

# Image Captioning with Unseen Objects

Berkan Demirel  
berkan.demirel@metu.edu.tr

Ramazan Gokberk Cinbis  
gcinbis@ceng.metu.edu.tr  
Nazli Ikizler-Cinbis  
nazli@cs.hacettepe.edu.tr

HAVELSAN Inc. & Middle East  
Technical University  
Ankara, Turkey

Middle East Technical University  
Ankara, Turkey  
Hacettepe University  
Ankara, Turkey

---

## Abstract

Image caption generation is a long standing and challenging problem at the intersection of computer vision and natural language processing. A number of recently proposed approaches utilize a fully supervised object recognition model within the captioning approach. Such models, however, tend to generate sentences which only consist of objects predicted by the recognition models, excluding instances of the classes without labelled training examples. In this paper, we propose a new challenging scenario that targets the image captioning problem in a fully zero-shot learning setting, where the goal is to be able to generate captions of test images containing objects that are not seen during training. The proposed approach jointly uses a novel zero-shot object detection model and a template-based sentence generator. Our experiments show promising results on the COCO dataset.

## 1 Introduction

Image captioning, the problem of generating a concise textual summary of a given image, is one of the most actively studied problems standing at the intersection of computer vision and natural language processing. Following the success of deep learning based object detection approaches [0, 18, 31, 32, 35, 37, 49, 51, 51, 52, 54], there have been a recent interest in generating visually grounded image captions by constructing captioning models that operate on the outputs of supervised object detectors [5, 29, 38, 57]. However, the success of such approaches are inherently limited by the set of classes provided in the detector training dataset, which is typically too small to construct a visually comprehensive model. Therefore, such models are arguably unsuitable for captioning in uncontrolled images, which are likely to contain object classes unseen at training time.

In the context of image classification, *zero-shot learning* (ZSL) has emerged as a promising alternative towards overcoming the practical limits in collecting labelled image datasets and constructing image classifiers with very large object vocabularies. ZSL approaches typically tackle the problem of transferring visual and/or semantic knowledge from the *seen* classes (training classes) to *unseen classes* (novel classes appearing at test time) by using a variety of techniques and information sources, such as class attributes [0, 23], class hierarchies [0, 52], attribute-to-class name mappings [02], label relation graphs [04]. While




Figure 1: Our zero-shot captioning framework, which consists of two components: (i) zero-shot object detector, and (ii) image caption generator. For inference, input image is fed into the object detector, which provides the candidate objects and their representations. The candidates are then given to the sentence generation RNN to obtain a sentence template, which consists of textual and visual words. The final sentence is obtained by setting the visual words according to the data obtained from the zero-shot detector.

ZSL remains as an unsolved problem, significant progress have been made in recent years in zero-shot image classification [62]. Very recently, a few methods have also been proposed for the *zero-shot object detection* (ZSD) problem [6, 13, 48].

In a similar manner, *zero-shot image captioning* (ZSC), where the goal is to create captions of images containing objects classes unseen at the training time, have the promise of overcoming the data collection bottleneck in image captioning. However, we observe that there is no benchmark directly tailored to study captioning in a true zero-shot setting, and, there is no prior work explicitly towards addressing true ZSC. While there are few very recent works on ZSC [58, 66], all these methods study the problem purely as a language modeling problem and presume the availability of a pre-trained fully-supervised object detector over all object classes of interest, to the best of our knowledge.

Following these observations, we argue that the true zero-shot captioning problem, where some of the test objects have no supervised visual or textual examples, needs to be studied towards (i) developing semantically scalable captioning methods, and, (ii) evaluating captioning approaches in a realistic setting where not all object classes have training examples. To tackle this problem, we propose a novel ZSC approach that consists of a novel ZSD model and a template-based [63] caption generator defined over the ZSD outputs. Our ZSD model uses the class-to-class similarities obtained over the distributed word representations [42] between each target class and all training classes. In order to address the generalized zero-shot detection (GZSD) problem, where test images contain a mixture seen and unseen class instances, we propose a scaling scheme to make scores of seen and unseen classes more comparable, as illustrated in Figure 1. For experimental evaluation, we use the MS-COCO dataset [63] with train and test splits where the model is trained over a strict subset of object classes.

To sum up, this work aims to make a number conceptual, technical and experimental contributions in image captioning, which can be summarized as follows: (i) we define a new paradigm for generating captions of unseen classes, (ii) we propose a novel ZSD approach that incorporates a probability scaling scheme for the generalized zero-shot object detection (GZSD) problem, (iii) we evaluate several caption evaluation metrics and discuss their suitability for the zero-shot image captioning scenario.

## 2 Related Work

**Zero-shot Learning.** Early work on ZSL focused on directly using attributes for transferring semantic information from seen to unseen classes [50]. More recent work studies other knowledge transfer mediums, such as distributed word representations and hierarchies [0], visually consistent word representations [12], as well as better knowledge transfer methods, such as synthesized classifiers [9], semantic autoencoders [27], hierarchy graphs [14], diffusion regularization [57], attribute regression [59] and latent space encoding [59]. There are also studies that aim to reduce the bias towards seen classes, *e.g.* [56]. More recently, data-generating models have been proposed for ZSL [43, 53, 60, 70].

**Object Detection.** Most recent object detection methods can be categorized into the following two groups: (i) regression based approaches, and, (ii) region proposal based approaches. Regression based approaches work on detecting the candidate objects and their locations directly without any intermediate step [17, 31, 36, 49, 50, 55]. Region proposal based approaches firstly generate region proposals, and then classify each region proposal into candidate classes with a confidence score [11, 18, 19, 32, 35, 51, 54]. Recently, zero-shot object detection methods are also emerging [6, 13, 47, 48]. Among these, Rahman *et al.* [48] use semantic clustering loss for clustering similar classes. Demirel *et al.* [13] propose a hybrid method to handle visually and semantically meaningful word vectors together. Bansal *et al.* [6] propose a background-aware approach to learn a generalized zero-shot object detector. Rahman *et al.* [47] reshape embeddings so that visual features are well-aligned with related semantics. As an alternative to zero-shot detection, [21, 22, 49] propose to transfer knowledge from learned image classifiers to detectors, similar to weakly supervised learning [8, 10]. However, such approaches require labeled training images for all classes of interest, which significantly reduces their suitability to zero-shot captioning.

The approach closest to our ZSD method is the one proposed by Demirel *et al.* [13]. Our approach differs by (i) leveraging class-to-class similarities measured in the word embedding space, instead of directly using the word embeddings, (ii) learning a class score scaling factor to improve generalized zero-shot detection performance.

**Image Captioning.** Most recent image captioning approaches are based on deep neural networks [25, 26, 29, 38, 40, 53, 68]. Mainstream methods can be categorized as (i) template-based techniques [16, 29, 38] and retrieval-based ones [21, 26, 44, 57]. Template-based captioning approaches generate templates with empty slots, and fill those slots by using attributes or detected objects. Kulkarni *et al.* [29] use conditional random fields (CRF) to push tight connections between the image content and sentence generation process before filling the empty slots. Farhadi *et al.* [16] use triplets of scene elements for filling the empty slots in generated template. Lu *et al.* [68] use recurrent neural network to generate sentence templates for slot filling. Retrieval-based image captioning methods rely on retrieving captions from the set of training examples. More specifically, a set of training images similar to the test example are retrieved and the captioning is performed over their captions.

In this work, we aim to generate captions that can include classes that are not seen within the supervised training set, where retrieval-based approaches are not directly suitable. For this reason, we adopt a template-based approach that generates sentence templates and fills the visual word slots with the ZSD outputs.

Dense captioning [24, 28, 65] appear to be similar to the zero-shot image captioning, but the focus is significantly different: while dense captioning aims to generate rich descriptions,

in ZSC our goal is to achieve captioning over novel object classes. Some captioning methods go beyond training with fully supervised captioning data and allow learning with a captioning dataset that covers only some of the object classes plus additional supervised examples for training object detectors and/or classifiers for all classes of interest [4, 5, 58, 61, 69]. Since these methods presume that all necessary visual information can be obtained from some pre-trained object recognition model, we believe they cannot be seen as true zero-shot captioning approaches.

### 3 Method

In this section, we describe the proposed ZSC model, and its components. It consists of a ZSD model that leverages the class-to-class similarities obtained over the distributed word representations, and a template-based image captioning model that generates sentence templates using a recurrent neural network.

In the rest of this section, we explain the details of the model components and then describe how we build the final ZSC model by using these components.

#### 3.1 Zero-Shot Object Detection

In zero-shot object detection, the goal is to learn a detection model over the examples given for the seen classes, denoted by  $Y_s$  such that the detector can recognize and localize the bounding boxes of the instances of all classes  $Y = Y_s \cup Y_u$ , where  $Y_u$  represents the set of unseen classes. For this purpose, we use the YOLO [60] architecture as our backend, and adapt it to the ZSD problem.

In the original YOLO approach, the loss function consists of three components: (i) the localization loss, which measures the error between ground truth locations and predicted bounding boxes, (ii) the objectness loss, and, (iii) the recognition loss. The original recognition loss  $\ell_{cls}$  is the squared error of class conditional probabilities at each cell, in a grid of size  $S \times S$  ( $S = 13$ , by default):

$$\ell_{cls}(x) = \sum_{i=0}^{S^2} \mathbb{1}_{\text{object}}^i \sum_{c \in Y_s} (t(x, c, i) - f(x, c, i))^2 \quad (1)$$

where the target indicator mapping  $t(x, c, i)$  is 1 if the cell  $i$  of image  $x$  contains an instance of class  $c$  and otherwise 0.  $f(x, c, i)$  is the prediction score corresponding two the class  $c$  and cell  $i$ .  $\mathbb{1}_{\text{object}}^i = 1$  if an object instance appears in cell  $i$ , and otherwise 0.

Following the approach in [43], we adapt the YOLO model to the ZSD problem by re-defining the prediction function as a compatibility estimator between the *cell embeddings* and *class embeddings*, so that novel test classes can be predicted at the test time based purely on their class embeddings:

$$f(x, c, i) = \frac{\Omega(x, i)^T \Psi(c)}{\|\Omega(x, i)\| \|\Psi(c)\|} \quad (2)$$

Here,  $\Omega(x, i)$  denotes the cell embedding obtained for the image cell  $i$  for the input image  $x$ , and,  $\Psi(c)$  represents the class embedding for the class  $c$ . Unlike the approach in [43], which uses word embeddings (or attributes) of class names directly as the class embeddings, we

use the vector of class-to-class similarities in the word embedding space to obtain the class embeddings to obtain more compact and descriptive class embeddings. More specifically, we define the class embedding for the class  $c$  as the vector of similarities between  $c$  and each reference training class  $\bar{c}$ :

$$\varphi(c)^T \varphi(\bar{c}) + 1 \quad (3)$$

where  $\varphi(c)$  denotes the distributed word representations [24] for class  $c$ .

In this work, we need to obtain a generalized zero-shot detector, where both training and test classes appear at the test images, as opposed to zero-shot detection as in [14]. However, there can still be a significant bias towards the seen classes, as the unseen test classes are not directly represented in the training set. We aim to overcome this problem by introducing a scaling coefficient  $\alpha$  for the unseen test classes as follows:

$$f(x, c, i) = \begin{cases} \alpha f(x, c, i), & \text{if } c \in \bar{Y}_s \\ f(x, c, i), & \text{otherwise} \end{cases} \quad (4)$$

Here,  $\bar{Y}_s$  represents the selected subset of classes from the training classes  $Y_s$  to simulate the unseen class scores. In order to learn the coefficient  $\alpha$ , we first train the ZSD model over all training classes without  $\alpha$ , freeze the network, select a subset of the training classes as unseen classes, and remove (*i.e.* set to zero) the corresponding similarities in the class embeddings. Then, in this setting, we train only  $\alpha$  to make the scores among all training classes comparable. At test time,  $\alpha$  is used as a scaling factor only for the unseen classes.

## 3.2 Template-based Image Captioning

We aim to generate accurate captions for images containing classes not seen during training. For this purpose, we use a template-based captioning method which provides the sentence templates whose visual word slots are to be filled in using the ZSD outputs. For this purpose, we use the slotted sentence template generation component of the Neural Baby Talk (NBT) approach [58].

The NBT method generates sentence templates which consist of the empty word slots by using a recurrent neural network. Moreover, NBT incorporates the pointer networks idea [59] to obtain a content-based attention mechanism over the grounding regions. There are two word types in NBT method: textual and visual words. Textual words are not related to any image region, therefore the model provides only dummy grounding for them. The template generation network uses the object detection outputs to fill empty visual word slots.

In our approach, the ZSD model and the sentence template generation component of NBT is trained over the seen classes. The ZSD outputs over all classes are used as input to the NBT sentence generator at test time.

## 4 Experiments

In this section, we present our experimental analysis for the ZSD component of our ZSC approach and the ZSC outputs themselves. For this purpose, we first define the COCO dataset splits that we use and the word embeddings used for computing the class-to-class similarities in Section 4.1. We present the experimental results for the ZSD component in Section 4.2, and, those for the complete ZSC model in Section 4.3. Finally, we present ablative studies in Section 4.4.

## 4.1 Dataset

We use the COCO image captioning dataset [43] in our experiments. As discussed in Section 2, there are several zero-shot captioning studies that assume that image-caption training example pairs are provided only for a subset of the object classes during training, while presuming the availability of visual training examples for all classes [4, 5, 58, 58, 56]. We refer to them as *partial zero-shot* methods and consider them as upper-bounds (golden baselines) for our true ZSC problem definition. In order to compare against these methods, we use the same COCO dataset splits used by the respective publications.

The details of the COCO splits are as follows: from the 80 classes, we select 8 of them as unseen classes (*bottle, bus, couch, microwave, pizza, tennis racket, suitcase, and zebra*). During training, we use the subset of the COCO training images that do not contain any instances of these unseen classes. For evaluation, we use the validation set split prepared by [5].

**Word embeddings.** To obtain the class-to-class similarities required by our zero-shot detector, we use the 300-dimensional word2vec embeddings of class names [44]. For the classes with multiple words, we use the average word embeddings. For the NBT component of our ZSC approach, we use 300-dimensional GloVe [46] word embeddings for template generation, as in the original NBT model.

## 4.2 Zero-Shot Object Detection

For training the ZSD component of our approach, we train the YOLO [50] based network for 160 epochs, with the learning rate set to 0.001 and a batch size of 32. The detector training is conducted on COCO object detection dataset using only the training examples of selected 72 seen classes. Then, we select 64 of these 72 classes as seen and remaining 8 of which as unseen, and learn the  $\alpha$  scaling factor.

We evaluate the proposed zero-shot detection approach in ZSD and Generalized ZSD (GZSD) settings. For the ZSD evaluation, we use the COCO validation images that consist of only unseen object instances and by feeding only unseen classes as target classes to the detector. The results are shown in the first column of Table 1. On average, the method yields 31.4% mean average precision (mAP) score, which can be interpreted as a promising result considering the difficulty of the zero-shot detection problem.

For the GZSD evaluation, we use the COCO val5k split, which contains both seen and unseen class instances. We separately compute mAP scores over the seen and unseen classes. As the final scalar performance metric, we use the harmonic mean (HM) of seen and unseen class mAP scores, following its use in generalized zero-shot classification evaluation [62]. The GZSD results without and with alpha scaling are presented in the last two rows of Table 1, respectively. It can be seen that, with alpha-scaling, the unseen class mAP increases from 0.3% to 7.3% mAP but the seen class mAP drops from 27.4% to 19.2% mAP. These results are as expected considering the fact that alpha-scaling acts as a prior that promotes the detection of unseen classes over the seen ones. Overall, the harmonic mean score increases from 0.7% to 10.6% mAP using alpha-scaling, which suggests that alpha scaled detector is more suitable for being used as a component of the ZSC model.

Finally, we note that these zero-shot detection results are not comparable to the previously reported COCO ZSD results of Bansal *et al.* [5] since the splits that we take from [5] (primarily for ZSC evaluation purposes) are different from those used in [5].

Exp. Type	Test	bottle	bus	couch	microwave	pizza	racket	suitcase	zebra	U-mAP(%)	S-mAP(%)	HM
ZSD	U	5.2	53.3	35.1	23.9	44.4	36.4	9.1	43.7	31.4	-	-
GZSD w/o $\alpha$	S+U	0	0	2.7	0	0	0	0	0	0.3	27.4	0.7
GZSD	S+U	0.8	21.4	4.9	1.2	4.8	0.7	9.1	15.8	7.3	19.2	10.6

Table 1: Our ZSD and Generalized ZSD (GZSD) results. The first row shows the ZSD results where the detector is evaluated over the images containing only unseen classes, without computing the seen class scores. The last two rows show the GZSD results where both seen and unseen class scores are computed on the full COCO val5k split. HM: Harmonic Mean.

Method	bottle	bus	couch	microwave	pizza	racket	suitcase	zebra	Avg. F1	METEOR	SPICE
NBT-baseline	0	0	0	0	0	0	0	0	0	18.2	12.7
Our Method	2.4	75.2	26.6	24.6	29.8	3.6	0.6	75.4	29.8	21.9	14.2
Partial zero-shot captioning methods (upper-bound)											
DCC [8]	4.6	29.8	45.9	28.1	64.6	52.2	13.2	79.9	39.8	21.0	13.4
NOC [58]	17.8	68.8	25.6	24.7	69.3	68.1	39.9	89.0	49.1	21.4	-
C-LSTM [60]	29.7	74.4	38.8	27.8	68.2	70.3	44.8	91.4	55.7	23.0	-
Base+T4 [9]	16.3	67.8	48.2	29.7	77.2	57.1	49.9	85.7	54.0	23.3	15.9
NBT+G [58]	14.0	74.8	42.8	63.7	74.4	19.0	44.5	92.0	53.2	23.9	16.6
DNOC [60]	33.0	77.0	54.0	46.6	75.8	33.0	59.5	84.6	57.9	21.6	-

Table 2: Our ZSC results with comparison to upper-bound partial ZSC methods.

### 4.3 Zero-Shot Image Captioning

For the ZSC experiments, we use our detection model in the GZSL setting. We obtain the candidate object regions by setting the class confidence threshold of the ZSD detector to 0.5, in order to retain only confident detections, as needed by the template-based caption generator. We prepare the same experimental setup described in [58], and exclude the image-sentence pairs which consist of unseen classes. To establish a fair comparison, we use the NBT method [58] for learning the language model.

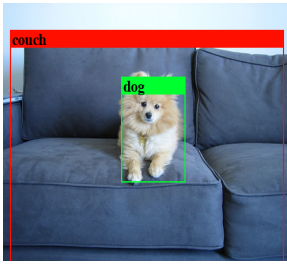
For comparison, we define the *NBT-baseline* by training a seen class only YOLO detector and feeding its outputs to the captioning model. For evaluation, we use the METEOR [15], SPICE [9] and Average F1 metrics to evaluate the NBT baseline, the proposed method and the upper-bound partial ZSC methods, on the validation split of [9].

In Table 2, we see that the proposed approach obtains 29.8% average F1 score, which is computed by averaging per-class F1 scores. For the F1 score, a generated sentence is counted as correct only if the sentence contains the name of the unseen class of interest in an image containing at least one instance of that class. Here, an image is considered as a positive example only if it contains at least one instance of the unseen object of interest. We observe that the obtained F1 score of the proposed ZSC approach is promising but relatively lower than those of the upper-bound partial ZSC methods, which is not a surprise considering that they are built upon supervised detectors over all classes.

Compared to the upper-bound techniques in terms of METEOR and SPICE metrics, our method obtains better results than DCC [8] and NOC [58] methods, however lower than the other upper-bound methods, which benefit from using supervised object classification or detection models over all classes. Compared to the NBT-baseline, we observe significant improvements in terms of all metrics, which highlights the importance of ZSD for the true ZSC problem.

Figure 2 gives qualitative ZSC results, where the seen (*italic* typeface) and unseen (**bold** typeface) class names are denoted. Overall, these results are encouraging as the single final model shows (partial) ability to generate captions with both seen and unseen classes, even in images where a group of seen and unseen class instances appear jointly.







A small white *dog* sitting on a *couch*.



A red *bus* is driving down the street.



A couple of *zebra* standing in a field.



A *tennis player* is about to hit a *tennis racket*.



A white plate topped with a piece of *pizza*.




A kitchen with a *microwave* and a counter.



A *bus* is parked on the side of the street.



A *bird* sitting on top of a metal pole.



A bunch of *banana* that are on a table.

Figure 2: Image captioning results of images which consist of *seen* and *unseen* classes.

## 4.4 Ablative Studies

In this section, we measure the effect of the ZSD method on ZSC success in a more detailed manner. For this purpose, we compare our method with the NBT-baseline method, using the same split as in Section 4.3, by evaluating in terms of METEOR [15], SPICE [3], ROUGE-L [8], BLEU [15] metrics and the F1 score.

We present the results in Table 3. According to these results, we get the average F1 = 0 in NBT-baseline as expected, as it is unable to caption with any unseen object instance. However, we observe that the results of the NBT-baseline and the proposed approach are very close to each other, especially in terms of BLEU and ROUGE-L metrics, despite the inability of the NBT-baseline to incorporate any unseen class in the captioning outputs. These results suggest that these metrics are not suitable for the ZSC scenario.

According to Table 3, METEOR and SPICE metrics give slightly more indicative results for ZSC. We note that the METEOR metric uses  $n$ -gram based synonym matching and



Method	BLEU1	BLEU2	BLEU3	BLEU4	METEOR	SPICE	ROUGE-L	F1 Scores
NBT-baseline	65.8	47.5	34.1	24.4	18.2	12.7	48.2	0
Our Method	<b>67.0</b>	<b>48.3</b>	<b>34.5</b>	<b>24.5</b>	<b>21.9</b>	<b>14.2</b>	<b>48.9</b>	<b>29.8</b>

Table 3: Zero-shot captioning evaluation results in terms of various performance metrics.



◆: A yellow and black **train** traveling down the road.

★: A yellow and black **bus** driving down a road.



◆: A couple of **elephants** standing next to each other.

★: A couple of **zebra** standing next to each other.



◆: A piece of **cake** on a white plate.

★: A piece of **pizza** on a white plate.

Figure 3: Zero-shot image captioning results using the NBT-baseline (◆) and the proposed approach (★). **Bold** typed words represent visual words given by the detectors.

the SPICE metric uses scene-graph based synonym matching. On the other hand, remaining metrics use  $n$ -gram precision,  $n$ -gram recall, tf-idf weighted  $n$ -gram similarities without synonym matching. Therefore, we may preliminary conclude that using synonym matching contributes to obtaining more meaningful ZSC evaluations.

Finally, we show a qualitative comparison between the NBT-baseline and our method in Figure 3. These results show how the NBT baseline generates captions with somewhat similar but inaccurate objects due to its inability to handle unseen class instances.

## 5 Conclusion

An important shortcoming of current image captioning methods aiming to learn with class-restricted caption annotations is their inability to operate in a fully zero-shot learning setting. These methods generate captions for images which consist of classes not seen in captioning datasets, but they still assume the availability of a fully supervised object classifier and/or detector over all classes of interest. To initiate research towards overcoming this important weakness in captioning research, we define the zero-shot image captioning problem, and propose a novel approach that consist of a zero-shot object detector and a detector-driven caption generator. Our experimental results based on the COCO dataset show that our method yields promising results towards achieving true ZSC goals. In addition, the qualitative results show that proposed architecture is able to generate natural looking and visually grounded captions in several challenging test cases.

## 6 Acknowledgements

This work was supported in part by the TUBITAK Grant 116E445.

## References

- [1] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for attribute-based classification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 819–826. IEEE, 2013.
- [2] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.* IEEE Computer Society, 2015.
- [3] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *Proc. European Conf. on Computer Vision*, pages 382–398. Springer, 2016.
- [4] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Guided open vocabulary image captioning with constrained beam search. In *Proc. of the Empirical Methods in Natural Language Processing*, pages 936–945, 2017.
- [5] Lisa Anne Hendricks, Subhashini Venugopalan, Marcus Rohrbach, Raymond Mooney, Kate Saenko, and Trevor Darrell. Deep compositional captioning: Describing novel object categories without paired training data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2016.
- [6] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. In *Proc. European Conf. on Computer Vision*, pages 384–400, 2018.
- [7] Sean Bell, C Lawrence Zitnick, Kavita Bala, and Ross Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2874–2883, 2016.
- [8] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2846–2854, 2016.
- [9] Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, and Fei Sha. Synthesized classifiers for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5327–5336, 2016.
- [10] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(1):189–203, 2016.
- [11] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 379–387, 2016.
- [12] Berkan Demirel, Ramazan Gokberk Cinbis, and Nazli Ikizler-Cinbis. Attributes2classname: A discriminative model for attribute-based unsupervised zero-shot learning. In *Proc. IEEE Int. Conf. on Computer Vision*, Oct 2017.
- [13] Berkan Demirel, Ramazan Gokberk Cinbis, and Nazli Ikizler-Cinbis. Zero-shot object detection by hybrid region embedding. In *BMVC*, 2018.

- [14] Jia Deng, Nan Ding, Yangqing Jia, Andrea Frome, Kevin Murphy, Samy Bengio, Yuan Li, Hartmut Neven, and Hartwig Adam. Large-scale object classification using label relation graphs. In *Proc. European Conf. on Computer Vision*, pages 48–64. Springer, 2014.
- [15] Michael Denkowski and Alon Lavie. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the ninth workshop on statistical machine translation*, pages 376–380, 2014.
- [16] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. Every picture tells a story: Generating sentences from images. In *Proc. European Conf. on Computer Vision*, pages 15–29. Springer, 2010.
- [17] Cheng-Yang Fu, Wei Liu, Ananth Ranga, Ambrish Tyagi, and Alexander C Berg. Dssd: Deconvolutional single shot detector. *arXiv preprint arXiv:1701.06659*, 2017.
- [18] Ross Girshick. Fast r-cnn. In *Proc. IEEE Int. Conf. on Computer Vision*, pages 1440–1448, 2015.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1904–1916, 2015.
- [20] Micah Hodosh, Peter Young, and Julia Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899, 2013.
- [21] Judy Hoffman, Sergio Guadarrama, Eric S Tzeng, Ronghang Hu, Jeff Donahue, Ross Girshick, Trevor Darrell, and Kate Saenko. Lsda: Large scale detection through adaptation. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 3536–3544, 2014.
- [22] Judy Hoffman, Deepak Pathak, Trevor Darrell, and Kate Saenko. Detector discovery in the wild: Joint multiple instance and representation learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2883–2891, 2015.
- [23] Dinesh Jayaraman and Kristen Grauman. Zero-shot recognition with unreliable attributes. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 3464–3472, 2014.
- [24] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4565–4574, 2016.
- [25] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3128–3137, 2015.
- [26] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [27] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3174–3183, 2017.

- [28] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 706–715, 2017.
- [29] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. Babytalk: Understanding and generating simple image descriptions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(12):2891–2903, 2013.
- [30] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 951–958, 2009.
- [31] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proc. European Conf. on Computer Vision*, pages 734–750, 2018.
- [32] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*, 2004.
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proc. European Conf. on Computer Vision*, pages 740–755. Springer, 2014.
- [34] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017.
- [35] Tsung-Yi Lin, Priyal Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2018.
- [36] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Proc. European Conf. on Computer Vision*, pages 21–37. Springer, 2016.
- [37] Yang Long, Li Liu, Fumin Shen, Ling Shao, and Xuelong Li. Zero-shot learning using synthesised unseen visual data with diffusion regularisation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(10):2498–2512, 2018.
- [38] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural baby talk. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7219–7228, 2018.
- [39] Changzhi Luo, Zhetao Li, Kaizhu Huang, Jiashi Feng, and Meng Wang. Zero-shot learning via attribute regression and class prototype rectification. *IEEE Trans. on Image Processing*, 27(2):637–648, 2018.
- [40] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn). *Proc. Int. Conf. Learn. Represent.*, 2015.
- [41] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

- [42] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 3111–3119, 2013.
- [43] Ashish Mishra, Shiva Krishna Reddy, Anurag Mittal, and Hema A Murthy. A generative model for zero shot learning using conditional variational autoencoders. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog. Workshops*, pages 2188–2196, 2018.
- [44] Vicente Ordonez, Girish Kulkarni, and Tamara L Berg. Im2text: Describing images using 1 million captioned photographs. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 1143–1151, 2011.
- [45] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318, 2002.
- [46] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proc. of the Empirical Methods in Natural Language Processing*, pages 1532–1543, 2014.
- [47] Shafin Rahman, Salman Khan, and Nick Barnes. Polarity loss for zero-shot object detection. *arXiv preprint arXiv:1811.08982*, 2018.
- [48] Shafin Rahman, Salman Khan, and Fatih Porikli. Zero-shot object detection: Learning to simultaneously recognize and localize novel concepts. *arXiv preprint arXiv:1803.06049*, 2018.
- [49] Joseph Redmon and Ali Farhadi. Yolo9000: Better, faster, stronger. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7263–7271, 2017.
- [50] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 779–788, 2016.
- [51] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 91–99, 2015.
- [52] Marcus Rohrbach, Michael Stark, and Bernt Schiele. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1641–1648. IEEE, 2011.
- [53] Mert Bulent Sariyildiz and Ramazan Gokberk Cinbis. Gradient matching generative networks for zero-shot learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2168–2178, 2019.
- [54] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Robert Fergus, and Yann Lecun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *Proc. Int. Conf. Learn. Represent.*, 2014.
- [55] Zhiqiang Shen, Zhuang Liu, Jianguo Li, Yu-Gang Jiang, Yurong Chen, and Xiangyang Xue. Dsod: Learning deeply supervised object detectors from scratch. In *Proc. IEEE Int. Conf. on Computer Vision*, pages 1919–1927, 2017.

- [56] Jie Song, Chengchao Shen, Yezhou Yang, Yang Liu, and Mingli Song. Transductive unbiased embedding for zero-shot learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1024–1033, 2018.
- [57] Chen Sun, Chuang Gan, and Ram Nevatia. Automatic concept discovery from parallel text and visual corpora. In *Proc. IEEE Int. Conf. on Computer Vision*, pages 2596–2604, 2015.
- [58] Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, Raymond Mooney, Trevor Darrell, and Kate Saenko. Captioning images with diverse objects. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5753–5761, 2017.
- [59] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 2692–2700, 2015.
- [60] Wenlin Wang, Yunchen Pu, Vinay Kumar Verma, Kai Fan, Yizhe Zhang, Changyou Chen, Piyush Rai, and Lawrence Carin. Zero-shot learning via class-conditioned deep generative models. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [61] Yu Wu, Linchao Zhu, Lu Jiang, and Yi Yang. Decoupled novel object captioner. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 1029–1037. ACM, 2018.
- [62] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2018.
- [63] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *Proc. Int. Conf. Mach. Learn.*, pages 2048–2057, 2015.
- [64] Junjie Yan, Zhen Lei, Longyin Wen, and Stan Li. The fastest deformable part model for object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2497–2504, 2014.
- [65] Linjie Yang, Kevin Tang, Jianchao Yang, and Li-Jia Li. Dense captioning with joint inference and visual context. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2193–2202, 2017.
- [66] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Incorporating copying mechanism in image captioning for learning novel objects. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6580–6588, 2017.
- [67] Xuwang Yin and Vicente Ordonez. Obj2text: Generating visually descriptive language from object layouts. In *Proc. of the Empirical Methods in Natural Language Processing*, pages 177–187, 2017.
- [68] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4651–4659, 2016.



- [69] Yunlong Yu, Zhong Ji, Jichang Guo, and Zhongfei Zhang. Zero-shot learning via latent space encoding. *IEEE transactions on cybernetics*, (99):1–12, 2018.
- [70] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1004–1013, 2018.