



Conditional Image-Text Embedding Networks

Bryan A. Plummer¹(✉), Paige Kordas¹, M. Hadi Kiapour², Shuai Zheng²,
Robinson Piramuthu², and Svetlana Lazebnik¹

¹ University of Illinois at Urbana-Champaign, Champaign, USA

{bplumme2, pkordas2, slazebni}@illinois.edu

² Ebay Inc., San Jose, USA

{mkiapour, shuzheng, rpiramuthu}@ebay.com

Abstract. This paper presents an approach for grounding phrases in images which jointly learns multiple text-conditioned embeddings in a single end-to-end model. In order to differentiate text phrases into semantically distinct subspaces, we propose a concept weight branch that automatically assigns phrases to embeddings, whereas prior works predefine such assignments. Our proposed solution simplifies the representation requirements for individual embeddings and allows the underrepresented concepts to take advantage of the shared representations before feeding them into concept-specific layers. Comprehensive experiments verify the effectiveness of our approach across three phrase grounding datasets, Flickr30K Entities, ReferIt Game, and Visual Genome, where we obtain a (resp.) 4%, 3%, and 4% improvement in grounding performance over a strong region-phrase embedding baseline (Code: <https://github.com/BryanPlummer/cite>).

Keywords: Natural language grounding · Phrase localization
Embedding methods · Conditional models

1 Introduction

Phrase grounding attempts to localize a given natural language phrase in an image. This constituent task has applications to image captioning [6, 12, 14, 19, 34], image retrieval [9, 26], and visual question answering [1, 7, 29]. Research on phrase grounding has been spurred by the release of several datasets, some of which primarily contain relatively short phrases [15, 18], while others contain longer queries, including entire sentences that can provide rich context [22, 25]. The difference in query length compounds the already challenging problem of generalizing to any (including never before seen) natural language input. Despite this, much of the recent attention has focused on learning a single embedding model between image regions and phrases [7, 10, 21, 22, 28, 31, 32, 35].

In this paper, we propose a Conditional Image-Text Embedding (CITE) network that jointly learns different embeddings for subsets of phrases (Fig. 1). This

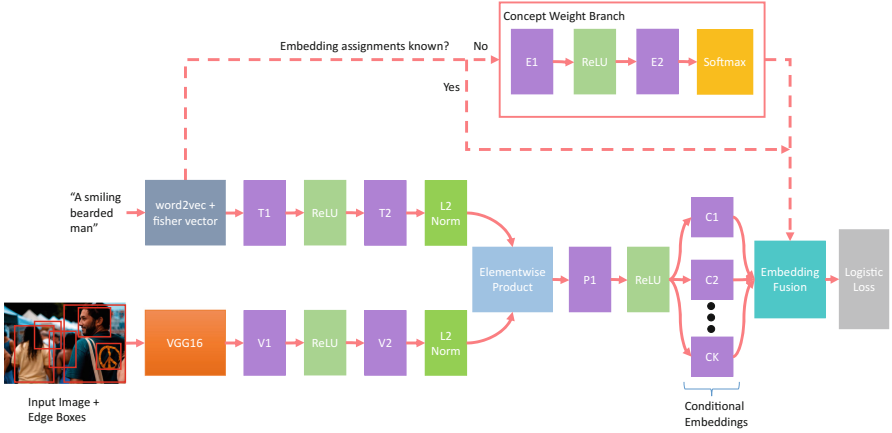


Fig. 1. Our CITE model separates phrases into different groups and learns conditional embeddings for these groups in a single end-to-end model. Assignments of phrases to embeddings can either be pre-defined (e.g. by separating phrases into distinct concepts like *people* or *clothing*), or can be jointly learned with the embeddings using the concept weight branch. Similarly colored blocks refer to layers of the same type, with purple blocks representing fully connected layers. Best viewed in color

enables our model to train separate embeddings for phrases that share a concept. Each conditional embedding can learn a representation specific to a subset of phrases while also taking advantage of weights that are shared across phrases. This is especially important for smaller groups of phrases that would be prone to overfitting if we were to train separate embeddings for them. In contrast to similar approaches that manually determine how to group concepts [20, 24, 30], we use a concept weight branch, trained jointly with the rest of the network, to do a soft assignment of phrases to learned embeddings automatically. The concept weight branch can be thought of producing a unique embedding for each region-phrase pair based on a phrase-specific linear combination of individual conditional embeddings. By training multiple embeddings our model also reduces variance akin to an ensemble of networks, but with far fewer parameters and lower computational cost.

Our idea of conditional embeddings was directly inspired by the conditional similarity networks of Veit et al. [30], although that work does not deal with cross-modal data and does not attempt to automatically assign different input items to different similarity subspaces. An earlier precursor of the idea of conditional similarity metrics can be found in [2]. Our work is also similar in spirit to Zhang et al. [37], who produced a linear classifier used to discriminate between image regions based on the textual input.

Our primary focus is on improving methods of associating individual image regions with individual phrases. Orthogonal to this goal, other works have focused on performing global inference for multiple phrases in a sentence and multiple regions in an image. Wang et al. [33] modeled the pronoun relationships

between phrases and forced each phrase prediction associated with a caption to be assigned to a different region. Chen *et al.* [3] also took into account the predictions made by other phrases when localizing phrases and incorporated bounding box regression to improve their region proposals. In their follow-up work [4], they introduced a region proposal network for phrases effectively reproducing the full Faster RCNN detection pipeline [27]. Yu *et al.* [36] took into account the visual similarity of objects in a single image when providing context for their predictions. Plummer *et al.* [24] performed global inference using a wide range of image-language constraints derived from attributes, verbs, prepositions, and pronouns. Yeh *et al.* [35] used a word prior in combination with segmentation masks, geometric features, and detection scores to select a region from all possible bounding boxes in an image. Many of these modifications could be used in combination with our approach to further improve performance.

The contributions of our paper are summarized below:

- By conditioning the embedding used by our model on the input phrase we simplify the representation requirements for each embedding, leading to a more generalizable model.
- We introduce a concept weight branch which enables our embedding assignments to be learned jointly with the image-text model.
- We introduce several improvements to the Similarity Network of Wang *et al.* [32] boosting the baseline model’s localization performance by 3.5% over the original paper.
- We perform extensive experiments over three datasets, Flickr30K Entities [25], ReferIt Game [15], and Visual Genome [18], where we report a (resp.) 4%, 3% and 4% improvement in phrase grounding performance over the baseline.

We begin Sect. 2.1 by describing the image-text Similarity Network [32] that we use as our baseline model. Section 2.2 describes our text-conditioned embedding model. Section 2.3 discusses three methods of assigning phrases to the trained embeddings. Lastly, Sect. 3 contains detailed experimental results and analysis of our proposed approach.

2 Our Approach

2.1 Image-Text Similarity Network

Given an image and a phrase, our goal is to select the most likely location of the phrase from a set of region proposals. To accomplish this, we build upon the image-text similarity network introduced in Wang *et al.* [32]. The image and text branches of this network each have two fully connected layers with batch normalization [11] and ReLUs. The final outputs of these branches are L2 normalized before performing an element-wise product between the image and text representations. This representation is then fed into a triplet of fully

connected layers using batch normalization and ReLUs. This is analogous to using the CITE model in Fig. 1 with a single conditional embedding.

The training objective for this network is a logistic regression loss computed over phrases P , the image regions R , and labels Y . The label y_{ij} for the i th input phrase and j th region is $+1$ where they match and -1 otherwise. Since this is a supervised learning approach, matching pairs of phrases and regions need to be provided in the annotations of each dataset. After producing some score x_{ij} measuring the affinity between the image region and text features using our network, the loss is given by

$$L_{sim}(P, R, Y) = \sum_{ij} \log(1 + \exp(-y_{ij}x_{ij})). \quad (1)$$

In this formulation, it is easy to consider multiple regions for a given phrase as positive examples and to use a variable number of region proposals per image. This is in contrast to competing methods which score regions with softmax with a cross entropy loss over a set number of proposals per image (*e.g.* [3, 7, 28]).

Sampling Phrase-Region Training Pairs. Following Wang *et al.* [32], we consider any regions with at least 0.6 intersection over union (IOU) with the ground truth box for a given phrase as a positive example. Negative examples are randomly sampled from regions of the same image with less than 0.3 IOU with the ground truth box. We select twice the number of negative regions as we have positive regions for a phrase. If too few negative regions occur for an image-phrase pair, then the negative example threshold is raised to 0.4 IOU.

Features. We represent phrases using the HGLMM fisher vector encoding [17] of word2vec [23] PCA reduced down to 6,000 dimensions. We generate region proposals using Edge Boxes [38]. Similarly to most state-of-the-art methods on our target datasets, we represent image regions using a Fast RCNN network [8] fine-tuned on the union of PASCAL 2007 and 2012 trainval sets [5]. The only exception is the experiment reported in Table 1(d), where we fine-tune the Fast RCNN parameters (corresponding to the VGG16 box in Fig. 1) on the Flickr30K Entities dataset.

Spatial Location. Following [3, 4, 28, 36], we experiment with concatenating bounding box location features to our region representation. This way our model can learn to bias predictions for phrases based on their location (*e.g.* that *sky* typically occurs in the top part of an image). For Flickr30K Entities we encode this spatial information as defined in [3, 4] for this dataset. For an image of height H and width W and a box with height h and width w is encoded as $[x_{min}/W, y_{min}/H, x_{max}/W, y_{max}/H, wh/WH]$. For a fair comparison to prior work [3, 4, 28], experiments on the ReferIt Game dataset encode the spatial information as an 8-dimensional feature vector $[x_{min}, y_{min}, x_{max}, y_{max}, x_{center}, y_{center}, w, h]$. For Visual Genome we adopt the same method of encoding spatial location as used for the ReferIt Game dataset.

2.2 Conditional Image-Text Network

Inspired by Veit *et al.* [30], we modify the image-text similarity model of the previous section to learn a set of conditional or concept embedding layers denoted C_1, \dots, C_K in Fig. 1. These are K parallel fully connected layers each with output dimensionality M . The outputs of these layers, in the form of a matrix of size $M \times K$, are fed into the embedding fusion layer, together with a K -dimensional concept weight vector U , which can be produced by several methods, as discussed in Sect. 2.3. The fusion layer simply performs a matrix-vector product, *i.e.*, $F = CU$. This is followed by another fully connected layer representing the final classifier (*i.e.*, the layer’s output dimension is 1).

2.3 Embedding Assignment

This section describes three possible methods for producing the concept weight vector U for combining the conditional embeddings as introduced in Sect. 2.2.

Coarse Categories. The Flickr30K Entities dataset comes with hand-constructed dictionaries that group phrases into eight coarse categories: *people, clothing, body parts, animals, vehicles, instruments, scene, other*. We use these dictionaries to map phrases to binary concept vectors representing their group membership. This is analogous to the approach of Veit *et al.* [30], which defines the concepts based on meta-data labels. Both the remaining approaches base their assignments on the training data rather than a hand-defined category label.

Nearest Cluster Center. A simple method of creating concept weights is to perform K-means clustering on the text features of the queries in the test set. Each cluster center becomes its own concept to learn. The concept weights U are then encoded as one-hot cluster membership vectors which we found to work better than alternatives such as similarity of a sample to each cluster center.

Concept Weight Branch. Creating a predefined set of concepts to learn, either using dictionaries or K-means clustering, produces concepts that don’t necessarily have anything to do with the difficulty or ease in localizing the phrases within them. An alternative is to let the model decide which concepts to learn. With this in mind, we feed the raw text features into a separate branch of the network consisting of two fully connected layers with batch normalization and a ReLU between them, followed by a softmax layer to ensure the output sums to 1 (denoted as the concept weight branch in Fig. 1). The output of the softmax is then used as the concept weights U . This can be seen as analogous to using soft attention [34] on the text features to select concepts for the final representation of a phrase. We use L1 regularization on the output of the last fully connected layer before being fed into the softmax to promote sparsity in our assignments. The training objective for our full CITE model then becomes

$$L_{CITE} = L_{sim}(P, R, Y) + \lambda \|\phi\|_1, \quad (2)$$

where ϕ are the inputs to the softmax layer and λ is a parameter controlling the importance of the regularization term. Note that we do not enforce diversity of assignments between different phrases, so it is possible that all phrases attend to a single embedding. However, we do not see this actually occur in practice. We also tried to use entropy minimization rather than L1 regularization for our concept weight branch as well as hard attention instead of soft attention, but found all worked similarly in our experiments.

3 Experiments

3.1 Datasets and Protocols

We evaluate the performance of our phrase-region grounding model on three datasets: Flickr30K Entities [25], ReferIt Game [15], and Visual Genome [18]. The metric we report is the proportion of correctly localized phrases in the test set. Consistent with prior work, a 0.5 IOU between the best-predicted box for a phrase and its ground truth is required for a phrase to be considered successfully localized. Similarly to [4, 24, 32], for phrases associated with multiple bounding boxes, the phrase is represented as the union of its boxes.

Training Procedure. We begin training our models with Adam [16]. After every epoch, we evaluate our model on the validation set. After it hasn't improved performance for 5 epochs, we fine-tune our model with stochastic gradient descent at 1/10th the learning rate and the same stopping criteria. We report test set performance for the model that performed best on the validation set.

Comparative Evaluation. In addition to comparing to previously published numbers of state-of-the-art approaches on each dataset, we systematically evaluate the following baselines and variants of our model:

- **Similarity Network.** Our first baseline is given by our own implementation of the model from Wang *et al.* [32], trained using the procedure described above. Phrases are pre-processed using stop word removal rather than part-of-speech filtering as done in the original paper. This change, together with a more careful tuning of the training settings, leads to a 2.5% improvement in performance over the reported results in [32]. The model is further enhanced by using the spatial location features (Sect. 2.1), resulting in a total improvement of 3.5%.
- **Individual Coarse Category Similarity Networks.** We train multiple Similarity Networks on different subsets of the data created according to the coarse category assignments as described in Sect. 2.3.

- **Individual K-means Similarity Networks.** We train multiple Similarity Networks on different subsets of the data created according to the nearest cluster center assignments as described in Sect. 2.3.
- **CITE, Coarse Categories.** No concept weight branch. Phrases are assigned according to their coarse category.
- **CITE, Random.** No concept weight branch. Phrases are randomly assigned to an embedding. At test time, phrases seen during training keep their assignments, while new phrases are randomly assigned.
- **CITE, K-means.** No concept weight branch. Phrases are matched to embeddings using nearest cluster center assignments.
- **CITE, Learned.** Our full model with the concept weight branch used to automatically produce concept weights as described in Sect. 2.3.

3.2 Flickr30K Entities

We use the same splits as Plummer *et al.* [25], which separates the images into 29,783 for training, 1,000 for testing, and 1,000 for validation. Models are trained with a batch size of 200 (128 if necessary to fit into GPU memory) and learning rate of $5e-5$. We set $\lambda = 5e-5$ in Eq. (2). We use the top 200 Edge Box proposals per image and embedding dimension $M = 256$ unless stated otherwise.

Grounding Results. Table 1 compares overall localization accuracies for a number of methods. The numbers for our Similarity Network baseline are reported in Table 1(b), and as stated above, they are better than the published numbers from [32]. Table 1(c) reports results for variants of conditional embedding models. From the first two lines, we can see that learning embeddings from subsets of the data without any shared weights leads to only a small improvement ($\leq 1\%$) over the Similarity Network baseline. The third line of Table 1(c) reports that separating phrases by manually defined high-level concepts only leads to a 1% improvement even when weights are shared across embeddings. This is likely due, in part, to the significant imbalance between different coarse categories, as a uniform random assignment shown in the fourth line of Table 1(c) lead to a 3% improvement. The fifth line of Table 1(c) demonstrates that grouping phrases based on their text features better reflects the needs of the data, resulting in just over 3% improvement over the baseline, only slightly better than random assignments. An additional improvement is reported in the eighth line of Table 1(c) by incorporating our concept weight branch, enabling our model to both determine what concepts are important to learn and how to assign phrases to them. We see in the last line of Table 1(c) that going from 200 to 500 bounding box proposals provides a small boost in localization accuracy. This results in our best performance using PASCAL-tuned features which is 3% better than the prior work reported in Table 1(a) and 4.5% better than the Similarity Network. We also note that the time to test an image-phrase pair is almost unaffected using our approach (the CITE, Learned, $K = 4$ model performs inference on 200

Table 1. Phrase localization performance on the Flickr30k Entities test set. (a) State-of-the-art results when predicting a single phrase at a time taken from published works. (b,c) Our baselines and variants using PASCAL-tuned features. (d) Results using Flickr30k-tuned features

Method	Accuracy
(a) Single phrase methods (PASCAL-tuned features)^a	
NonlinearSP [31]	43.89
GroundeR [28]	47.81
MCB [7]	48.69
RtP [25]	50.89
Similarity Network [32]	51.05
IGOP [35]	53.97
SPC [24]	55.49
MCB + Reg + Spatial [3]	51.01
MNN + Reg + Spatial [3]	55.99
(b) Our implementation	
Similarity Network	53.45
Similarity Network + Spatial	54.52
(c) Conditional models + spatial	
Individual Coarse Category Similarity Networks, $K = 8$	55.32
Individual K-means Similarity Networks, $K = 8$	54.95
CITE, Coarse Categories, $K = 8$	55.42
CITE, Random, $K = 16$	57.58
CITE, K-means, $K = 16$	57.89
CITE, Learned, $K = 4$	58.69
CITE, Learned, $K = 4$, 500 Edge Boxes	59.27
(d) Flickr30K-tuned features + spatial	
PGN + QRN [4]	60.21
CITE, Learned, $K = 4$, 500 Edge Boxes	61.89

^aPerformance on this task can be further improved by taking into account the predictions made for other phrases in the same sentence [3, 4, 24, 33], with the best result using Pascal-tuned features of 57.53% achieved by Chen *et al.* [3] and 65.14% using Flickr30K-tuned features [4].

Edge Boxes at 0.182 s per pair using a NVIDIA Titan X GPU with our implementation) compared with the baseline Similarity Network (0.171 s per pair). Finally, Table 1(d) gives results for models whose visual features were fine-tuned for localization on the Flickr30K Entities dataset. Our model still obtains a 1.5% improvement over the approach of Chen *et al.* [4], which used bounding box regression as well as a region proposal network. In principle, we could also incorporate these techniques to further improve the model.

Table 2. Comparison of phrase grounding performance over coarse categories on the Flickr30K Entities dataset. Our models were tested with 500 Edge Box proposals

	People	Clothing	Body parts	Animals	Vehicles	Instruments	Scene	Other
PASCAL-tuned Features								
GroundedR [28]	61.00	38.12	10.33	62.55	68.75	36.42	58.18	29.08
RtP [25]	64.73	46.88	17.21	65.83	68.75	37.65	51.39	31.77
IGOP [35]	68.71	56.83	19.50	70.07	73.75	39.50	60.38	32.45
MCB + Reg + Spatial [3]	62.75	43.67	14.91	65.44	65.25	24.74	64.10	34.62
MNN + Reg + Spatial [3]	67.38	47.57	20.11	73.75	72.44	29.34	63.68	37.88
CITE, Learned, $K = 4 +$ Spatial	73.20	52.34	30.59	76.25	75.75	48.15	55.64	42.83
Flickr30K-tuned Features								
PGN + QRN + Spatial [4]	75.05	55.90	20.27	73.36	68.95	45.68	65.27	38.80
CITE, Learned, $K = 4 +$ Spatial	75.95	58.50	30.78	77.03	79.25	48.15	58.78	43.24


**Fig. 2.** Effect of the number of learned embeddings (K) on Flickr30K Entities localization accuracy using PASCAL-tuned features

Table 2 breaks down localization accuracy by coarse category. Of particular note are our results on the challenging *body part* category, which are typically small and represent only 3.5% of the phrases in the test set, improving over the next best model as well as the Similarity Network trained on just body part phrases by 10% when using Flickr30K-tuned features. We also see a substantial improvement in the *vehicles* and *other* categories, seeing a 5–9% improvement over the previous state-of-the-art. The only category where we perform worse are phrases referring to scenes, which commonly cover the majority (or entire) image. Here, incorporating a bias towards selecting larger proposals, as in [24, 25], can lead to significant improvements.

Parameter Selection. In addition to reporting the localization performance, we also provide some insight into the effect of different parameter choices and what information our model is capturing. In Fig. 2 we show how the number K of learned embeddings affects performance. Using our concept weight branch consistently outperforms K-means cluster assignments. Table 3 shows how the embedding dimensionality M affects performance. Here we see that reducing the output dimension from 256 to 64 (*i.e.*, by 1/4th) leads to a minor (1%) decrease in performance. This result is particularly noteworthy as the CITE network with $K = 4, M = 64$ has 4 million parameters compared the 14 million

the baseline Similarity Network has with $M = 256$ while still maintaining a 3% improvement in performance. We also experimented with different ways of altering the Similarity Network to have the same number of parameters to ours at similar points (*e.g.* increasing the last fully connected layer to be K times larger or adding K additional layers), but found they performed comparably to the baseline Similarity Network (*i.e.* their performance was about 4% worse than our approach). In addition to experiments on how many layers to use and the size of each layer, we also explored the effect the number of Edge Boxes has on performance in Table 4. In contrast to some prior work which performed best using 200 candidates (*e.g.* [24, 25]), our model’s increased discriminate power enables us to still be able to obtain a benefit from using up to 500 proposals.

Concept Weight Branch Examination. To analyze what our model is learning, Fig. 3 shows the means and standard deviations of the weights over the different embeddings broken down by coarse categories. Interestingly, *people* end up being split between two embeddings. We find that people phrases tend to be split by plural vs. singular. Table 5 gives a closer look at the conditional embeddings by listing the ten phrases with the highest weight for each embedding. While most phrases give the first embedding little weight, it appears to provide the most benefit for finding very specific references to people rather than generic terms (*e.g.* *little curly hair girl* instead of *girl* itself). These patterns generally hold through multiple runs of the model, indicating they are important concepts to learn for the task.

Qualitative Results. Figure 4 gives a look into areas where our model could be improved. Of the phrases that occur at least 100 times in the test set, the lowest performing phrases are *street* and *people* at (resp.) 60% and 64% accuracy. The highest performing of these common phrases is *man* at 81% accuracy, which also happens to be the most common phrase with 1065 instances in the test set. In the top-left example of Fig. 4, the word *people*, which is not correctly localized, refers to partially visible background pedestrians. Analyzing the saliency of a phrase in the context of the whole caption may lead to treating these phrases differently. Global inference constraints, for example, a requirement that predictions for *a man* and *a woman* must be different, would be useful for the top-center example. Performing pronoun resolution, as attempted in [24], would help in the top-right example. In the test set, the pronoun *one* is correctly localized around 36% of the time, whereas *the blond woman* is correctly localized 81% of the time. Having an understanding of relationships between entities may help in cases such as the bottom-left example of Fig. 4, where the extent of the table could be refined by knowing that the groceries are “on” it. Our model also performs relatively poorly on phrases referring to classic “stuff” categories, as shown in the bottom-center and bottom-right examples. The *water* and *street* phrases in these examples are only partly localized. Using pixel-level predictions may help to recover the full extent of these types of phrases since the parts of the images they refer to are relatively homogeneous.

Table 3. Localization accuracy with different embedding sizes using the CITE, Learned, $K = 4$ model on Flickr30K Entities with PASCAL-tuned features. Embedding size refers to M , the output dimensionality of layers P1 and the conditional embeddings in Fig. 1. The remaining fully connected layers’ output dimensions (excluding those that are part of the VGG16 network) are four times the embedding size

Embedding size (M)	64	128	256	512
Validation set accuracy	56.32	57.51	57.53	57.42
Test set accuracy	57.77	58.48	58.69	58.64

Table 4. Localization accuracy with different numbers of proposals using the CITE, Learned, $K = 4$ model on Flickr30K Entities with PASCAL-tuned features

#Edge Box proposals	100	200	500	1000
Validation set accuracy	49.61	57.53	58.48	57.87
Test set accuracy	51.32	58.69	59.27	58.63

3.3 ReferIt Game

We use the same splits as Hu *et al.* [10], which consist of 10,000 images combined for training and validation with the remaining 10,000 images for testing. Models are trained with a batch size of 128, learning rate of $5e-4$, and $\lambda = 5e-4$ in Eq. (2). We generate 500 Edge Box proposals per image.

Results. Table 6 reports the localization accuracy across the ReferIt Game test set. The first line of Table 6(b) shows that our model using the nearest cluster center assignments results in a 2.5% improvement over the baseline Similarity Network. Using our concept weight branch in order to learn assignments yields an additional small improvement.

We note that we do not outperform the approach of Yeh *et al.* [35] on this dataset. This can likely be attributed to the failures of Edge Boxes to produce adequate proposals on the ReferIt Game dataset. Oracle performance using the top 500 proposals is 93% on Flickr30K Entities, while it is only 86% on this dataset. As a result, the specialized bounding box methods used by Yeh *et al.* as well as Chen *et al.* [3] may play a larger role here. Our model would also likely benefit from these improved bounding boxes.

As with the Flickr30K Entities dataset, we show the effect of the number K of embeddings on localization performance in Fig. 5. While the concept weight branch provides a small performance improvement across many different choices of K , when $K = 2$ the clustering assignments actually perform a little better. However, this behavior is atypical in our experiments across all three datasets, and may simply be due to the small size of the ReferIt Game training data, as it has far fewer ground truth phrase-region pairs to train our models with.




Fig. 3. The mean weight for each embedding (left) along with the standard deviation of those weights (right) broken down by coarse category for the Flickr30K Entities dataset using Flickr30K-tuned features

Table 5. The ten phrases with the highest weight per embedding on the Flickr30K Entities dataset using Flickr30K-tuned features

Embedding 1	Soldiers (0.08), male nun (0.07), rather angry looking woman (0.07), skinny dark complected boy (0.07), little curly hair girl (0.07), middle eastern woman (0.07), first man’s leg (0.07), statue athletic man (0.07), referee (0.07), woman drink wine (0.07)
Embedding 2	Red scooter (0.97), blue clothes (0.97), yellow bike (0.97), red bike (0.97), red buckets (0.97), yellow backpack (0.97), street window shops (0.97), red blue buckets (0.97), red backpack (0.97), purple red backpack (0.97)
Embedding 3	Two people (0.94), two men (0.93), two young kids (0.93), two kids (0.93), two white-haired women (0.93), two women (0.93), group three boys (0.93), two young people (0.93), three people (0.92), crowd people (0.92)
Embedding 4	Blond-haired woman (0.91), dark-skinned woman (0.91), gray-haired man (0.91), one-armed man (0.91), dark-haired man (0.91), red-haired man (0.91), boy young man (0.91), man (0.91), well-dressed man (0.91), dark-skinned man (0.91)

3.4 Visual Genome

We use the same splits as Zhang *et al.* [37], consisting of 77,398 images for training and 5,000 each for testing and validation. Models are trained with a learning rate of $5e-5$, and $\lambda = 5e-4$ in Eq. (2). We generate 500 Edge Box proposals per image, and use a batch size of 128.

Results. Table 7 reports the localization accuracy across the Visual Genome dataset. Table 7(a) lists published numbers from several recent methods. The current state of the art performance belongs to Zhang *et al.* [37], who fine-tuned visual features on this dataset and created a cleaner set during training by pruning ambiguous phrases. We did not perform either fine-tuning or phrase



Fig. 4. Examples demonstrating some common failure cases on the Flickr30K Entities dataset. See Sect. 3.2 for discussion




Fig. 5. Effect of the number K of embeddings on localization accuracy on the ReferIt Game dataset

pruning, so the most comparable reference number for our methods is their 17.5% accuracy without these steps.

The baseline accuracies for our Similarity Network with and without spatial features are given in the last two lines of Table 7(a). We can see that including the spatial features gives only a small improvement. This is likely due to the denser annotations in this dataset as compared to Flickr30K Entities. For example, a phrase like *a man* in Flickr30K Entities would typically refer to a relatively large region towards the center since background instances are commonly not mentioned in an image-level caption. However, entities in Visual Genome include both foreground and background instances.

In the first line of Table 7(b), we see our K-means model is 3.5% better than the Similarity Network baseline, and over 6% better than the 17.5% accuracy of [37]. According to the second line of Table 7(b), using the concept weight branch obtains a further improvement. In fact, our full model with pre-trained PASCAL features has better performance than [37] with fine-tuned features.

Table 6. Localization performance on the ReferIt Game test set. (a) Published results and our Similarity Network baseline. (b) Our best-performing conditional models

Method	Accuracy
(a) State-of-the-art	
SCRC [10]	17.93
GroundeR + Spatial [28]	26.93
MCB + Reg + Spatial [3]	26.54
CGRE [21]	31.85
MNN + Reg + Spatial [3]	32.21
IGOP [35]	34.70
Similarity Network + Spatial	31.26
(b) Conditional Models + Spatial	
CITE, K-Means, $K = 2$	34.01
CITE, Learned, $K = 12$	34.13

Table 7. Phrase localization performance on Visual Genome. (a) Published results and our Similarity Network baselines. APP refers to ambiguous phrase pruning (see [37] for details). (b) Our best-performing conditional models

Method	Accuracy
(a) State-of-the-art	
Densecap [13]	10.1
SCRC [10]	11.0
DBNet [37]	17.5
DBNet (with APP) [37]	21.2
DBNet (with APP, V. Genome-tuned Features) [37]	23.7
Similarity Network	19.76
Similarity Network + Spatial	20.08
(b) Conditional Models + Spatial	
CITE, K-Means, $K = 12$	23.67
CITE, Learned, $K = 12$	24.43

As with the other two datasets, Fig. 6 reports performance as a function of the number of learned embeddings. Echoing most of the earlier results, we see a consistent improvement for the learned embeddings over the K-means ones. The large size of this dataset (>250,000 instances in the test set) helps to reinforce the significance of our results.




Fig. 6. Effect of the number of learned embeddings on performance on the Visual Genome with models trained on 1/3 of the available training data

4 Conclusion

This paper introduced a method of learning a set of conditional embeddings and phrase-to-embedding assignments in a single end-to-end network. The effectiveness of our approach was demonstrated on three popular and challenging phrase-to-region grounding datasets. In future work, our model could be further improved by including a term to enforce that distinct concepts are being learned by each embedding.

Our experiments focused on localizing individual phrases to a fixed set of category-independent region proposals. As such, our absolute accuracies could be further improved by incorporating a number of orthogonal techniques used in competing work. By jointly predicting multiple phrases in an image our model could take advantage of relationships between multiple entities (*e.g.* [3, 4, 24, 33]). Including bounding box regression and a region proposal network as done in [3, 4] would also likely lead to a better model. In fact, tying the regression parameters to a specific concept embedding may further improve performance since it would simplify our prediction task as a result of needing to learn parameters for just the phrases assigned to that embedding.

Acknowledgements. This material is based upon work supported in part by the National Science Foundation under Grants No. 1563727 and 1718221, Amazon Research Award, AWS Machine Learning Research Award, and Google Research Award.

References

1. Antol, S., et al.: VQA: visual question answering. In: ICCV (2015)
2. Babenko, B., Branson, S., Belongie, S.: Similarity metrics for categorization: from monolithic to category specific. In: ICCV (2009)
3. Chen, K., Kovvuri, R., Gao, J., Nevatia, R.: MSRC: multimodal spatial regression with semantic context for phrase grounding. In: ICMR (2017)
4. Chen, K., Kovvuri, R., Nevatia, R.: Query-guided regression network with context policy for phrase grounding. In: ICCV (2017)
5. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL visual object classes challenge 2012 (VOC2012) results (2012). <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>

6. Fang, H., et al.: From captions to visual concepts and back. In: CVPR (2015)
7. Fukui, A., Park, D.H., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M.: Multi-modal compact bilinear pooling for visual question answering and visual grounding. In: EMNLP (2016)
8. Girshick, R.: Fast R-CNN. In: ICCV (2015)
9. Gordo, A., Almazán, J., Revaud, J., Larlus, D.: Deep image retrieval: learning global representations for image search. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9910, pp. 241–257. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46466-4_15
10. Hu, R., Xu, H., Rohrbach, M., Feng, J., Saenko, K., Darrell, T.: Natural language object retrieval. In: CVPR (2016)
11. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
12. Johnson, J., Karpathy, A., Fei-Fei, L.: DenseCap: fully convolutional localization networks for dense captioning. In: CVPR (2016)
13. Johnson, J., et al.: Image retrieval using scene graphs. In: CVPR (2015)
14. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: CVPR (2015)
15. Kazemzadeh, S., Ordonez, V., Matten, M., Berg, T.: ReferitGame: referring to objects in photographs of natural scenes. In: EMNLP (2014)
16. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: International Conference for Learning Representations (2015)
17. Klein, B., Lev, G., Sadeh, G., Wolf, L.: Associating neural word embeddings with deep image representations using fisher vector. In: CVPR (2015)
18. Krishna, R.: Visual genome: connecting language and vision using crowdsourced dense image annotations. IJCV **123**, 32–73 (2017)
19. Liu, C., Mao, J., Sha, F., Yuille, A.: Attention correctness in neural image captioning. In: AAAI (2017)
20. Liu, J., Wang, L., Yang, M.H.: Referring expression generation and comprehension via attributes. In: ICCV (2017)
21. Luo, R., Shakhnarovich, G.: Comprehension-guided referring expressions. In: CVPR (2017)
22. Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A., Murphy, K.: Generation and comprehension of unambiguous object descriptions. In: CVPR (2016)
23. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. [arXiv:1301.3781](https://arxiv.org/abs/1301.3781) (2013)
24. Plummer, B.A., Mallya, A., Cervantes, C.M., Hockenmaier, J., Lazebnik, S.: Phrase localization and visual relationship detection with comprehensive image-language cues. In: ICCV (2017)
25. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: collecting region-to-phrase correspondences for richer image-to-sentence models. IJCV **123**(1), 74–93 (2017)
26. Radenović, F., Toliás, G., Chum, O.: CNN image retrieval learns from BoW: unsupervised fine-tuning with hard examples. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 3–20. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_1
27. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: NIPS (2015)

28. Rohrbach, A., Rohrbach, M., Hu, R., Darrell, T., Schiele, B.: Grounding of textual phrases in images by reconstruction. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 817–834. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_49
29. Tommasi, T., Mallya, A., Plummer, B.A., Lazebnik, S., Berg, A.C., Berg, T.L.: Solving visual madlibs with multiple cues. In: BMVC (2016)
30. Veit, A., Belongie, S., Karaletsos, T.: Conditional similarity networks. In: CVPR (2017)
31. Wang, L., Li, Y., Lazebnik, S.: Learning deep structure-preserving image-text embeddings. In: CVPR (2016)
32. Wang, L., Li, Y., Lazebnik, S.: Learning two-branch neural networks for image-text matching tasks. [arXiv:1704.03470](https://arxiv.org/abs/1704.03470) (2017)
33. Wang, M., Azab, M., Kojima, N., Mihalcea, R., Deng, J.: Structured matching for phrase localization. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9912, pp. 696–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_42
34. Xu, K., et al.: Show, attend and tell: neural image caption generation with visual attention. In: ICML (2015)
35. Yeh, R.A., Xiong, J., Hwu, W.M., Do, M.N., Schwing, A.G.: Interpretable and globally optimal prediction for textual grounding using image concepts. In: NIPS (2017)
36. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 69–85. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_5
37. Zhang, Y., Yuan, L., Guo, Y., He, Z., Huang, I.A., Lee, H.: Discriminative bimodal networks for visual localization and detection with natural language queries. In: CVPR (2017)
38. Zitnick, C.L., Dollár, P.: Edge boxes: locating object proposals from edges. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 391–405. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_26