrases. We use a subset (25%) for validation. **VG-Split-2**

contains 17K images and 23K query phrases in the unbalanced set, 10K images and 12K query phrases in the balanced set. **VG-Split-3** contains 41K images with 68K query phrases in the unbalanced set, 23K images and 25K query phrases in the balanced set.

4.3. Experimental Setup

Evaluation Metric: We use the same metric as in [7]. For each query phrase, we assume that there is only one ground truth bounding box. Given an image and a query phrase if the *IoU* of our predicted bounding box and the ground truth box is more than 0.5 we mark it as correct. However, in the case of Visual Genome splits, the annotations are not precise so we use 0.3 as the threshold. The final accuracy is averaged over all image query phrase pairs.

Baselines: To explicitly compare the performance of dense proposals, we create a new baseline **QRG** based on QRC [7] which uses GloVe embeddings instead of embeddings learned from the data. We benchmark it on Flickr30k to show there is no drop in performance compared to QRC. We further use it as a strong baseline on the unseen splits. In all cases, we use a fasterR-CNN [38] pretrained on Pascal-VOC [12] and fine-tune it on the target dataset. For **Flickr-Split-0** we fine-tune on all the entities, for **Flickr-Split-1** we fine-tune on all entities except “other”. We use the top-100 box predictions provided by the fine-tuned network after applying non-maxima suppression to be consistent with implementation in [7]. For **VG-Split**, we train on all the seen-classes, *i.e.* union of all seen objects in every cluster ($\cup_k S_k$). In this case, we don’t use non-maxima suppression and instead consider all the 300 boxes provided by the fine-tuned region-proposal network.

Implementation details: We train ZSGNet and baseline models till validation accuracy saturates and report our values on the test set. We found Adam [21] with learning rate $1e^{-4}$ for 20 epochs to be sufficient in most cases. For ZSGNet, to generate image feature maps at different scales, we use two variations: (i) SSD [28] with VGG network (ii) RetinaNet [26] with Resnet-50 [16] network. Note that these are not pretrained on any detection dataset. Initially, we resize the image to 300×300 for faster training and later retrain with image sizes 600×600 which gives a consistent 2 – 3% improvement. We note that while image augmentations (like horizontal flipping) are crucial for object detectors it is harmful for grounding as the query phrases often have location information (like “person standing on the left”, “person to the right of the tree”).

4.4. Results on Existing Grounding datasets

Table 2 compares ZSGNet with prior works on Flickr30k Entities [34] and ReferIt [19]. We use “det” and “cls” to denote models using Pascal VOC [12] detection weights and ImageNet [10, 40] classification weights. Networks marked

Method	Net	Flickr30k	ReferIt
SCRC [18]	VGG	27.8	17.9
GroundeR (cls) [39]	VGG	42.43	24.18
GroundeR (det) [39]	VGG	48.38	28.5
MCB (det) [13]	VGG	48.7	28.9
Li (cls) [23]	VGG	-	40
QRC* (det) [7]	VGG	60.21	44.1
CITE* (cls) [33]	VGG	61.89	34.13
QRG* (det)	VGG	60.1	-
ZSGNet (cls)	VGG	60.12	53.31
ZSGNet (cls)	Res50	63.39	58.63

Table 2. Comparison of our model with other state of the art methods. We denote those networks which use classification weights from ImageNet [40] using “cls” and those networks which use detection weights from Pascal VOC [12] using “det”. The reported numbers are all $Accuracy@IoU = 0.5$ or equivalently $Recall@1$. Models marked with “*” fine-tune their detection network on the entities in the Flickr30k.

with “*” fine-tune their object detector pretrained on Pascal-VOC [12] on the fixed entities of Flickr30k [34].

However, such information is not available in ReferIt dataset which explains $\sim 9\%$ increase in performance of ZSGNet over other methods. This shows that our model learns about the entities directly from the grounding dataset.

For Flickr30k we also note entity-wise accuracy in Table 3 and compare against [7, 33]. We don’t compare to the full model in [7] since it uses additional context queries from the sentence for disambiguation. As these models use object detectors pretrained on Pascal-VOC [12], they have somewhat higher performance on classes that are common to both Flickr30k and Pascal-VOC (“animals”, “people” and “vehicles”). However, on the classes like “clothing” and “bodyparts” our model shows much better performance; likely because both “clothing” and “bodyparts” are present along with “people” category and so the other methods choose the “people” category. Such biases are not exhibited in our results as our model is category agnostic.

4.5. Results on ZSG Datasets

Table 4 shows the performance of our ZSGNet model compared to QRG on the four unseen splits described in Section 4.1 and 4.2. Across all splits, ZSGNet shows 4–8% higher performance than QRG even though the latter has seen more data (the object detector is pretrained on Pascal VOC [12]). Next, we observe that the accuracy obtained on Flickr-Split-0,1 are higher than the VG-split likely due to larger variation of the referred objects in Visual Genome. Finally, the accuracy remains the same across the balanced and unbalanced sets indicating the model performs well across all clusters as our training set is balanced.

We also study the relationship between accuracy and the

Method	Overall	people	clothing	bodyparts	animals	vehicles	instruments	scene	other
QRC - VGG(det)	60.21	75.08	55.9	20.27	73.36	68.95	45.68	65.27	38.8
CITE - VGG(det)	61.89	75.95	58.50	30.78	77.03	79.25	48.15	58.78	43.24
ZSGNet - VGG (cls)	60.12	72.52	60.57	38.51	63.61	64.47	49.59	64.66	41.09
ZSGNet - Res50 (cls)	63.39	73.87	66.18	45.27	73.79	71.38	58.54	66.49	45.53

Table 3. Category-wise performance with the default split of Flickr30k Entities.

Method	Net	Flickr-Split-0	Flickr-Split-1	VG-2B		VG-2UB		VG-3B		VG-3UB	
				0.3	0.5	0.3	0.5	0.3	0.5	0.3	0.5
QRC	VGG	35.62	24.42	13.17	7.64	12.39	7.15	14.21	8.35	13.03	7.52
ZSGNet	VGG	39.32	29.35	17.09	11.02	16.48	10.55	17.63	11.42	17.35	10.97
	Res50	43.02	31.23	19.95	12.90	19.12	12.37	20.77	13.77	19.72	12.82

Table 4. Accuracy across various unseen splits. For Flickr-Split-0,1 we use Accuracy with IoU threshold of 0.5. Since Visual Genome annotations are noisy we additionally report Accuracy with IoU threshold of 0.3. The second row denotes the IoU threshold at which the Accuracy is calculated. “B” and “UB” denote the balanced and unbalanced sets.

	Method	Semantic Distances				
		3-4	4-5	5-6	6-7	7-8
VG 2B	# I-P	310	1050	3543	5321	1985
	QRC (Vgg)	25.16	16.67	15.16	10.96	12.54
	ZSGNet (Vgg)	28.71	21.52	19.02	15.37	14.76
	ZSGNet (Res50)	31.94	25.14	21.99	17.89	17.98
VG 3B	# I-P	974	3199	7740	9873	3765
	QRC (Vgg)	23.1	20.13	14.73	12.19	11.24
	ZSGNet (Vgg)	23.82	25.73	17.16	16	14.56
	ZSGNet (Res50)	29.57	27.85	21.3	18.77	16.71

Table 5. Accuracy of various models on the balanced VG-Splits-2,3 w.r.t the semantic distance of the referred object (A) to the closest object seen at train time. VG-2B and VG-3B refer to the balanced test set for Case2, 3. #I-P denotes the number of image-pharse pairs in the given semantic distance range.

distance of the referred object at test time (A) from the training set. We re-use the clusters obtained while creating the **VG-Split** and consider the closest-seen object which lies in the same cluster as that of the referred object. For every unseen object (o_i) in a particular cluster k ($o_i \in U_k$) we find the closest seen object in the same cluster $o_j = \operatorname{argmin}_{o_j \in S_k} \operatorname{dist}(o_i, o_j)$. For dist calculation, we use the GloVe embeddings [32] corresponding to the objects and take the L2-norm of their difference. We group the distances into 5 intervals of unit length and report the accuracy of the subset where the distance of the referred object from the training set lies in that interval in Table 5.

We note few examples of various intervals. We use the

Model	Accuracy on RefClef
BM + Softmax	48.54
BM + BCE	55.20
BM + FL	57.13
BM + FL + Img-Resize	61.75

Table 6. Ablation study: BM=Base Model, softmax means we classify only one candidate box as foreground, BCE = Binary Cross Entropy means we classify each candidate box as the foreground or background, FL = Focal Loss, Img-Resize: use images of dimension 600×600

notation (A, b) *i.e.* a tuple of the referred object and the closest object in the same cluster. Examples for the interval 3-4 are {(bouquet, flower), (cupcake, cake), (tablecloth, curtain)} and for 7-8 are {(printer, paper), (notebook, pen), (tattoo, poster)}. As expected, the accuracy declines with the semantic distance but smoothly *i.e.* there is no sudden drop in performance.

4.6. Ablation Study

We show the performance of our model with different loss functions using the base model of ZSGNet on the validation set of ReferIt [19] in Table 6. Note that using softmax loss by itself places us higher than the previous methods. Further using Binary Cross Entropy Loss and Focal loss [26] give a significant (7%) performance boost which is expected in a single shot framework. Finally, image resizing gives another 4% increase.

4.7. Visualization

To qualitatively analyze our model we show a few visualizations in Fig 4. The first row shows grounding results on **Flickr30k** (first, second column) and **ReferIt** (third, fourth column). Our model learns about the attribute(s) (“red”), location (“leftmost”) and entities (“cap”, “nightstand”) and



Figure 4. Few grounding visualizations. In all cases, red denotes the ground truth box; green is the box predicted by ZSGNet. Row-1: Flickr30k, ReferIt; Row-2: Flickr-Split-0, 1; Row-3: VG-Split-2,3. In Row-3, the query word Q is emphasized and the closest seen object is provided in parenthesis. The last column shows incorrect predictions.

predicts very tight bounding box. In the last column our model incorrectly predicts only one “microphone”.

The second row shows **Flickr-Split-0** (first, second column) and **Flickr-Split-1** (second, third column) predictions. The query phrases “cigar”, “kayak” are never encountered in the training set though close synonyms like “cigarettes”, “kayakers” are *i.e.* our model generalizes to unseen words even if they haven’t been explicitly seen before. This generalization can be attributed to pre-trained GloVe [32] embedding. On **Flickr-Split-1** the model predicts a good bounding box even though referred objects lies in a different category. However, when there are too many objects in the image the model gets confused (last column).

The third row shows predictions on **VG-Split-2** (first, second column) and **VG-Split-3** (third, fourth column). Additionally, we italicize the query word Q which refers to the object A in the image and mention the closest object encountered during training in parenthesis. In **VG-Split-2** our model effectively utilizes word embedding knowledge and performs best when the closest object seen during training is visually similar to the referred object. In **VG-Split-3** our model additionally needs to disambiguate between a seen object and the referred object and performs well when they are visually distinct like “shorts” and “person”. However, when the two are visually similar as in the case of “planter” and “plants” our model incorrectly grounds the seen object.

5. Conclusion

In this work, we introduce the task of Zero-Shot grounding (ZSG) which aims at localizing novel objects from a query phrase. We outline four cases of zero-shot grounding to perform finer analysis. We address the limitations posed by previous systems and propose a simple yet effective architecture ZSGNet. Finally, we create new datasets by sub-sampling existing datasets to evaluate each of the four grounding cases. We verify that our proposed model ZSGNet performs significantly better than existing baseline in the zero-shot setting.

Acknowledgment: We thank Jiyang Gao for helpful discussions and the anonymous reviewers for helpful suggestions. This paper is based , in part, on research sponsored by the Air Force Research Laboratory and the Defense Advanced Research Projects Agency under agreement number FA8750-16-2-0204. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory and the Defense Advanced Research Projects Agency or the U.S. Government.

References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiaseen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 1
- [2] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. *ECCV*, 2018. 2
- [3] Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. Mutan: Multimodal tucker fusion for visual question answering. In *ICCV*, 2017. 2
- [4] Hedi Ben-younes, Rémi Cadene, Nicolas Thome, and Matthieu Cord. Block: Bilinear superdiagonal fusion for visual question answering and visual relationship detection. *AAAI*, 2019. 2
- [5] Kan Chen, Trung Bui, Chen Fang, Zhaowen Wang, and Ram Nevatia. AMC: Attention guided multi-modal correlation learning for image search. *CVPR*, 2017. 1
- [6] Kan Chen, Jiyang Gao, and Ram Nevatia. Knowledge aided consistency for weakly supervised phrase grounding. In *CVPR*, 2018. 2
- [7] Kan Chen, Rama Kovvuri, and Ram Nevatia. Query-guided regression network with context policy for phrase grounding. In *ICCV*, 2017. 1, 2, 3, 5, 6
- [8] Harm De Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron Courville. Guess-what?! visual object discovery through multi-modal dialogue. In *CVPR*, 2017. 2
- [9] Chaorui Deng, Qi Wu, Qingyao Wu, Fuyuan Hu, Fan Lyu, and Mingkui Tan. Visual grounding via accumulated attention. In *CVPR*, 2018. 1, 2, 3
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [11] Hugo Jair Escalante, Carlos A Hernández, Jesus A Gonzalez, Aurelio López-López, Manuel Montes, Eduardo F Morales, L Enrique Sucar, Luis Villaseñor, and Michael Grubinger. The segmented and annotated iapr tc-12 benchmark. *CVIU*, 2010. 5
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 6
- [13] 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. *EMNLP*, 2016. 2, 6
- [14] Sergio Guadarrama, Erik Rodner, Kate Saenko, Ning Zhang, Ryan Farrell, Jeff Donahue, and Trevor Darrell. Open-vocabulary object retrieval. In *Robotics: science and systems*. Citeseer, 2014. 2
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4, 6
- [17] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997. 4
- [18] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *CVPR*, 2016. 1, 2, 5, 6
- [19] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara L. Berg. Referit game: Referring to objects in photographs of natural scenes. In *EMNLP*, 2014. 2, 3, 5, 6, 7
- [20] Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Hadamard product for low-rank bilinear pooling. *ICLR*, 2017. 2
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 6
- [22] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 2, 5
- [23] Jianan Li, Yunchao Wei, Xiaodan Liang, Fang Zhao, Jianshu Li, Tingfa Xu, and Jiashi Feng. Deep attribute-preserving metric learning for natural language object retrieval. In *ACM Multimedia*, 2017. 6
- [24] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. 2
- [25] Tsung-Yi Lin, Piotr Dollár, Ross B Girshick, Kaiming He, Bharath Hariharan, and Serge J Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 2, 4
- [26] Tsung-Yi Lin, Priyanka Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *TPAMI*, 2018. 2, 3, 4, 6, 7
- [27] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *ICCV*, 2015. 2
- [28] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single shot multibox detector. In *ECCV*, 2016. 2, 3, 6
- [29] Jiaseen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural baby talk. In *CVPR*, 2018. 2
- [30] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *CVPR*, 2016. 2
- [31] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 1995. 5
- [32] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 4, 5, 7, 8
- [33] Bryan A Plummer, Paige Kordas, M Hadi Kiapour, Shuai Zheng, Robinson Piramuthu, and Svetlana Lazebnik. Conditional image-text embedding networks. In *ECCV*, 2018. 1, 2, 3, 6
- [34] 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. 2, 5, 6

- [35] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 2
- [36] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. *CVPR*, 2017. 2
- [37] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 2, 4, 6
- [39] 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
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 6
- [41] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE TSP*, 1997. 4
- [42] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *ECCV*, 2018. 2
- [43] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *TACL*, 2014. 2
- [44] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mattnet: Modular attention network for referring expression comprehension. In *CVPR*, 2018. 1, 2, 3
- [45] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *ECCV*, 2016. 2
- [46] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multimodal factorized bilinear pooling with co-attention learning for visual question answering. In *ICCV*, 2017. 2
- [47] Zhou Yu, Jun Yu, Chenchao Xiang, Jianping Fan, and Dacheng Tao. Beyond bilinear: generalized multimodal factorized high-order pooling for visual question answering. *IEEE TNLS*, 2018. 2
- [48] Zhou Yu, Jun Yu, Chenchao Xiang, Zhou Zhao, Qi Tian, and Dacheng Tao. Rethinking diversified and discriminative proposal generation for visual grounding. *IJCAI*, 2018. 1, 2, 3
- [49] Hanwang Zhang, Yulei Niu, and Shih-Fu Chang. Grounding referring expressions in images by variational context. In *CVPR*, 2018. 2
- [50] Yundong Zhang, Juan Carlos Niebles, and Alvaro Soto. Interpretable visual question answering by visual grounding from attention supervision mining. *WACV*, 2019. 1
- [51] Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, and Anton van den Hengel. Parallel attention: A unified framework for visual object discovery through dialogs and queries. In *CVPR*, 2018. 2