# Learning like a Child:
# Fast Novel Visual Concept Learning from Sentence Descriptions of Images

Junhua Mao[1]     Wei Xu[2]     Yi Yang[2]     Jiang Wang[2]     Zhiheng Huang[2]     Alan L. Yuille[1]
[1]University of California, Los Angeles     [2]Baidu Research

mjhustc@ucla.edu, {wei.xu,yangyi05,wangjiang03,huangzhiheng}@baidu.com, yuille@stat.ucla.edu

## Abstract

*In this paper, we address the task of learning novel visual concepts, and their interactions with other concepts, from a few images with sentence descriptions. Using linguistic context and visual features, our method is able to efficiently hypothesize the semantic meaning of new words and add them to its word dictionary so that they can be used to describe images which contain these novel concepts. Our method has an image captioning module based on [38] with several improvements. In particular, we propose a transposed weight sharing scheme, which not only improves performance on image captioning, but also makes the model more suitable for the novel concept learning task. We propose methods to prevent overfitting the new concepts. In addition, three novel concept datasets are constructed for this new task, and are publicly available on the project page. In the experiments, we show that our method effectively learns novel visual concepts from a few examples without disturbing the previously learned concepts. The project page is: www.stat.ucla.edu/~junhua.mao/projects/child_learning.html.*

## 1. Introduction

Recognizing, learning and using novel concepts is one of the most important cognitive functions of humans. When we were very young, we learned new concepts by observing the visual world and listening to the sentence descriptions of our parents. The process was slow at the beginning, but got much faster after we accumulated enough learned concepts [4]. In particular, it is known that children can form quick and rough hypotheses about the meaning of new words in a sentence based on their knowledge of previous learned words [5, 21], associate these words to the objects or their properties, and describe novel concepts using sentences with the new words [4]. This phenomenon has been researched for over 30 years by the psychologists and linguists who study the process of word learning [52].

For the computer vision field, several methods are proposed [16, 46, 53, 31] to handle the problem of learning new



Figure 1: An illustration of the Novel Visual Concept learning from Sentences (NVCS) task. We start with a model (i.e. model-base) trained with images that do not contain the concept of "quidditch" [1]. Using a few "quidditch" images with sentence descriptions, our method is able to learn that "quidditch" is played by people with a ball.

categories of objects from a handful of examples. This task is important in practice because we sometimes do not have enough data for novel concepts and hence need to transfer knowledge from previously learned categories. Moreover, we do not want to retrain the whole model every time we add a few images with novel concepts, especially when the amount of data or model parameters is very big.

However, these previous methods concentrate on learning classifiers, or mappings, between single words (e.g. a novel object category) and images. We are unaware of any computer vision studies into the task of learning novel visual concepts from a few sentences and then using these concepts to describe new images – a task that children seem to do effortlessly. We call this the *Novel Visual Concept learning from Sentences (NVCS)* task (see Figure 1).

In this paper, we present a novel framework to address the NVCS task. We start with a model that has already been trained with a large amount of visual concepts. We propose a method that allows the model to enlarge its word dictionary to describe the novel concepts using a few examples and without extensive retraining. In particular, we do not need to retrain models from scratch on all of the data (all the previously learned concepts and the novel concepts). We propose three datasets for the NVCS task to validate our

---

[1]"quidditch" is a sport created in "Harry Potter". It is played by teams of people holding brooms with a ball (see Figure 1).

model, which are available on the project page.

Our method requires a *base model* for image captioning which will be adapted to perform the NVCS task. We choose the m-RNN model [38], which performs at the state of the art, as our base model. Note that we could use most of the current image captioning models as the base model in our method. But we make several changes to the model structure of m-RNN partly motivated by the desire to avoid overfitting, which is a particular danger for NVCS because we want to learn from a few new images. We note that these changes also improve performance on the original image captioning task, although this improvement is not the main focus of this paper. In particular, we introduce a transposed weight sharing (TWS) strategy (motivated by auto-encoders [3]) which reduces, by a factor of one half, the number of model parameters that need to be learned. This allows us to increase the dimension of the word-embedding and multimodal layers, without overfitting the data, yielding a richer word and multimodal dense representation. We train this image captioning model on a large image dataset with sentence descriptions. This is the base model which we adapt for the NVCS task.

Now we address the task of learning the new concepts from a small new set of data that contains these concepts. There are two main difficulties. Firstly, the weights for the previously learned concepts may be disturbed by the new concepts. Although this can be solved by fixing these weights. Secondly, learning the new concepts from positive examples can introduce bias. Intuitively, the model will assign a baseline probability for each word, which is roughly proportional to the frequency of the words in the sentences. When we train the model on new data, the baseline probabilities of the new words will be unreliably high. We propose a strategy that addresses this problem by fixing the baseline probability of the new words.

We construct three datasets to validate our method, which involves new concepts of man-made objects, animals, and activities. The first two datasets are derived from the MS-COCO dataset [34]. The third new dataset is constructed by adding three uncommon concepts which do not occur in MS-COCO or other standard datasets. These concepts are: quidditch, t-rex and samisen (see section 5)[2]. The experiments show that training our method on only a few examples of the new concepts gives us as good performance as retraining the entire model on all the examples.

## 2. Related Work

**Deep neural network** Recently there have been dramatic progress in deep neural networks for natural language and computer vision. For natural language, Recurrent Neural Networks (RNNs [14, 40]) and Long-Short Term Memories (LSTMs [22]) achieve the state-of-the-art performance for many NLP tasks such as machine translation [23, 8, 51] and speech recognition [40]. For computer vision, deep Convolutional Neural Networks (CNN [33]) outperform previous methods by a large margin for the tasks of object classification [27, 48] and detection [19, 43, 59]. The success of these methods for language and vision motivate their use for multimodal learning tasks (e.g. image captioning and sentence-image retrieval).

**Multimodal learning of language and vision** The methods of image-sentence retrieval [17, 50], image description generation [28, 42, 20] and visual question-answering [18, 36, 1] have developed very fast in recent years. Very recent works of image captioning includes [39, 25, 24, 55, 10, 15, 7, 32, 37, 26, 57, 35]. Many of them (e.g. [38, 55]) adopt an RNN-CNN framework that optimizes the log-likelihood of the caption given the image, and train the networks in an end-to-end way. An exception is [15], which incorporates visual detectors, language models, and multimodal similarity models in a high-performing pipeline. The evaluation metrics of the image captioning task is also discussed [13, 54]. All of these image captioning methods use a pre-specified and fixed word dictionary, and train their model on a large dataset. Our method can be directly applied to any captioning models that adopt an RNN-CNN framework, and our strategy to avoid overfitting is useful for most of the models in the novel visual concept learning task.

**Zero-shot and one-shot learning** For zero-shot learning, the task is to associate dense word vectors or attributes with image features [49, 17, 12, 2, 31]. The dense word vectors in these papers are pre-trained from a large amount of text corpus and the word semantic representation is captured from co-occurrence with other words [41]. [31] developed this idea by only showing the novel words a few times. In addition, [47] adopted auto-encoders with attribute representations to learn new class labels and [56] proposed a method that scales to large datasets using label embeddings.

Another related task is one-shot learning task of new categories [16, 29, 53]. They learn new objects from only a few examples. However, these work only consider words or attributes instead of sentences and so their learning target is different from that of the task in this paper.

## 3. The Image Captioning Model

We need an image captioning as the base model which will be adapted in the NVCS task. The base model is based on the m-RNN model [38]. Its architecture is shown in Figure 2(a). We make two main modifications of the architecture to make it more suitable for the NVCS task which, as a side effect, also improves performance on the original image captioning task. Firstly and most importantly, we propose a transposed weight sharing strategy which sig-

---

[2]The dataset is available at www.stat.ucla.edu/~junhua.mao/projects/child_learning.html. We are adding more novel concepts in this dataset. The latest version of the dataset contains 8 additional novel concepts: tai-ji, huangmei opera, kiss, rocket gun, tempura, waterfall, wedding dress, and windmill.

Figure 2: (a). Our image captioning model. For each word in a sentence, the model takes the current word index and the image as inputs, and outputs the next word index. The weights are shared across the sub-models for the words in a sentence. The number on the top right of each layer denotes its dimension. As in the m-RNN model [38], we add a start sign $w_{start}$ and an end sign $w_{end}$ to each training sentence. (b). The transposed weight sharing of $\mathbf{U}_D$ and $\mathbf{U}_M$. (Best viewed in color)

nificantly reduces the number of parameters in the model (see section 3.2). Secondly, we replace the recurrent layer in [38] by a Long-Short Term Memory (LSTM) layer [22]. LSTM is a recurrent neural network which is designed to solve the gradient explosion and vanishing problems. We briefly introduce the framework of the model in section 3.1 and describe the details of the transposed weight sharing strategy in section 3.2.

### 3.1. The Model Architecture

As shown in Figure 2(a), the input of our model for each word in a sentence is the index of the current word in the word dictionary as well as the image. We represent this index as a one-hot vector (a binary vector with only one non-zero element indicating the index). The output is the index of the next word. The model has three components: the language component, the vision component and the multi-modal component. The language component contains two word embedding layers and a LSTM layer. It maps the index of the word in the dictionary into a semantic dense word embedding space and stores the word context information in the LSTM layer. The vision component contains a 16-layer deep convolutional neural network (CNN [48]) pre-trained on the ImageNet classification task [45]. We remove the final SoftMax layer of the deep CNN and connect the top fully connected layer (a 4096 dimensional layer) to our model. The activation of this 4096 dimensional layer can be treated as image features that contain rich visual attributes for objects and scenes. The multimodal component contains a one-layer representation where the information from the language part and the vision part merge together. We build a SoftMax layer after the multimodal layer to predict the index of the next word. The weights are shared across the sub-models of the words in a sentence. As in the m-RNN model [38], we add a start sign $w_{start}$ and an end sign $w_{end}$ to each training sentence. In the testing stage for image captioning, we input the start sign $w_{start}$ into the

model and pick the $K$ best words with maximum probabilities according to the SoftMax layer. We repeat the process until the model generates the end sign $w_{end}$.

### 3.2. The Transposed Weight Sharing (TWS)

For the original m-RNN model [38], most of the weights (i.e. 98.49%) are contained in the following two weight matrices: $\mathbf{U}_D \in \mathbb{R}^{512 \times N}$ and $\mathbf{U}_M \in \mathbb{R}^{N \times 1024}$ where $N$ represents the size of the word dictionary.

The weight matrix $\mathbf{U}_D$ between the one-hot layer and first word embedding layer is used to compute the input of the first word embedding layer $\mathbf{w}(t)$:

$$\mathbf{w}(t) = f(\mathbf{U}_D \mathbf{h}(t)) \qquad (1)$$

where $f(.)$ is an element-wise non-linear function, $\mathbf{h}(t) \in \mathbb{R}^{N \times 1}$ is the one-hot vector of the current word. Note that it is fast to calculate Equation 1 because there is only one non-zero element in $\mathbf{h}(t)$. In practice, we do not need to calculate the full matrix multiplication operation since only one column of $\mathbf{U}_D$ is used for each word in the forward and backward propagation.

The weight matrix $\mathbf{U}_M$ between the multimodal layer and the SoftMax layer is used to compute the activation of the SoftMax layer $\mathbf{y}(t)$:

$$\mathbf{y}(t) = g(\mathbf{U}_M \mathbf{m}(t) + \mathbf{b}) \qquad (2)$$

where $\mathbf{m}(t)$ is the activation of the multimodal layer and $g(.)$ is the SoftMax non-linear function.

Intuitively, the role of the weight matrix $\mathbf{U}_D$ in Equation 1 is to encode the one-hot vector $\mathbf{h}(t)$ into a dense semantic vector $\mathbf{w}(t)$. The role of the weight matrix $\mathbf{U}_M$ in Equation 2 is to decode the dense semantic vector $\mathbf{m}(t)$ back to a pseudo one-hot vector $\mathbf{y}(t)$ with the help of the Soft-Max function, which is very similar to the inverse operation of Equation 1. The difference is that $\mathbf{m}(t)$ is in the dense multimodal semantic space while $\mathbf{w}(t)$ is in the dense word semantic space.

Figure 3: Illustration of training novel concepts. We only update the sub-matrix $\mathbf{U}_{D_n}$ (the green parts in the figure) in $\mathbf{U}_D$ that is connected to the node of new words in the One-Hot layer and the SoftMax layer during the training for novel concepts. (Best viewed in color)

To reduce the number of the parameters, we decompose $\mathbf{U}_M$ into two parts. The first part maps the multimodal layer activation vector to an intermediate vector in the word semantic space. The second part maps the intermediate vector to the pseudo one-hot word vector, which is the inverse operation of Equation 1. The sub-matrix of the second part is able to share parameters with $\mathbf{U}_D$ in a transposed manner, which is motivated by the tied weights strategy in autoencoders for unsupervised learning tasks [3]. Here is an example of linear decomposition: $\mathbf{U}_M = \mathbf{U}_D^T \mathbf{U}_I$, where $\mathbf{U}_I \in \mathbb{R}^{512 \times 1024}$. Equation 2 is accordingly changed to:

$$\mathbf{y}(t) = g[\mathbf{U}_D^T f(\mathbf{U}_I \mathbf{m}(t)) + \mathbf{b}] \qquad (3)$$

where $f(.)$ is a element-wise function. If $f(.)$ is an identity mapping function, it is equivalent to linearly decomposing $\mathbf{U}_M$ into $\mathbf{U}_D^T$ and $\mathbf{U}_I$. In our experiments, we find that setting $f(.)$ as the scaled hyperbolic tangent function leads to a slightly better performance than linear decomposition. This strategy can be viewed as adding an intermediate layer with dimension 512 between the multimodal and SoftMax layers as shown in Figure 2(b). The weight matrix between the intermediate and the SoftMax layer is shared with $\mathbf{U}_D$ in a transposed manner. This Transposed Weight Sharing (TWS) strategy enables us to use a much larger dimensional word-embedding layer than the m-RNN model [38] without increasing the number of parameters. We also benefit from this strategy when addressing the novel concept learning task.

## 4. The Novel Concept Learning (NVCS) Task

Suppose we have trained a model based on a large amount of images and sentences. Then we meet with images of novel concepts whose sentence annotations contain words not in our dictionary, what should we do? It is time-consuming and unnecessary to re-train the whole model from scratch using all the data. In many cases, we cannot even access the original training data of the model. But fine-tuning the whole model using only the new data causes severe overfitting on the new concepts and decrease the performance of the model for the originally trained ones.

To solve these problems, we propose the following strategies that learn the new concepts with a few images without losing the accuracy on the original concepts.

### 4.1. Fixing the originally learned weights

Under the assumption that we have learned the weights of the original words from a large amount of data and that the amount of the data for new concepts is relatively small, it is straightforward to fix the originally learned weights of the model during the incremental training. More specifically, the weight matrix $\mathbf{U}_D$ can be separated into two parts: $\mathbf{U}_D = [\mathbf{U}_{D_o}, \mathbf{U}_{D_n}]$, where $\mathbf{U}_{D_o}$ and $\mathbf{U}_{D_n}$ associate with the original words and the new words respectively. E.g., as shown in Figure 3, for the novel visual concept "cat", $\mathbf{U}_{D_n}$ is associated with 29 new words, such as cat, kitten and pawing. We fix the sub-matrix $\mathbf{U}_{D_o}$ and update the sub-matrix $\mathbf{U}_{D_n}$ as illustrated in Figure 3.

### 4.2. Fixing the baseline probability

In Equation 3, there is a bias term $\mathbf{b}$. Intuitively, each element in $\mathbf{b}$ represents the tendency of the model to output the corresponding word. We can think of this term as the baseline probability of each word. Similar to $\mathbf{U}_D$, $\mathbf{b}$ can be separated into two parts: $\mathbf{b} = [\mathbf{b}_o, \mathbf{b}_n]$, where $\mathbf{b}_o$ and $\mathbf{b}_n$ associate with the original words and the new words respectively. If we only present the new data to the network, the estimation of $\mathbf{b}_n$ is unreliable. The network will tend to increase the value of $\mathbf{b}_n$ which causes overfitting to the new data.

The easiest way to solve this problem is to fix $\mathbf{b}_n$ during the training for novel concepts. But this is not enough. Because the average activation $\bar{\mathbf{x}}$ of the intermediate layer across all the training samples is not $\mathbf{0}$, the weight matrix $\mathbf{U}_D$ plays a similar role to $\mathbf{b}$ in changing the baseline probability. To avoid this problem, we centralize the activation of the intermediate layer $\mathbf{x}$ and turn the original bias term $\mathbf{b}$ into $\mathbf{b}'$ as follows:

$$\mathbf{y}(t) = g[\mathbf{U}_D^T(\mathbf{x} - \bar{\mathbf{x}}) + \mathbf{b}']; \quad \mathbf{b}'_o = \mathbf{b}_o + \mathbf{U}_{D_o}^T \bar{\mathbf{x}} \qquad (4)$$

After that, we set every element in $\mathbf{b}'_n$ to be the average value of the elements in $\mathbf{b}'_o$ and fix $\mathbf{b}'_n$ when we train on the new images. We call this strategy Baseline Probability Fixation (BPF).

In the experiments, we adopt a stochastic gradient descent algorithm with an initial learning rate of 0.01 and use AdaDelta [58] as the adaptive learning rate algorithm for both the base model and the novel concept model.

### 4.3. The Role of Language and Vision

In the novel concept learning (NVCS) task, the sentences serve as a weak labeling of the image. The language part of the model (the word embedding layers and the LSTM layer) hypothesizes the basic properties (e.g. the parts of speech) of the new words and whether the new words are closely related to the content of the image. It also hypothesizes

Figure 4: Organization of the novel concept datasets

which words in the original dictionary are semantically and syntactically close to the new words. For example, suppose the model meets a new image with the sentence description "A woman is playing with a cat". Also suppose there are images in the original data containing sentence description such as "A man is playing with a dog". Then although the model has not seen the word "cat" before, it will hypothesize that the word "cat" and "dog" are close to each other.

The vision part is pre-trained on the ImageNet classification task [45] with 1.2 million images and 1,000 categories. It provides rich visual attributes of the objects and scenes that are useful not only for the 1,000 classification task itself, but also for other vision tasks [11].

Combining cues from both language and vision, our model can effectively learn the new concepts using only a few examples as demonstrated in the experiments.

## 5. Datasets

### 5.1. Strategies to Construct Datasets

We use the annotations and images from the MS COCO [34] to construct our Novel Concept (NC) learning datasets. The current release of COCO contains 82,783 training images and 40,504 validation images, with object instance annotations and 5 sentence descriptions for each image. To construct the NC dataset with a specific new concept (e.g. "cat"), we remove all images containing the object "cat" according to the object annotations. We also check whether there are some images left with sentences descriptions containing cat related words. The remaining images are treated as the *Base Set* where we will train, validate and test our base model. The removed images are used to construct the *Novel Concept set (NC set)*, which is used to train, validate and test our model for the task of novel concept learning.

### 5.2. The Novel Visual Concepts Datasets

We construct three datasets involving five different novel visual concepts:

**NewObj-Cat** and **NewObj-Motor** The corresponding new concepts of these two datasets are *"cat"* and *"motorcycle"* respectively. The model need to learn all the related

| | Train | NC Test | Validation |
|---|---|---|---|
| NewObj-Cat | 2840 | 1000 | 490 |
| NewObj-Motor | 1854 | 600 | 349 |
| NC-3 | 150 (50 × 3) | 120 (40 × 3) | 30 (10 × 3) |

Table 1: The number of images for the three datasets.



a quidditch player kneeling with ball and broom while looking at players running at him

a close up of a black statue of a t-rex faces right and roars as it holds up a valentine's day sign

a woman wearing a light pink kimono kneels on a cushion while playing the samisen

Figure 5: Sample images and annotations from Novel Concept-3 (NC-3) dataset (see more on the project page).

words that describe these concepts and their activities.

**NC-3 dataset**[3] The two datasets mentioned above are all derived from the MS COCO dataset. To further verify the effectiveness of our method, we construct a new dataset contains three novel concepts: *"quidditch"* (a recently created sport derived from "Harry Potter"), *"t-rex"* (a dinosaur), and *"samisen"* (an instrument). It contains not only object concepts (e.g. t-rex and samisen), but also activity concepts (e.g. quidditch). We labeled 100 images for each concept with 5 sentence annotations for each image. To diversify the labeled sentences for different images in the same category, the annotators are instructed to label the images with different sentences by describing the details in each image. It leads to a different style of annotation from that of the MS COCO dataset. The average length of the sentences is also 26% longer than that of the MS COCO (13.5 v.s. 10.7). We construct this dataset for two reasons. Firstly, the three concepts are not included in the 1,000 categories of the ImageNet Classification task [45] where we pre-trained the vision component of our model. Secondly, this dataset has richer and more diversified sentence descriptions compared to NewObj-Cat and NewObj-Motor. We denote this dataset as Novel Concept-3 dataset (NC-3). Some samples images and annotations are shown in Figure 5.

We randomly separate the above three datasets into training, testing and validation sets. The number of images for the three datasets are shown in Table 1. To investigate the possible overfitting issues on these datasets, in the testing stage, we randomly picked images from the testing set of the Base Set and treated them as a separate set of testing images. The number of added images is equal to the size of the original test set (e.g. 1000 images are picked for NewObj-Cat testing set). We denote the original new concept testing images as *Novel Concept (NC) test* set and the added base testing images as *Base test* set. A good novel visual concept learning method should perform better than the base model on NC test set and comparable on Base test set. The organization of NC datasets is illustrated in Figure 4.

|  | B-1 | B-2 | B-3 | B-4 | METEOR | CIDEr | ROUGE_L |
|---|---|---|---|---|---|---|---|
| m-RNN [38] | 0.680 | 0.506 | 0.369 | 0.272 | 0.225 | 0.791 | 0.499 |
| ours-TWS | **0.685** | **0.512** | **0.376** | **0.279** | **0.229** | **0.819** | **0.504** |

Table 2: The performance comparisons of our model and m-RNN [38] for the standard image captioning task.

# 6. Experiments

## 6.1. Evaluation Metrics

To evaluate the output sentence descriptions for novel visual concepts, we adopt two evaluation metrics that are widely used in recent image captioning work: BLEU scores [44] (BLEU score for n-gram is denoted as as B-n in the paper) and METEOR [30].

Both BLEU scores and METEOR target on evaluating the overall quality of the generated sentences. In the NVCS task, however, we focus more on the accuracy for the new words than the previously learned words in the sentences. Therefore, to conduct a comprehensive evaluation, we also calculate the $f$ score for the words that describe the new concepts. E.g. for the cat dataset, there are 29 new words such as cat, cats, kitten, and pawing. The precision $p$ and recall $r$ for each new word in the dictionary $(w_n^d)$ are calculated as follows:

$$p = \frac{\mathbf{N}(w_n^d \in \mathbb{S}_{gen} \wedge w_n^d \in \mathbb{S}_{ref})}{\mathbf{N}(w_n^d \in \mathbb{S}_{gen})}; r = \frac{\mathbf{N}(w_n^d \in \mathbb{S}_{gen} \wedge w_n^d \in \mathbb{S}_{ref})}{\mathbf{N}(w_n^d \in \mathbb{S}_{ref})}$$

where $\mathbb{S}_{gen}$ denotes generated sentence, $\mathbb{S}_{ref}$ denotes reference sentences, $\mathbf{N}(condition)$ represents number of testing images that conform to the condition. Note that $p$ and $r$ are calculated on the combined testing set of the NC test set and the base test set (i.e. All test).

A high $r$ with a low $p$ indicates that the model overfits the new data (We can always get $r = 1$ if we output the new word every time) while a high $p$ with a low $r$ indicates underfitting. We use the $f = \frac{2}{p^{-1}+r^{-1}}$ as a balanced measurement between $p$ and $r$. Best $f$ score is 1. Note that $f = 0$ if either $p = 0$ or $r = 0$. Compared to METEOR and BLEU, the $f$ score show the effectiveness of the model to learn new concepts more explicitly.

## 6.2. Effectiveness of TWS and BPF

We test our base model with the Transposed Weight Sharing (TWS) strategy in the original image captaining task on the MS COCO [6] and compare to m-RNN [38], which does not use TWS. Our model performs better than m-RNN in this task as shown in Table 2. We choose the layer dimensions of our model so that the number of parameters matches that of [38]. Models with different hyperparameters, features or pipelines might lead to better performance, which is beyond the scope of this paper. E.g. [38, 55, 15] further improve their results after the submission of this draft and achieve a B-4 score of 0.302, 0.309 and 0.308 respectively using, e.g., fine-tuned image features on COCO or consensus reranking [9, 38], which are complementary with TWS.

|  | BiasFix | Centralize | TWS | $f$ |
|---|---|---|---|---|
| Deep-NVCS-UnfixedBias | × | × | √ | 0.851 |
| Deep-NVCS-FixedBias | √ | × | √ | 0.860 |
| Deep-NVCS-NoBPF-NoTWS | × | × | × | 0.839 |
| Deep-NVCS-BPF-NoTWS | √ | √ | × | 0.850 |
| Deep-NVCS-BPF-TWS | √ | √ | √ | **0.875** |

Table 3: Performance of Deep-NVCS models with different novel concept learning strategies on NewObj-Cat. TWS and BPF improve the performance.

We also validate the effectiveness of our Transposed Weight Sharing (TWS) and Baseline Probability Fixation (BPF) strategies for the novel concept learning task on the NewObj-Cat dataset. We compare the performance of five Deep-NVCS models. Their properties and performance in terms of $f$ score for the word "cat" are summarized in Table 3. "BiasFix" means that we fix the bias term $\mathbf{b}_n$ in Equation 3. "Centralize" means that we centralize the intermediate layer activation $\mathbf{x}$ (see Equation 4) so that $\mathbf{U}_D$ will not affect the baseline probability.

We achieve 2.5% increase of performance in terms of $f$ using TWS (Deep-NVCS-BPF-TWS v.s. Deep-NVCS-BPF-noTWS[4]), and achieves 2.4% increase using BPF (Deep-NVCS-BPF-TWS v.s. Deep-NVCS-UnfixedBias). We use Deep-NVCS to represent Deep-NVCS-BPF-TWS in short for the rest of the paper.

## 6.3. Results on NewObj-Motor and NewObj-Cat

### 6.3.1 Using all training samples

We show the performance of our Deep-NVCS models compared to strong baselines on the NewObj-Cat and NewObj-Motor datasets in Table 4. For Deep-NVCS, we only use the training data from the novel concept set. For Deep-NVCS-Inc1:1, we add training data randomly sampled from the training set of the base set. The number of added training images is the same as that of the training images for novel concepts. *Model-base* stands for the model trained only on the base set (no novel concept images). We implement a baseline model, *Model-word2vec*, where the weights of new words ($\mathbf{U}_{D_n}$) are calculated using a weighted sum of the weights of 10 similar concepts measured by the unsupervised learned word-embeddings from word2vec [41]. We also implement a strong baseline, *Model-retrain*, by retraining the whole model from scratch on the combined training set (training images from both the base set and the NC set).

The results show that compared to the Model-base which is only trained on base set, the Deep-NVCS models perform much better on the novel concept test set while reaching comparable performance on the base test set. Deep-NVCS also performs better than the Model-word2vec model. The

---

[4]We tried two versions of the model without TWS: (I). the model with multimodal layer directly connected to softmax layer like [38], (II). the model with an additional intermediate layer like TWS but does not share the weights. In our experiments, (I) performs slightly better than (II) so we report the performance of (I) here.

| Evaluation Metrics | $f$ | All test | | | | | NC test | | | | | Base test | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | B-1 | B-2 | B-3 | B-4 | METEOR | B-1 | B-2 | B-3 | B-4 | METEOR | B-1 | B-2 | B-3 | B-4 | METEOR |
| NewObj-Cat | | | | | | | | | | | | | | | | |
| Model-retrain | 0.866 | 0.689 | 0.531 | 0.391 | 0.291 | 0.228 | 0.696 | 0.549 | 0.403 | 0.305 | 0.227 | 0.683 | 0.513 | 0.379 | 0.277 | 0.229 |
| Model-base | 0.000 | 0.645 | 0.474 | 0.339 | 0.247 | 0.201 | 0.607 | 0.436 | 0.303 | 0.217 | 0.175 | 0.683 | 0.511 | 0.375 | 0.277 | 0.227 |
| Model-word2vec | 0.183 | 0.642 | 0.471 | 0.341 | 0.245 | 0.200 | 0.610 | 0.432 | 0.307 | 0.217 | 0.176 | 0.674 | 0.510 | 0.375 | 0.273 | 0.224 |
| Deep-NVCS | 0.875 | 0.682 | 0.521 | 0.382 | 0.286 | 0.224 | 0.684 | 0.534 | 0.392 | 0.299 | 0.224 | 0.680 | 0.508 | 0.372 | 0.274 | 0.225 |
| Deep-NVCS-1:1Inc | 0.881 | 0.683 | 0.523 | 0.385 | 0.288 | 0.226 | 0.686 | 0.538 | 0.398 | 0.303 | 0.226 | 0.679 | 0.507 | 0.371 | 0.273 | 0.225 |
| NewObj-Motor | | | | | | | | | | | | | | | | |
| Model-retrain | 0.797 | 0.697 | 0.526 | 0.386 | 0.284 | 0.240 | 0.687 | 0.512 | 0.368 | 0.263 | 0.244 | 0.707 | 0.539 | 0.404 | 0.305 | 0.236 |
| Model-base | 0.000 | 0.646 | 0.460 | 0.327 | 0.235 | 0.218 | 0.586 | 0.380 | 0.245 | 0.160 | 0.203 | 0.705 | 0.536 | 0.401 | 0.301 | 0.235 |
| Model-word2vec | 0.279 | 0.663 | 0.476 | 0.338 | 0.243 | 0.226 | 0.624 | 0.423 | 0.279 | 0.183 | 0.223 | 0.701 | 0.530 | 0.397 | 0.303 | 0.229 |
| Deep-NVCS | 0.784 | 0.688 | 0.512 | 0.373 | 0.276 | 0.236 | 0.677 | 0.494 | 0.349 | 0.252 | 0.241 | 0.698 | 0.530 | 0.398 | 0.299 | 0.231 |
| Deep-NVCS-1:1Inc | 0.790 | 0.687 | 0.512 | 0.374 | 0.280 | 0.235 | 0.672 | 0.492 | 0.347 | 0.256 | 0.237 | 0.702 | 0.532 | 0.401 | 0.303 | 0.234 |

Table 4: Results on the NewObj-Cat and NewObj-Motor dataset using all the training samples. The Deep NVCS models outperform the simple baselines. They achieve comparable performance with the strong baseline (i.e. Model-retrain) but only need ≤2% of the time. Model-base and Model-retrain stand for the model trained on base set (no novel concepts) and the model retrained on the combined data (all the images of base set and novel concept set) respectively. Model-word2vec is a baseline model based on word2vec [41]. Deep-NVCS stands for the model trained only with the new concept data. Deep-NVCS-1:1Inc stands for the Deep-NVCS model trained by adding equal number of training images from the base set.
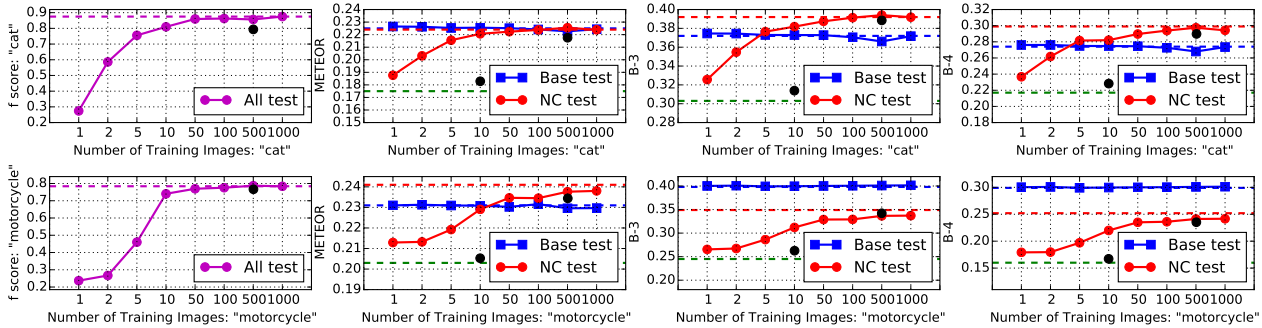


Figure 6: Performance comparison of our model with different number of training images on NewObj-cat and NewObj-Motor datasets. The red, blue, and magenta dashed line indicates the performance of Deep-NVCS using all the training images on the base test set, the NC test set and the all test set respectively. The green dashed line indicates the performance of Model-base. The black dots stand for the performance of Model-retrain for NC test. We show that our model trained with 10 to 50 images achieves comparable performance with the model trained on the full training set. (Best viewed in color)

performance of our Deep-NVCS models is very close to that of the strong baseline Model-retrain but needs *only less than 2% of the time*. This demonstrates the effectiveness of our novel concept learning strategies. The model learns the new words for the novel concepts without disturbing the previous learned words.

The performance of Deep-NVCS is also comparable with, though slightly lower than Deep-NVCS-1:1Inc. Intuitively, if the image features can successfully capture the difference between the new concepts and the existing ones, it is sufficient to learn the new concept only from the new data. However, if the new concepts are very similar to some previously learned concepts, such as cat and dog, it is helpful to present the data of both novel and existing concepts to make it easier for the model to find the difference.

### 6.3.2 Using a few training samples

We also test our model under the one or few-shot scenarios. Specifically, we randomly sampled $k$ images from the train-

ing set of NewObj-Cat and NewObj-Motor, and trained our Deep-NVCS model only on these images ($k$ ranges from 1 to 1000). We conduct the experiments 10 times and average the results to avoid the randomness of the sampling.

We show the performance of our model with different number of training images in Figure 6. We only show the results in terms of $f$ score, METEOR, B-3 and B-4 because of space limitation. The results of B-1 and B-2 and consistent with the shown metrics. The performance of the model trained with the full NC training set in the last section is indicated by the blue (Base test), red (NC test) or magenta (All test) dashed lines in Figure 6. These lines represent the experimental upper bounds of our model under the one or few-shot scenario. The performance of the Model-base is shown by a green dashed line. It serves as an experimental lower bound. We also show the results of Model-retrain for NC test with black dots in Figure 6 trained with 10 and 500 novel concepts images.

The results show that using about 10 to 50 training im-

| | | | | |
|---|---|---|---|---|
| Model-base: | a window with a window | a street with a lot of cars and a bus | a group of people playing frisbee in a park | a man in a blue shirt is riding a horse | a woman is holding a child |
| Deep-NVCS: | a cat standing in front of a window | a motorcycle parked in a parking lot with a car in the background | a group of quidditch players playing quidditch on a field | a t-rex is standing in a room with a man | a woman is holding a samisen in her hand |

Figure 7: The generated sentences for the test images from novel concept datasets. In these examples, cat, motorcycle, quidditch, t-rex and samisen are the novel concepts respectively.

ages, the model achieves comparable performance with the Deep-NVCS model trained on the full novel concept training set. In addition, using about 5 training images, we observe a nontrivial increase of performance compared to the base model. Our deep-NVCS also better handles the case for a few images and runs much faster than Model-retrain.

## 6.4. Results on NC-3

The NC-3 dataset has three main difficulties. Firstly, the concepts have very similar counterparts in the original image set, such as Samisen v.s. Guitar, Quidditch v.s. football. Secondly, the three concepts rarely appear in daily life. They are not included in the ImageNet 1,000 categories where we pre-trained our vision deep CNN. Thirdly, the way we describe the three novel concepts is somewhat different from that of the common objects included in the base set. The requirement to diversify the annotated sentences makes the difference of the style for the annotated sentences between NC-3 and MS COCO even larger. The effect of the difference in sentence style leads to decreased performance of the base model compared to that on the NewObj-Cat and NewObj-Motor dataset (see Model-base in Table 5 compared to that in Table 4 on NC test). Furthermore, it makes it harder for the model to hypothesize the meanings of new words from a few sentences.

Faced with these difficulties, our model still learns the semantic meaning of the new concepts quite well. The $f$ scores of the model shown in Table 5 indicate that the model successfully learns the new concepts with a high accuracy from only 50 examples.

It is interesting that Model-retrain performs very badly

| New Word | Five nearest neighbours |
|---|---|
| cat | kitten; tabby; puppy; calico; doll; |
| motorcycle | motorbike; moped; vehicle; motor; motorbikes; |
| quidditch | soccer; football; softball; basketball; frisbees; |
| t-rex | giraffe's; bull; pony; goat; burger; |
| samisen | guitar; wii; toothbrushes; purse; contents; |

Table 6: The five nearest neighbors of the new words as measured by the activation of the word-embedding layer.

on this dataset. It does not output the word "quidditch" and "samisen" in the generated sentences. The BLEU scores and METEOR are also very low. This is not surprising since there are only a few training examples (i.e. 50) for these three novel concepts and so it is easy to be overwhelmed by other concepts from the original MS COCO dataset.

## 6.5. Qualitative Results

In Table 6, we show the five nearest neighbors of the new concepts using the activation of the word-embedding layer learned by our Deep-NVCS model. It shows that the learned novel word embedding vectors captures the semantic information from both language and vision. We also show some sample generated sentence descriptions of the base model and our Deep-NVCS model in Figure 7.

## 7. Conclusion

In this paper, we propose the Novel Visual Concept learning from Sentences (NVCS) task. In this task, methods need to learn novel concepts from sentence descriptions of a few images. We describe a method that allows us to train our model on a small number of images containing novel concepts. This performs comparably with the model retrained from scratch on all of the data if the number of novel concept images is large, and performs better when there are only a few training images of novel concepts available. We construct three novel concept datasets where we validate the effectiveness of our method. These datasets have been released to encourage future research in this area.

| Evaluation Metrics | $f$ | B-3 | B-4 | MET. | $f$ | B-3 | B-4 | MET. |
|---|---|---|---|---|---|---|---|---|
| | quidditch | | | | t-rex | | | |
| Model-retrain | 0.000 | 0.196 | 0.138 | 0.120 | 0.213 | 0.224 | 0.141 | 0.105 |
| Model-base | 0.000 | 0.193 | 0.139 | 0.122 | 0.000 | 0.166 | 0.102 | 0.088 |
| Deep-NVCC | 0.854 | 0.237 | 0.167 | 0.168 | 0.861 | 0.247 | 0.144 | 0.187 |
| Deep-NVCC-1:1Inc | 0.863 | 0.244 | 0.170 | 0.170 | 0.856 | 0.242 | 0.132 | 0.186 |
| | samisen | | | | Base Test | | | |
| Model-retrain | 0.000 | 0.209 | 0.133 | 0.122 | - | 0.412 | 0.328 | 0.234 |
| Model-base | 0.000 | 0.177 | 0.105 | 0.122 | - | 0.414 | 0.325 | 0.240 |
| Deep-NVCC | 0.630 | 0.229 | 0.140 | 0.161 | - | 0.414 | 0.326 | 0.239 |
| Deep-NVCC-1:1Inc | 0.642 | 0.233 | 0.144 | 0.164 | - | 0.414 | 0.327 | 0.239 |

Table 5: Results of our model on the NC-3 Datasets.

# References

[1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. Vqa: Visual question answering. *arXiv*, 2015. 2

[2] S. Antol, C. L. Zitnick, and D. Parikh. Zero-shot learning via visual abstraction. In *ECCV*, pages 401–416. 2014. 2

[3] Y. Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009. 2, 4

[4] P. Bloom. *How children learn the meanings of words*. MIT press, 2002. 1

[5] S. Carey and E. Bartlett. Acquiring a single new word. 1978. 1

[6] X. Chen, H. Fang, T. Lin, R. Vedantam, S. Gupta, P. Dollr, and C. L. Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6

[7] X. Chen and C. L. Zitnick. Learning a recurrent visual representation for image caption generation. *arXiv preprint arXiv:1411.5654*, 2014. 2

[8] K. Cho, B. van Merrienboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 2

[9] J. Devlin, S. Gupta, R. Girshick, M. Mitchell, and C. L. Zitnick. Exploring nearest neighbor approaches for image captioning. *arXiv preprint arXiv:1505.04467*, 2015. 6

[10] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *arXiv preprint arXiv:1411.4389*, 2014. 2

[11] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. *arXiv preprint arXiv:1310.1531*, 2013. 5

[12] M. Elhoseiny, B. Saleh, and A. Elgammal. Write a classifier: Zero-shot learning using purely textual descriptions. In *ICCV*, pages 2584–2591, 2013. 2

[13] D. Elliott and F. Keller. Comparing automatic evaluation measures for image description. In *ACL*, volume 2, pages 452–457, 2014. 2

[14] J. L. Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. 2

[15] H. Fang, S. Gupta, F. Iandola, R. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. Platt, et al. From captions to visual concepts and back. *arXiv preprint arXiv:1411.4952*, 2014. 2, 6

[16] L. Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *TPAMI*, 28(4):594–611, 2006. 1, 2

[17] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NIPS*, pages 2121–2129, 2013. 2

[18] H. Gao, J. Mao, J. Zhou, Z. Huang, L. Wang, and W. Xu. Are you talking to a machine? dataset and methods for multilingual image question answering. In *NIPS*, 2015. 2

[19] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 2

[20] A. Gupta and P. Mannem. From image annotation to image description. In *ICONIP*, 2012. 2

[21] T. H. Heibeck and E. M. Markman. Word learning in children: An examination of fast mapping. *Child development*, pages 1021–1034, 1987. 1

[22] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2, 3

[23] N. Kalchbrenner and P. Blunsom. Recurrent continuous translation models. In *EMNLP*, pages 1700–1709, 2013. 2

[24] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. *arXiv preprint arXiv:1412.2306*, 2014. 2

[25] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014. 2

[26] B. Klein, G. Lev, G. Sadeh, and L. Wolf. Fisher vectors derived from hybrid gaussian-laplacian mixture models for image annotation. *arXiv preprint arXiv:1411.7399*, 2014. 2

[27] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 2

[28] G. Kulkarni, V. Premraj, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Baby talk: Understanding and generating image descriptions. In *CVPR*, 2011. 2

[29] B. M. Lake, R. Salakhutdinov, J. Gross, and J. B. Tenenbaum. One shot learning of simple visual concepts. In *CogSci*, volume 172, 2011. 2

[30] A. Lavie and A. Agarwal. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgements. In *Workshop on Statistical Machine Translation*, pages 228–231, 2007. 6

[31] A. Lazaridou, E. Bruni, and M. Baroni. Is this a wampimuk? cross-modal mapping between distributional semantics and the visual world. In *ACL*, pages 1403–1414, 2014. 1, 2

[32] R. Lebret, P. O. Pinheiro, and R. Collobert. Simple image description generator via a linear phrase-based approach. *arXiv preprint arXiv:1412.8419*, 2014. 2

[33] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012. 2

[34] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. *arXiv preprint arXiv:1405.0312*, 2014. 2, 5

[35] L. Ma, Z. Lu, L. Shang, and H. Li. Multimodal convolutional neural networks for matching image and sentence. *arXiv preprint arXiv:1504.06063*, 2015. 2

[36] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *NIPS*, pages 1682–1690, 2014. 2

[37] M. Malinowski and M. Fritz. A pooling approach to modelling spatial relations for image retrieval and annotation. *arXiv preprint arXiv:1411.5190*, 2014. 2

[38] J. Mao, W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn). In *ICLR*, 2015. 1, 2, 3, 4, 6

[39] J. Mao, W. Xu, Y. Yang, J. Wang, and A. L. Yuille. Explain images with multimodal recurrent neural networks. In *NIPS Deep Learning Workshop*, 2014. 2

[40] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur. Recurrent neural network based language model. In *INTERSPEECH*, pages 1045–1048, 2010. 2

[41] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013. 2, 6, 7

[42] M. Mitchell, X. Han, J. Dodge, A. Mensch, A. Goyal, A. Berg, K. Yamaguchi, T. Berg, K. Stratos, and H. Daumé III. Midge: Generating image descriptions from computer vision detections. In *EACL*, 2012. 2

[43] W. Ouyang, P. Luo, X. Zeng, S. Qiu, Y. Tian, H. Li, S. Yang, Z. Wang, Y. Xiong, C. Qian, et al. Deepid-net: multi-stage and deformable deep convolutional neural networks for object detection. *arXiv preprint arXiv:1409.3505*, 2014. 2

[44] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318, 2002. 6

[45] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge, 2014. 3, 5

[46] R. Salakhutdinov, J. Tenenbaum, and A. Torralba. One-shot learning with a hierarchical nonparametric bayesian model. 2010. 1

[47] V. Sharmanska, N. Quadrianto, and C. H. Lampert. Augmented attribute representations. In *ECCV*, pages 242–255. 2012. 2

[48] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 3

[49] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, pages 935–943, 2013. 2

[50] R. Socher, Q. Le, C. Manning, and A. Ng. Grounded compositional semantics for finding and describing images with sentences. In *TACL*, 2014. 2

[51] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112, 2014. 2

[52] D. Swingley. Fast mapping and slow mapping in children's word learning. *Language learning and Development*, 6(3):179–183, 2010. 1

[53] T. Tommasi, F. Orabona, and B. Caputo. Learning categories from few examples with multi model knowledge transfer. *TPAMI*, 36(5):928–941, 2014. 1, 2

[54] R. Vedantam, C. L. Zitnick, and D. Parikh. Cider: Consensus-based image description evaluation. *arXiv preprint arXiv:1411.5726*, 2014. 2

[55] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. *arXiv preprint arXiv:1411.4555*, 2014. 2, 6

[56] J. Weston, S. Bengio, and N. Usunier. Large scale image annotation: learning to rank with joint word-image embeddings. *Machine learning*, 81(1):21–35, 2010. 2

[57] K. Xu, J. Ba, R. Kiros, C. A. Cho, Kyunghyun, R. Salakhutdinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044*, 2015. 2

[58] M. D. Zeiler. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012. 4

[59] J. Zhu, J. Mao, and A. L. Yuille. Learning from weakly supervised data by the expectation loss svm (e-svm) algorithm. In *NIPS*, pages 1125–1133, 2014. 2