

Object Captioning and Retrieval with Natural Language

Anh Nguyen¹, Quang D. Tran², Thanh-Toan Do³, Ian Reid⁴, Darwin G. Caldwell⁵, Nikos G. Tsagarakis⁵

¹Imperial College London, ²AIOZ Pte Ltd, ³University of Liverpool, ⁴The University of Adelaide, ⁵IIT

Abstract

We address the problem of jointly learning vision and language to understand the object in a fine-grained manner. The key idea of our approach is the use of object descriptions to provide the detailed understanding of an object. Based on this idea, we propose two new architectures to solve two related problems: object captioning and natural language-based object retrieval. The goal of the object captioning task is to simultaneously detect the object and generate its associated description, while in the object retrieval task, the goal is to localize an object given an input query. We demonstrate that both problems can be solved effectively using hybrid end-to-end CNN-LSTM networks. The experimental results on our new challenging dataset show that our methods outperform recent methods by a fair margin, while providing a detailed understanding of the object and having fast inference time. The source code will be made available.

1. INTRODUCTION

Over the past few years, deep learning has become a popular approach to solve visual problems, with traditional problems in computer vision such as image classification [1], instance segmentation [2], and object detection [3] experiencing many mini-revolutions. In spite of these remarkable results, the way these problems are defined prevents them from being widely used in robotic applications. For example, the problem of instance segmentation is formed as a binary segment mask inside a rectangle bounding box. While this is reasonable for computer vision applications, we usually need more information (e.g., object part understanding, grasping frame, etc.) for real-world robotic applications [4].

In this paper, we extend the traditional object detection problem to make it more realistic and suitable for robotic applications. We argue that although recent successful object detectors can achieve reasonable results on a dataset with a few thousand classes [5], they are still limited by the predefined classes presented during the training. Furthermore, the object detector is also unable to provide more use-




Figure 1. A comparison between the traditional object detection problem and our approach. **Top row:** The groundtruth of the traditional object detection problem which is restricted to the predefined object categories. **Bottom row:** We propose to use textual caption as the object description, providing a more detailed understanding of each object.

ful information about the object. On the other hand, humans are able to distinguish more than 30,000 basic categories, and we not only recognize the object based on its category, but also are able to describe the object based on its properties and attributes [6]. Motivated by these limitations, we propose to integrate *natural language* into the object detector. Compared to the traditional object detection approaches that output only the category information, our approach provides a better way to understand the objects by outputting its fine-grained textual description. From this observation, we propose a new method to simultaneously detect the object and generate its caption. Moreover, we show that by using natural language, we can easily adapt an object captioning architecture to a retrieval system, which has excellent potential in many real-world robotic applications [7].

In particular, we first define a small set of *superclasses* (e.g., *animals*, *vehicles*, etc.), then each object has the caption as its description. This is the main difference between our approach and the traditional object detection problem. The superclass in our usage is a general class,

which can contain many (unlimited) fine-grained classes (e.g., the `animals` class contains all the sub-classes such as `dog`, `cat`, `horse`, etc.). While the superclass provides only the general information, the object descriptions provide the fine-grained understanding of the object (e.g., “a black dog”, “a little Asian girl”, etc.). Fig. 1 shows a comparison between the tradition object detection problem and our approach.

Based on the aforementioned definition, we consider two problems: (1) object captioning and (2) object retrieval using natural language. Our goal in the first problem is to simultaneously detect the object and generate its description. In this way, we have a new object detector that can localize the object while providing more meaningful information. In the second problem, the goal is to localize an object given an input query. This is particularly useful in the human-robot interaction applications since it allows humans to use natural language to communicate with the robot [7]. We show that both problems can be solved effectively with the deep networks, providing a detailed understanding of the object while achieving the state-of-the-art performance.

The rest of this paper is organized as follows. We first review related work in Section 2. We then describe two end-to-end architectures for two tasks: object captioning and object retrieval with natural language in Section 3. In Section 4, we present the extensive experimental results on our new challenging dataset. Finally, the conclusion and future work are presented in Section 5.

2. Related Work

In computer vision, detecting objects from RGB images is a well-known task. With the rise of deep learning, recent works design different deep architectures to solve this problem. These architectures can be divided into two groups: region-based [3] and non region-based approaches [5] [8]. While non region-based methods can achieve real-time performance, region-based architectures provide a more flexible way to adapt the object detection problem to other scenarios [2]. However, the drawback of the methods in [3] [5] [8] is their reliance on a fixed set of classes for both training and inference. Therefore, they are unable to deal with a new class or provide more detailed understanding of the object.

Along with the object detection problem, image captioning is also an active field in computer vision. Current research has shown recurrent neural networks such as LSTM [9] to be effective on this task. Recently, the authors in [10] proposed to fuse deep reinforcement learning with LSTM and achieved competitive results. While we retain the concept of image captioning for object description, our goal here is more closely related to the dense captioning task [11] since we want to generate the caption for each object, not for the entire image. However, unlike [11] that

generates the captions for all possible regions, our network only generates captions for objects in the superclasses. This allows us to have a more detailed understanding of each object, while still being able to distinguish objects based on their superclasses.

In the robotics community, the work in [7] introduced a method to localize an object based on a text query. Recently, Hu et al. [12] improved on this by fusing the query text, the spatial, and the global context of the image into three recurrent neural networks. Similarly, the authors in [13] introduced a new discriminative training method for this task. Although these methods are able to localize the object based on the input query, their architectures are not end-to-end and unable to run in real-time since the object proposals are generated offline and not trained with the network. With a different viewpoint, Plummer et al. [14] proposed to localize objects that correspond to the noun phrases of a textual image description. Although our goal in the retrieval task is similar to [7] [12] [13], the key difference with our approach is the use of an end-to-end architecture, which has a fast inference time and does not depend on any external bounding box extraction method [12] [13].

3. Object Captioning and Retrieval with Natural Language

We start by briefly describing three basic building blocks used in our architecture: Convolutional backbone with Region Proposal Network (RPN) as proposed in Faster R-CNN [3], Long-Short Term Memory (LSTM) network [9] and the embedding of word representations. We then present in details the network architectures for two sub-problems in section 3.2 and 3.3.

3.1. Background

Convolutional Backbone Inspired by [2] [4], we build an end-to-end network with two branches: the first branch localizes and classifies the object based on its superclass, while the second branch handles the object caption. This architectural methodology was proposed in Faster R-CNN [3] and is now widely used since it provides a robust way to extract both the image feature map and the region features. In particular, given an input image, the image feature map is first extracted by a convolutional backbone (e.g., VGG16 [15]). An RPN that shares the weights with the convolutional backbone is then used to generate candidate bounding boxes (RoIs). For each RoI, unlike Faster R-CNN that uses the RoIPool layer, we use the RoIAlign [2] layer to robustly pool its corresponding features from the image feature map into a fixed size feature map.

LSTM In this work, we use LSTM to model the sequential relationship in the object caption. The robustness of the LSTM network is due to the gate mechanism that helps the

network encodes the sequential knowledge for long periods of time, while is still remaining sturdy against the vanishing gradient problem. In particular, the LSTM takes an input \mathbf{x}_t at each time step t , and computes the hidden state \mathbf{h}_t and the memory cell state \mathbf{c}_t as follows:

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}_{xi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{b}_i) \\ \mathbf{f}_t &= \sigma(\mathbf{W}_{xf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1} + \mathbf{b}_f) \\ \mathbf{o}_t &= \sigma(\mathbf{W}_{xo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{b}_o) \\ \mathbf{g}_t &= \phi(\mathbf{W}_{xg}\mathbf{x}_t + \mathbf{W}_{hg}\mathbf{h}_{t-1} + \mathbf{b}_g) \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t \\ \mathbf{h}_t &= \mathbf{o}_t \odot \phi(\mathbf{c}_t) \end{aligned} \quad (1)$$

where \odot represents element-wise multiplication, σ is the sigmoid non-linearity, and ϕ is the hyperbolic tangent non-linearity function. The weight \mathbf{W} and bias \mathbf{b} are the parameters of the network.

Word Embedding For simplicity, we choose the one-hot encoding technique as our word representation. The one-hot vector $\mathbf{y} \in \mathbb{R}^{|D|}$ is a binary vector with only one non-zero entry indicating the index of the current word in the vocabulary. Formally, each value in the one-hot vector \mathbf{y} is defined by:

$$\mathbf{y}^j = \begin{cases} 1, & \text{if } j = \text{ind}(\mathbf{y}) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where $\text{ind}(\mathbf{y})$ is the index of the current word in the dictionary D . In practice, we add two extra words to the dictionary (i.e., EOC word to denote the end of the caption, and UNK word to denote the unknown word).

3.2. Object Captioning

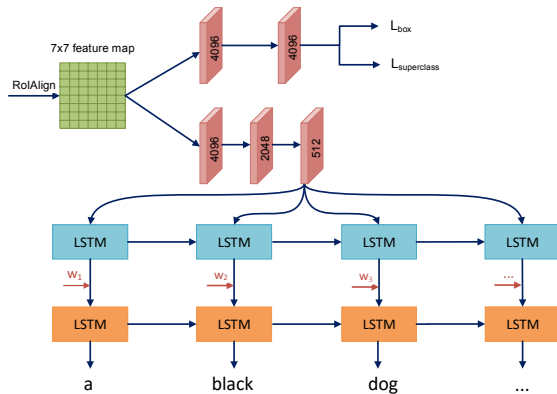


Figure 2. An overview of our object captioning network. In the second branch, each positive RoI is first cloned and fed into the $LSTM_{roi}$ layer (in light blue), then the input words and the outputs of $LSTM_{roi}$ layer are combined to become the input for the $LSTM_{word}$ layer (in dark yellow).

In this task, our goal is to simultaneously find the object location, its superclass, and the object caption. While the object location and its superclass are learned using the first branch of the network as in [2] [4], inspired by [16] [17] we propose to use two LSTM layers to handle the object caption. The first LSTM layer ($LSTM_{roi}$) encodes the visual information from each RoI provided by the RPN network, while the second layer ($LSTM_{word}$) combines the visual information from the first $LSTM_{roi}$ layer with the input words to generate the object caption. Fig. 2 shows an overview of our object captioning network.

In particular, we first use the convolutional backbone to extract the image feature map from the input image. From this feature map, the RoIAlign layer is used to pool the variable-sized RoIs to a fixed-size region feature map (i.e., 7×7). In the first branch, this region feature map is fed into two fully connected layers, each having 4096 neurons, to localize the object location and classify its superclass. In the second branch, this region feature map is fed into three fully connected layers to gradually squeeze the region feature map into a smaller map that fits with the LSTM input. The $LSTM_{roi}$ layer uses the feature from the last fully connected layer to encode the visual information for each RoI, while the $LSTM_{word}$ layer then encodes both the visual features provided by the $LSTM_{roi}$ layer and the input words to generate the object caption.

In practice, we use three fully connected layers with 4096, 2048, 512 neurons, respectively in the second branch. The feature from the last fully connected layer is used as the input for the $LSTM_{roi}$ layer. More formally, the input of the $LSTM_{roi}$ layer is a visual feature vector: $\mathbf{X}^r = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_\ell})$, where n_ℓ is the number of LSTM time steps. We note that all the $\mathbf{x}_i \in \mathbf{X}^r$ are identical since they are cloned from one RoI. The $LSTM_{roi}$ layer then encodes its visual input into a list of hidden state vectors $\mathbf{H}^r = (\mathbf{h}_1^r, \mathbf{h}_2^r, \dots, \mathbf{h}_{n_\ell}^r)$. Each hidden state vector \mathbf{h}_i^r is combined with one input word to become the input for the second LSTM layer, i.e., the input for the $LSTM_{word}$ layer is a vector: $\mathbf{X}^w = (\mathbf{h}_1^r \oplus \mathbf{w}_1, \mathbf{h}_2^r \oplus \mathbf{w}_2, \dots, \mathbf{h}_{n_\ell}^r \oplus \mathbf{w}_{n_\ell})$, where \oplus denotes the concatenation operation and \mathbf{w} is the input word of the object caption. In this way, the network is able to learn the sequential information in the input object caption, while is aware of which object the current caption belongs to (via the concatenation operation).

Multi-task Loss We train the network end-to-end with a multi-task loss L function as follows:

$$L = L_{loc} + L_{superclass} + L_{caption} \quad (3)$$

where L_{loc} and $L_{superclass}$ are defined on the first branch to regress the object location and classify its superclass, $L_{caption}$ is defined on the second branch to generate the object caption. We refer readers to [3] for a full description of the L_{loc} and $L_{superclass}$ loss, while we present the

$L_{caption}$ loss in details here.

Let \mathbf{z}_t denote the output of each cell at each time step t in the $LSTM_{word}$ layer. Similar to [18], this output is then passed through a linear prediction layer $\hat{\mathbf{y}}_t = \mathbf{W}_z \mathbf{z}_t + \mathbf{b}_z$, and the predicted distribution $P(\mathbf{y}_t)$ is computed by taking the softmax of $\hat{\mathbf{y}}_t$ as follows:

$$P(\mathbf{y}_t = \mathbf{w} | \mathbf{z}_t) = \frac{\exp(\hat{\mathbf{y}}_{t,w})}{\sum_{\mathbf{w}' \in D} \exp(\hat{\mathbf{y}}_{t,w'})} \quad (4)$$

where \mathbf{W}_z and \mathbf{b}_z are learned parameters, \mathbf{w} is a word in the dictionary D . The $L_{caption}$ loss is then computed as follows:

$$L_{caption} = \sum_{i=0}^{n_{RoI}} \sum_{t=0}^{n_\ell} P_i(\mathbf{y}_t = \mathbf{w} | \mathbf{z}_t) \quad (5)$$

where n_{RoI} is the number of positive RoIs, n_ℓ is the number of LSTM time steps. Intuitively, Equation 5 computes the loss at each word of the current outputted caption for each positive RoI provided by the RNP network.

Training and Inference The network is trained end-to-end for 200k iterations using stochastic gradient descent with 0.9 momentum and 0.0005 weight decay. The learning rate is empirically set to 0.001 and keeps unchanging during the training. We select 2000 RoIs from RPN to compute the multi-task loss. A RoI is considered as positive if it has the IoU with a ground-truth box by at least 50%, and negative otherwise. We note that the L_{box} is calculated from both the positive and negative RoIs, while the $L_{superclass}$ and $L_{caption}$ losses are computed only from the positive RoIs. In the second branch, each positive RoI is cloned and fed into the first $LSTM_{roi}$ layer, then the word embedding and the hidden states of the first LSTM layer are combined and fed into the second $LSTM_{word}$ layer. This layer converts the inputs into a sequence of outputted words as the predicted object caption. This process is performed sequentially for each word in the predicted caption until the network generates the end-of-caption (EOC) token.

During the inference phase, only the input image is given to the network. We first select top 300 RoIs produced by RPN as the object candidates. The object detection branch uses these RoIs to assign the object superclass. The results are then suppressed by the non-maximum suppression method [19]. In the captioning branch, all RoIs are also forwarded into two LSTM layers in order to generate the caption for each RoI. The outputted boxes that have the classification score higher than 0.5 and its associated caption are chosen as the final detected objects. In case there are no boxes satisfying this condition, we select the one with highest classification score as the only detected object.

3.3. Object Retrieval

In this task, rather than generating the object caption, we want to retrieve the desired object in the image given a natu-

ral language query. For simplicity, the object is also defined as a rectangle bounding box. The general idea is to select the “best” (i.e., with the highest probability) bounding box from all region proposals. To this end, our goal is similar to [7] [12], however we note that while the authors in [12] focus more on finding the local and global relationship of the query object and other parts of the image, we propose to perform the retrieval task within the concept of the object superclass. In this way, we can train the network end-to-end, while still be able to select the desired object.

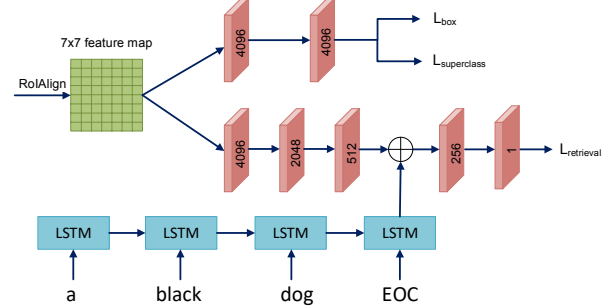


Figure 3. An overview of our object retrieval network (ORN). Each input query is encoded by an LSTM layer, then this feature is combined with the RoI feature in order to regress the retrieval score for this RoI.

Since we need a system that can generate the object proposals, the RPN network is well suitable for our purpose. We first encode the input query using one LSTM layer. Similar to the object captioning task, we feed each positive RoI into a sequence of fully connected layers (with 4096, 2048, 512 neurons, respectively). The feature map of the last fully connected layer with 512 neurons is combined with the output of the LSTM layer to create a visual word. This visual word then goes into another fully connected layer with 256 neurons, then finally the retrieval score is regressed at the last layer with only 1 neuron. Intuitively, this whole process computes a retrieval score for each positive RoI given the input text query. We note that in parallel with the retrieval branch, the object detection branch is also trained as in the object captioning task. Fig. 3 illustrates the details of our object retrieval network.

Multi-task Loss Similar to the object captioning task, we train the network end-to-end with a multi-task loss L function as follows:

$$L = L_{loc} + L_{superclass} + L_{retrieval} \quad (6)$$

where L_{loc} and $L_{superclass}$ loss are identical to the ones in Equation 5. $L_{retrieval}$ is the sigmoid cross entropy loss of the retrieval branch and is defined on the positive RoIs as follows:

$$L_{retrieval} = \sum_{i=0}^{n_{RoI}} \log(1 + \exp(-y_i f(\mathbf{x}_i))) \quad (7)$$

where y is the groundtruth label (retrieval label) of the current positive RoI, and $f(\mathbf{x})$ is predicted output of the retrieval branch of the network.

Training and Inference We generally follow the training procedure in the object captioning task to train our object retrieval network. The key difference between these two networks relies on the second branch. In particular, in the object retrieval task, at each iteration, we randomly select one groundtruth object caption as the input query and feed it into the LSTM network. The output of LSTM is then combined with each positive RoI to compute the $L_{retrieval}$ loss for this RoI. Note that, the retrieval groundtruth score of each RoI is automatically reconstructed during training since we know the current input text query belongs to which object (i.e., the positive RoIs associated with the object of the current query has the retrieval score 1, otherwise 0).

During the testing phase, the inputs for the network are an image and a text query. Our goal is to select the outputted box with the highest retrieval score. To evaluate the result, each object caption groundtruth of the test image is considered as one input query. This query and the test image are forwarded through the network to generate the object bounding boxes and their retrieval scores. Similar to the object captioning task, we select top 300 RoIs as the object candidates. These candidates are pruned by the non-maximum suppression process [19], then the one with highest retrieval score is selected as the desired object. We notice that along with the retrieval score, the network also provides the object superclass and its location via its first branch.

4. EXPERIMENTS

4.1. Dataset

Currently, there are many popular datasets for object detection such as Pascal VOC [20] and MS COCO [21]. However, these datasets only provide the bounding boxes groundtruth [20] or the caption for the entire image [21]. In the field of referring expressions, we also have the ReferIt dataset [22] and the G-Ref dataset [13]. Although these datasets can be used in the object retrieval task, they focus mostly on the context of the object in the image, while we focus on the fine-grained understanding of each object alone. Motivated by these limitations, we introduce a new dataset (*Flick5k*) which contains both the object superclass and its descriptions for the fine-grained object understanding. With our new dataset, we can train the network end-to-end to detect the object, generate its caption, or retrieve the object from an input query.

Table 1. Object Captioning Results on Flick5k Dataset

	Bleu_1	Bleu_2	Bleu_3	Bleu_4	METEOR	ROUGE_L	CIDEr
OCN1_VGG16	0.564	0.366	0.150	0.000	0.213	0.559	0.845
OCN1_ResNet50	0.566	0.361	0.150	0.000	0.215	0.563	0.847
OCN1_ResNet101	0.571	0.369	0.158	0.000	0.221	0.570	0.852
OCN2_VGG16	0.682	0.555	0.469	0.405	0.344	0.675	1.628
OCN2_ResNet50	0.693	0.568	0.468	0.403	0.345	0.681	1.630
OCN2_ResNet101	0.696	0.570	0.469	0.400	0.342	0.685	1.632

In particular, we select 5000 images from the Flick30k dataset [14]. We only reuse the bounding boxes that come with the Flick30k dataset then manually assign the superclass and annotate the object captions. Note that, one bounding box only has one specific superclass, while it can have many captions. Totally, our new Flick5k dataset has 4 object superclasses (people, instruments, animals, vehicles), 7979 object bounding boxes, and 18,214 object captions. The number of bounding boxes for each superclass are 3839, 475, 2387, and 1278 for the people, instruments, animals, vehicles, respectively. We randomly use 50% of the dataset for training and the remaining 50% for testing.

4.2. Implementation

For both two sub-problems, we use the LSTM network with 512 hidden units. The number of LSTM timestep n_t is empirically set to 6. Subsequently, the longer captions are cut from the beginning to the sixth word, while the shorter captions are padded with the EOC word until they reach 6 words. From all the captions, we build a dictionary from the words that appear at least twice, resulting in a dictionary with 866 words. We use the strategy in [2] [4] to resize the input image to (600, 1000) size. The object proposals are generated by the RPN network with 12 anchors (4 scales and 3 aspect ratios). We use three popular convolutional backbones: VGG16 [15], ResNet50 and ResNet101 [23] in our experiments. All the networks are trained end-to-end for 200k iterations. The training time is approximately 2 days on an NVIDIA Titan X GPU.

4.3. Object Captioning Results

Evaluation Protocol Although both the traditional object detection and image captioning tasks have the standard metrics to evaluate the results, the object captioning task is more challenging to evaluate since its output contains many information (i.e., the object class, the location, and its caption). In [13], the authors proposed to use human evaluation, however this approach is not scalable. Unlike the system in [13] which is not end-to-end and only provides the caption for each bounding box without the super-

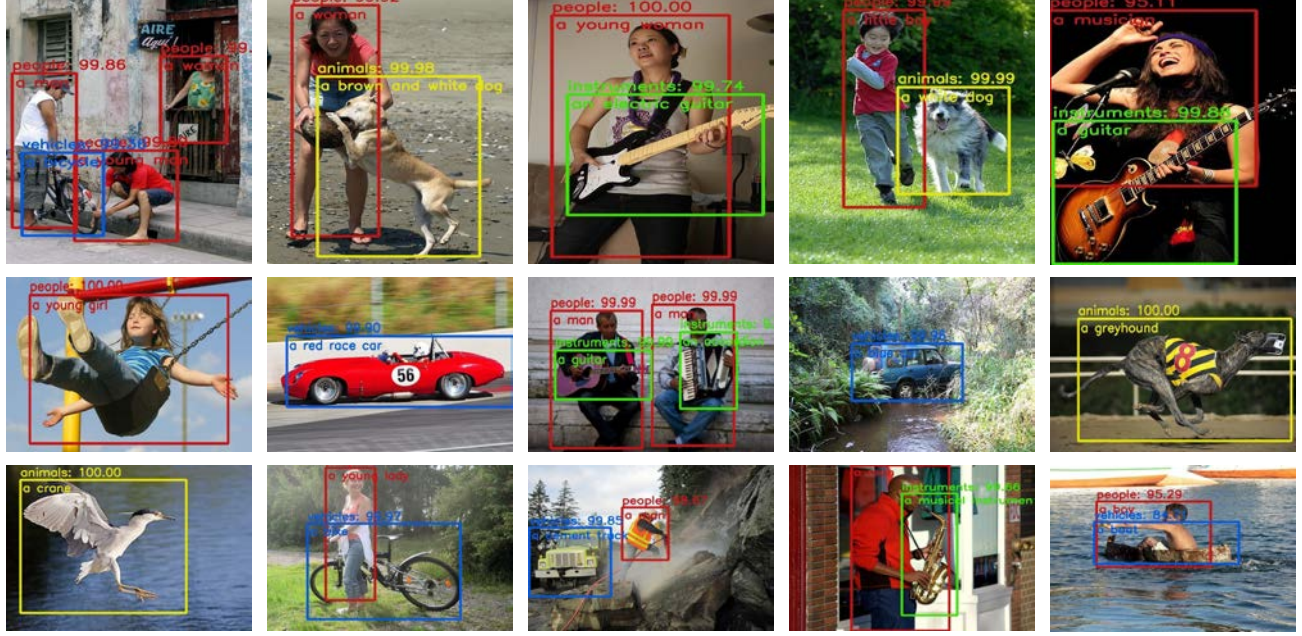


Figure 4. Example of prediction results using our OCN2_ResNet101 network. Besides the correct object detection, the network is able to generate an accurate caption for each object based on its properties (e.g., “a red race car” vs. “a blue car”, etc.).

class information, our network provides all this information. Therefore, we propose to use the standard metrics of image captioning [24] to evaluate the outputted caption of the bounding boxes that have high classification score. This protocol is also widely used in other problems when the network provides both the detected object and its other information [4] [25].

Table 1 summarizes the object captioning results. We compare our object captioning network with two LSTM layers (denotes as OCN2) with the baseline that uses only one LSTM layer (denotes as OCN1, see details in Appendix). Overall, OCN2 clearly outperforms OCN1 by a substantial margin in both three backbones VGG16, ResNet50, and ResNet101. This demonstrates that the way we combine the region feature map with the input caption plays an important role in this task. Intuitively, the approach in OCN2 is more robust than OCN1 since in OCN2 the feature of each positive RoI is combined at every word of the input caption, while in OCN1 the feature is only combined once with the first word. Table 1 also shows the ResNet101 backbone achieves the highest performance and outperforms VGG16 backbone. However, this improvement is not very significant. Fig. 4 shows some example results from our OCN2_ResNet101 network. It is worth noting that the network is able to generate accurate captions within a superclass based on the object properties (e.g., “a red race car” vs. “a blue car”, etc.).

Table 2. Object Retrieval Results on Flickr5k Dataset

	$R@1$
SCRC [12]	68.23%
ORN_VGG16 (ours)	70.61%
ORN_ResNet50 (ours)	74.08%
ORN_ResNet101 (ours)	76.36%

4.4. Object Retrieval Results

Similar to [12], we use the $R@1$ score to evaluate the object retrieval results. The $R@1$ score is the percentage of the predicted box with highest retrieval score being correct. We notice that the predicted box is considered as correct if it has the overlap with the groundtruth box by at least 50% IoU. Table 2 summarizes the object captioning results on the Flickr5k dataset. Overall, our ORN_ResNet101 achieves the highest performance with 76.36% of the input queries has the correct retrieval bounding box. This is a significant improvement over SCRC [12]. While we employ an end-to-end architecture to jointly train both the bounding box location and the input query, in SCRC the bounding boxes are pre-extracted and not trained with the network, hence there are many cases the network does not have the reliable bounding box candidates to regress the retrieval score. Furthermore, the inference time of our ORN network is only around 150ms per query, which is significantly faster than the non end-to-end SCRC approach. Fig. 5 shows some exam-



Figure 5. Example of retrieval results using our ORN_ResNet101 network. The network is able to retrieve the correct object in the challenging cases (e.g., “a black dog” vs. “a spotted dog”). The text query (Q) is in purple at the bottom of the image. The retrieval score (RS) is denoted inside each detected object. Best viewed in color.

ple of retrieval results using our ORN_ResNet101 network. It is worth noting that the network successfully retrieves the correct object in challenging scenarios such as when there are two dogs (“a black dog” vs. “a spotted dog”) in the image.

4.5. Ablation Studies

Object Superclass Unlike the traditional object detection methods [3] [8] which use the normal object categories (e.g., dog, cat, etc.), we train the detection branch using the superclass (e.g., animals, etc.). With this setup, the object detection branch only provides the location and general knowledge of the object, while the fine-grained understanding of the object is given by the second branch. In the main experiment, we classify all objects into 4 superclasses in order to keep the basic knowledge of the object categories. However, in applications that do not require the object category understanding, we can consider all the objects belong to one unique superclass (i.e., the object superclass). To this end, we group all the objects of 4 superclasses into only one object superclass, and then train the captioning and retrieval networks with the ResNet101 backbone as usual.

We follow the same testing procedure as described above. The Bleu_1, Bleu_2, Bleu_3, Bleu_4, METEOR, ROUGE_L and CIDEr scores of the OCN2_ResNet101 network in this experiment are 0.673, 0.544, 0.454, 0.395, 0.330, 0.666, and 1.572, respectively. While the $R@1$ score of the ORN_ResNet101 network is 73.06%. As we ex-

pected, the accuracy of the networks is slightly dropped in comparison with the 4 superclasses setup, but it is still very reasonable. This demonstrates that the object detection branch can be used to just localize the object location, while the fine-grained knowledge of the object can be effectively learned in the captioning/retrieval branch. More importantly, from this experiment we can conclude that the captioning/retrieval results do not strongly depend on the object classification results of the detection branch, but are actually learned by the captioning/retrieval branch. Compared to the dense captioning framework [11] that does not take the object category knowledge into account, or the non end-to-end object retrieval methods [7] [12], our approach provides a flexible yet detailed understanding of the object, while still is able to complete both the captioning and retrieval tasks effectively with fast inference time.

Generalization Although we train both of the OCN and ORN networks on a relatively small training set (i.e., there are only 2500 images in the training set), they still generalize well under challenging testing environments. Fig. 6-a shows a qualitative result when our OCN2_ResNet101 network successfully detects the object and generates its caption from an artwork image. In Fig. 6-b, the ORN_ResNet101 is able to localize the desired object in an image from Gazebo simulation. Besides the generalization ability, the inference time of both networks is only around 150ms per image (or query) on an NVIDIA Titan X GPU, which makes them well suitable for the real-time robotic applications.

Failure Cases Since we use an end-to-end network to simultaneously train the object detection and the captioning/retrieval branch, the outputted results of the second branch strongly depend on the object location given by the object detection branch. Therefore, a typical failure case in our networks is when the object detection branch outputs the incorrect object location. Fig. 7-a and Fig. 7-b show two examples when the detection branch misrecognizes the object (i.e., the dog) or is unable to detect the object (i.e., the bird). Similarly, Fig. 7-c shows a case when the detection branch is unable to provide the object location for the retrieval branch. We notice that, although in this case the object location is wrong, the retrieval branch is able to assign a very low retrieval score to the wrong object, which shows that it is not confident about the final result.



Figure 6. Qualitative results. Our networks generalize well under different testing environments. (a) The OCN2_ResNet101 successfully generates the caption for an object in an artwork image. (b) The ORN_ResNet101 retrieves the correct object in a simulation image.




Figure 7. Some failure cases from the OCN2_ResNet101 network (a, b), and ORN_ResNet101 network (c).

5. Conclusions and Future Work

In this paper, we address the problem of jointly learn vision and language to understand objects in the fine-grained manner. By integrating natural language, we provide a detailed understanding of the object through its caption, while still is able to have the category knowledge from its superclass. Based on the proposed definition, we introduce two deep architectures to tackle two problems: object caption-

ing and object retrieval using natural language. We show that both problems can be effectively solved with the end-to-end hybrid CNN-LSTM networks. The extensive experimental results on our new dataset show that our proposed methods not only achieve the state-of-the-art performance but also generalize well under challenging testing environments and have fast inference time. We plan to release a new large-scale version of our dataset and the full source code of this paper in the future. We hope that these resources can further improve the development of the object captioning and retrieval tasks, making them ready for the real-world robotic applications.

Appendix

The architecture of OCN1 network is as follows:

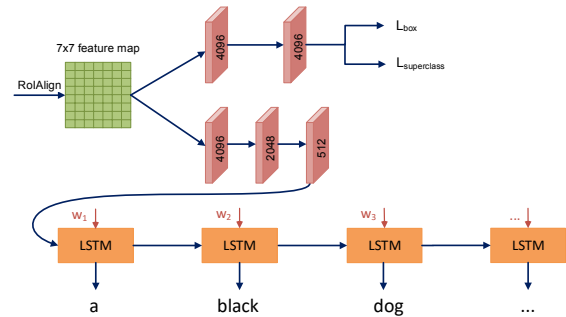


Figure 8. An illustration of the OCN1 network with one LSTM layer.

While our proposed object captioning network with two LSTM layers (Fig. 2) combines each input word with the visual feature, the OCN1 network only combines the first word with the visual feature. The experimental results from Table 1 show that the OCN1 network has poor performance and cannot effectively generate the long caption.

References

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” in *Advances in Neural Information Processing Systems (NIPS)*, 2012. 1
- [2] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, “Mask R-CNN,” in *International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 3, 5
- [3] S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems (NIPS)*, 2015. 1, 2, 3, 7
- [4] T.-T. Do, A. Nguyen, and I. Reid, “Affordancenet: An end-to-end deep learning approach for object

- affordance detection,” in *International Conference Robotics and Automation (ICRA)*, 2018. 1, 2, 3, 5, 6
- [5] J. Redmon and A. Farhadi, “YOLO9000: better, faster, stronger,” *CoRR*, vol. abs/1612.08242, 2016. 1, 2
- [6] I. Biederman, “Recognition-by-components: a theory of human image understanding,” *Psychological review*, 1987. 1
- [7] S. Guadarrama, E. Rodner, K. Saenko, N. Zhang, R. Farrell, J. Donahue, and T. Darrell, “Open-vocabulary object retrieval,” in *Robotics: Science and Systems (RSS)*, 2014. 1, 2, 4, 7
- [8] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, “Ssd: Single shot multi-box detector,” in *European Conference on Computer Vision (ECCV)*, 2016. 2, 7
- [9] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computing*, 1997. 2
- [10] Z. Ren, X. Wang, N. Zhang, X. Lv, and L.-J. Li, “Deep reinforcement learning-based image captioning with embedding reward,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [11] J. Johnson, A. Karpathy, and L. Fei-Fei, “Densecap: Fully convolutional localization networks for dense captioning,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 7
- [12] R. Hu, H. Xu, M. Rohrbach, J. Feng, K. Saenko, and T. Darrell, “Natural language object retrieval,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 4, 6, 7
- [13] J. Mao, J. Huang, A. Toshev, O. Camburu, A. L. Yuille, and K. Murphy, “Generation and comprehension of unambiguous object descriptions,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5
- [14] 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,” *International Journal of Computer Vision (IJCV)*, 2017. 2, 5
- [15] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” *CoRR*, vol. abs/1409.1556, 2014. 2, 5
- [16] V. Ramanishka, A. Das, J. Zhang, and K. Saenko, “Top-down visual saliency guided by captions,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [17] A. Nguyen, D. Kanoulas, L. Muratore, D. G. Caldwell, and N. G. Tsagarakis, “Translating videos to commands for robotic manipulation with deep recurrent neural networks,” in *International Conference Robotics and Automation (ICRA)*, 2018. 3
- [18] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, “Long-term recurrent convolutional networks for visual recognition and description,” in *Computer Vision and Pattern Recognition (CVPR)*, 2014. 4
- [19] R. Girshick, F. Iandola, T. Darrell, and J. Malik, “Deformable part models are convolutional neural networks,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 4, 5
- [20] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes (voc) challenge,” *International Journal of Computer Vision (IJCV)*, 2010. 5
- [21] 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 *European Conference on Computer Vision (ECCV)*, 2014. 5
- [22] S. Kazemzadeh, V. Ordonez, M. Matten, and T. Berg, “Referitgame: Referring to objects in photographs of natural scenes,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014. 5
- [23] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [24] J. Xu, T. Mei, T. Yao, and Y. Rui, “Msr-vtt: A large video description dataset for bridging video and language,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [25] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, “Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes,” *arXiv preprint arXiv:1711.00199*, 2017. 6