Decoupled Novel Object Captioner

Yu Wu¹, Linchao Zhu¹, Lu Jiang², Yi Yang^{1,3}

¹CAI, University of Technology Sydney, ²Google Inc.

³ State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences
yu.wu-3@student.uts.edu.au, yi.yang@uts.edu.au

ABSTRACT

Image captioning is a challenging task where the machine automatically describes an image by sentences or phrases. It often requires a large number of paired image-sentence annotations for training. However, a pre-trained captioning model can hardly be applied to a new domain in which some novel object categories exist, *i.e.*, the objects and their description words are unseen during model training. To correctly caption the novel object, it requires professional human workers to annotate the images by sentences with the novel words. It is labor expensive and thus limits its usage in real-world applications.

In this paper, we introduce the zero-shot novel object captioning task where the machine generates descriptions without extra training sentences about the novel object. To tackle the challenging problem, we propose a Decoupled Novel Object Captioner (DNOC) framework that can fully decouple the language sequence model from the object descriptions. DNOC has two components. 1) A **Sequence Model with the Placeholder** (SM-P) generates a sentence containing placeholders. The placeholder represents an unseen novel object. Thus, the sequence model can be decoupled from the novel object descriptions. 2) A **key-value object memory** built upon the freely available detection model, contains the visual information and the corresponding word for each object. A query generated from the SM-P is used to retrieve the words from the object memory. The placeholder will further be filled with the correct word, resulting in a caption with novel object descriptions. The experimental results on the held-out MSCOCO dataset demonstrate the ability of DNOC in describing novel concepts.

CCS CONCEPTS

• Computing methodologies → Natural language generation;

KEYWORDS

image captioning, novel object, novel object captioning

ACM Reference Format:

Yu Wu, Linchao Zhu, Lu Jiang, and Yi Yang. 2018. Decoupled Novel Object Captioner. In *2018 ACM Multimedia Conference (MM '18)*, October 22–26, 2018, Seoul, Republic of Korea. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3240508.3240640>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '18, October 22–26, 2018, Seoul, Republic of Korea

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5665-7/18/10...\$15.00

<https://doi.org/10.1145/3240508.3240640>

Figure 1: An example of the novel object captioning. The colored bounding boxes show the object detection results. The novel object “zebra” is not present in the training data. LRCN [6] fails to describe the image with the novel object “zebra”. NOC [34] can generate the correct caption but requires extra text training data containing the word “zebra” to learn this concept. Our algorithm can generate correct captions, and *more importantly*, we do not need any extra sentence data. Specifically, we first generate the caption template with a placeholder “<PL>” that represents the novel object. We then fill in the placeholder with the word “zebra” from the object detection model.

1 INTRODUCTION

Image captioning is an important task in vision and language research [15, 27, 37, 40]. It aims at automatically describing an image by natural language sentences or phrases. Recent encoder-decoder architectures have been successful in many image captioning tasks [5, 6, 15, 23, 27, 37, 40], in which the Convolutional Neural Network (CNN) is usually used as the image encoder, and the decoder is usually a Recurrent Neural Network (RNN) to sequentially predict the next word given the previous words. The captioning networks need a large number of image-sentence paired data to train a meaningful model.

These captioning models fail in describing the *novel objects* which are unseen words in the paired training data. For example, as shown in Figure 1, the LRCN [6] model cannot correctly generate captions for the novel object “zebra”. As a result, to apply the model in a new domain where novel objects can be visually detected, it requires professional annotators to caption new images in order to generate

paired training sentences. This is labor expensive and thus limits the applications of captioning models.

A few works have been proposed recently to address the novel object captioning problem [3, 34, 41]. Essentially, these methods attempt to incorporate the class label produced by the pre-trained object recognition model. The class label can be used as the novel object description in language generation. To be specific, Hendricks *et al.* [3] trained a captioning model by leveraging a pre-trained image tagger model and a pre-trained language sequence model from external text corpora. Yao *et al.* [41] exploited a pre-trained sequence model to copy the object detection result into the output sentence.

However, to feed the novel object description into the generated captions, existing approaches either employ the pre-trained language sequence model [3, 34] or require extra unpaired training sentences of the novel object [41]. In both cases, the *novel* objects have been used in training and, hence, is not really novel. A more precise meaning of *novel* in existing works is *unseen* in the *paired* training sentences. However, the word is *seen* in the *unpaired* training sentences. For example, in Figure 1, existing methods require extra training sentences containing the word “zebra” to produce the caption, even though the object zebra is confidently detected in the image. The assumption that such training sentences of the novel object always exist may not hold in many real-world scenarios. It is considerably difficult to collect sentences about brand new products in a timely manner, *e.g.*, self-balancing scooters, robot vacuums, drones, and emerging topics trending on the social media like fidget spinners. Someone may collect these training sentences about novel objects. However, it can still introduce language biases into the captioning model. For example, if training sentences are all about bass (a sea fish), the captioning model will never learn to caption the instrument bass and may generate awkward sentences like “A man is eating a bass with a guitar amplifier.”

In this paper, we tackle the image captioning for novel objects, where we do not need any training sentences containing the object. We utilize a pre-trained object detection model about the novel object. We call it *zero-shot* novel object captioning to distinguish it from the traditional problem setting [3, 34, 41]. In the traditional setting, in addition to the pre-trained object detection model, extra training sentences of the novel object are provided. In the zero-shot novel object captioning, there are *zero training sentences* about the novel object, *i.e.*, there is no information about the semantic meaning, sense, and context of the object. As a result, existing approaches of directly training the word embedding and sequence model become infeasible.

To address this problem, we propose a Decoupled Novel Object Captioner (DNOC) framework that is able to generate natural language descriptions without extra training sentences of the novel object. DNOC follows the standard encoder-decoder architecture but with an improved decoder. Specifically, we first design a sequence model with the placeholder (SM-P) to generate captions with placeholders. The placeholder represents the unseen word for a novel object. Then we build a key-value object memory for each image, which contains the visual information and the corresponding words for objects. Finally, a query is generated to retrieve a value from the key-value object memory and the placeholder is filled by the corresponding word. In this way, the sequence model

is fully decoupled from the novel object descriptions. Our DNOC is thus capable of dealing with the unseen novel object. For example, in Fig. 1, our method first generates the captioning sentence by generating a placeholder “<PL>” to represent any novel object. Then it learns to fill in the placeholder with “zebra” based on the visual object detection result.

In summary, the main contributions of this work are listed as follows:

- We introduce the zero-shot novel object captioning task, which is an important research direction but has been largely ignored.
- To tackle the challenging task, we design the sequence model with the placeholder (SM-P). SM-P generates a sentence with placeholders to fully decouple the sequence model from the novel object descriptions.
- A key-value object memory is introduced to incorporate external visual knowledge. It interactively works with SM-P to tackle the zero-shot novel object captioning problem. Extensive experimental results on the held-out MSCOCO dataset show that our DNOC is effective for zero-shot novel object captioning. Without extra training data, our model significantly outperforms state-of-the-art methods (with additional training sentences) on the F1-score metric.

2 RELATED WORK

The problem of generating natural language descriptions from visual data has been a popular research area in the computer vision field. This section first reviews recent works for image captioning and then extends the discussion to the zero-shot novel object captioning task.

2.1 Image Captioning

Image captioning aims at automatically describing the content of an image by a complete and natural sentence. This is a fundamental problem in artificial intelligence that connects computer vision and natural language processing [38]. Some early works such as template-based approaches [18, 25] and search-based approaches [8, 26] generate captioning by the sentence template and the sentence pool. Recently, language-based models have achieved promising performance. Most of them are based on the encoder-decoder architecture to learn the probability distribution of both visual embedding and textual embeddings [5–7, 15, 17, 23, 27, 37, 40, 43]. In this architecture, the encoder is a CNN model which processes and encodes the input image into an embedding representation, while the decoder is a RNN model that takes the CNN representation as the initial input and sequentially predicts the next word given the previous word. Among recent contributions, Kiros *et al.* [17] proposed a multi-modal log-bilinear neural language model to jointly learn word representations and image feature embeddings. Vinyals *et al.* [37] proposed an end-to-end neural network consisting of a vision CNN followed by a language generating RNN. Xu *et al.* [40] improved [37] by incorporating the attention mechanism into captioning. The attention mechanism focuses on the salient image regions when generating corresponding words. You *et al.* [42] further utilized an independent high-level concepts/attribute detector to improve the attention mechanism for image captioning.

Tavakoliy *et al.* [33] proposed a saliency-boosted image captioning model in order to investigate benefits from low-level cues in language models.

2.2 Zero-Shot Novel Object Captioning

Zero-shot learning aims to recognize objects whose instances may not have been seen during training [12, 13, 19, 29, 39]. Zero-shot learning bridge the gap between the visual and the textual semantics by learning a dictionary of concept detectors on external data source.

Novel object captioning is a challenging task where there is no paired visual-sentence data for the novel object in training. Only a few works have been proposed to address this captioning problem. Mao *et al.* [22] tried to describe novel objects with only a few paired image-sentence data. Hendricks *et al.* [3] proposed the Deep Compositional Captioner (DCC), a pilot work to address the task of generating descriptions of novel objects which are not present in paired image-sentence datasets. DCC leverages a pre-trained image tagger model from large object recognition datasets and a pre-trained language sequence model from external text corpora. The captioning model is trained on the paired image-sentence data with above pre-trained models. Venugopalan *et al.* [34] discussed a Novel Object Captioner (NOC) to further improve the DCC to an end-to-end system by jointly training the visual classification model, language sequence model, and the captioning model. Anderson *et al.* [2] leveraged an approximate search algorithm to forcibly guarantee the inclusion of selected words during the evaluation stage of a caption generation model. Yao *et al.* [41] exploited a mechanism to copy the detection results to the output sentence with a pre-trained language sequence model. Concurrently to us, Lu *et al.* [21] also proposed to generate a sentence template with slot locations, which are then filled in by visual concepts from object detectors. However, when captioning the novel objects, they have to manually defined category mapping list to replace the novel object's word embedding with an existing one.

Note that all of the above methods have to use extra data of the novel object to train their word embedding. Different from existing methods, our method focuses on zero-shot novel object captioning task in which there are no additional sentences or pre-trained models to learn such embeddings for novel objects. Our method thus needs to introduce a new approach to exploit the object word's meaning, sense, and embedding in the zero-shot training condition.

3 THE PROPOSED METHOD

In this section, we first introduce the preliminaries and further show the two key parts of the proposed Decoupled Novel Object Captioner, *i.e.*, the sequence model with the placeholder (Section 3.2) and the key-value object memory (Section 3.3). The overview of the DNOC framework and training details are illustrated in Section 3.4 and Section 3.5, respectively.

3.1 Preliminaries

We first introduce the notations for image captioning. In image captioning, given an input image I , the goal is to generate an associated natural language sentence s of length n_I , denoted as $s =$

$(w_1, w_2, \dots, w_{n_I})$. Each w represents a word and the length n_I is usually varied for different sentences. Let $\mathcal{P} = \{(I_1, s_1), \dots, (I_{n_p}, s_{n_p})\}$ be the set with n_p image-sentences pairs. The vocabulary of $\mathcal{P}$ is $\mathcal{W}_{paired} = \{w_1, w_2, \dots, w_{N_I}\}$ which contains N_I words. Each word $w_i \in \{0, 1\}^{N_I}$ is a one-hot (1 of N_I) encoding vector. The one-hot vector is then embedded into a D_w -dimensional real-valued vector $x_i = \phi_w(w_i) \in \mathbb{R}^{D_w}$. The embedding function $\phi_w(\cdot)$ is usually a trainable linear transformation $x_i = T_w w_i$, where $T_w \in \mathbb{R}^{D_w \times N_I}$ is the embedding matrix. The typical architecture for captioning is the encoder-decoder model. In the followings, we show the encoding and the decoding procedures during the testing phase.

The encoder. We obtain the representation for an input image I by $\phi_e(I)$, where $\phi_e(\cdot)$ is the embedding function for encoder. The function ϕ_e is usually an ImageNet pre-trained CNN model with the classification layer removed. It extracts the top-layer outputs as the visual features.

The decoder. The decoder is a word-by-word sequence model designed to generate the sentence given the encoder outputs. In specific, at the first time step $t = 0$, a special token w_0 ($\langle GO \rangle$) is the input to the sequence model, which indicates the start of the sentence. At time step t , the decoder generates a word w_t given the visual content $\phi_e(I)$ and previous words $(w_0, \dots, w_{t-1})$. Therefore, we formulate the probability of generating the sentence s as

$$p(s|I) = \prod_{t=1}^{n_I} p(w_t | w_0, \dots, w_{t-1}, \phi_e(I)). \quad (1)$$

The Long Short-Term Memory (LSTM) [10] is commonly used as the decoder in visual captioning and natural language processing tasks [3, 35, 41]. The core of the LSTM model is the memory cell, which encodes the knowledge of the input that has been observed at every time step. There are three gates that modulate the memory cell updating, *i.e.*, the input gate, the forget gate, and the output gate. These three gates are all computed by the current input x_t and the previous hidden state h_{t-1} . The input gate controls how the current input should be added to the memory cell; the forget gate is used to control what the cell should forget from the previous memory; the output gate controls whether the current memory cell should be passed as output. Given inputs x_t, h_{t-1} , we get the predicted output word o_t by updating the LSTM unit at time step t as follows:

$$o_t, h_t = \text{LSTM}(x_t, h_{t-1}). \quad (2)$$

Combining Eqn (1) and Eqn (2), at t -th time step, the previous word w_{t-1} is taken as the decoder input x_t . In the training stage, we feed the ground-truth word of the previous step as the model input. In the evaluation stage, we take the output o_{t-1} of the model at $(t-1)$ -th time step as model input x_t .

Zero-Shot Novel Object Captioning. This paper studies the zero-shot novel object captioning task, where the model needs to caption novel objects *without* additional training sentence data about the object. The novel object words are shown neither in the paired image-sentence training data $\mathcal{P}$ nor unpaired sentence training data. We denote $\mathcal{W}_{unseen}$ as the vocabulary for the novel object words which are unseen in training. Given an input image I_n containing novel objects, the captioning model should generate a sentence with the corresponding unseen word $\tilde{w} \in \mathcal{W}_{unseen}$ to describe the novel objects.

Figure 2: The comparison of the typical sequence model and the proposed SM-P. In this example, “zebra” is an unseen word during training. The bottom are the sentences generated by the two models. The left is the classical sequence model, which cannot handle the input out-of-vocabulary (OOV) word “zebra”. The right is our sequence model with the placeholder (SM-P). It generates the special token word “<PL>” (placeholder) to represent the novel object, and is able to continue to output the subsequent word given the input “<PL>”.

A notable challenge for this task is to deal with the out-of-vocabulary (OOV) words. The learned word embedding function ϕ_w is unable to encode the unseen words, since these words cannot simply be found in $\mathcal{W}_{paired}$. As a result, these unseen words cannot be fed into the decoder for caption generation. Previous works [3, 34, 41] circumvented this problem by learning the word embeddings of unseen words using additional sentences that contain the words. We denote these extra training sentences as $\mathcal{S}_{unpaired}$. Since the words of “novel” objects have been used in training, the “novel” objects are not really novel. However, in our zero-shot novel object task, we do not assume the availability of additional training sentences $\mathcal{S}_{unpaired}$ of the novel object. Therefore, we propose a novel approach to deal with the OOV words in the sequence model.

3.2 Sequence Model with the Placeholder

We propose the Sequence Model with the placeholder (SM-P) to fully decouple the sequence model from novel object descriptions. As discussed above, the classical sequence model cannot take an out-of-vocabulary word as input. To solve this problem, we design a new token, denoted as “<PL>”. “<PL>” is the *placeholder* that represents any novel words $\tilde{w} \in \mathcal{W}_{unseen}$. It is used in the decoder similarly to other tokens, such as “<GO>”, “<PAD>”, “<EOS>”, “<UNKNOWN>” in most natural language processing works. We add the token “<PL>” into the paired vocabulary $\mathcal{W}_{paired}$ to learn the embedding. The training details for the placeholder are discussed in Section 3.5.

A new embedding ϕ_w is learned for “<PL>”, which encodes all unknown words with a compact representation. The new representation could be jointly learned with known words. We carefully

designed the new token “<PL>” in both the input and the output of the decoder, which enables us to handle the out-of-vocabulary words. When the decoder outputs “<PL>”, our model utilizes the external knowledge from the object detection model to replace it with novel description. Our SM-P is flexible that can be readily incorporated in the sequence to sequence model. Without loss of generality, we use the LSTM as the backbone of our SM-P. When the SM-P decides to generate a word about a novel object at time step t , it will output a special word w_t , the token “<PL>”, as the output. At time step $t + 1$, the SM-P takes the previous output word w_t “<PL>” as input instead of the novel word $\tilde{w}$. In this way, regardless of the existence of unseen words, the word embedding function ϕ_w is able to encode all the input tokens. For example, in Figure 2, the classical sequence model cannot handle the out-of-vocabulary word “zebra” as input. Instead, the SM-P model outputs the “<PL>” token when it needs to generate a word about the novel object “zebra”. This token is further fed to the decoder at the next time step. Thus, the subsequent words can be generated. Finally, the SM-P generates the sentence with the placeholder “A <PL> is standing ...”. The “<PL>” token will be replaced by the novel word generated by the key-value object memory.

3.3 Key-Value Object Memory

To incorporate the novel words into the generated sentences with the placeholder, we exploit a pre-trained object detection model to build the key-value object memory.

A freely available object detection model is applied on the input images to find novel objects. For the i -th detected object obj_i , the CNN feature representations $\mathbf{f}_i \in \mathbb{R}^{1 \times N_f}$ and the predicted semantic class label $\mathbf{l}_i \in \mathbb{R}^{1 \times N_D}$ form a key-value pair, with the feature as key and the semantic label as the value. N_f is the feature dimension of CNN representation, N_D is the number of total detection classes. The key-value pairs associate the semantic class labels (descriptions of the novel objects) with their appearance feature. Following [14], we extract the CNN feature $\mathbf{f}_i$ for obj_i from the ROI pooling layer of the detection model. Among all the detected results, the top N_{det} key-value pairs are selected according to their confidence scores, which form the key-value object memory $\mathcal{M}_{obj}$. For each input image, the memory $\mathcal{M}_{obj}$ is initialized to be empty.

Let $\mathcal{W}_{det}$ be the vocabulary of the detection model, which consists of N_D detection class labels. Note that each word in $\mathcal{W}_{det}$ is the detection class label in the *one-hot* format, since we cannot obtain trained word embedding function $\phi_w(\cdot)$ for the new word. To generate the novel word and replace the placeholder in the sentence at time step t , we define the query $\mathbf{q}_t$ to be a linear transformation of previous hidden state $\mathbf{h}_{t-1}$ when the model meets the special token “<PL>” at time step t :

$$\mathbf{q}_t = \mathbf{w}\mathbf{h}_{t-1}, \quad (3)$$

where $\mathbf{h}_{t-1} \in \mathbb{R}^{N_h}$ is the previous hidden state at $(t-1)$ -th step from the sequence model, and $\mathbf{w} \in \mathbb{R}^{N_f \times N_h}$ is the linear transformation to convert the hidden state from semantic feature space to CNN appearance feature space. We have the following operations on the key-value memory $\mathcal{M}_{obj}$:

$$\mathcal{M}_{obj} \leftarrow \text{WRITE}(\mathcal{M}_{obj}, (\mathbf{f}_i, \mathbf{l}_i)), \quad (4)$$

$$\mathbf{w}_{obj} \leftarrow \text{READ}(\mathbf{q}, \mathcal{M}_{obj}). \quad (5)$$

Figure 3: The overview of the DNOC framework. We design a novel sequence model with the placeholder (SM-P) to handle the unseen objects by replacing them with the special token “<PL>”. The SM-P first generates a sentence with placeholders, which refer to the unknown objects in an image. For example, in this figure, the “dog” and “cake” are unseen in the training set. The SM-P generates the sentence “a <PL> is looking at a <PL>”. Meanwhile, we exploit a freely available object detection model to build a key-value object memory, which associate the semantic class labels (descriptions of the novel objects) with their appearance feature. When SM-P generates a placeholder, we take the linear transformation of the previous hidden state as a query to read the memory and output the correct object description, e.g., “dog” and “cake”. Finally, we replace the placeholders by the query results and generate the sentence with novel words.

- **WRITE** operation is to write the input key-value pair (f_i, l_i) into the existing memory $\mathcal{M}_{obj}$. The input key-value pair is written to a new slot of the memory. The key-value object memory is similar to the support set widely used in many few-shot works [9, 30, 36]. The difference is that in their work the key-value memory is utilized for long-term memorization, while our motivation is to build a structured mapping from the detection bounding-box-level feature to its semantic label.
- **READ** operation takes the query q as input, and conducts content-based addressing on the object memory $\mathcal{M}_{obj}$. It aims to find related object information according to the similarity metric, qK^T . The output of READ operation is,

$$w_{obj} = (qK^T)V, \quad (6)$$

where $K^T \in \mathbb{R}^{N_f \times N_{det}}$, $V \in \mathbb{R}^{N_{det} \times N_D}$ are the vertical concatenations of all keys and values in the memory, respectively. The output $w_{obj} \in \mathbb{R}^{N_D}$ is the combination of all semantic labels. In evaluation, the word with the max prediction is used as the query result.

3.4 Framework Overview

With the above two components, we propose the DNOC framework to caption images with novel objects. The framework is based on the encoder-decoder architecture with the SM-P and key-value object memory. For an input image with novel objects, we have the following steps to generate the captioning sentence:

- We first exploit the SM-P to generate a captioning sentence with some placeholders. Each placeholder represents an unseen word/phrase for a novel object;
- We then build a key-value object memory $\mathcal{M}_{obj}$ for each input based on the detection feature-label pairs $\{f_i, l_i\}$ on the image;
- Finally, we replace the placeholders of the sentence by corresponding object descriptions. For the placeholder generated at time step t , we take the previous hidden state h_{t-1} from SM-P as a query to read the object memory $\mathcal{M}_{obj}$, and replace the placeholder by the query results w_{obj} .

In the example shown in Figure 3, the “dog” and “cake” are the novel objects which are not present in training. The SM-P first generates a sentence “a <PL> is looking at a <PL>”. Meanwhile, we build the key-value object memory $\mathcal{M}_{obj}$ based on the detection results, which contains both the visual information and the corresponding word (the detection class label). The hidden state at the step before each placeholder is used as the query to read from the memory. The memory will then return the correct object description, i.e., “dog” and “cake”. Finally, we replace the placeholders by the query results and thus generate the sentence with novel words “a dog is looking at a cake”.

3.5 Training

To learn how to exploit the “out-of-vocabulary” words, we modify the input and target for SM-P in training. We define $\mathcal{W}_{pd}$ as the intersection set of the vocabulary $\mathcal{W}_{paired}$ and vocabulary $\mathcal{W}_{det}$,

$$\mathcal{W}_{pd} = \mathcal{W}_{paired} \cap \mathcal{W}_{det}. \quad (7)$$

$\mathcal{W}_{pd}$ contains the words in the paired visual-sentence training data and the labels of the pre-trained detection model. For the i -th paired visual-sentence input (I_i, s_i) , we first encode the visual input by the encoder $\phi_e(\cdot)$. We modify the input annotation sentence $s_i = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{n_i})$ of the sequence model SM-P by replacing each word $\mathbf{w}_i \in \mathcal{W}_{pd}$ with the token " $\langle PL \rangle$ ". The new input word $\hat{\mathbf{w}}_t$ at t -th time step is,

$$\hat{\mathbf{w}}_t = \begin{cases} \langle PL \rangle, & \mathbf{w}_t \in \mathcal{W}_{pd} \\ \mathbf{w}_t, & \text{otherwise.} \end{cases} \quad (8)$$

We replace some known objects with the placeholder to help the SM-P learn to output the placeholder token at the correct place, and train the key-value object memory to output the correct word given the query. The actual input sentence for the SM-P is,

$$\hat{\mathbf{s}}_i = (\hat{\mathbf{w}}_0, \hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_{n_i}). \quad (9)$$

We take the $\hat{\mathbf{s}}_i$ as the optimizing target for SM-P. Let $F_{SM}(\cdot)$ denote the function of SM-P, and θ_{SM-P} denote its parameters. The output of function $F_{SM}(\cdot)$ is the probability of next word prediction. Each step of word generation is a word classification on existing vocabulary $\mathcal{W}_{paired}$. SM-P is trained to predict the next word $\hat{\mathbf{w}}_t$ given $\phi_e(I)$ and sequence of words $(\hat{\mathbf{w}}_0, \hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_{t-1})$. The optimizing loss for $\mathcal{L}_{SM-P}$ is,

$$\begin{aligned} \mathcal{L}_{SM-P}(\hat{\mathbf{w}}_0, \hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_{t-1}, \phi_e(I); \theta_{SM-P}) = \\ - \sum_t \log(\text{softmax}_t(F_{SM}(\hat{\mathbf{w}}_0, \hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_{t-1}, \phi_e(I); \theta_{SM-P}))), \end{aligned} \quad (10)$$

where the softmax_t denotes the softmax operation on the t -th step.

For the key-value object memory $\mathcal{M}_{obj}$, we define the optimizing loss by comparing the query result $\mathbf{w}_{obj_t}$ from object memory and the word $\mathbf{w}_t$ from annotation,

$$\mathcal{L}_{\mathcal{M}_{obj}} = \sum_t a_t CE(\mathbf{w}_{obj_t}, \mathbf{w}_t), \quad (11)$$

where $CE(\cdot)$ is the cross-entropy loss function, and a_t is the weight at time step t that is calculated by,

$$a_t = \begin{cases} 1, & \mathbf{w}_t \in \mathcal{W}_{pd} \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

There are two trainable components in optimizing Eqn. (11). One is the query q , the hidden state from the LSTM model. The other is the linear transformation on detection features in the computation of the memory key. We simultaneously minimize the two loss functions. The final objective function for the DNOC framework is,

$$\mathcal{L} = \mathcal{L}_{SM-P} + \mathcal{L}_{\mathcal{M}_{obj}}. \quad (13)$$

4 EXPERIMENTS

We first discuss the experimental setups and then compare DNOC with the state-of-the-art methods on the held-out MSCOCO dataset. Ablation studies and qualitative results are provided to show the effectiveness of DNOC.

4.1 The held-out MSCOCO dataset

The MSCOCO dataset [20] is a large scale image captioning dataset. For each image, there are five human-annotated paired sentence descriptions. Following [3, 34, 41], we employ the subset of the MSCOCO dataset, but excludes all image-sentence paired captioning annotations which describe at least one of eight MSCOCO

objects. The eight objects are chosen by clustering the vectors from the word2vec embeddings over all the 80 objects in MSCOCO segmentation challenge [3]. It results in the final eight novel objects for evaluation, which are "bottle", "bus", "couch", "microwave", "pizza", "racket", "suitcase", and "zebra". These novel objects are held-out in the training split and appear only in the evaluation split. We use the same training, validation and test split as in [3].

4.2 Experimental Settings

The Object Detection Model. We employ a freely available pre-trained object detection model to build the key-value object memory. Specifically, we use Faster R-CNN [28] model with Inception-ResNet-V2 [32] to generate detection bounding boxes and scores. The object detection model is pre-trained on all the MSCOCO training images of 80 objects, including the eight novel objects. We use the pre-trained models released by [11] which are publicly available. For each image, we write the top $N_{det} = 4$ detection results to the key-value object memory.

Evaluation Metrics. Metric for Evaluation of Translation with Explicit Ordering (METEOR) [4] is an effective machine translation metric which relies on the use of stemmers, WordNet [24] synonyms and paraphrase tables to identify matches between candidate sentence and reference sentences. However, as pointed in [3, 34, 41], the METEOR metric is not well designed for the novel object captioning task. It is possible to achieve high METEOR scores even without mentioning the novel objects. Therefore, to better evaluate the description quality, we also use the F1-score as an evaluation metric following [3, 34, 41]. F1-score considers false positives, false negatives, and true positives, indicating whether a generated sentence includes a new object.

Implementation Details. Following [3, 34, 41], we use a 16-layer VGG [31] pre-trained on the ImageNet ILSVRC12 dataset [25] as the visual encoder. The CNN encoder is fixed during model training. The decoder is an LSTM with cell size 1,024 and 15 sequence steps. For each input image, we take the output of the fc7 layer from the pre-trained VGG-16 model with 4,096 dimensions as the image representation. The representations are processed by a fully-connected layer and then fed to the decoder SM-P as the initial state. For the word embedding, unlike [3, 41], we do not exploit the pre-trained word embeddings with additional knowledge data. Instead, we learn the word embedding ϕ_w with 1,024 dimensions for all words including the placeholder token. We implement our DNOC model with TensorFlow [1]. Our DNOC is optimized by ADAM [16] with the learning rate of 1×10^{-3} . The weight decay is set to 5×10^{-5} . We train the DNOC for 50 epochs and choose the model with the best validation performance for testing.

4.3 Comparison to the state-of-the-art results

We compare our DNOC with the following state-of-the-art methods on the held-out MSCOCO dataset.

(1). *Long-term Recurrent Convolutional Networks (LRCN)* [6]. LRCN is one of the basic RNN-based image captioning models. Since it has no mechanism to deal with novel objects, we train LRCN only on the paired visual-sentence data.

(2). *Deep Compositional Captioner (DCC)* [3]. DCC leverages a pre-trained image tagger model from large object recognition datasets

Table 1: The comparison with the state-of-the-art methods on the eight novel objects in the held-out MSCOCO dataset. Per-object F1-score, averaged F1-score and METEOR score are reported. All the results are reported using VGG-16 [31] feature and without beam search. Note that we adopt the zero-shot novel object captioning setting where no additional language data is used in training. With few training data, we achieve a higher average F1-score than all the previous methods with external sentence data. All F1-score values are reported as percentage (%).

Settings	Methods	F _{bottle}	F _{bus}	F _{couch}	F _{microwave}	F _{pizza}	F _{racket}	F _{suitcase}	F _{zebra}	F _{average}	METEOR
With External Semantic Data	DCC [3]	4.63	29.79	45.87	28.09	64.59	52.24	13.16	79.88	39.78	21
	NOC [34]										
	–(One hot)	16.52	68.63	42.57	32.16	67.07	61.22	31.18	88.39	50.97	20.7
	–(One hot+Glove)	14.93	68.96	43.82	37.89	66.53	65.87	28.13	88.66	51.85	20.7
	LSTM-C[41]										
	–(One hot)	29.07	64.38	26.01	26.04	75.57	66.54	55.54	92.03	54.40	22
	–(One hot+Glove)	29.68	74.42	38.77	27.81	68.17	70.27	44.76	91.4	55.66	23
Zero-shot	NBT+G [34]	7.1	73.7	34.4	61.9	59.9	20.2	42.3	88.5	48.5	22.8
	LRCN [6]	0	0	0	0	0	0	0	0	0	19.33
	DNOC w/o object memory	26.91	57.14	46.05	41.88	58.50	18.41	48.04	75.17	46.51	20.41
	DNOC (ours)	33.04	76.87	53.97	46.57	75.82	32.98	59.48	84.58	57.92	21.57

Table 2: The comparison of our method and the baseline LRCN [6] on the six known objects in the held-out MSCOCO dataset. Per-object F1-scores and averaged F1-scores are reported as percentage (%).

Methods	F _{bear}	F _{cat}	F _{dog}	F _{elephant}	F _{horse}	F _{motorcycle}	F _{average}
LRCN [6]	66.23	75.73	53.62	65.49	55.20	71.45	64.62
DNOC	62.86	87.28	71.57	77.46	71.20	77.59	74.66

and a pre-trained language sequence model from external text corpora. The captioning model is trained on the paired image-sentence data with the two pre-trained models.

(3). *Novel Object Captioner (NOC)* [34]. NOC improves the DCC to an end-to-end system by jointly training the visual classification model, language sequence model, and the captioning model.

(4). *LSTM-C* [41]. LSTM-C leverages a copying mechanism to copy the detection results to the output sentence with a pre-trained language sequence model.

(5). *Neural Baby Talk (NBT)* [21]. NBT incorporates visual concepts from object detectors to the sentence template. They manually define category mapping list to replace the novel object’s word embedding with an existing one to incorporate the novel words.

Table 1 summarizes the F1 scores and METEOR scores of the above methods and our DNOC on the held-out MSCOCO dataset. All the state-of-the-art methods except LRCN use additional semantic data containing the words of the eight novel objects. Nevertheless, without external sentence data, our method achieves competitive performance to the state of the art. Our model on average yields a higher F1-score than the best state-of-the-art result (57.92% versus 55.66%). The improvement is significant considering our model uses no additional training sentences. Our METEOR score is slightly worse than the LSTM-C with GloVe [41]. One possible reason is that much more training sentences containing the novel words are used to learn the LSTM-C model.

Table 2 shows the comparison of our DNOC and the baseline method LRCN on the known objects in the held-out MSCOCO dataset. Our method achieves much higher F1-scores on the known

Table 3: Ablation studies in terms of Averaged F1-score and METEOR score on the held-out MSCOCO. “LRCN” is the baseline method. “DNOC w/o detection model” indicates the DNOC framework without any detected objects as input. “DNOC w/o object memory” indicates the DNOC framework with SM-P but without the key-value object memory.

Model	F _{average}	METEOR
LRCN [6]	0	19.33
DNOC w/o detection model	0	17.52
DNOC w/o object memory	46.51	20.41
DNOC	57.92	21.57

objects than LRCN, indicating that DNOC also benefits the captioning on the known objects. DNOC enhances the ability of the model to generate sentences with the objects shown in the image. The results strongly support the validity of the proposed model in both the novel objects and the known objects.

4.4 Ablation Studies

The ablation studies are designed to evaluate the effectiveness of each component in DNOC.

The effectiveness of SM-P and key-value object memory.

We conduct the ablation studies on the held-out MSCOCO dataset. The results are shown in Table 3. The “DNOC w/o detection model” indicates the DNOC framework with SM-P but without any detection objects as input. All the objects in the visual inputs will not be detected. Thus, the placeholder token remains in the final generated sentence. We observe a performance drop compared to LRCN. “DNOC w/o object memory” indicates the DNOC framework with SM-P but without the key-value object memory. In this experiment, we take the top N_{det} confident detected labels and randomly feed them into the placeholder. Only with the SM-P component, “DNOC w/o object memory” outperforms LRCN by 46.51% in F1-score and 1.08% in METEOR score, which shows the effectiveness of the SM-P component. Our full DNOC framework outperforms “DNOC w/o

Figure 4: Qualitative results for the held-out MSCOCO dataset. The words in pink are not present during training. “Detected” shows the object detection results. “GT” and “LRCN” are the human-annotated sentences and the sentences generated by LRCN, respectively. “SM-P” indicates the sentence generated by SM-P (the first step of DNOC). The SM-P first generates a sentence template with placeholder, and DNOC further feeds the detection results into the placeholder.

Figure 5: The performance curve with different number of selected detected objects N_{det} .

object memory” by 11.41% in F1-score and 1.16% in METEOR score. It validates that the key-value object memory can enhance the semantic understanding of the visual content. From the experimental results, we can conclude that our full DNOC framework with SM-P and the key-value object memory greatly improves the performance in both the F1-score and the METEOR score. It shows that the two components are effective in exploiting the external detection knowledge.

Analysis of the number of selected detection objects N_{det} . N_{det} is the number of selected top detection results. We show the performance curves with different N_{det} values in Figure 5. When N_{det} varies in a range from two to ten, we can see that the curves of F1-score and METEOR are relatively smooth. When we only adopt one detected object to build the object memory $\mathcal{M}_{obj}$, the F1-score significantly drops. If we do not write any detection results into

the memory, the F1-score is zero. This curve also demonstrates the effectiveness of the key-value object memory.

4.5 Qualitative Result

In Figure 4, we show some qualitative results on the held-out MSCOCO dataset. We take the “zebra” image in the first row as an example. The classic captioning model LRCN [6] could only describe the image with a wrong word “giraffe”, where “zebra” does not exist in the vocabulary. SM-P first generates the sentence with the placeholder, where each placeholder represents a novel object. The DNOC then replaces the placeholders with the detection results by querying the key-value object memory. It generates the sentence with the novel word “zebra” in the correct place. As can be seen, zero-shot novel object captioning is a very challenging task since the evaluation examples contain unseen objects and no additional sentence data is available.

5 CONCLUSIONS

In this paper, we tackle the novel object captioning under a challenging condition where none sentence of the novel object is available. We propose a novel Decoupled Novel Object Captioner (DNOC) framework to generate natural language descriptions of the novel object. Our experiments validate its effectiveness on the held-out MSCOCO dataset. The comprehensive experimental results demonstrate that DNOC outperforms the state-of-the-art methods for captioning novel objects.

Acknowledgment. This work is partially supported by an Australian Research Council Discovery Project. We acknowledge the Data to Decisions CRC (D2D CRC) and the Cooperative Research Centres Programme for funding this research.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, Vol. 16. 265–283.
- [2] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2017. Guided open vocabulary image captioning with constrained beam search. In *EMNLP*.
- [3] Lisa Anne Hendricks, Subhashini Venugopalan, Marcus Rohrbach, Raymond Mooney, Kate Saenko, Trevor Darrell, Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, et al. 2016. Deep compositional captioning: Describing novel object categories without paired training data. In *CVPR*.
- [4] Satantjeet Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *ACL-W*. 65–72.
- [5] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*. 1171–1179.
- [6] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*. 2625–2634.
- [7] Xuanyi Dong, Linchao Zhu, De Zhang, Yi Yang, and Fei Wu. 2018. Fast Parameter Adaptation for Few-shot Image Captioning and Visual Question Answering. In *ACM on Multimedia*.
- [8] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. 2010. Every picture tells a story: Generating sentences from images. In *ECCV*. 15–29.
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML*. 1126–1135.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [11] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, et al. 2017. Speed/accuracy trade-offs for modern convolutional object detectors. In *CVPR*.
- [12] Lu Jiang, Shou-I Yu, Deyu Meng, Teruko Mitamura, and Alexander G Hauptmann. 2015. Bridging the ultimate semantic gap: A semantic search engine for internet videos. In *ICMR*. 27–34.
- [13] Lu Jiang, Shou-I Yu, Deyu Meng, Yi Yang, Teruko Mitamura, and Alexander G Hauptmann. 2015. Fast and accurate content-based semantic search in 100m internet videos. In *ACM on Multimedia*. 49–58.
- [14] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. 2016. Densecap: Fully convolutional localization networks for dense captioning. In *CVPR*. 4565–4574.
- [15] Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *CVPR*. 3128–3137.
- [16] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- [17] Ryan Kiros, Ruslan Salakhutdinov, and Rich Zemel. 2014. Multimodal neural language models. In *ICML*. 595–603.
- [18] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. 2013. Babytalk: Understanding and generating simple image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 12 (2013), 2891–2903.
- [19] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. 2014. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 3 (2014), 453–465.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *ECCV*. 740–755.
- [21] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2018. Neural Baby Talk. In *CVPR*. 7219–7228.
- [22] Junhua Mao, Xu Wei, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan L Yuille. 2015. Learning like a child: Fast novel visual concept learning from sentence descriptions of images. In *ICCV*. 2533–2541.
- [23] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. 2015. Deep Captioning with Multimodal Recurrent Neural Networks (m-RNN). *ICLR* (2015).
- [24] George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J Miller. 1990. Introduction to WordNet: An on-line lexical database. *International journal of lexicography* 3, 4 (1990), 235–244.
- [25] Margaret Mitchell, Xufeng Han, Jesse Dodge, Alyssa Mensch, Amit Goyal, Alex Berg, Kota Yamaguchi, Tamara Berg, Karl Stratos, and Hal Daumé III. 2012. Midge: Generating Image Descriptions From Computer Vision Detections. In *EACL*. 747–756.
- [26] Vicente Ordonez, Girish Kulkarni, and Tamara L Berg. 2011. Im2text: Describing images using 1 million captioned photographs. In *NIPS*. 1143–1151.
- [27] Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. In *ICLR*.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*. 91–99.
- [29] Marcus Rohrbach, Michael Stark, and Bernt Schiele. 2011. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *CVPR*. 1641–1648.
- [30] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. One-shot learning with memory-augmented neural networks. *NIPS-W* (2016).
- [31] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *ICLR*.
- [32] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*.
- [33] Hamed R Tavakoli, Rakshit Shetty, Ali Borji, and Jorma Laaksonen. 2017. Paying Attention to Descriptions Generated by Image Captioning Models. In *ICCV*. 2506–2515.
- [34] Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2017. Captioning Images with Diverse Objects. In *CVPR*.
- [35] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence-video to text. In *ICCV*. 4534–4542.
- [36] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *NIPS*. 3630–3638.
- [37] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *CVPR*. 3156–3164.
- [38] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. 2017. Show and Tell: Lessons Learned from the 2015 MSCOCO Image Captioning Challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 4 (April 2017), 652–663.
- [39] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata. 2018. Zero-Shot Learning - A Comprehensive Evaluation of the Good, the Bad and the Ugly. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018), 1–1. <https://doi.org/10.1109/TPAMI.2018.2857768>
- [40] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*. 2048–2057.
- [41] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. 2017. Incorporating copying mechanism in image captioning for learning novel objects. In *CVPR*. 5263–5271.
- [42] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *CVPR*. 4651–4659.
- [43] Linchao Zhu, Zhongwen Xu, Yi Yang, and Alexander G. Hauptmann. 2017. Uncovering the Temporal Context for Video Question Answering. *International Journal of Computer Vision* 124, 3 (01 Sep 2017), 409–421. <https://doi.org/10.1007/s11263-017-1033-7>