

Learning Robust Visual-Semantic Embeddings

Yao-Hung Hubert Tsai Liang-Kang Huang Ruslan Salakhutdinov

School of Computer Science, Machine Learning Department, Carnegie Mellon University

yaohungt@cs.cmu.edu liangkah@andrew.cmu.edu rsalakhu@cs.cmu.edu

Abstract

Many of the existing methods for learning joint embedding of images and text use only supervised information from paired images and its textual attributes. Taking advantage of the recent success of unsupervised learning in deep neural networks, we propose an end-to-end learning framework that is able to extract more robust multi-modal representations across domains. The proposed method combines representation learning models (i.e., auto-encoders) together with cross-domain learning criteria (i.e., Maximum Mean Discrepancy loss) to learn joint embeddings for semantic and visual features. A novel technique of unsupervised-data adaptation inference is introduced to construct more comprehensive embeddings for both labeled and unlabeled data. We evaluate our method on *Animals with Attributes* and *Caltech-UCSD Birds 200-2011* dataset with a wide range of applications, including zero and few-shot image recognition and retrieval, from inductive to transductive settings. Empirically, we show that our framework improves over the current state of the art on many of the considered tasks.

1. Introduction

Over the past few years, due to the availability of large amount of data and the advancement of the training techniques, learning effective and robust representations directly from images or text becomes feasible [19, 27, 32]. These learned representations have facilitated a number of high-level tasks, such as image recognition [40], sentence generation [17], and object detection [35]. Despite useful representations being developed for specific domains, learning more comprehensive representations across different data modalities remains challenging. In practice, more complex tasks, such as image captioning [45] and image tagging [23] often involve data from different modalities (i.e., images and text). Additionally, the learning process would be faster, requiring fewer labeled examples, and hence more scalable to handling a large number of categories if we could transfer cross-domain knowledge more effectively [9]. This motivates learning multi-modal em-




Figure 1. Illustration of our proposed ReViSE (Robust sEmi-supervised Visual-Semantic Embeddings) architecture.

beddings. In this paper, we consider learning robust joint embeddings across visual and textual modalities in an end-to-end fashion under zero and few-shot setting.

Zero-shot learning aims at performing specific tasks, such as recognition and retrieval of novel classes, when no label information is available during training [16]. On the other hand, few-shot learning enables us to have few labeled examples in our of-interest categories [38]. In order to compensate the missing information under the zero and few-shot setting, the model should learn to associate novel concepts in image examples with textual attributes and transfer knowledge from training to test classes. A common strategy for deriving the visual-semantic embeddings is to make use of images and textual attributes in a supervised way [41, 2, 48, 49, 50, 22, 7]. Specifically, one can learn transformations of images and textual attributes under the objective that the transformed visual and semantic vectors of the same class should be similar in the joint embeddings space. Despite good performance, this common strategy basically boils down to a supervised learning setting, learning from labeled or paired data only. In this paper, we show that to learn better joint embeddings across different data modalities, it is beneficial to combine supervised and unsupervised learning from both labeled and unlabeled data.

Our contributions in this work are as follows. First, to extract meaningful feature representations from both labeled and unlabeled data, one possible option is to train an

auto-encoder [33, 5]. In this way, instead of learning representations directly to align the visual and textual inputs, we choose to learn representations in an auto-encoder using reconstruction objective. Second, we impose a cross-modality distribution matching constraint to require the embeddings learned by the visual and textual auto-decoders to have similar distributions. By minimizing the distributional mismatch between visual and textual domain, we show improved performance on recognition and retrieval tasks. Finally, to achieve better adaptation on the unlabeled data, we perform a novel unsupervised-data adaptation inference technique. We show that by adopting this technique, the accuracy increases significantly not only for our method but also for many of the existing other models. Fig. 1 illustrates our overall end-to-end differentiable model.

To summarize, our proposed method successfully combines supervised and unsupervised learning objectives, and learns from both labeled and unlabeled data to construct joint embeddings of visual and textual data. We demonstrate improved performance on Animals with Attributes (AwA) [21] and Caltech-UCSD Birds 200-2011 [47] datasets on both image recognition and image retrieval tasks under zero and few-shot setting.

2. Related Work

In this section, we provide an overview of learning multi-modal embeddings across visual and textual domain.

Zero and Few-Shot Learning

Zero-shot [7, 1, 2] and few-shot learning [8, 39, 20] are related problems, but somewhat different in the setting of the training data. While few-shot learning aims to learn specific classes through one or few examples, zero-shot learning aims to learn even when no examples of the classes are presented. In this setting, zero-shot learning should rely on the side information provided by other domains. In the case of image recognition, this often comes in the form of textual descriptions. Thus, the focus of zero-shot image recognition is to derive joint embeddings of visual and textual data, so that the missing information of specific classes could be transferred from the textual domain.

Since the relation between raw pixels and text descriptions is non-trivial, most of the previous work relied on learning the embeddings through a large amount of data. Witnessing the success of deep learning in extracting useful representations, much of the existing work mostly applies deep neural networks to first transform raw pixels and text into more informative representations, followed by using various techniques to further identify the relation between them. For example, Socher *et al.* [41] used deep architectures [13] to learn representations for both images and text, and then used a Bayesian framework to perform classification. Norouzi *et al.* [29] introduced a simple idea

that treated classification scores output by the deep network [19] as weights in convex combination of word vectors. Fu *et al.* [10] proposed a method that learns projections from low-level visual and textual features to form a hypergraph in the embedding space and performed label propagation for recognition. A number of similar methods learn transformations from input image representations to the semantic space for the recognition or retrieval purposes [2, 49, 1, 48, 7, 50, 51, 30, 37, 6, 12].

A number of recent approaches also attempt to learn the entire task with deep models in an end-to-end fashion. Frome *et al.* [9] constructed a deep model that took visual embeddings extracted by CNN [19] and word embeddings as input, and trained the model with the objective that the visual and word embeddings of the same class should be well aligned under linear transformations. Ba *et al.* [22] predicted the output weights of both the convolutional and fully connected layers in a deep convolutional neural network. Instead of using textual attributes or word embeddings model, Reed *et al.* [34] proposed to train a neural language model directly from raw text with the goal of encoding only the relevant visual concepts for various categories.

Visual and Semantic Knowledge Transfer

Liu *et al.* [24] developed multi-task deep visual-semantic embeddings model for selecting video thumbnails based on side semantic information (i.e., title, description, and query). By incorporating knowledge about objects similarities between visual and semantic domains, Tang *et al.* [44] improved object detection in a semi-supervised fashion. Kottur *et al.* [18] proposed to learn visually grounded word embeddings (vis-w2v) and showed improvements over text only word embeddings (word2vec) on various challenging tasks. Reed *et al.* designed a text-conditional convolutional GAN architecture to synthesize an image from text. Recently, Wang *et al.* [46] introduced structure-preserving constraints in learning joint embeddings of images and text for image-to-sentence and sentence-to-image retrieval tasks.

Unsupervised Multi-modal Representations Learning

One of our key contributions is to effectively combine supervised and unsupervised learning tasks for learning multi-modal embeddings. This is inspired and supported by several previous works that provided evidence of how unsupervised learning tasks could benefit cross-modal feature learning.

Ngiam *et al.* [28] proposed various models based on Restricted Boltzmann Machine, Deep Belief Network, and Deep Auto-encoder to perform feature learning over multiple modalities. The derived multi-modal features demonstrated an improved performance over single-modal features on the audio-visual speech classification tasks. Srivastava and Salakhutdinov [42] developed a Multimodal Deep

Boltzmann Machine for fusing together multiple diverse modalities even when some of them are absent. Providing inputs of images and text, their generative model manifested noticeable performance improvement on classification and retrieval tasks.

3. Proposed Method

First, we define the problem setting and the corresponding notation. Let $\mathbf{V}_{tr} = \{v_i^{(tr)}\}_{i=1}^{N_{tr}}$ denote labeled training images from C_{tr} classes, $\mathbf{V}_{ut} = \{v_i^{(ut)}\}_{i=1}^{N_{ut}}$ denote unlabeled training images from C_{ut} possibly different classes, and $\mathbf{V}_{te} = \{v_i^{(te)}\}_{i=1}^{N_{te}}$ denote test images from C_{te} novel classes. For each class, following [49, 50, 51, 2, 48, 7], its textual attributes are either provided from human annotated attributes [21] or learned from unsupervised text corpora (Wikipedia) [32]. We denote these class-specific textual attributes as $\mathbf{T}_{tr} = \{t_c^{(tr)}\}_{c=1}^{C_{tr}}$, $\mathbf{T}_{ut} = \{t_c^{(ut)}\}_{c=1}^{C_{ut}}$, and $\mathbf{T}_{te} = \{t_c^{(te)}\}_{c=1}^{C_{te}}$ for labeled training, unlabeled training, and test classes, respectively.

Under zero-shot setting, our goal is to predict labels of the test images coming from *novel*, previously unseen, classes given textual attributes. That is, for a given test image $v_i^{(te)}$, its label is determined by

$$\arg \max_{c \in \{1, \dots, C_{te}\}} P_{\theta} \left(c \mid t_c^{(te)}, v_i^{(te)} \right), \quad (1)$$

where θ denotes model parameters. We will also consider a few-shot learning, where a few labeled training images are available in each of the test classes. In the following, we omit the model parameters θ for brevity.

3.1. Basic Formulation

The goal of learning multi-modal embeddings can be formulated as learning transformation functions f_v and f_t , such that given an image v and a textual attribute t , $f_v(v)$ should be similar to $f_t(t)$ if v and t are of the same class. Much of the previous work for learning multi-modal embeddings can be generalized to this formulation. For instance, in Cross-Modal Transfer (CMT) [41], $f_v(\cdot)$ can be viewed as a pre-defined feature extraction model followed by a two-layer neural network, while $f_t(\cdot)$ is set to an identity matrix. To be more specific, [41] aim at learning a non-linear projection directly from visual features to semantic word vectors.

Over the past few years, deep architectures have been shown to learn useful representations that could embed high-level semantics for both visual and textual data. This gives rise to the attempts of applying successful deep architectures to learn $f_v(\cdot)$ and $f_t(\cdot)$. For example, DeViSE [9] designed $f_v(\cdot)$ as a CNN model followed by a linear transformation matrix. On the other hand, they adopted the well known skip-gram text modeling architecture [27] for learning $f_t(\cdot)$ from raw text on Wikipedia. It is worth noting




Figure 2. Illustration of traditional visual-semantic embeddings model.

that, to further take advantage of previous success, these deep models are often pre-trained on large datasets where they have shown to learn effective representations.

Figure 2 shows the basic formulation of the visual-semantic embeddings model. Our method is built on top of this basic architecture by adding additional components as well as modifying existing ones.

3.2. Reconstructing Features from Auto-Encoder

Although the basic architecture provides a way to utilize label information during training, the learning process could further benefit if unlabeled data are provided at the same time. To be more specific, we propose to combine supervised and unsupervised learning objectives together by incorporating auto-encoders [4] for both image and text data. Typical setting of auto-encoders consists of a symmetric encoder-decoder architecture, with the hidden representations in the middle being compact representations that could be used to reconstruct the original input data. In our model, the auto-encoders are added after the image and text data are processed by the pre-trained networks. For learning visual embeddings, we use contractive auto-encoder proposed by [36], which is able to learn more robust visual codes for the images of same class. Given a visual feature vector \tilde{v} , the contractive auto-encoder maps \tilde{v} to a hidden representation \tilde{v}_h , and seeks to reconstruct \tilde{v} from \tilde{v}_h . Let us denote the reconstructed vector by $r_v(\cdot)$. Model parameters are thus learned by minimizing the regularized reconstruction error

$$\mathcal{L}_v = \frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} \|\tilde{v}_i - r_v(\tilde{v}_i)\|^2 + \gamma \|J(\tilde{v}_i)\|_F^2, \quad (2)$$

where $J(\cdot)$ is the Jacobian matrix [36].

On the other hand, for a given semantic feature vector or textual attribute t , we use a vanilla auto-encoder to first encode and then reconstruct from its hidden representation t_h . We hence minimize the reconstruction error

$$\mathcal{L}_t = \frac{1}{C_{tr}} \sum_{c=1}^{C_{tr}} \|t_c - r_t(t_c)\|^2. \quad (3)$$

Combining (2) and (3) gives us the reconstruction loss

$$\mathcal{L}_{reconstruct} = \mathcal{L}_v + \mathcal{L}_t. \quad (4)$$

In practice, if we have access to a large unlabeled set or a set of test inputs (without labels), we can easily incorporate them into the reconstruction loss. In our experimental results, we find that with the visual and textual auto-encoders, image and text data are transformed into visual and textual embeddings with more meaningful representations. In order to further transfer the knowledge across modalities, we impose discriminative constraints on the hidden representations (\tilde{v}_h and t_h) learned by these auto-encoders, as we discuss next.

3.3. Cross-Modality Distributions Matching

Distributions matching technique has been proven to be effective for transferring knowledge from one modality to another [31, 25, 14]. A common nonparametric way to analyze and compare distributions is to use Maximum Mean Discrepancy (MMD) [11] criterion. We can view MMD as a two-sample test on \tilde{v}_h and t_h , and thus its loss can be formulated as

$$\mathcal{L}_{MMD} = \|\mathbf{E}_p[\phi(\tilde{v}_h)] - \mathbf{E}_q[\phi(t_h)]\|_{\mathcal{H}_k}^2, \quad (5)$$

where p, q are the distributions of visual and textual embeddings (i.e., $\tilde{v}_h \sim p$ and $t_h \sim q$), ϕ is the feature map with canonical form $\phi(\mathbf{x}) = k(\mathbf{x}, \cdot)$, and \mathcal{H}_k is the reproducing kernel Hilbert space (RKHS) endowed with a characteristic kernel k . Note that the kernel in the MMD criterion must be a universal kernel, and thus we empirically choose a Gaussian kernel:

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\kappa \|\mathbf{x} - \mathbf{x}'\|^2\right). \quad (6)$$

We can now minimize the MMD criterion between visual and textual embeddings by minimize eq. (5). This can be further regarded as shrinking the gap between information across two data modalities. In our experiments, we find that the MMD loss helps improve model performance on both recognition and retrieval tasks in zero and few-shot setting.

3.4. Learning

After we derive the hidden representations \tilde{v}_h and t_h , the transformation functions $f_v(\cdot)$ and $f_t(\cdot)$ can be reformulated as

$$f_v(v) = f'_v(\tilde{v}_h) \text{ and } f_t(t) = f'_t(t_h), \quad (7)$$

where $f'_v(\cdot)$ and $f'_t(\cdot)$ are the mapping functions from the hidden representations to the visual and textual output.

To leverage the supervised information from labeled training images \mathbf{V}_{tr} and the corresponding textual attributes \mathbf{T}_{tr} , we minimize the binary prediction loss:

$$\mathcal{L}_{supervised} = -\frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} \sum_{c=1}^{C_{tr}} I_{i,c} \left\langle f'_v(\tilde{v}_{h,i}^{(tr)}), f'_t(t_{h,c}^{(tr)}) \right\rangle, \quad (8)$$

where $I_{i,c}$ indicates a $\{0, 1\}$ encoding of positive and negative classes and $\langle \cdot \rangle$ denotes a dot-product. It is worth noting that we can adopt different loss functions, including binary cross-entropy loss or multi-class hinge loss. However, empirically, using the simple binary prediction loss results in the best performance of our model.

Similar to eq. (8), we adopt the binary prediction loss for unlabeled training images \mathbf{V}_{ut} and the attributes \mathbf{T}_{ut} :

$$\mathcal{L}_{unlab}^{unsup} = -\frac{1}{N_{ut}} \sum_{i=1}^{N_{ut}} \sum_{c=1}^{C_{ut}} \tilde{I}_{i,c}^{(ut)} \left\langle f'_v(\tilde{v}_{h,i}^{(ut)}), f'_t(t_{h,c}^{(ut)}) \right\rangle, \quad (9)$$

where

$$\tilde{I}_{i,c}^{(ut)} = \begin{cases} 1 & \text{if } c = \arg \max_{c \in \{1, \dots, C_{ut}\}} \left\langle f'_v(\tilde{v}_{h,i}^{(ut)}), f'_t(t_{h,c}^{(ut)}) \right\rangle \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

We refer to eq. (9) as *unsupervised-data adaptation inference*, which acts as a self-reinforcing strategy using the unsupervised data with unknown labels. The intuition is that by minimizing eq. (9), we can further adapt our unlabeled data into the learning of $f'_v(\cdot)$ and $f'_t(\cdot)$ based on the empirical predictions. The choice of λ does influence its effectiveness. However, we find that setting $\lambda = 1.0$ works quite well for many methods we considered in this work.

In sum, our model learns by minimizing the total loss from both supervised and unsupervised objectives:

$$\mathcal{L}_{Total} = \mathcal{L}_{supervised} + \alpha \mathcal{L}_{unsupervised}, \quad (11)$$

where

$$\mathcal{L}_{unsupervised} = \mathcal{L}_{reconstruct} + \lambda \mathcal{L}_{unlab}^{unsup} + \beta \mathcal{L}_{MMD}, \quad (12)$$

with α, λ , and β representing the trade-off parameters for different components. Note that we can also view the unsupervised objective here as a regularizer for learning more robust visual and textual representations (see Figure 1 for our overall model architecture).

Before computing the loss, we find it useful to perform ℓ_2 normalization on the output scores $f_v(v)$ and $f_t(t)$ along the batch-wise direction. It can be viewed as a mixture of Batch Normalization [15] and Layer Normalization [3]. The idea is simple, we encourage the competence between different instances in the data batch, rather than across different categories.

Table 1. The statistics of AwA and CUB datasets. Images and classes are disjoint across *training(+ validation) / test* split.

Dataset	<i>training (+ validation)</i>		<i>test</i>	
	# of images	# of classes	# of images	# of classes
AwA [21]	24293	40	6180	10
CUB [47]	8855	150	2931	50

4. Experiments

In the experiments, we denote our proposed method as ReViSE (Robust sEmi-supervised Visual-Semantic Embeddings). Extensive experiments on zero and few-shot image recognition and retrieval tasks are conducted using two benchmark datasets: Animals with Attributes (AwA) [21] and Caltech-UCSD Birds 200-2011 (CUB) [47]. CUB is a fine-grained dataset in which the objects are both visually and semantically very similar, while AwA is a more general concept dataset. We use the same *training (+validation)/ test* splits as in [2, 48]. Table 1 lists the statistics of the datasets.

To verify the performance of our method, we consider two state-of-the-art deep-embeddings methods: CMT [41] and DeViSE [9]. CMT and DeViSE can be viewed as a special case of our proposed method with $\alpha = 0$ (without using unsupervised objective in eq. (11)). The difference between them is that DeViSE learns a nonlinear transformation on raw visual images and textual attributes for the alignment purpose, while CMT only learns the nonlinear transformation from visual to semantic embeddings.

We choose GoogLeNet [43] as the CNN model in DeViSE, CMT, and our architecture. For the textual attributes of classes, we consider three alternatives: human annotated attributes (*att*) [21], Word2Vec attributes (*w2v*) [27], and Glove attributes (*glo*) [32]. *att* are continuous attributes judged by humans: CUB contains 312 attributes and AwA contains 85 attributes. *w2v* and *glo* are unsupervised methods for obtaining distributed text representations of words. We use the pre-extracted Word2Vec and Glove vectors from Wikipedia provided by [2, 48]. Both *w2v* and *glo* are 400-dim. features.

4.1. Network Design and Training Procedure

Please see Supplementary for the design details of ReViSE and its parameters. Note that we report results averaged over 10 random trials.

4.2. Zero-Shot Learning

Following the partitioning strategy of [2, 48], we split AwA dataset into 30/10/10 classes and CUB dataset into 100/50/50 classes for *labeled training/ unlabeled training/ test* data. We adopt *att* attributes as a textual description of each class. For zero-shot learning, not only the labels of images are unknown in the *unlabeled training* and *test* set, but classes are also disjoint across *labeled training/ unlabeled training/ test* splits.

Table 2. Zero-shot recognition using top-1 classification accuracy (%). *att* attributes are used to describe each category.

Dataset <i>recognition for</i>	AwA		CUB		average top-1 acc.
	V_{ut}	V_{te}	V_{ut}	V_{te}	
using only labeled training data					
DeViSE [9]	60.8	63.0	38.9	36.8	49.9
CMT [41]	59.3	61.6	41.1	40.6	50.7
ReViSE ^a	60.3	61.2	46.4	45.0	53.2
ReViSE ^b	64.5	65.0	49.6	47.3	56.6
using both labeled and unlabeled training data					
DeViSE* [9]	76.0	63.7	37.2	36.2	53.3
CMT* [41]	77.8	58.5	39.9	39.8	54.0
ReViSE ^c	64.7	67.6	52.2	48.2	58.2
ReViSE	78.0	68.6	56.6	49.6	63.2

Table 3. Zero-shot retrieval using mean Average Precision (mAP) (%). *att* attributes are used to describe each category.

Dataset <i>retrieval for</i>	AwA		CUB		average mAP
	V_{ut}	V_{te}	V_{ut}	V_{te}	
using only labeled training data					
DeViSE [9]	60.5	61.6	32.6	31.5	46.6
CMT [41]	58.8	61.4	35.8	37.0	48.3
ReViSE ^a	60.2	60.1	33.1	32.0	46.4
ReViSE ^b	64.4	63.2	36.0	37.4	50.3
using both labeled and unlabeled training data					
DeViSE* [9]	65.6	58.9	35.4	31.3	47.8
CMT* [41]	63.5	57.1	39.7	38.0	49.6
ReViSE ^c	66.8	63.6	39.4	37.5	51.8
ReViSE	74.2	68.1	47.6	40.4	57.6

To verify how unlabeled training data could benefit the learning of ReViSE, we provide four variants: ReViSE^a, ReViSE^b, ReViSE^c, and ReViSE. ReViSE^a is when we only consider supervised objective. That is, $\alpha = 0$ in eq. (11). ReViSE^b is when we further take unsupervised objective in *labeled training data* into account; that is, only $\mathcal{L}_{reconstruct}$ and \mathcal{L}_{MMD} are considered in $\mathcal{L}_{unsupervised}$ (see eq. (12)) for labeled training data. Next, for ReViSE^c, we consider *unlabeled training data* in $\mathcal{L}_{unsupervised}$ without unsupervised-data adaptation technique (setting $\beta = 0$). Last, ReViSE denotes our complete training architecture.

For completeness, we also consider the technique of unsupervised-data adaptation inference (see section 3.4) for DeViSE [9] and CMT [41]. In other words, we also evaluate how DeViSE and CMT benefit from the unlabeled training data. We adopt the same procedure as in eq. (9) and report results as DeViSE* and CMT*, respectively.

Similar to [49, 50, 51], the results and comparisons are reported using top-1 classification accuracy (**top-1**) (see Table 2) and mean average precision (**mAP**) (see Table 3) for recognition and retrieval tasks, respectively, on the unlabeled training and test images. To be more specific, we define the prediction score as $\hat{y}_{i,c}^{(\cdot)} = \langle f'_v(\tilde{v}_{h,i}^{(\cdot)}), f'_t(t_{h,c}^{(\cdot)}) \rangle$ for a given image $v_i^{(\cdot)}$ and textual attributes $t_c^{(\cdot)}$ for class c . Results are provided after ranking $\hat{y}_{i,c}^{(\cdot)}$ on all unlabeled training or test classes.

Table 2 and 3 list the results for our zero-shot recognition and retrieval experiments. We first observe that **NOT** all the methods benefit from using unlabeled training data dur-

Table 4. Transductive zero-shot recognition using top-1 classification accuracy (%).

Dataset attributes	AwA			CUB			average top-1 acc.
	<i>att</i>	<i>w2v</i>	<i>glo</i>	<i>att</i>	<i>w2v</i>	<i>glo</i>	
test data not available during training							
DeViSE [9]	67.4	67.0	66.7	40.8	28.8	25.6	49.3
CMT [41]	67.6	69.5	68.0	42.4	29.6	25.7	50.5
test data available during training							
DeViSE* [9]	90.7	84.8	88.0	41.4	31.6	26.9	60.6
CMT* [41]	89.4	87.8	81.8	43.1	31.8	28.9	60.5
ReViSE ^{††}	92.1	92.3	90.3	62.4	30.0	27.5	65.8
ReViSE [†]	92.8	92.6	91.7	62.7	31.8	28.9	66.8
ReViSE ^c	73.0	67.0	73.4	53.7	26.4	28.2	53.6
ReViSE	93.4	93.5	92.2	65.4	32.4	31.5	68.1

ing training. For example, in AwA dataset for test images \mathbf{V}_{te} , there is a 2.7% retrieval deterioration from DeViSE to DeViSE* and a 3.1% recognition deterioration from CMT to CMT*. On the other hand, our proposed method enjoys 2.6% recognition improvement and 0.4% retrieval improvement from ReViSE^b to ReViSE^c. This shows that the learning method of our proposed architecture can actually benefit from unlabeled training data \mathbf{V}_{ut} and \mathbf{T}_{ut} .

Next, we examine different variants in our proposed architecture. Comparing the average results from ReViSE^a to ReViSE^b, we observe 3.4% recognition improvement and 3.9% retrieval improvement. This indicates that taking unsupervised objectives $\mathcal{L}_{reconstruct}$ and \mathcal{L}_{MMD} into account results in learning better feature representations and thus yields a better recognition/ retrieval performance. Moreover, when unsupervised-data adaptation technique is introduced, we enjoy 5.0% average recognition improvement and 5.8% average retrieval improvement from ReViSE^c to ReViSE. It is worth noting that the significant performance improvement for unlabeled training images \mathbf{V}_{ut} further verifies that our unsupervised-data adaptation technique leads to a more accurate prediction on \mathbf{V}_{ut} .

4.3. Transductive Zero-Shot Learning

In this subsection, we extend our experiments to a transductive setting, where test data are available during training. Therefore, the test data can now be regarded as the unlabeled training data ($\mathbf{V}_{tr} = \mathbf{V}_{ut}$ and $\mathbf{T}_{tr} = \mathbf{T}_{ut}$). To perform the experiments, as in Table 1, we split AwA dataset into 40/10 disjoint classes and CUB dataset into 150/50 disjoint classes for *labeled training/ test* data.

In order to evaluate different components in ReViSE, we further provide two variants: ReViSE[†] and ReViSE^{††}. ReViSE[†] is when we consider no distributional matching between the codes across modalities ($\beta = 0$). ReViSE^{††} is when we further consider no contractive loss in our visual auto-encoder ($\beta = \gamma = 0$). Similar to previous subsection, we also consider DeViSE*, CMT*, and ReViSE^c to evaluate the effect of our unsupervised-data adaptation inference.

Zero-Shot Recognition: Table 4 reports top-1 classification accuracy. Observe that ReViSE clearly outperforms

other state-of-the-art methods by a large margin. On average, we have at least 17% gain compared to the methods without using unsupervised objective and 7.5% gain compared to DeViSE* and CMT*. Note that all the methods work better on human annotated attributes (*att*) than on unsupervised attributes (*w2v* and *glo*) in CUB dataset. One possible reason is that for visually and semantically similar classes in a fine-grained dataset (CUB), attributes obtained in an unsupervised way (*glo* word vectors) cannot fully differentiate between them. Nonetheless, for the more general concept dataset AwA, using either supervised or unsupervised textual attributes, the performance does not differ by that much. For instance, our method achieves comparable performance using *att*, *w2v*, and *glo* (93.4%, 93.5%, and 92.2% top-1 classification accuracy) on AwA dataset.

The recognition performance for DeViSE* and CMT* (60.6% and 60.5% on average) compared to DeViSE and CMT (49.3% and 50.5% on average) further verifies that using unsupervised-data adaptation inference technique does benefit transductive zero-shot recognition. Furthermore, all of the variants of ReViSE using unsupervised-data adaptation inference (ReViSE^{††}, ReViSE[†], and ReViSE itself) have noticeable improvement over DeViSE* and CMT*. This shows that the proposed model succeeds in leveraging unsupervised information in test data for constructing more effective cross-modal embeddings.

Next, we evaluate the effects of different components designed in our architecture. First of all, we compare the results between ReViSE[†] (set $\beta = 0$) and ReViSE. The performance gain (66.8% to 68.1% on average) indicates that minimizing MMD distance between visual and textual codes enables our model to learn more robust visual-semantic embeddings. In other words, we can better associate cross-modal information when we match the distributions across visual and textual domains (please refer to Supplementary for the study of MMD distance). Second, we observe that, without contractive loss, performance slightly drop from 66.8% (ReViSE[†]) to 65.8% (ReViSE^{††}). This is not surprising since the contractive auto-encoder aims at learning less varying features/codes with similar visual input, and therefore we can expect to learn more robust visual codes. Finally, similar to the observations found in comparing DeViSE/CMT to DeViSE*/CMT*, the *unsupervised-data adaptation inference* in ReViSE substantially improves the average top-1 classification accuracy from 53.6% (ReViSE^c) to 68.1% (ReViSE). Please see Supplementary material for more detailed comparisons to the following non-deep-embeddings methods: SOC [30], ConSE [29], SSE [49], SJE [2], ESZSL [37], JLSE [50], LatEm [48], Sync [7], MTE [6], TMV [10], and SMS [12].

Zero-Shot Retrieval: In Table 5, we report zero-shot retrieval results by measuring the retrieval performance by mean average precision (mAP). On average, methods that




Figure 3. Images-retrieval experiments for CUB with *att* attributes.

Table 5. Transductive zero-shot retrieval using mean Average Precision (mAP) (%).

Dataset <i>attributes</i>	AwA			CUB			average mAP
	<i>att</i>	<i>w2v</i>	<i>glo</i>	<i>att</i>	<i>w2v</i>	<i>glo</i>	
test data not available during training							
DeViSE [9]	67.5	67.6	66.2	31.9	26.6	24.5	47.4
CMT [41]	66.3	70.6	69.5	39.3	25.2	21.9	48.8
test data not available during training							
DeViSE* [9]	82.3	78.0	84.4	36.9	25.8	21.3	54.8
CMT* [41]	85.8	77.3	73.0	44.1	28.9	28.1	56.2
ReViSE ^{††}	96.7	96.8	95.1	60.7	29.4	27.2	67.7
ReViSE [†]	97.2	96.9	96.4	62.0	29.8	28.2	68.4
ReViSE ^c	73.0	67.0	73.4	53.7	26.4	28.2	53.6
ReViSE	97.4	97.4	96.7	68.9	30.5	30.9	70.3

leverage unsupervised information yield better performance compared to the methods using no unsupervised objective. However, in few cases, the performance drops when we take unsupervised information into account. For example, on CUB dataset, DeViSE* performs unfavorably compared to DeViSE when *w2v* and *glo* word embeddings are used as textual attributes.

Overall, our method does help improve zero-shot retrieval by at least 14.1% compared to CMT*/DeViSE* and 21.5% compared to CMT/DeViSE. It clearly demonstrates the effectiveness of leveraging unsupervised information for improving zero-shot retrieval (please see Supplementary for the plot of precision-recall curves).

In addition to quantitative results, we also provide qualitative results of ReViSE. Fig. 3 is the image retrieval experiments for classes *Chestnut_sided_Warbler* and *White_eyed_Vireo*. Given a class embedding, the nearest image neighbors are retrieved based on the cosine similarity between transformed visual and textual features. We consider two conditions: images from the same class and images from all test classes. In *Chestnut_sided_Warbler*, most of the images (71.7%) are correctly classified, and we also observe that three nearest image neighbors are also in *Chestnut_sided_Warbler*. On the other hand, only 43.3% images are correctly classified in *White_eyed_Vireo*, and two of the three nearest image neighbors are from wrong class *Wilson_Warbler*.

Availability of Unlabeled Test Images: We next evaluate the performance of our method w.r.t. to the availability of




Figure 4. Fraction p of test images used for training ReViSE on transductive (a) zero-shot recognition (b) zero-shot retrieval for CUB dataset with *att* attributes.

test images for unsupervised objective (see Fig. 4) on CUB dataset with *att* attributes. We alter the fraction p of unlabeled test images used in the training stage from 0% to 100% by a step size of 10%. That is, in eq. (12), only p portion (randomly chosen) of test images contributes to $\mathcal{L}_{unsupervised}$. Fig. 4 clearly indicates the performance increases when p increases. That is, with more unsupervised information (test images) available, our model can better associate the supervised and unsupervised data. Another interesting observation is that with only 40% test images available, ReViSE achieves favorable performance on both transductive zero-shot recognition and retrieval.

Expand the test-time search space: Note that most of the methods [9, 41, 29, 49, 2, 50, 48, 7, 10] consider that, at test time, queries come from only test classes. For AwA dataset with *att* attributes, we expand the test-time search space to all training and test classes and perform transductive zero-shot recognition for DeViSE*, CMT*, and ReViSE. We discover severe performance drops from 90.7%, 89.4%, and 93.4% to 47.4%, 45.8%, and 42.5%. Similar results can also be observed in other non-deep-embeddings methods. Although challenging, it remains interesting to consider this generalized zero-/few-shot learning setting in our future work.

4.4. From Zero to Few-Shot Learning

In this subsection, we extend our experiments from transductive zero-shot to transductive few-shot learning. Compared to zero-shot learning, few-shot learning allows us to have a few labeled images in our test classes. Here, 3 images are randomly chosen to be labeled per test category. We use the same performance comparison metrics as in Sec. 4.2 to report the results.

Transductive Few-Shot Recognition and Retrieval: Tables 6 and 7 list the results of transductive few-shot recognition and retrieval tasks. Generally speaking, ReViSE achieves the best performance compared to its variants and other methods. Moreover, as expected, when we compare the results with transductive zero-shot recognition (Table 4) and retrieval (Table 5), every methods perform better when few (i.e., 3) labeled images are observed in the test classes. For example, for CUB dataset with *w2v* attributes, there is a 22.5% recognition improvement for CMT* and a 32.3%




Figure 5. (a) Original CNN features (b) Reconstructed features (c) Visual codes for AWA dataset in ReViSE under transductive zero-shot setting. We use *glo* as our textual attributes for classes. Different colors denote different classes. Best viewed in colors.




Figure 6. Output features of : (a) DeVISE* (b) CMT* (c) ReViSE. *glo* attributes are used on AWA dataset under transductive zero-shot setting. Different colors denote different classes. Best view in colors.

Table 6. Few-shot recognition comparison using top-1 classification accuracy (%). For each test class, 3 images are randomly labeled, while the rest are unlabeled.

Dataset <i>attributes</i>	AwA			CUB			average top-1 acc.
	<i>att</i>	<i>w2v</i>	<i>glo</i>	<i>att</i>	<i>w2v</i>	<i>glo</i>	
test data not available during training							
DeViSE [9]	80.9	75.3	79.4	54.0	45.7	46.0	63.6
CMT [41]	85.1	83.4	84.3	56.7	53.4	52.0	69.2
test data available during training							
DeViSE* [9]	92.6	91.1	91.3	57.5	50.7	52.9	72.7
CMT* [41]	90.6	90.2	91.1	62.5	54.3	55.4	74.0
ReViSE ^{††}	93.3	93.3	93.1	66.9	57.6	59.0	77.2
ReViSE [†]	93.3	93.8	93.5	67.7	59.6	60.0	78.0
ReViSE ^c	87.8	88.7	90.2	61.1	55.3	55.0	73.0
ReViSE	94.2	94.1	94.4	68.4	59.9	61.7	78.8

retrieval improvement for ReViSE.

We also observe that the performance gap between our proposed ReViSE and other methods is reduced compared to transductive zero-shot learning. For instance, in average retrieval performance, ReViSE has 15.5% mAP improvement over DeVISE* under zero-shot experiments, while only 9.3% improvement under few-shot experiments.

4.5. t-SNE Visualization

Figure 5 further shows the t-SNE [26] visualization for the original CNN features, the reconstructed visual features $r_v(\tilde{v}^{(te)})$, and the visual codes $\tilde{v}_h^{(te)}$ on AWA dataset with *glo* attributes under transductive zero-shot setting. First of all, observe that both the reconstructed features and the visual codes have more separate clusters over different classes, which suggest ReViSE has learned useful representations. Another interesting observation is that the affinities between classes might change after learning visual codes. For example, “leopard” images (green dots) are near “humpback whale” images (light purple dots) in the original CNN feature space. However, in the visual code space, leopard images are far from humpback whale images. One possible explanation is that we know leopard is semantically

Table 7. Few-shot retrieval comparison using mean Average Precision (mAP) (%). For each test class, 3 images are randomly labeled, while the rest are unlabeled.

Dataset <i>attributes</i>	AwA			CUB			average mAP
	<i>att</i>	<i>w2v</i>	<i>glo</i>	<i>att</i>	<i>w2v</i>	<i>glo</i>	
test data not available during training							
DeViSE [9]	85.0	79.3	84.9	46.4	42.6	42.9	63.5
CMT [41]	88.4	88.2	89.2	58.5	54.0	52.7	71.8
test data available during training							
DeViSE* [9]	96.7	95.5	95.8	47.5	49.2	51.6	72.7
CMT* [41]	95.3	94.8	95.8	60.0	54.7	56.4	76.2
ReViSE ^{††}	97.2	97.1	97.1	71.2	59.4	61.4	80.6
ReViSE [†]	97.3	97.5	97.4	72.5	61.4	62.5	81.4
ReViSE ^c	92.3	93.0	94.6	60.8	55.0	57.1	75.5
ReViSE	97.8	97.7	97.8	72.9	62.8	63.0	82.0

distinct from humpback whale, and thus their semantic attributes must also be very different. This leads to different image clusters in our designed framework.

Next, we provide the t-SNE visualization on the output visual test scores $f_v(\mathbf{V}_{te})$ for DeVISE*, CMT*, and ReViSE in Fig. 6. Clearly, ReViSE can better separate instances from different classes.

5. Conclusion

In this paper, we showed how we can augment a typical supervised formulation with unsupervised techniques for learning joint embeddings of visual and textual data. We empirically evaluate our proposed method on both general and fine-grained image classification datasets, with comparisons against the state-of-the-art methods in zero-shot and few-shot recognition and retrieval tasks, from inductive to transductive setting. In all the experiments, our method consistently outperforms other methods, substantially improving performance in some cases. We believe that this work sheds light on the advantages of combining supervised and unsupervised learning techniques, and makes a step towards learning more useful representations from multimodal data.

References

- [1] Z. Akata, M. Malinowski, M. Fritz, and B. Schiele. Multi-cue zero-shot learning with strong supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 59–68, 2016. 2
- [2] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, 2015. 1, 2, 3, 5, 6, 7
- [3] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [4] Y. Bengio. Learning deep architectures for ai. *Foundations and trends in Machine Learning*, 2(1):1–127, 2009. 3
- [5] Y. Bengio, L. Yao, G. Alain, and P. Vincent. Generalized denoising auto-encoders as generative models. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 899–907. Curran Associates, Inc., 2013. 2
- [6] M. Bucher, S. Herbin, and F. Jurie. Improving semantic embedding consistency by metric learning for zero-shot classification. In *European Conference on Computer Vision*, pages 730–746. Springer, 2016. 2, 6
- [7] S. Changpinyo, W.-L. Chao, B. Gong, and F. Sha. Synthesized classifiers for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5327–5336, 2016. 1, 2, 3, 6, 7
- [8] L. Fei-fei, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 28:2006, 2006. 2
- [9] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NIPS*, 2013. 1, 2, 3, 5, 6, 7, 8
- [10] Y. Fu, T. M. Hospedales, T. Xiang, and S. Gong. Transductive multi-view zero-shot learning. *IEEE transactions on pattern analysis and machine intelligence*, 37(11):2332–2345, 2015. 2, 6, 7
- [11] A. Gretton, K. M. Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola. A kernel method for the two-sample-problem. In *NIPS*, 2006. 4
- [12] Y. Guo, G. Ding, X. Jin, and J. Wang. Transductive zero-shot recognition via shared model space learning. In *AAAI*, volume 3, page 8, 2016. 2, 6
- [13] E. H. Huang, R. Socher, C. D. Manning, and A. Y. Ng. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 873–882. Association for Computational Linguistics, 2012. 2
- [14] Y.-H. Hubert Tsai, Y.-R. Yeh, and Y.-C. Frank Wang. Learning cross-domain landmarks for heterogeneous domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5081–5090, 2016. 4
- [15] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 448–456, 2015. 4
- [16] D. Jayaraman and K. Grauman. Zero-shot recognition with unreliable attributes. In *Advances in Neural Information Processing Systems*, pages 3464–3472, 2014. 1
- [17] R. Kiros, Y. Zhu, R. R. Salakhutdinov, R. Zemel, R. Urtasun, A. Torralba, and S. Fidler. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302, 2015. 1
- [18] S. Kottur, R. Vedantam, J. M. Moura, and D. Parikh. Visual word2vec (vis-w2v): Learning visually grounded word embeddings using abstract scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4985–4994, 2016. 2
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, page 2012. 1, 2
- [20] B. M. Lake, R. R. Salakhutdinov, and J. Tenenbaum. One-shot learning by inverting a compositional causal process. In *Advances in neural information processing systems*, pages 2526–2534, 2013. 2
- [21] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):453–465, 2014. 2, 3, 5
- [22] J. Lei Ba, K. Swersky, S. Fidler, et al. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *CVPR*, 2015. 1, 2
- [23] X. Li, S. Liao, W. Lan, X. Du, and G. Yang. Zero-shot image tagging by hierarchical semantic embedding. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 879–882. ACM, 2015. 1
- [24] W. Liu, T. Mei, Y. Zhang, C. Che, and J. Luo. Multi-task deep visual-semantic embedding for video thumbnail selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3707–3715, 2015. 2
- [25] M. Long, Y. Cao, J. Wang, and M. I. Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105, 2015. 4
- [26] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. 8
- [27] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013. 1, 3, 5
- [28] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011. 2
- [29] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. S. Corrado, and J. Dean. Zero-shot learning by convex combination of semantic embeddings. In *ICLR*, 2014. 2, 6, 7

- [30] M. Palatucci, D. Pomerleau, G. E. Hinton, and T. M. Mitchell. Zero-shot learning with semantic output codes. In *Advances in neural information processing systems*, pages 1410–1418, 2009. 2, 6
- [31] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010. 4
- [32] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–43, 2014. 1, 3, 5
- [33] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *CoRR*, abs/1511.06434, 2015. 2
- [34] S. Reed, Z. Akata, H. Lee, and B. Schiele. Learning deep representations of fine-grained visual descriptions. In *29th IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–58, 2016. 2
- [35] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems (NIPS)*, 2015. 1
- [36] S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio. Contractive auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 833–840, 2011. 3
- [37] B. Romera-Paredes and P. H. Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, pages 2152–2161, 2015. 2, 6
- [38] R. Salakhutdinov, J. B. Tenenbaum, and A. Torralba. One-shot learning with a hierarchical nonparametric bayesian model. In *ICML Unsupervised and Transfer Learning*, pages 195–206, 2012. 1
- [39] R. Salakhutdinov, J. B. Tenenbaum, and A. Torralba. Learning with hierarchical-deep models. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1958–1971, 2013. 2
- [40] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [41] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, 2013. 1, 2, 3, 5, 6, 7, 8
- [42] N. Srivastava and R. R. Salakhutdinov. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*, pages 2222–2230, 2012. 2
- [43] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. 5
- [44] Y. Tang, J. Wang, B. Gao, E. Dellandréa, R. Gaizauskas, and L. Chen. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2016)*. IEEE, 2016. 2
- [45] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3156–3164, 2015. 1
- [46] L. Wang, Y. Li, and S. Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5005–5013, 2016. 2
- [47] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010. 2, 5
- [48] Y. Xian, Z. Akata, G. Sharma, Q. Nguyen, M. Hein, and B. Schiele. Latent embeddings for zero-shot classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 69–77, 2016. 1, 2, 3, 5, 6, 7
- [49] Z. Zhang and V. Saligrama. Zero-shot learning via semantic similarity embedding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4166–4174, 2015. 1, 2, 3, 5, 6, 7
- [50] Z. Zhang and V. Saligrama. Zero-shot learning via joint latent similarity embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6034–6042, 2016. 1, 2, 3, 5, 6, 7
- [51] Z. Zhang and V. Saligrama. Zero-shot recognition via structured prediction. In *European Conference on Computer Vision*, pages 533–548. Springer, 2016. 2, 3, 5