

Grounding Referring Expressions in Images by Variational Context

Hanwang Zhang[†], Yulei Niu[‡], Shih-Fu Chang[†]

[†]Columbia University, [‡]Renmin University of China

hanwangzhang@gmail.com; niu@ruc.edu.cn; shih.fu.chang@columbia.edu

Abstract

We focus on grounding (i.e., localizing or linking) referring expressions in images, e.g., “largest elephant standing behind baby elephant”. This is a general yet challenging vision-language task since it does not only require the localization of objects, but also the multimodal comprehension of context — visual attributes (e.g., “largest”, “baby”) and relationships (e.g., “behind”) that help to distinguish the referent from other objects, especially those of the same category. Due to the exponential complexity involved in modeling the context associated with multiple image regions, existing work oversimplifies this task to pairwise region modeling by multiple instance learning. In this paper, we propose a variational Bayesian method, called Variational Context, to solve the problem of complex context modeling in referring expression grounding. Our model exploits the reciprocal relation between the referent and context, i.e., either of them influences estimation of the posterior distribution of the other, and thereby the search space of context can be greatly reduced, resulting in better localization of referent. We develop a novel cue-specific language-vision embedding network that learns this reciprocity model end-to-end. We also extend the model to unsupervised setting where no annotation for the referent is available. Extensive experiments on various benchmarks show consistent improvement over state-of-the-art methods in both supervised and unsupervised settings.

1. Introduction

Grounding natural language in visual data is a hallmark of AI, since it establishes a communication channel between humans, machines, and the physical world, underpinning a variety of multimodal AI tasks such as robotic navigation [35], visual Q&A [1], and visual chatbot [6]. Thanks to the rapid development in deep learning-based CV and NLP, we have witnessed promising results not only in grounding nouns (e.g., object detection [28]), but also short phrases

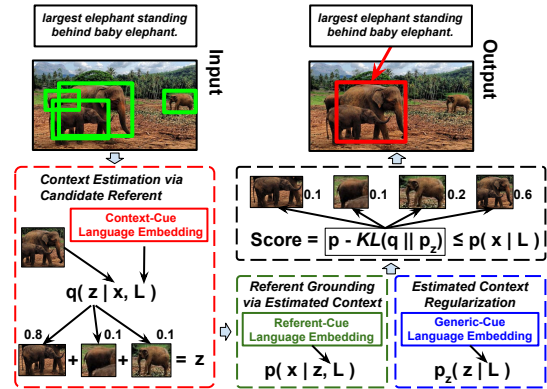


Figure 1. The proposed model of Variational Context. Given an input referring expression and an image with extracted regions, we localize the referent as output. We develop a grounding score function, with the variational lower-bound consisting of three multimodal modules, each of which can be interpreted as a cue-specific grounding task (indicated by the description in the dashed color boxes).

(e.g., noun phrases [26] and relations [42]). However, the more general task: grounding referring expressions [22], is still far from resolved due to the challenges in understanding of both language and scene compositions [10]. As illustrated in Figure 1, given an input referring expression “largest elephant standing behind baby elephant” and an image with region proposals, a model that can only localize “elephant” is not satisfactory as there are multiple elephants. Therefore, the key for referring expression grounding is to comprehend and model the context. Here, we refer to *context* as the visual objects (e.g., “elephant”), attributes (e.g., “largest” and “baby”), and relationships (e.g., “behind”) mentioned in the expression that help to distinguish the referent from other objects.

One straightforward way of modeling the relations between the referent and context is to: 1) use external syntactic parsers to parse the expression into entities, modifiers, and relations [32], and then 2) apply visual relation detectors to localize them [42]. However, this two-stage approach is not practical due to the limited generalization ability of the detectors applied in the highly unrestricted

language and scene compositions [19]. To this end, recent approaches use multimodal embedding networks that jointly comprehend language and model the visual relations [24, 11, 30]. Due to the high cost of annotating both referent and context of referring expressions in images [43], multiple instance learning (MIL) [7] is usually adopted to handle the weak supervision of the unannotated context objects, by maximizing the joint likelihood of every region pair. However, for a referent, the MIL framework essentially oversimplifies the number of context configurations of N regions from $\mathcal{O}(2^N)$ to $\mathcal{O}(N)$. For example, to localize the “elephant” in Figure 1, we may need to consider the other three elephants all together as a multinomial subset for modeling the context such as “largest”, “behind” and “baby elephant”.

In this paper, we propose a novel model called *Variational Context* for grounding referring expressions in images. Compared to the previous MIL-based approaches [24, 11], our model approximates the combinatorial context configurations with weak supervision using a variational Bayesian framework [15]. Intuitively, it exploits the reciprocity between referent and context, given either of which can help to localize the other. As shown in Figure 1, for each region x , we first estimate a coarse context z , which will help refine the true localizations of the referent. This reciprocity is formulated into the variational lower-bound of the grounding likelihood $p(x|L)$, where L is the text expression and the context is considered as a hidden variable z (cf. Section 3). Specifically, the model consists of three multimodal modules: context posterior $q(z|x, L)$, referent posterior $p(x|z, L)$, and context prior $p_z(z|L)$, each of which performs a grounding task (cf. Section 4.3) that aligns image regions with a cue-specific language feature; each cue dynamically encodes different subsets of words in the expression L that help accomplish localization of the corresponding cues (cf. Section 4.2).

Thanks to the reciprocity between referent and context, our model can not only be used in the conventional supervised setting, where there is annotation for referent, but also in the challenging unsupervised setting, where there is no instance-level annotation (e.g., bounding boxes) of both referent and context. We perform extensive experiments on four benchmark referring expression datasets: RefCLEF [14], RefCOCO [40], RefCOCO+ [40], and RefCOCOg [22]. Our results can consistently outperform previous methods in both supervised and unsupervised settings. We also qualitatively show that our model can ground the context in the expression to the corresponding image regions (cf. Section 5).

2. Related Work

Grounding Referring Expressions. Grounding referring expression is also known as referring expression com-

prehension, whose inverse task is called referring expression generation [22]. Different from grounding phrases [27, 26] and descriptive sentences [12, 30], the key for grounding referring expression is to use the context (or pragmatics in linguistics [34]) to distinguish the referent from other objects, usually of the same category [10]. However, most previous works resort to use holistic context such as the entire image [22, 12, 30] or visual feature difference between regions [40, 41]. Our model is similar to works on explicitly modeling the referent and context region pairs [11, 24], however, due to the lack of context annotation, they reduce the grounding task into a multiple instance learning framework [7]. As we will discuss later, this framework is not a proper approximation to the original task. There are also studies on visual relation detection that detect objects and their relationships [19, 5, 42, 43]. However, they are limited to a fixed-vocabulary set of relation triplets and hence are difficult to be applied in natural language grounding. Our cue-specific language feature is similar to the language modular network [11] that learns to decompose a sentence into referent/context-related words, which are different from other approaches that treat the expression as a whole [22, 21, 41, 17].

Variational Bayesian Model vs. Multiple Instance Learning. Our proposed variational context model is in a similar vein of the deep neural network based variational autoencoder (VAE) [15], which uses neural networks to approximate the posterior distribution of the hidden value $q(z|x)$, i.e., encoder, and the conditional distribution of the observation $p(x|z)$, i.e., decoder. VAE shows efficient and effective end-to-end optimization for the *intractable* log-sum likelihood $\log \sum_z p(x, z)$ that is widely used in generative processes such as image synthesis [39] and video dynamics [38]. Considering the unannotated context as the hidden variable z , the referring expression grounding task can also be formulated into the above log-sum marginalization (cf. Eq. (2)). The MIL framework [7] is essentially a sum-log approximation of the log-sum, i.e., $\sum_z \log p(x, z)$. To see this, the max-pooling function $\log \max_z p(x, z)$ used in [11] can be viewed as the sum-log $\sum_z \log p(x|z)p(z)$, where $p(z) = 1$ if z is the correct context and 0 otherwise, indicating there is only one positive instance; maximizing the noisy-or function $\log(1 - \prod_z (1 - p(x, z)))$ used in [24] is equivalent to maximize $\sum_z \log p(x, z)$, assuming there is at least one positive instance. However, due to the numerical property of the log function, this sum-log approximation will unnecessarily force every (x, z) pair to explain the data [8]. Instead, we use the variational Bayesian upper-bound to obtain a better sum-log approximation. Note that visual attention models [2, 37] simplify the variational lower bound by assuming $p(z) = q(z|x)$; however, we explicitly use the KL divergence $KL(q(z|x)||p(z))$ in the lower bound to regularize the approximate posterior $q(z|x)$

not being too far from the prior $p(z)$.

3. Variational Context

In this section, we derive the variational Bayesian formulation of the proposed variational context model and the objective function for training and test.

3.1. Problem Formulation

The task of grounding a referring expression L in an image I , represented by a set of regions $x \in \mathcal{X}$, can be viewed as a region retrieval task with the natural language query L . Formally, we maximize the log-likelihood of the conditional distribution to localize the referent region $x^* \in \mathcal{X}$:

$$x^* = \arg \max_{x \in \mathcal{X}} \log p(x|L), \quad (1)$$

where we omit the image I in $p(x|I, L)$.

As there is usually no annotation for the context, we consider it as a hidden variable z . Therefore, Eq. (1) can be rewritten as the following maximization of the log-likelihood of the conditional marginal distribution:

$$x^* = \arg \max_{x \in \mathcal{X}} \log \sum_z p(x, z|L). \quad (2)$$

Generally, z is NOT necessary to be one region as assumed in recent MIL approaches [11, 24], *i.e.*, $z \in \mathcal{X}$. For example, the contextual objects ‘‘surrounding elephants’’ in ‘‘a bigger elephant than the surrounding elephants’’ should be composed by a multinomial subset of \mathcal{X} , resulting in an extremely large sample space that requires $\mathcal{O}(2^{|\mathcal{X}|})$ search complexity. Therefore, the marginalization in Eq (2) is intractable.

To this end, we use the variational lower-bound [15] to approximate the marginal distribution in Eq. (2) as:

$$\begin{aligned} \log p(x|L) &= \log \sum_z p(x, z|L) \geq \mathcal{Q}(x, L) = \\ &\underbrace{\mathbb{E}_{z \sim q_\phi(z|x, L)} \log p_\theta(x|z, L)}_{\text{Localization}} - \underbrace{KL(q_\phi(z|x, L) || p_\omega(z|L))}_{\text{Regularization}}, \end{aligned} \quad (3)$$

where $KL(\cdot)$ is the Kullback-Leibler divergence, ϕ , θ , and ω are independent parameter sets for the respective distributions. As shown in Figure 1, the lower bound $\mathcal{Q}(x, L)$ offers a new perspective for exploiting the reciprocal nature of referent and context in referring expression grounding:

Localization. This term calculates the localization score for x given an estimated context z , using the referent-cue of L parameterized by θ . In particular, we design a new posterior $q_\phi(z|x, L)$ that approximates the true context prior $p(z|x, L)$, which models the context z using the context-cue of L parameterized by ϕ . In the view of variational auto-encoder [15, 33], this term works in an encoding-decoding

fashion: q_ϕ is the encoder from x to z , and p_θ is the decoder from z to x .

Regularization. As KL is non-negative, maximizing $\mathcal{Q}(x, L)$ would encourage that the posterior q_ϕ is similar to the prior p_ω , *i.e.*, the estimated context z sampled from $q_\phi(z|x, L)$ should not be too far from the referring expression, which is modeled by $p_\omega(z|L)$ with the generic-cue of L parameterized by ω . This term is necessary as the estimated z could be overfitted to region features that are inconsistent with the visual context described in the expression.

3.2. Training and Test

The lower-bound $\mathcal{Q}(x, L)$ transforms the intractable log-sum in Eq. (2) into the efficient sum-log in Eq. (3), where SGD is applicable. One approach is to sample $z = \sum_i \alpha_i x_i$, where α_i is the parameter of the Bernoulli distribution indicating whether x_i is the context. However, the resultant Monte Carlo gradient estimator of $\mathcal{Q}(x, L)$ requires complicated variance reduction techniques [2, 23] and may lead to unstable training¹. Instead, we adopt a deterministic approximation to obtain z :

$$z = \sum_{x' \in \mathcal{X}} x' \cdot q_\phi(x'|x, L), \quad (4)$$

where q_ϕ is considered as a soft attention function over \mathcal{X} [3, 37], and transforms Eq. (3) into a function of only x and L , which can be rewritten as:

$$\mathcal{Q}(x, L) = \log p_\theta(x|z, L) - \log q_\phi(z|x, L) + \log p_\omega(z|L), \quad (5)$$

which is smooth and can be maximized using standard back propagation.

In supervised setting where the ground truth of the referent is known, to distinguish the referent from other objects, we need to train a model that outputs a high $p(x|L)$ (*i.e.*, $\mathcal{Q}(x, L)$), while maintaining a low $p(x'|L)$ (*i.e.*, $\mathcal{Q}(x', L)$), whenever $x' \neq x$. Therefore, we use the so-called Maximum Mutual Information loss as in [22] $-\log\{\mathcal{Q}(x, L)/\sum_{x'} \mathcal{Q}(x', L)\}$, where we do not need to explicitly model the distributions with normalizations; we use the following score function:

$$\mathcal{Q}(x, L) \propto \mathcal{S}(x, L) = s_\theta(x, L) - s_\phi(x, L) + s_\omega(x, L), \quad (6)$$

where s_θ , s_ϕ , and s_ω are the score functions (*e.g.*, $p_\theta \propto s_\theta$) for p_θ , q_ϕ , and p_ω , respectively. These functions will be detailed in Section 4.3. In this way, maximizing Eq. (5) is equivalent to minimizing the following softmax loss:

$$\mathcal{L}_s = -\log \text{softmax } \mathcal{S}(x_{gt}, L), \quad (7)$$

¹This is mainly due to the varying number of regions in an image.

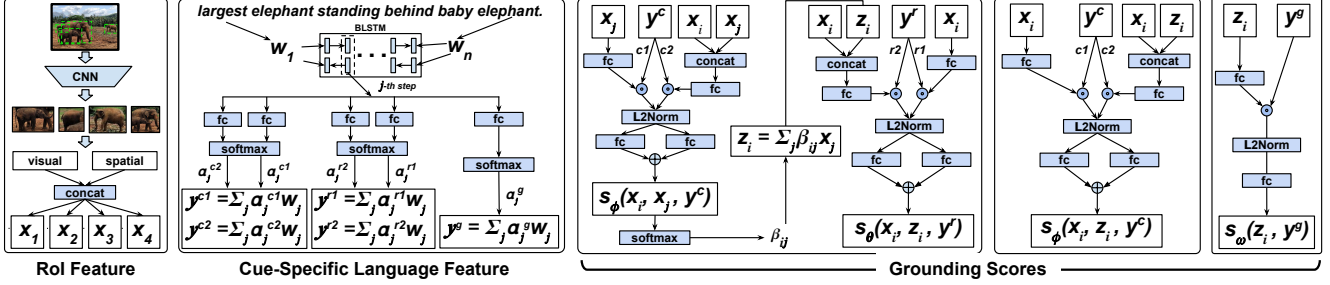


Figure 2. The architecture of the proposed Variational Context model. It consists of a region feature extraction module (Section 4.1, and a language feature extraction module (Section 4.2), and three grounding modules (Section 4.3). It can be trained in an end-to-end fashion with the input of a set of image regions and a referring expression, using the supervised loss (Eq. (7)) or the unsupervised loss (Eq. (8)). fc: fully-connected layer. concat: vector concatenation. L2Norm: L2 normalization layer. \odot : element-wise vector multiplication. \oplus : add.

where the softmax is over $x \in \mathcal{X}$ and x_{gt} is the ground truth referent region.

Note that the reciprocity between referent and context can be extended to unsupervised learning, where neither of the referent and context has annotation. In this setting, we adopt the max-pooled MIL loss functions for unsupervised referring expression grounding:

$$\mathcal{L}_u = - \max_{x \in \mathcal{X}} \log \text{softmax} \mathcal{S}(x, L), \quad (8)$$

where the softmax is over $x \in \mathcal{X}$. Note that the max-pooled MIL function is reasonable since there is only one ground truth referent given an expression and image training pair.

At test stage, in both supervised and unsupervised settings, we predict the referent region x^* by selecting the region $x \in \mathcal{X}$ with the highest score:

$$x^* = \arg \max_{x \in \mathcal{X}} \mathcal{S}(x, L), \quad (9)$$

4. Model Architecture

The overall architecture of the proposed variational context model is illustrated in Figure 2. Thanks to the deterministic approximation in Eq. (4), the five modules in our model can be integrated into an end-to-end fashion. Next, we will detail the implementation of each module.

4.1. RoI Features

Given an image with a set of Region of Interests (RoIs) \mathcal{X} , obtained by any off-the-shelf proposal generator [44] or object detectors [18], this module extracts the feature vector \mathbf{x}_i for every RoI. In particular, \mathbf{x}_i is the concatenation of visual feature \mathbf{v}_i and spatial feature \mathbf{p}_i . For \mathbf{v}_i , we can use the output of a pre-trained convolutional network (cf. Section 5). If the object category of each RoI is available, we can further utilize the comparison between the referent and other objects to capture the visual difference such as ‘the largest/baby elephant’ and ‘the blue/red shirt’. Specifically, we append the visual difference feature [40] $\delta \mathbf{v}_i = \frac{1}{n} \sum_{j \neq i} \frac{\mathbf{v}_i - \mathbf{v}_j}{\|\mathbf{v}_i - \mathbf{v}_j\|}$ to the original \mathbf{v}_i visual

feature, where n is the number of objects chosen for comparison (e.g., the number of RoI in the same object category). For spatial feature, we use the 5-d spatial attributes $\mathbf{p}_i = [\frac{x_{tl}}{W}, \frac{y_{tl}}{H}, \frac{x_{br}}{W}, \frac{y_{br}}{H}, \frac{w \cdot h}{W \cdot H}]$, where x and y are the coordinates the top left (tl) and bottom right (br) RoI of the size $w \times h$, and the image is of the size $W \times H$.

4.2. Cue-Specific Language Features

The cue-specific language feature representation for a referring expression is inspired by the attention weighted sum of word vectors [11, 20, 3], where the weights are parameterized by context-cue ϕ , referent-cue θ , and generic-cue ω . The context-cue language feature $\mathbf{y}^c = [\mathbf{y}^{c1}, \mathbf{y}^{c2}]$ is a concatenation of \mathbf{y}^{c1} for language-vision association between *single* RoI and the expression, and \mathbf{y}^{c2} for the association between *pairwise* RoIs; the referent-cue language feature \mathbf{y}^r can be represented in a similar way to \mathbf{y}^c ; the generic-cue language feature \mathbf{y}^g is only for single RoI association as it is an independent prior. The weights of each cue are calculated from the hidden state vectors of a 2-layer bi-directional LSTM (BLSTM) [31], scanning through the expression. The hidden states encode forward and backward compositional semantic meanings of the sentences, beneficial for selecting words that are useful for single and pairwise associations. Specifically, suppose \mathbf{h}_j as the 4,000-d concatenation of forward and backward hidden vectors of the j -th word, without loss of generality, the word attention weight α_j and the language feature \mathbf{y} for single/pairwise association of any cue can be calculated as:

$$\mathbf{m}_j = \text{fc}(\mathbf{h}_j), \alpha_j = \text{softmax}_j(\mathbf{m}_j), \mathbf{y} = \sum_j \alpha_j \mathbf{w}_j, \quad (10)$$

where \mathbf{w}_j is a 300-d vector. Note that the BLSTM module can be jointly trained with the entire model.

Figure 3 shows that the cue-specific language features dynamically weight words in different expressions. We can have two interesting observations. First, $c1$ almost weights every word equally, while $c2$ is highly informative; although $r2$ is more informative than $c1$, it is still less informative




Figure 3. Two qualitative examples of the cue-specific language feature word weights. Darker color indicates higher weights. $c/r+1/2$: context/referent-cue + single/pairwise.

than $r1$. This demonstrates the reciprocity between referent and context, that is, at the stage of context estimation $q_\phi(z|x, L)$ from scratch, context should be defined by the pairwise relationships (e.g., “left”) with other objects (e.g., “frisbee”); at the stage of referent localization $p_\theta(x|z, L)$, given the estimated context, the referent can be easily found by single descriptions (e.g., “dog lying” and “black white dog”). Second, g is adaptive to the number of object categories in the expression, *i.e.*, if the context object is of the same category as the referent, g weights descriptive or relationship words higher (e.g., “lying, standing, left”), and nouns higher (e.g., “frisbee”), otherwise.

4.3. Score Functions

For any image and expression pair, given the RoI feature \mathbf{x}_i , and the cue-specific language feature \mathbf{y}^c , \mathbf{y}^r , and \mathbf{y}^g , we implement the final grounding score in Eq. (6) as follows.

Context Estimation Score: $s_\phi(\mathbf{x}_i, \mathbf{x}_j, \mathbf{y}^c)$. It is a score function for modeling the context posterior $q_\phi(z|x, L)$, *i.e.*, given an RoI \mathbf{x}_i as the candidate referent, we calculate the likelihood of any RoI \mathbf{x}_j to be the context. We can also use this function to estimate the final context posterior score $s_\phi(\mathbf{x}_i, \mathbf{z}_i, \mathbf{y}^c)$. Specifically, the context estimation score is a sum of the single and pairwise vision-language association scores: \mathbf{x}_j and \mathbf{y}^{c1} , $[\mathbf{x}_i, \mathbf{x}_j]$ and \mathbf{y}^{c2} . Each associate score is an fc output from the input of a normalized multimodal feature:

$$\begin{aligned} \mathbf{m}_i^1 &= \mathbf{y}^{c1} \odot \text{fc}(\mathbf{x}_j), \quad \mathbf{m}_j^2 = \mathbf{y}^{c2} \odot \text{fc}([\mathbf{x}_i, \mathbf{x}_j]), \\ \tilde{\mathbf{m}}_j^1 &= \text{L2Norm}(\mathbf{m}_j^1), \quad \tilde{\mathbf{m}}_j^2 = \text{L2Norm}(\mathbf{m}_j^2), \\ s_\phi(\mathbf{x}_i, \mathbf{x}_j, \mathbf{y}^c) &= \text{fc}(\tilde{\mathbf{m}}_j^1) + \text{fc}(\tilde{\mathbf{m}}_j^2), \end{aligned} \quad (11)$$

where the element-wise multiplication \odot is an effective way for multimodal features [2]. According to Eq. (4), we can obtain the estimated context z as $\mathbf{z}_i = \sum_j \beta_j \mathbf{x}_j$, where $\beta_j = \text{softmax}_j(s_\phi(\mathbf{x}_i, \mathbf{x}_j, \mathbf{y}^c))$

Referent Grounding Score: $s_\theta(\mathbf{x}_i, \mathbf{z}_i, \mathbf{y}^r)$. After obtaining the context feature \mathbf{z}_i , we can use this score function to calculate how likely a candidate RoI \mathbf{x}_i is the referent given the context \mathbf{z}_i . This function follows a similar way to Eq. (11).

Context Regularization Score: $s_\omega(\mathbf{z}_i, \mathbf{y}^g)$. This function scores how likely the estimated context feature \mathbf{z}_i is consistent with the content mentioned in the expression.

Compared to the above scores, this score is only dependent on single RoI and language association:

$$\mathbf{m}_i = \mathbf{y}_i^g \odot \text{fc}(\mathbf{z}_i), \quad \tilde{\mathbf{m}}_i = \text{L2Norm}(\mathbf{m}_i), \quad s_\omega(\mathbf{z}_i, \mathbf{y}_i^g) = \text{fc}(\tilde{\mathbf{m}}_i). \quad (12)$$

Note that the final regularization score is $s_\omega(\mathbf{z}_i, \mathbf{y}^g) - s_\phi(\mathbf{x}_i, \mathbf{z}_i, \mathbf{y}^c)$ as in Eq. (6).

5. Experiment

5.1. Datasets

We used four popular benchmarks for the referring expression grounding task.

RefCOCO [40]. It has 142,210 referring expressions for 50,000 referents (e.g., object instances) in 19,994 images from MSCOCO [16]. The expressions are collected in an interactive way [14]. The dataset is split into train, validation, Test A, and Test B, which has 120,624, 10,834, 5,657 and 5,095 expression-referent pairs, respectively. An image contains multiple people in Test A and multiple objects in Test B.

RefCOCO+ [40]. It has 141,564 expressions for 49,856 referents in 19,992 images from MSCOCO. The difference from RefCOCO is that it only allows appearances but no locations to describe the referents. The split is 120,191, 10,758, 5,726 and 4,889 expression-referent pairs for train, validation, Test A, and Test B respectively.

RefCOCOg [22]. It has 95,010 referring expressions for 49,822 objects in 25,799 images from MSCOCO. Different from RefCOCO and RefCOCO+, this dataset not collected in an interactive way and contains longer sentences containing both appearance and location expressions. The split is 85,474 and 9,536 expression-referent pairs for training and validation. Note that there is no open test split for RefCOCOg, so we used the hyper-parameters cross-validated on RefCOCO and RefCOCO+.

RefCLEF [14]. It contains 20,000 images with annotated image regions. It has some ambiguous (e.g. anywhere) phrases and mistakenly annotated image regions that are not described in the expressions. For fair comparison, we used the split released by [12, 30], *i.e.*, 58,838, 6,333 and 65,193 expression-referent pairs for training, validation and test, respectively.

5.2. Settings and Metrics

We used an English vocabulary of 72,704 words contained in the GloVe pre-trained word vectors [25], which was also used for the initialization of our word vectors. We used a “unk” symbol for the input word of the BLSTM if the word is out of the vocabulary; we set the sentence length to 20 and used “pad” symbol to pad expression sentence < 20 . For RoI visual features on RefCOCO, RefCOCO+, and RefCOCOg which have MSCOCO annotated regions with object categories, we used the concatenation of the 4,096-d

Table 1. Supervised grounding performances (Acc%) of comparing methods on RefCOCO, RefCOCO+, and RefCOCOg. Note that [41] reports slightly higher accuracies using ensemble models of Listener and Speaker. For fair comparison, we only report their single models.

| Dataset | Split | State-of-The-Arts | | | | | | Our Baselines | | |
|---------------|--------|-------------------|-------------|-----------|----------|--------------|---------------|---------------|-----------------|--------------|
| | | MMI [22] | NegBag [24] | Attr [17] | CMN [11] | Speaker [41] | Listener [41] | VC w/o reg | VC w/o α | VC |
| RefCOCO | Test A | 71.72 | 75.6 | 78.85 | 75.94 | 78.95 | 78.45 | 75.59 | 74.03 | 78.98 |
| | Test B | 71.09 | 78.0 | 78.07 | 79.57 | 80.22 | 80.10 | 79.69 | 78.27 | 82.39 |
| RefCOCO+ | Test A | 58.42 | — | 61.47 | 59.29 | 64.60 | 63.34 | 60.76 | 57.61 | 62.56 |
| | Test B | 51.23 | — | 57.22 | 59.34 | 59.62 | 58.91 | 60.14 | 54.37 | 62.90 |
| RefCOCOg | Val | 62.14 | 68.4 | 69.83 | 69.30 | 72.63 | 72.25 | 71.05 | 65.13 | 73.98 |
| RefCOCO(det) | Test A | 64.90 | 58.6 | 72.08 | 71.03 | 72.95 | 72.95 | 70.78 | 70.73 | 73.33 |
| | Test B | 54.51 | 56.4 | 57.29 | 65.77 | 63.43 | 62.98 | 65.10 | 64.63 | 67.44 |
| RefCOCO+(det) | Test A | 54.03 | — | 57.97 | 54.32 | 60.43 | 59.61 | 56.82 | 53.33 | 58.40 |
| | Test B | 42.81 | — | 46.20 | 47.76 | 48.74 | 48.44 | 51.30 | 46.88 | 53.18 |
| RefCOCOg(det) | Val | 45.85 | 39.5 | 52.35 | 57.47 | 59.51 | 58.32 | 60.95 | 55.72 | 62.30 |

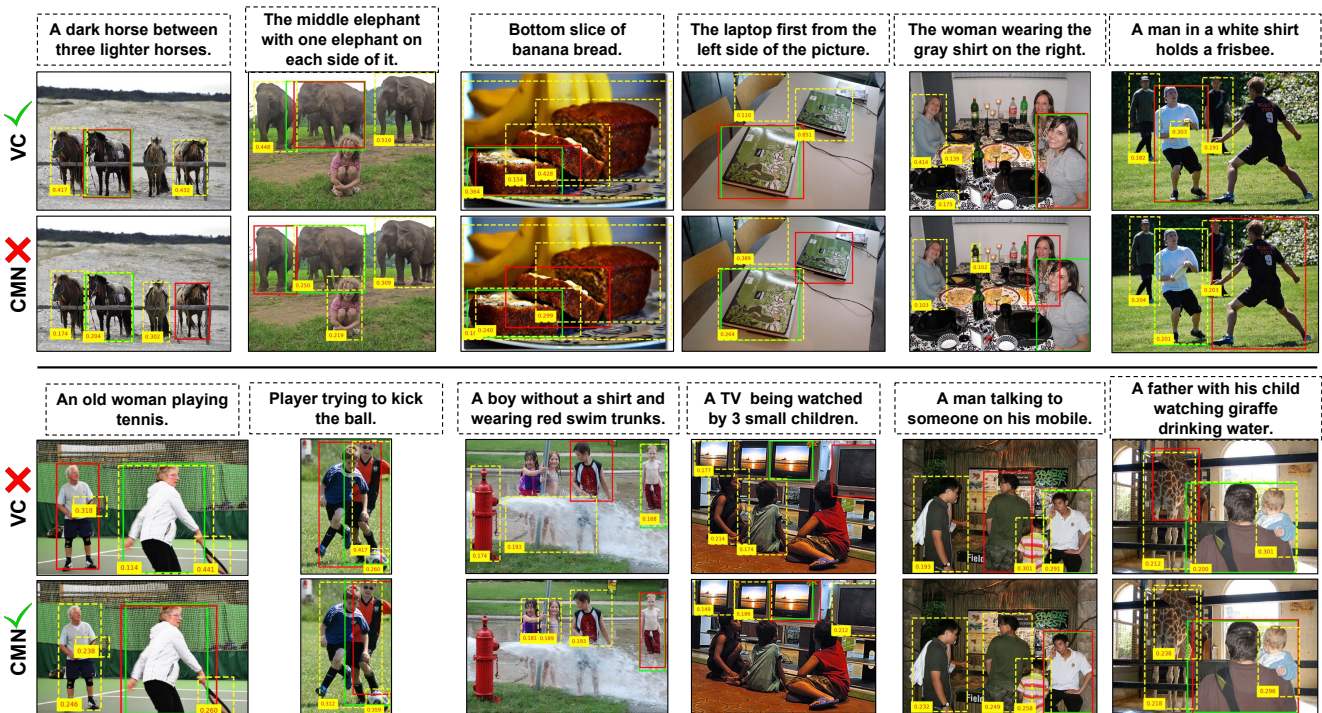


Figure 4. Qualitative results on RefCOCOg (det) showing comparisons between correct (green tick) and wrong referent grounds (red cross) by VC and CMN. The denotations of the bounding box colors are as follows. Solid red: referent ground; solid green: ground truth; dashed yellow: context ground. We only display top 3 context objects with the context ground probability > 0.1 . We can observe that VC has more reasonable context localizations than CMN, even in cases when the referent ground of VC fails.

fc7 output of a VGG-16 based Faster-RCNN network [29] trained on MSCOCO and its corresponding 4,096-d visdiff feature [40]; although RefCLEF regions also have object categories, for fair comparison with [30], we did not use the visdiff feature.

The model training is single-image based, with all referring expressions annotated. We applied SGD of 0.95-momentum with initial learning rate of 0.01, multiplied by 0.1 after every 120,000 iterations, up to 160,000 iterations. Parameters in BILSTM and fc-layers were initialized by

Xavier [9] with 0.0005 weight decay. Other settings were default in TensorFlow. Note that our model is trained without bells and whistles, therefore, other optimization tricks such as batch normalization [13] and GRU [4] are expected to further improve the results reported here. Besides the ground truth annotations, grounding to automatically detected objects is a more practical setting. Therefore, we also evaluated with the SSD-detected bounding boxes [18] on the four datasets provided by [41]. A grounding is considered as correct if the intersection-over-union (IoU) of the

top-1 scored region and the ground-truth object is larger than 0.5. The grounding accuracy (a.k.a, P@1) is the fraction of correctly grounded test expressions.

5.3. Evaluations of Supervised Grounding

We compared our variational context model (VC) with state-of-the-art referring expression methods published in recent years, which can be categorized into: 1) generation-comprehension based such as MMI [22], Attr [17], Speaker [41], Listener [41], and SCRC [12]; 2) localization based such as GroundR [30], NegBag [24], CMN [11]. Note that NegBag and CMN are MIL-based models. In particular, we used the author-released code to obtain the results of CMN on RefCLEF, RefCOCO, and RefCOCO+.

From the results on RefCOCO, RefCOCO+, and RefCOCOg in Table 1 and that on RefCLEF in Table 2, we can see that VC achieves the state-of-the-art performance. We believe that the improvement is attributed to the variational Bayesian modeling of context. First, on all datasets, except for the most recent reinforcement learning based [41], VC outperforms all the other sentence generation-comprehension methods that do not model context. Second, compared to VC without the regularization term in Eq. (3) (VC w/o reg), VC can boost the performance by around 2% on all datasets. This demonstrates the effectiveness of the KL divergence for the prevention of the overfitted context estimation.

In particular, we further demonstrate the superiority of VC over the most recent MIL-based method CMN. As illustrated in Figure 4, VC has better context comprehension in both of the language and image regions than CMN. For example, in the top two rows where VC is correct and CMN is wrong, for the grounding in the second column, CMN unnecessarily considers the “girl” as context but the expression only describes using “elephant”; in the last column, CMN misses the key context “frisbee”. Even in the failure cases where VC is wrong and CMN is correct, VC still localizes reasonable context. For example, in the fourth column, although CMN grounds the correct TV, but it is based on incorrect context of other TVs; while VC can predict the correct context “children”. In addition, we observed that most of the cases that CMN is better than VC involves multiple humans. This demonstrates that VC is better at grounding objects of different categories.

VC is also effective in images with more objects. Figure 4 shows the performances of VC and CMN with various number of bounding boxes. We can observe that VC considerably outperforms CMN over all bounding boxes numbers. Recall that context is the key to distinguish objects of the same category. In particular, on the Test A sets of RefCOCO and RefCOCO+ where the grounding is only about people, *i.e.*, the same object category, the gap between VC and CMN is becoming larger as the box number increases.

Table 2. Performances (Acc%) of supervised and unsupervised methods on RefCLEF.

| | Sup. | Sup. (det) | Unsup. (det) |
|-----------------|--------------|--------------|--------------|
| SCRC [12] | 72.74 | 17.93 | — |
| GroundR [30] | — | 26.93 | 10.70 |
| CMN [11] | 81.52 | 28.33 | — |
| VC | 82.43 | 31.13 | 14.11 |
| VC w/o α | 79.60 | 27.40 | 14.50 |

Table 3. Unsupervised grounding performances (Acc%) of comparing methods on RefCOCO, RefCOCO+, and RefCOCOg.

| Dataset | Split | VC w/o reg | VC | VC w/o α |
|---------------|--------|------------|--------------|-----------------|
| RefCOCO | Test A | 13.59 | 17.34 | 33.29 |
| | Test B | 21.65 | 20.98 | 30.13 |
| RefCOCO+ | Test A | 18.79 | 23.24 | 34.60 |
| | Test B | 24.14 | 24.91 | 31.58 |
| RefCOCOg | Val | 25.14 | 33.79 | 30.26 |
| RefCOCO(det) | Test A | 17.14 | 20.91 | 32.68 |
| | Test B | 22.30 | 21.77 | 27.22 |
| RefCOCO+(det) | Test A | 19.74 | 25.79 | 34.68 |
| | Test B | 24.05 | 25.54 | 28.10 |
| RefCOCOg(det) | Val | 28.14 | 33.66 | 29.65 |

This demonstrates that MIL is ineffective in modeling context, especially when the number of image regions is large.

5.4. Evaluations of Unsupervised Grounding

We follow the unsupervised setting in GroundR [30]. To our best knowledge, it is the only work on unsupervised referring expression grounding. Note that it is also known as “weakly supervised” detection [43] as there is still image-level ground truth (*i.e.*, the referring expression). Table 2 reports the unsupervised results on the RefCLEF. We can see that VC outperforms the state-of-the-art GroundR, which is a generation-comprehension based method. This demonstrates that using context also helps unsupervised grounding. As there is no published unsupervised results on RefCOCO, RefCOCO+, and RefCOCOg, we only compared our baselines on them in Table 3. We can have the following three key observations which highlight the challenges of unsupervised grounding:

Context Prior. VC w/o reg is the baseline without the KL divergence as a context regularization in Eq. (3). We can see that in most of the cases, VC considerably outperforms VC w/o reg by over 2%, even over 5% on RefCOCO+ (det) and RefCOCOg (det). Note that this improvement is significantly higher than that in supervised setting (*e.g.*, < 3% as reported in Table 1). The reason is that the context estimation in Eq. (4) would be easier to be stuck in image regions that are irrelevant to the expression in unsupervised setting,




Figure 5. Performances of VC and CMN with different number of object bounding boxes on RefCOCO Test A & B, RefCOCO+ Test A & B, and RefCOCOg Val. Compared to CMN, we can see that VC is more effective in context modeling when the number of objects is large.

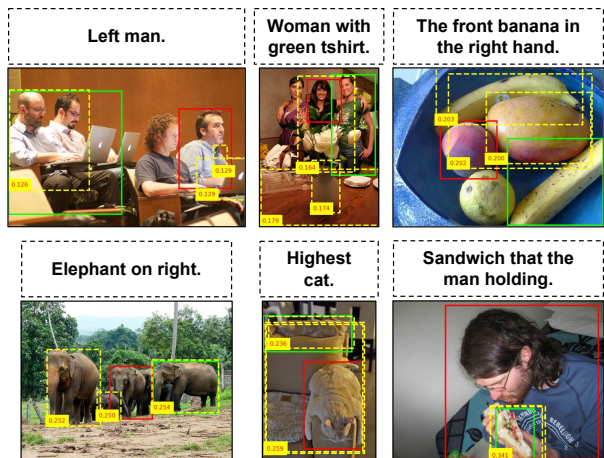


Figure 6. Common failure cases in unsupervised grounding with detected bounding boxes. From left to right: RefCOCO, RefCOCO+, and RefCOCOg. The failure is mainly to the challenging unsupervised relation modeling between referent and context.

therefore, context prior is necessary.

Language Feature. Except on RefCOCOg, we consistently observed the *ineffectiveness* of the cue-specific language feature in unsupervised setting, *i.e.*, VC w/o α outperforms VC in Table 2 and 3. Here α represents the cue-specific word attention. This is contrary to the observation in the supervised setting as listed in Table 1, where VC w/o α is consistently lower than VC. Note that without the cue-specific word attention α in Eq. (10), the language feature is merely the average value of the word embedding vectors in the expression. In this way, VC w/o α does not encode any structural language composition as illustrated in Figure 3, thus, it is better for short expressions. However, when the expression is long in RefCOCOg, discarding the language structure still degrades the performance on RefCOCOg.

Unsupervised Relation Discovery. Although we demonstrated that VC improves the unsupervised grounding by modeling context, we believe that there is still a large space for improving the quality of modeling the context. As the failure examples shown in Figure (6), 1) many context estimations are still out of the scope of the expression, *e.g.*, we may localize the “cup” and “table” as context even though the expression is “woman with green t-shirt”; 2) we may mistake due to the wrong comprehension of the relations, *e.g.*, “right” as “left”, even if the objects belong to the same category, *e.g.*, “elephant”. For further investiga-

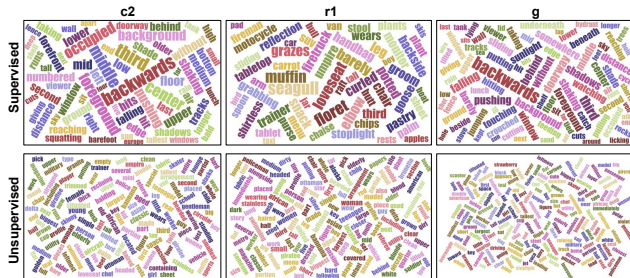


Figure 7. Word cloud visualizations of cue-specific word attention α in Eq. 10 of context-cue (c2), referent-cue (r1), and generic-cue (g) using supervised (top row) and unsupervised training (bottom row) on RefCOCOg. Without supervision, it is difficult to discover meaningful language compositions.

tion, Figure 7 visualizes the cue-specific word attentions in supervised and unsupervised settings. The almost identical word attentions in unsupervised setting reflect the fact that the relation modeling between referent and context is not as successful as in supervised setting. This inspires us to exploit stronger prior knowledge such as language structure [36] and spatial configurations [43].

6. Conclusions

We focused on the task of grounding referring expressions in images and discussed that the key problem is how to model the complex context, which is not effectively resolved by the multiple instance learning framework used in prior works. Towards this challenge we introduced the Variational Context model, where the variational lower-bound can be interpreted by the reciprocity between the referent and context: given any of which can help to localize the other, and hence is expected to significantly reduce the context complexity in a principled way. We implemented the model using cue-specific language-vision embedding network that can be efficiently trained end-to-end. We validated the effectiveness of this reciprocity by promising supervised and unsupervised experiments on four benchmarks. Moving forward, we are going to 1) incorporate expression language generation in the variational framework, 2) use more structural features of language rather than word attentions, and 3) further investigate the potential of our model in the unsupervised referring expression grounding.

7. Supplementary Material

7.1. Derivation of Eq (3) and (5)

Recall that we are going to transform the log-sum objective function in Eq (2) to sum-log for tractable training. Without loss of generality, we omit the conditional L in the derivation. By using the concavity of the log function:

$$\log((1 - \alpha)x + \alpha y) \geq (1 - \alpha) \log(x) + \alpha \log(y), \quad (13)$$

to the following rewriting of Eq (2):

$$\begin{aligned} \log \sum_z p(x, z) &= \log \left(\sum_z \frac{q(z|x)p(x, z)}{q(z|x)} \right) \\ &\geq \sum_z q(z|x) \log \frac{p(x, z)}{q(z|x)} \\ &= \sum_z q(z|x) \log p(x|z)p(z) - \sum_z q(z|x) \log q(z|x) \\ &= \sum_z q(z|x) \log p(x|z) + \sum_z q(z|x) \log p(z) \\ &\quad - \sum_z q(z|x) \log q(z|x) \\ &= \sum_z q(z|x) \log p(x|z) - \sum_z q(z|x) \frac{q(z|x)}{p(z)} \\ &= \mathbb{E}_{z \sim q(z|x)} \log p(x|z) - KL(q(z|x) || p(z)) = \mathcal{Q}(x) \end{aligned} \quad (14)$$

By substituting $z = \sum_{x'} x' \cdot q_\phi(x'|x)$ in Eq (4) into $\sum_z q(z|x) \log p(x|z)p(z) - \sum_z q(z|x) \log q(z|x)$, we have

$$\begin{aligned} \mathcal{Q}(x) &= \log p(x, z) - \log q(z|x) \\ &= \log p(x|z) - \log q(z|x) + \log p(z), \end{aligned} \quad (15)$$

which is Eq (5).

7.2. More Examples on Language Features

Figure 9 shows more cue-specific language features on RefCOCOg.

7.3. More Results on Unsupervised Grounding

7.3.1 External Parsers

As discussed in Section 5.4 that the language feature in unsupervised VC is not as good as that in the supervised setting. An alternative is to use external NLP parsers to obtain the compositions. However, conventional parsers (*e.g.*, Stanford Dependency) are observed to be suboptimal to the visual grounding task [11]. Therefore, we adopt the parser jointly trained on the referring expression grounding task [11]. As illustrated in Figure 8, this parser assigns word-level attention weights of subject, relation, and object.

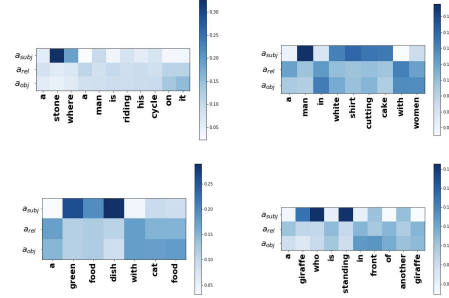


Figure 8. Some parsing examples from [11]. We can see that the composition of subject, object and predicate cannot be clearly parsed.

In particular, we consider the language features of c1, r2 as the average word embeddings, c2 as the relation weights, r1 as the subject weights, g as the object weights. Table 4 shows the performances on unsupervised grounding. We can see that there is no significant improvement of VC w/ parser over VC w/o α .

7.3.2 More Qualitative Results

Figure 10 shows more qualitative results on supervised and unsupervised grounding results on RefCOCO, RefCOCO+, and RefCOCOg.

References

- [1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Zitnick, and D. Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 1
- [2] J. Ba, V. Mnih, and K. Kavukcuoglu. Multiple object recognition with visual attention. In *ICLR*, 2015. 2, 3, 5
- [3] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. 2015. 3, 4
- [4] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014. 6
- [5] B. Dai, Y. Zhang, and D. Lin. Detecting visual relationships with deep relational networks. In *CVPR*, 2017. 2
- [6] A. Das, S. Kottur, J. M. Moura, S. Lee, and D. Batra. Learning cooperative visual dialog agents with deep reinforcement learning. In *ICCV*, 2017. 1
- [7] T. G. Dietterich, R. H. Lathrop, and T. Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 1997. 2
- [8] C. W. Fox and S. J. Roberts. A tutorial on variational bayesian inference. *Artificial intelligence review*, 2012. 2
- [9] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *ICAI*, 2010. 6
- [10] D. Golland, P. Liang, and D. Klein. A game-theoretic approach to generating spatial descriptions. In *EMNLP*, 2010. 1, 2
- [11] R. Hu, M. Rohrbach, J. Andreas, T. Darrell, and K. Saenko. Modeling relationships in referential expressions with compositional modular networks. In *CVPR*, 2017. 2, 3, 4, 6, 7, 9
- [12] R. Hu, H. Xu, M. Rohrbach, J. Feng, K. Saenko, and T. Darrell. Natural language object retrieval. In *CVPR*, 2016. 2, 5, 7

Table 4. Unsupervised grounding performances (%) of VC w/ or w/o parsers on the four datasets.

| GT | RefCLEF | RefCOCO TestA | RefCOCO TestB | RefCOCO+TestA | RefCOCO+TestB | RefCOCOg |
|-----------------|---------|---------------|---------------|----------------|----------------|----------|
| VC w/ parser | 22.57 | 23.00 | 27.50 | 24.69 | 28.96 | 30.73 |
| VC | 21.06 | 17.34 | 20.98 | 23.24 | 24.91 | 33.79 |
| VC w/o α | 20.72 | 33.29 | 30.13 | 34.60 | 31.58 | 30.26 |
| DET | RefCLEF | RefCOCO TestA | RefCOCO TestB | RefCOCO+ TestA | RefCOCO+ TestB | RefCOCOg |
| VC w/ parser | 14.90 | 24.07 | 25.12 | 26.18 | 26.92 | 30.64 |
| VC | 14.11 | 20.91 | 21.77 | 25.79 | 25.54 | 33.66 |
| VC w/o α | 14.50 | 32.68 | 27.22 | 34.68 | 28.10 | 29.65 |

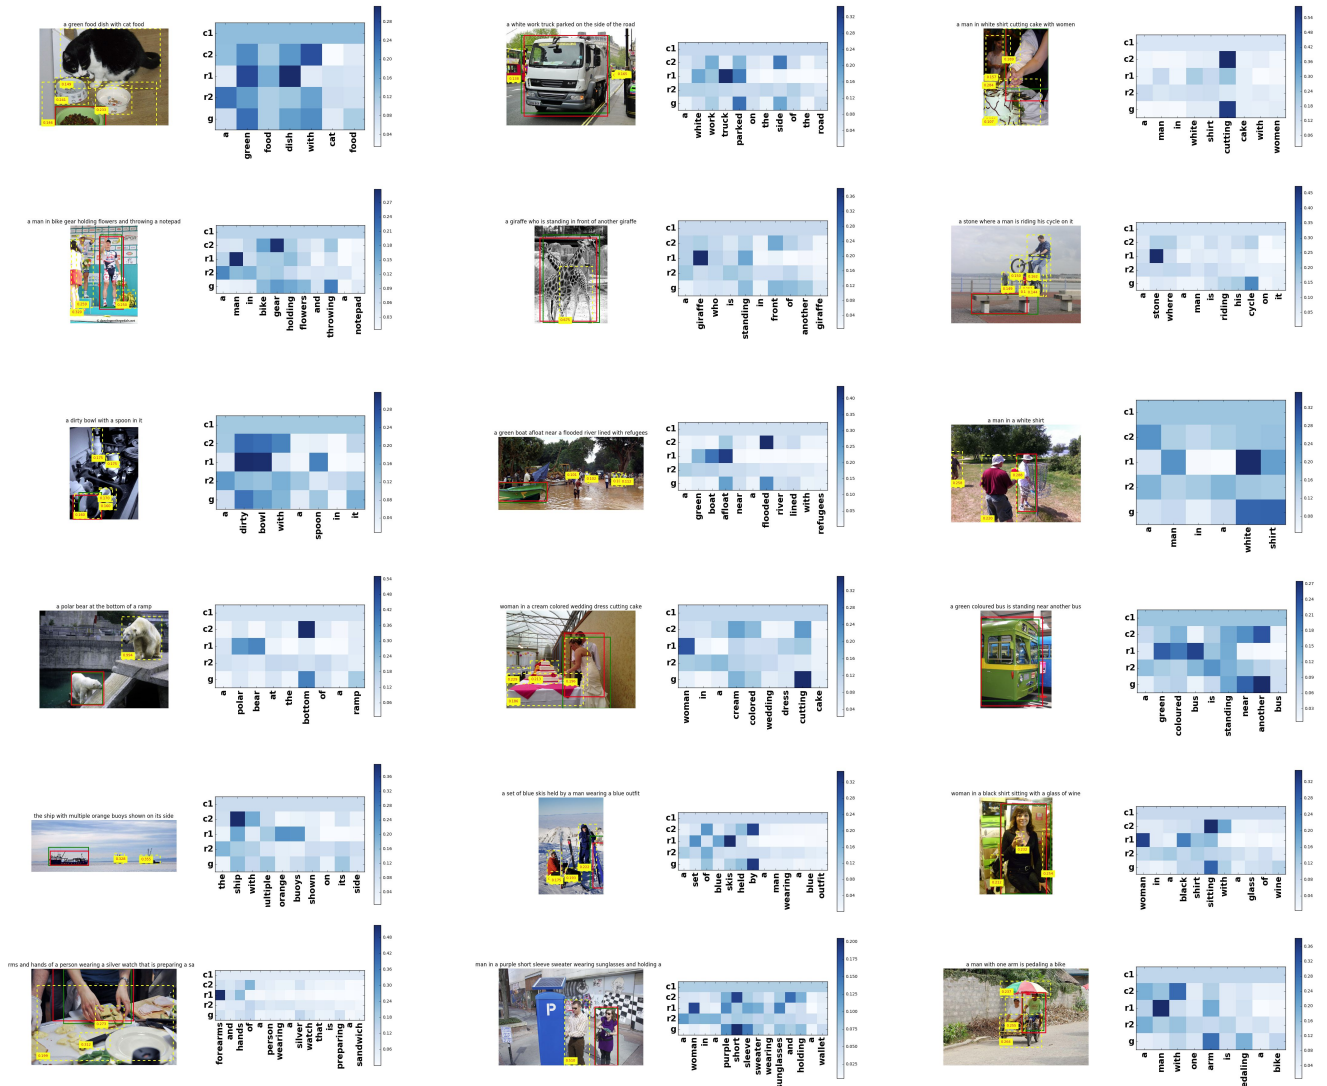


Figure 9. Qualitative results of the cue-specific language features on RefCOCOg.

[13] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

6

[14] S. Kazemzadeh, V. Ordonez, M. Matten, and T. L. Berg. Refer-

itgame: Referring to objects in photographs of natural scenes. In *EMNLP*, 2014. 2, 5

[15] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 2, 3

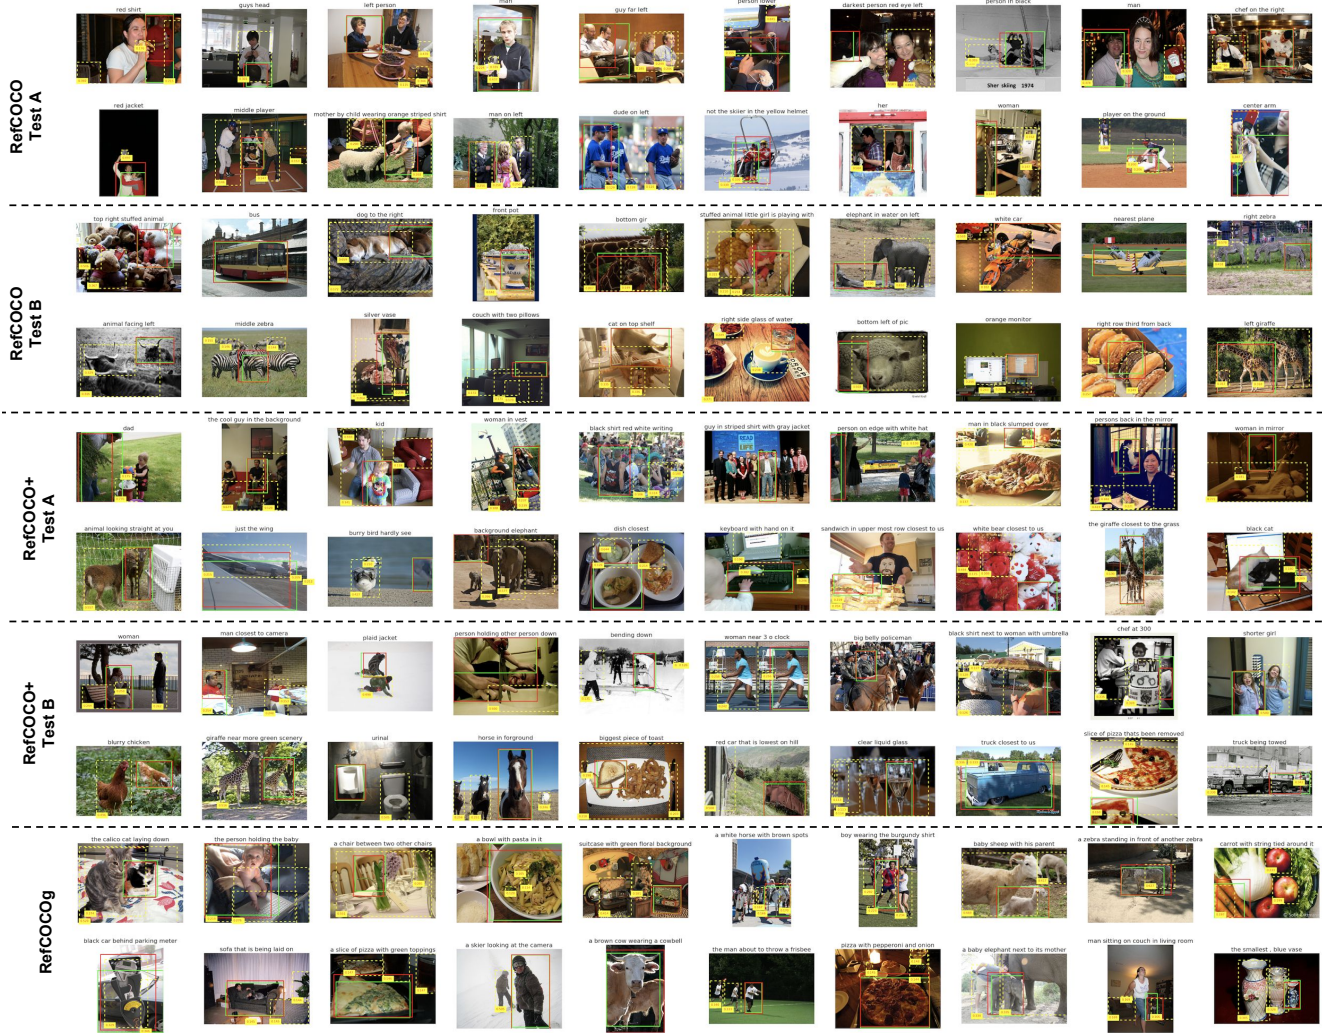


Figure 10. Qualitative results on RefCOCO, RefCOCO, RefCOCO+, and RefCOCOg. All of them are from det bounding boxes. The first line denotes the supervised results and the second line denotes the unsupervised results. Solid red: referent grounding; solid green: ground truth; dashed yellow: context grounding. We only display top 3 context objects with the context grounding probability > 0.1 . Please zoom in.

[16] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5

[17] J. Liu, L. Wang, and M.-H. Yang. Referring expression generation and comprehension via attributes. In *ICCV*, 2017. 2, 6, 7

[18] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 4, 6

[19] C. Lu, R. Krishna, M. Bernstein, and L. Fei-Fei. Visual relationship detection with language priors. In *ECCV*, 2016. 2

[20] J. Lu, J. Yang, D. Batra, and D. Parikh. Hierarchical question-image co-attention for visual question answering. In *NIPS*, 2016. 4

[21] R. Luo and G. Shakhnarovich. Comprehension-guided referring expressions. In *CVPR*, 2017. 2

[22] J. Mao, J. Huang, A. Toshev, O. Camburu, A. L. Yuille, and K. Murphy. Generation and comprehension of unambiguous object descriptions. In *CVPR*, 2016. 1, 2, 3, 5, 6, 7

[23] V. Mnih, N. Heess, A. Graves, et al. Recurrent models of visual attention. In *NIPS*, 2014. 3

[24] V. K. Nagaraja, V. I. Morariu, and L. S. Davis. Modeling context between objects for referring expression understanding. In *ECCV*, 2016. 2, 3, 6, 7

[25] J. Pennington, R. Socher, and C. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 5

[26] B. A. Plummer, A. Mallya, C. M. Cervantes, J. Hockenmaier, and S. Lazebnik. Phrase localization and visual relationship detection with comprehensive linguistic cues. In *ICCV*, 2017. 1, 2

[27] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015. 2

[28] J. Redmon and A. Farhadi. Yolo9000: better, faster, stronger. In *CVPR*, 2017. 1

[29] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 6

- [30] A. Rohrbach, M. Rohrbach, R. Hu, T. Darrell, and B. Schiele. Grounding of textual phrases in images by reconstruction. In *ECCV*, 2016. [2](#), [5](#), [6](#), [7](#)
- [31] M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *TSP*, 1997. [4](#)
- [32] S. Schuster, R. Krishna, A. Chang, L. Fei-Fei, and C. D. Manning. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *Workshop on Vision and Language*, 2015. [1](#)
- [33] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In *NIPS*, 2015. [3](#)
- [34] J. A. Thomas. *Meaning in interaction: An introduction to pragmatics*. Routledge, 2014. [2](#)
- [35] J. Thomason, J. Sinapov, and R. Mooney. Guiding interaction behaviors for multi-modal grounded language learning. In *Proceedings of the First Workshop on Language Grounding for Robotics*, pages 20–24, 2017. [1](#)
- [36] F. Xiao, L. Sigal, and Y.-J. Lee. Weakly-supervised visual grounding of phrases with linguistic structures. In *CVPR*, 2017. [8](#)
- [37] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. [2](#), [3](#)
- [38] T. Xue, J. Wu, K. Bouman, and B. Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *NIPS*, 2016. [2](#)
- [39] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2image: Conditional image generation from visual attributes. In *ECCV*, 2016. [2](#)
- [40] L. Yu, P. Poirson, S. Yang, A. C. Berg, and T. L. Berg. Modeling context in referring expressions. In *ECCV*, 2016. [2](#), [4](#), [5](#), [6](#)
- [41] L. Yu, H. Tan, M. Bansal, and T. L. Berg. A joint speaker-listener-reinforcer model for referring expressions. In *ICCV*, 2017. [2](#), [6](#), [7](#)
- [42] H. Zhang, Z. Kyaw, S.-F. Chang, and T.-S. Chua. Visual translation embedding network for visual relation detection. In *CVPR*, 2017. [1](#), [2](#)
- [43] H. Zhang, Z. Kyaw, J. Yu, and S.-F. Chang. Ppr-fcn: Weakly supervised visual relation detection via parallel pairwise r-fcn. In *ICCV*, 2017. [2](#), [7](#), [8](#)
- [44] C. L. Zitnick and P. Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, 2014. [4](#)