Zero-shot learning via category-specific visual-semantic mapping and label refinement

Niu, Li; Cai, Jianfei; Veeraraghavan, Ashok; Zhang, Liqing

2018

Niu, L., Cai, J., Veeraraghavan, A., & Zhang, L. (2019). Zero-shot learning via category-specific visual-semantic mapping and label refinement. *IEEE Transactions on Image Processing*, 28(2), 965-979. doi:10.1109/tip.2018.2872916

<https://hdl.handle.net/10356/142785>

<https://doi.org/10.1109/TIP.2018.2872916>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/TIP.2018.2872916>.

Downloaded on 26 Apr 2023 23:01:53 SGT

Zero-Shot Learning via Category-Specific Visual-Semantic Mapping

Li Niu, Jianfei Cai, and Ashok Veeraraghavan

Abstract—Zero-Shot Learning (ZSL) aims to classify a test instance from an unseen category based on the training instances from seen categories, in which the gap between seen categories and unseen categories is generally bridged via visual-semantic mapping between the low-level visual feature space and the intermediate semantic space. However, the visual-semantic mapping (*i.e.*, projection) learnt based on seen categories may not generalize well to unseen categories, which is known as the projection domain shift in ZSL. To address this projection domain shift issue, we propose a method named Adaptive Embedding ZSL (AEZSL) to learn an adaptive visual-semantic mapping for each unseen category, followed by progressive label refinement. Moreover, to avoid learning visual-semantic mapping for each unseen category in the large-scale classification task, we additionally propose a deep adaptive embedding model named Deep AEZSL (DAEZSL) sharing the similar idea (*i.e.*, visual-semantic mapping should be category-specific and related to the semantic space) with AEZSL, which only needs to be trained once, but can be applied to arbitrary number of unseen categories. Extensive experiments demonstrate that our proposed methods achieve the state-of-the-art results for image classification on three small-scale benchmark datasets and one large-scale benchmark dataset.

I. INTRODUCTION

Conventional classification tasks require the training and test categories to be identical, but collecting annotated data for all categories is time-consuming and expensive in large-scale or fine-grained classification tasks. Therefore, Zero-Shot Learning (ZSL) [1], [2], [3], which aims to recognize the test instances from the categories previously unseen in the training stage, has become increasingly popular in the field of computer vision. In ZSL, the gap between seen (*i.e.*, training) categories and unseen (*i.e.*, testing) categories is generally bridged based on the intermediate semantic representations. Semantic representation of each category means the representation of this category in the semantic embedding space. One popular semantic representation is manually designed attribute vector [1], [4], which is a high-level description for each category, such as the shape (*e.g.*, cylindrical), material (*e.g.*, cloth), and color (*e.g.*, white). Besides, there also exist automatically generated semantic representations [5], [6], [7], [8] including the textual features extracted from online corpus corresponding to this category or the word vector of this category name.

Recently, based on the intermediate semantic representation (*i.e.*, semantic embedding), many ZSL approaches have been

proposed, which can be roughly categorized into Semantic Relatedness (SR) and Semantic Embedding (SE) methods according to [9], [10]. In particular, the first category SR methods [11], [12], [13] tend to learn the visual classifiers for unseen categories using the similarities between unseen categories and seen categories while the second category SE methods attract more research interest [1], [14], [15], [16], [17], [18], [19], [6], [20], [21], [22], [23], [5], [10], [24], [25], [26], [27], which aim to build the semantic link between visual features and semantic representations using mapping functions. More details of SR and SE methods can be found in Section II. For Semantic Embedding (SE) methods, a mapping function needs to be learnt from the seen categories in the training stage, and then applied to the unseen categories in the testing stage. However, the visual-semantic mappings between the seen categories and unseen categories might be considerably different. In this sense, the learnt mapping (*i.e.*, projection) may not generalize well to the test set and thus results in poor performance, leading to the recently highlighted projection domain shift, which was first mentioned in [20] and then theoretically proved in [24].

To cope with the projection domain shift, some semi-supervised or transductive (the two terminologies semi-supervised and transductive are often interchangeably used in the field of ZSL) ZSL approaches have been proposed by utilizing unlabeled test instances from unseen categories in the training stage. In general, these methods either learn a common mapping function for both seen and unseen categories, or adapt the mapping function learnt based on the seen categories to the unseen categories. Nevertheless, this may be insufficient to tackle the projection domain shift because the mapping function of each individual category varies significantly. For example, as shown in [20], the categories “Zebra” and “Pig” share the same “hasTail” attribute, but in fact the visual appearances of their tails differ greatly. Another example will be given in our experiments, where many categories have the common attributes “cloth” and “cluttered space” while the visual appearances of these two attributes are dramatically diverse for these categories (see Section V for more details).

To this end, we focus on the projection domain shift and tend to address this problem by learning a category-specific mapping function for each unseen category, which has never been studied before. Since we do not have labeled instances for the unseen categories, the mapping functions of unseen categories can only be learnt by transferring from those of seen categories. However, it is hard to tell which seen categories are semantically closer to a given unseen category *w.r.t.* certain entry in the semantic representation. Thus, we

L. Niu is with Electric and Computer Engineering (ECE) department in Rice University, Houston, TX 77005, USA (e-mail:ln7@rice.edu).

J. Cai is with the School of Computer Engineering, Nanyang Technological University, Singapore (e-mail:asjfc@ntu.edu.sg).

A. Veeraraghavan is with Electric and Computer Engineering (ECE) department in Rice University, Houston, TX 77005, USA (e-mail:vashok@rice.edu).

make a simple yet effective assumption that when the semantic representations of two categories are similar, the common non-zero entries of these two semantic representations should be semantically similar, and thus the mapping functions of these two categories should also be similar. Specifically, for each unseen category, we calculate the similarities between this unseen category and all the seen categories based on their semantic representations, and assign higher weights to the classification tasks corresponding to the more similar seen categories. This idea can be unified with many existing ZSL works with various forms of classification losses. As an instantiation, we build our method upon ESZSL [24] considering its simplicity and effectiveness, which is named as Adaptive Embedding ZSL (AEZSL). Note that for AEZSL, we only utilize the semantic representations of unseen categories to learn the mapping functions. In order to further utilize the unlabeled instances from unseen categories like previous semi-supervised or transductive ZSL works [20], [15], [18], [19], [10], we propose a label refinement strategy following AEZSL, which can refine the predicted test labels and update the visual classifiers for unseen categories alternatively by exploiting the relations among unseen categories and among test instances.

One problem of AEZSL is its inefficiency for large-scale classification task with a large number of unseen categories because one visual-semantic mapping needs to be learnt for each unseen category. Thus, we aim to design a model which only needs to be trained once on seen categories but can be applied to any unseen category without retraining the model. With this aim, we develop a Deep Adaptive Embedding ZSL (DAEZSL) model, which only utilizes seen categories in the training stage but has the generalization ability to arbitrary number of unseen categories. Specifically, instead of learning one mapping for each unseen category based on its semantic representation as for AEZSL, we target at learning a projection function from semantic representation to visual-semantic mapping. In this sense, given a new unseen category, its visual-semantic mapping can be easily generated based on its semantic representation. Nevertheless, the size of visual-semantic mapping matrix is usually very large. To further simplify the task, we opt for learning the projection from semantic representation to feature mask which can be applied on visual features via elementwise product, considering that the size of feature mask is much smaller than that of visual-semantic mapping matrix. In the meantime, transforming visual features by virtue of category-specific feature masks can implicitly adapt the visual-semantic mapping to different categories. In the testing stage, given a set of unseen categories associated with semantic representations, our model is able to generate category-specific feature masks for each unseen category, which can better fit the classification tasks for unseen categories.

Our contributions are threefold: 1) this is the first work to address the projection domain shift by learning category-specific visual-semantic mappings with the idea that higher weights should be assigned to the classification tasks of more relevant seen categories for each unseen category. As an instantiation, we build our AEZSL method upon ESZSL [24], followed by a progressive label refinement strategy; 2) we ad-

ditionally propose a deep adaptive embedding model DAEZSL for large-scale ZSL, which needs to be trained only once on seen categories but can be applied to arbitrary unseen category; 3) comprehensive experiments are conducted on three small-scale datasets and one large-scale dataset to demonstrate the effectiveness of our approaches.

II. RELATED WORK

As discussed in Section I, the existing Zero-Shot Learning (ZSL) methods can be categorized into Semantic Relatedness (SR) approaches and Semantic Embedding (SE) approaches. From another perspective, ZSL methods can be categorized into standard ZSL and transductive/semi-supervised ZSL based on whether to use the unlabeled test instances from unseen categories in the training stage. Moreover, several deep learning models have been proposed for ZSL. Additionally, since the terminology domain shift is commonly used in the field of domain adaptation, we briefly describe the difference of this terminology used in ZSL and domain adaptation.

Semantic Relatedness (SR) and Semantic Embedding (SE) ZSL: For SR approaches, the methods in [11], [13], [12] construct the visual classifiers for unseen categories from those for seen categories based on the semantic similarities between seen categories and unseen categories. Then, an extension was made in [28], which assumes that the visual classifiers for both seen and unseen categories can be represented by a set of base classifiers. The above SR approaches do not exhibit obvious projection domain shift problem, but they cannot take full advantage of the semantic representations. Moreover, the above works did not discuss how to utilize the unlabeled test instances in the training stage. In contrast, our methods can fully exploit the semantic representations and leverage the unlabeled test instances during the training procedure.

For SE approaches, three different strategies are employed to build the semantic link between the visual feature space and the semantic representation space as discussed in Section I. Firstly, the methods in [1], [14], [15], [16] are proposed for projecting the visual features to semantic representations based on the learnt attribute classifiers. Secondly, the approaches in [18], [19] target at projecting the semantic representations to visual features based on the learnt dictionary. Thirdly, the works in [6], [20], [21], [22], [23], [5], [10], [24], [25], [26], [27] tend to map the visual feature space and the semantic representation space into a common space, or learn a mapping function which measures the compatibility between visual features and semantic representations. Among the above SE approaches, the approaches in [22], [23], [14], [5], [1], [26] do not address the projection domain shift while the remaining transductive or semi-supervised approaches can somehow account for the projection domain shift problem, which will be detailed next.

Transductive or Semi-Supervised ZSL: Some transductive or semi-supervised Semantic Embedding (SE) methods [20], [15], [18], [19], [10], [25], [21], [29] can alleviate the projection domain shift by utilizing the unlabeled test instances from unseen categories in the training phase. Specifically, the projection domain shift is rectified in [20] by projecting visual

features and multiple semantic embeddings of test instances into a common subspace via Canonical Correlation Analysis (CCA). Label propagation for zero-shot learning is used in both [20] and [29]. The approach in [18] learns dictionaries for the seen categories and the unseen categories separately while enforcing these two dictionaries to be close. The methods in [10], [19], [25], [21] learn the mapping function and simultaneously infer the test labels. In [21], [15], a Laplacian regularizer is employed based on the semantic representations or the label representations of test instances. The above transductive or semi-supervised ZSL approaches are able to account for the projection domain shift problem. However, the approach in [20] requires multi-view semantic embeddings, which are often unavailable. Besides, the methods in [18], [10], [19], [25], [21], [15] either learn a common mapping function for all the categories or adapt the mapping function learnt from the seen categories to the unseen categories, which is likely to be improper for the real-world applications since the mapping function of each category may vary significantly. In contrast, our methods learn a category-specific mapping for each unseen category, which can capture diverse semantic meanings of the same attribute for different categories.

Deep ZSL: Compared with traditional ZSL methods based on extracted deep learning features, there are fewer end-to-end deep ZSL models [6], [7], [30], [31], [32]. These works learn the mapping from visual features to semantic representations [7], map visual feature space and semantic space to a common space [6], [30], [32], or directly learn classifiers for different categories based on their semantic representations [31]. However, none of the above deep ZSL methods mentions the projection domain shift issue. In contrast, we focus on tackling the projection domain shift in this paper and propose a deep adaptive embedding model DAEZSL which can adapt the visual-semantic mapping to different categories implicitly by learning category-specific feature mask.

Domain Adaptation: Domain adaptation methods aim to address the domain shift issues, *i.e.*, to reduce the domain distribution mismatch between the source domain (*i.e.*, training set) and the target domain (*i.e.*, test set). Regarding the domain shift problem, domain adaptation methods focus on the difference of marginal probability or class conditional probability between the source domain and the target domain while ZSL methods focus on the difference of visual-semantic mappings (*i.e.*, projections) between seen categories and unseen categories, which is referred to as projection domain shift. Thus, the terminology of domain shift has quite different meanings in these two fields. In this paper, we focus on the projection domain shift problem in ZSL.

III. BACKGROUND

In this paper, for ease of presentation, a vector/matrix is denoted by a lowercase/uppercase letter in boldface. The transpose of a vector/matrix is denoted using the superscript $'$. We use $\mathbf{A} \circ \mathbf{B}$ to denote the dot product of two matrices. Moreover, we use $\mathbf{I}$ to denote the identity matrix and $\mathbf{A}^{-1}$ to denote the inverse matrix of $\mathbf{A}$. We use upperscript s (*resp.*, t) to indicate seen (*resp.*, unseen) categories while omitting

the upperscript for not being specific with seen or unseen categories.

A. Problem Definition

Assume we have C^s seen categories and C^t unseen categories. Let us denote the training (*resp.*, test) data from seen (*resp.*, unseen) categories as $\mathbf{X}^s \in \mathcal{R}^{d \times n^s}$ (*resp.*, $\mathbf{X}^t \in \mathcal{R}^{d \times n^t}$), where d is the dimensionality of visual features and n^s (*resp.*, n^t) is the number of training (*resp.*, test) instances from seen (*resp.*, unseen) categories. We assume each category is associated with an a -dim semantic representation and thus the semantic representations of seen (*resp.*, unseen) categories can be stacked as $\mathbf{A}^s \in \mathcal{R}^{a \times C^s}$ (*resp.*, $\mathbf{A}^t \in \mathcal{R}^{a \times C^t}$). In order to bridge the gap between visual features and semantic representations, we develop our method using the third strategy of SE discussed in Section II, which can exploit the discriminative capacity of semantic representations as claimed in [24]. Specifically, inspired by [24], [25], [33], we use the mapping matrix $\mathbf{W} \in \mathcal{R}^{d \times a}$ to match visual features with semantic representations, *i.e.*, $\mathbf{X}^{s'}\mathbf{W}\mathbf{A}^s$, which measures the compatibility between instances and categories. Ideally, the most compatible semantic representation of each training instance should be from its ground-truth category. In the testing stage, a test instance $\mathbf{x}^t$ is assigned to the category corresponding to the maximum value in the C^t -dim vector $\mathbf{x}^{t'}\mathbf{W}\mathbf{A}^t$, in which $\mathbf{W}\mathbf{A}^t$ is essentially the stacked visual classifiers for unseen categories.

B. Embarrassingly Simple Zero-Shot Learning (ESZSL)

Before introducing our method, we briefly introduce Embarrassingly Simple Zero-Shot Learning (ESZSL) [24], based on which we build our own method considering its simplicity and effectiveness. ESZSL is formulated as

$$\min_{\mathbf{W}} \quad \|\mathbf{X}^{s'}\mathbf{W}\mathbf{A}^s - \mathbf{Y}^s\|_F^2 + \gamma\|\mathbf{W}\mathbf{A}^s\|_F^2 \quad (1)$$

$$+ \lambda\|\mathbf{X}^{s'}\mathbf{W}\|_F^2 + \beta\|\mathbf{W}\|_F^2, \quad (2)$$

in which γ , λ , and β are trade-off parameters, and $\mathbf{Y}^s \in \mathcal{R}^{n^s \times C^s}$ is a binary label matrix with the i -th row being the label vector of the i -th training instance. Note that in (1), $\|\mathbf{X}^{s'}\mathbf{W}\mathbf{A}^s - \mathbf{Y}^s\|_F^2$ can fully exploit the discriminability of semantic representations by assigning each instance to the category with the most compatible semantic representation, and $\|\mathbf{W}\mathbf{A}^s\|_F^2$ (*resp.*, $\|\mathbf{X}^{s'}\mathbf{W}\|_F^2$) is used to control the complexity of the projection of semantic embeddings (*resp.*, visual features) onto the visual feature (*resp.*, semantic embedding) space. By setting the derivative of (1) *w.r.t.* $\mathbf{W}$ as zeros, we can have

$$(\mathbf{X}^s\mathbf{X}^{s'} + \gamma\mathbf{I})\mathbf{W}\mathbf{A}^s\mathbf{A}^{s'} + \lambda(\mathbf{X}^s\mathbf{X}^{s'} + \frac{\beta}{\lambda}\mathbf{I})\mathbf{W} = \mathbf{X}^s\mathbf{Y}^s\mathbf{A}^{s'}. \quad (3)$$

when setting $\beta = \gamma\lambda$, the problem in (3) has a close-form solution:

$$\mathbf{W} = (\mathbf{X}^s\mathbf{X}^{s'} + \gamma\mathbf{I})^{-1}\mathbf{X}^s\mathbf{Y}^s\mathbf{A}^{s'}(\mathbf{A}^s\mathbf{A}^{s'} + \lambda\mathbf{I})^{-1}. \quad (4)$$

IV. OUR METHODS

In this section, we first build our AEZSL method upon ESZSL to learn category-specific mapping matrix in Section IV-A1 followed by a label refinement strategy in Section IV-A2. Moreover, we propose a deep adaptive embedding model named DAEZSL for large-scale ZSL in Section IV-B, which shares the similar idea with AEZSL.

A. Adaptive Embedding Zero-Shot Learning

In this section, we first introduce how to learn category-specific visual-semantic mapping, followed by describing our label refinement strategy.

1) *Category-Specific Visual-Semantic Mapping*: Since the visual-semantic mappings of different categories could be largely different, we aim to learn a category-specific mapping matrix for each unseen category (*i.e.*, $\mathbf{W}^c$ for the c -th unseen category) to address the projection domain shift. Because no labeled instances are provided for the unseen categories, we can only transfer from the mapping matrices of seen categories. However, it is a challenging task to determine which seen categories are more semantically similar to a given unseen category *w.r.t.* certain entry in the semantic representation. To facilitate the transfer, we assume that when the semantic representations of two categories are similar, the common non-zero entries shared by these two semantic representations should be semantically similar, and thus the mapping matrices of these two categories should also be similar.

Specifically, recall that in (1), the $\tilde{c}$ -th column of $\mathbf{X}^{s'}\mathbf{W}\mathbf{A}^s - \mathbf{Y}^s$ corresponds to the classification task for the $\tilde{c}$ -th seen category, and the tasks for different seen categories are dependent on one another with a common $\mathbf{W}$. Given the c -th unseen category, in order to ensure that $\mathbf{W}^c$ is close to the mapping matrices of those similar seen categories, we assign higher weights to the classification tasks for those more similar seen categories. In particular, we formulate this idea as $(\mathbf{X}^{s'}\mathbf{W}^c\mathbf{A}^s - \mathbf{Y}^s)\mathbf{S}^c$, where $\mathbf{S}^c \in \mathcal{R}^{C^s \times C^s}$ is a diagonal matrix with the $\tilde{c}$ -th diagonal element being the cosine similarity $s_{\tilde{c}}^c = \frac{\mathbf{a}_{\tilde{c}}^{t'} \mathbf{a}_{\tilde{c}}^s}{\|\mathbf{a}_{\tilde{c}}^{t'}\| \|\mathbf{a}_{\tilde{c}}^s\|}$, which measures the similarity between the semantic representation $\mathbf{a}_{\tilde{c}}^t$ of the c -th unseen category and the semantic representation $\mathbf{a}_{\tilde{c}}^s$ of the $\tilde{c}$ -th seen category. Note that $(\mathbf{X}^{s'}\mathbf{W}^c\mathbf{A}^s - \mathbf{Y}^s)\mathbf{S}^c$ is equivalent to multiplying the $\tilde{c}$ -th column of $\mathbf{X}^{s'}\mathbf{W}^c\mathbf{A}^s - \mathbf{Y}^s$ by $s_{\tilde{c}}^c$. For better explanation, assuming that the c -th unseen category is similar to the $\tilde{c}$ -th seen category and their semantic representations share the common non-zero entry indices $\{j_1, j_2, \dots, j_l\}$, then a higher weight $s_{\tilde{c}}^c$ should be assigned to $\mathbf{X}^{s'}\mathbf{W}^c\mathbf{a}_{\tilde{c}}^s - \mathbf{y}_{\tilde{c}}^s$ with $\mathbf{y}_{\tilde{c}}^s$ being the $\tilde{c}$ -th column of $\mathbf{Y}^s$. In this way, the learnt $\mathbf{W}^c$ should be closer to the mapping matrix of the $\tilde{c}$ -th seen category *w.r.t.* the $\{j_1, j_2, \dots, j_l\}$ -th columns. To this end, we tend to solve $\mathbf{W}^c$'s for all unseen categories simultaneously using the following formulation:

$$\min_{\mathbf{W}^c} \frac{1}{2} \sum_{c=1}^{C^t} \|(\mathbf{X}^{s'}\mathbf{W}^c\mathbf{A}^s - \mathbf{Y}^s)\mathbf{S}^c\|_F^2 + \frac{\lambda_1}{2} \sum_{c=1}^{C^t} \|\mathbf{X}^{s'}\mathbf{W}^c\|_F^2 + \frac{\lambda_2}{2} \sum_{c=1}^{C^t} \|\mathbf{W}^c\|_F^2 + \frac{\lambda_3}{2} \sum_{c < \tilde{c}} \|\mathbf{W}^c - \mathbf{W}^{\tilde{c}}\|_F^2, \quad (5)$$

where λ_1 , λ_2 , and λ_3 are trade-off parameters, which can be obtained using cross-validation, and $\sum_{c < \tilde{c}} \|\mathbf{W}^c - \mathbf{W}^{\tilde{c}}\|_F^2$ is a co-regularizer which encourages $\mathbf{W}^c$'s for different unseen categories to share some common parts. Note that we omit the regularizer $\|\mathbf{W}^c\mathbf{A}^s\|_F^2$ in (1) for ease of optimization. When λ_3 approaches Infinity, the mappings of all categories are enforced to be the same $\mathbf{W}$. In this case, the problem in (5) reduces to

$$\min_{\mathbf{W}} \frac{1}{2} \|(\mathbf{X}^{s'}\mathbf{W}\mathbf{A}^s - \mathbf{Y}^s)\tilde{\mathbf{S}}^c\|_F^2 + \frac{\lambda_1}{2} \|\mathbf{X}^{s'}\mathbf{W}\|_F^2 + \frac{\lambda_2}{2} \|\mathbf{W}\|_F^2,$$

in which the $\tilde{\mathbf{S}}^c$ is a diagonal matrix with the $\tilde{c}$ -th diagonal element being $\sqrt{\frac{\sum_{c=1}^{C^t} (s_{\tilde{c}}^c)^2}{C^t}}$. Compared with (1), we assign higher weights on the classification tasks corresponding to the seen categories which are closer to the overall unseen categories based on $\tilde{\mathbf{S}}^c$.

To solve the problem in (5), we update each $\mathbf{W}^c$ by fixing all the other $\mathbf{W}^{\tilde{c}}$'s for $\tilde{c} \neq c$ in an alternating fashion. Specifically, by setting the derivative of (5) *w.r.t.* $\mathbf{W}^c$ to zeros, we obtain the following equation:

$$(\mathbf{X}^s\mathbf{X}^{s'})\mathbf{W}^c(\mathbf{A}^s\mathbf{S}^c\mathbf{S}^{c'}\mathbf{A}^{s'} + \lambda_1\mathbf{I}) + ((C^t - 1)\lambda_3 + \lambda_2)\mathbf{W}^c = \mathbf{X}^s\mathbf{Y}^s\mathbf{S}^c\mathbf{S}^{c'}\mathbf{A}^{s'} + \lambda_3 \sum_{\tilde{c} \neq c} \mathbf{W}^{\tilde{c}}. \quad (6)$$

By denoting $\mathbf{L} = \mathbf{X}^s\mathbf{X}^{s'}$, $\mathbf{T} = \mathbf{A}^s\mathbf{S}^c\mathbf{S}^{c'}\mathbf{A}^{s'} + \lambda_1\mathbf{I}$, $\mathbf{N} = \mathbf{X}^s\mathbf{Y}^s\mathbf{S}^c\mathbf{S}^{c'}\mathbf{A}^{s'} + \lambda_3 \sum_{\tilde{c} \neq c} \mathbf{W}^{\tilde{c}}$, and $\mu = (C^t - 1)\lambda_3 + \lambda_2$, the problem in (6) becomes a special case of Sylvester equation *w.r.t.* $\mathbf{W}^c$ as follows,

$$\mathbf{L}\mathbf{W}^c\mathbf{T} + \mu\mathbf{W}^c = \mathbf{N}. \quad (7)$$

Since $\mathbf{L}$ and $\mathbf{T}$ are symmetric real matrices, the problem in (7) has an efficient solution. Inspired by [34], we perform Singular Value Decomposition (SVD) on $\mathbf{L}$ and $\mathbf{T}$, *i.e.*, $\mathbf{L} = \hat{\mathbf{U}}\Sigma^L\hat{\mathbf{U}}'$ and $\mathbf{T} = \hat{\mathbf{V}}\Sigma^T\hat{\mathbf{V}}'$ with Σ^L (*resp.*, Σ^T) being a diagonal matrix, in which the i -th diagonal element is σ_i^L (*resp.*, σ_i^T). Then, we can rewrite (7) as

$$\hat{\mathbf{U}}\Sigma^L\hat{\mathbf{U}}'\mathbf{W}^c\hat{\mathbf{V}}\Sigma^T\hat{\mathbf{V}}' + \mu\mathbf{W}^c = \mathbf{N}. \quad (8)$$

After multiplying (8) by $\hat{\mathbf{U}}'$ (*resp.*, $\hat{\mathbf{V}}$) on the left (*resp.*, right) and considering the orthonormality of $\hat{\mathbf{U}}$ (*resp.*, $\hat{\mathbf{V}}$), we have

$$\Sigma^L\hat{\mathbf{U}}'\mathbf{W}^c\hat{\mathbf{V}}\Sigma^T + \mu\hat{\mathbf{U}}'\mathbf{W}^c\hat{\mathbf{V}} = \hat{\mathbf{U}}'\mathbf{N}\hat{\mathbf{V}}. \quad (9)$$

By denoting $\hat{\mathbf{W}}^c = \hat{\mathbf{U}}'\mathbf{W}^c\hat{\mathbf{V}}$ and $\hat{\mathbf{N}} = \hat{\mathbf{U}}'\mathbf{N}\hat{\mathbf{V}}$, we can arrive at

$$\Sigma^L\hat{\mathbf{W}}^c\Sigma^T + \mu\hat{\mathbf{W}}^c = \hat{\mathbf{N}}. \quad (10)$$

Based on (10), we can easily solve each entry in $\hat{\mathbf{W}}^c$ as $\hat{W}_{ij}^c = \frac{\hat{N}_{ij}}{\sigma_i^L\sigma_j^T + \mu}$. Note that $\mathbf{L}$ is positive semi-definite, $\mathbf{T}$ is positive definite, and $\mu > 0$, so $\sigma_i^L\sigma_j^T + \mu \neq 0$, $\forall i, j$, and the problem in (10) has unique solution. Then, we can recover $\mathbf{W}^c$ by using $\mathbf{W}^c = \hat{\mathbf{U}}\hat{\mathbf{W}}^c\hat{\mathbf{V}}'$. Because SVD on $\mathbf{L}$ and $\mathbf{T}$ can be precomputed, our algorithm is quite efficient. We update each $\mathbf{W}^c$ alternatively until the objective of (5) converges. We name this method as Adaptive Embedding ZSL (AEZSL) and the algorithm to solve (5) is listed in Algorithm 1.

After learning $\mathbf{W}^c$'s, we can obtain the visual classifier for the c -th unseen category as $\mathbf{p}^c = \mathbf{W}^c\mathbf{a}_{\tilde{c}}^t$. In the testing stage,

Algorithm 1 The algorithm to solve AEZSL (5)

```

1: Input:  $\mathbf{X}^s, \mathbf{Y}^s, \mathbf{A}^s, \mathbf{S}^c, \lambda_1, \lambda_2, \lambda_3$ 
2: Initialize all  $\mathbf{W}^c$ 's equally using ESZSL [24].
3: repeat
4:   for  $c = 1 : C^t$  do
5:     Update  $\mathbf{W}^c$  by solving (7).
6:   end for
7: until The objective of (5) converges.
8: Output:  $\mathbf{W}^c$ 's.

```

given a test instance $\mathbf{x}^t$ from unseen categories, we can obtain its decision value for the c -th unseen category as $\mathbf{p}^c \mathbf{x}^t$.

Discussion: The projection domain shift problem has been theoretically proved in [24], in which the seen (*resp.*, unseen) categories are referred to as the source (*resp.*, target) domain. Then, two extreme cases are presented in [24] when the semantic representations of two domains are identical or orthogonal. Specifically, when the semantic representations of two domains are identical (*i.e.*, $\mathbf{a}_c^t = \mathbf{a}_c^s$), the error bound of the learnt classifier is approximate to that of a standard classifier without projection domain shift. In this case, we have $s_c^c = 1$, which allows the maximum transfer. When their semantic representations are orthogonal (*i.e.*, $\mathbf{a}_c^t \mathbf{a}_c^s = 0$), the error bound is vacuous and no transfer can be done. In this case, we have $s_c^c = 0$, which means no transfer. So the analysis for our problem accords with that in [24], which verifies that it is reasonable to control how much to transfer from seen categories based on the cosine similarities.

In this paper, we build our AEZSL method upon ESZSL due to its simplicity and effectiveness. However, it is worth mentioning that the idea of AEZSL, *i.e.*, assigning higher weights on the classification tasks of more similar seen categories for each unseen category, can be incorporated into many existing ZSL frameworks with slight modification based on their learning paradigms and used losses.

2) *Progressive Label Refinement:* Note that in Section IV-A1, we only utilize the semantic representations of unseen categories to learn the adaptive mapping matrices. In order to further adapt the mapping matrices to the unseen categories by utilizing the unlabeled test instances, we propose a progressive approach to update visual classifiers and refine predicted test labels alternatively, similar to some progressive semi-supervised learning approaches like [35]. Unlike traditional semi-supervised learning which requires the labeled and unlabeled instances to be from the same set of categories, zero-shot learning does not have any labeled instances from the unseen categories. So we divide the test set into a confident set $\mathcal{L}$ and an unconfident set $\mathcal{U}$, in which the labels in $\mathcal{L}$ (*i.e.*, $\mathbf{Y}^l$) are expected to be relatively more accurate than those in $\mathcal{U}$ (*i.e.*, $\mathbf{Y}^u$). Note that $\mathbf{Y}^l$ and $\mathbf{Y}^u$ are both binary label matrices, similar to $\mathbf{Y}^s$ in (5). Initially, $\mathcal{L}$ is an empty set and $\mathcal{U}$ is the entire test set with initial labels predicted by the initial classifiers $\mathbf{p}^c$'s, which are obtained based on $\mathbf{W}^c$'s from Section IV-A1. Then, we move k most confident instances from $\mathcal{U}$ into $\mathcal{L}$, and update the visual classifiers based on the new confident and unconfident sets. We repeat the above process iteratively until $\mathcal{U}$ becomes empty and output $\mathbf{Y}^l$ as final predicted test labels. The details of each iteration will be

described as follows, in which we stack $\mathbf{p}^c$'s for all unseen categories as $\mathbf{P}$, and split $\mathbf{X}^t$ into $\mathbf{X}^l$ and $\mathbf{X}^u$ corresponding to $\mathcal{L}$ and $\mathcal{U}$.

In each iteration, we first use the visual classifiers $\mathbf{P}$ from the previous iteration to select k most confident instances from $\mathcal{U}$ based on the confidence score, which is defined as a softmax function $\text{conf}(\mathbf{x}^t) = \frac{\exp(\mathbf{p}^{c(\mathbf{x}^t)} \mathbf{x}^t)}{\sum_c \exp(\mathbf{p}^c \mathbf{x}^t)}$ with $c(\mathbf{x}^t)$ being the assigned label of $\mathbf{x}^t$ corresponding to $\mathbf{p}^{c(\mathbf{x}^t)}$ which achieves the highest prediction score on $\mathbf{x}^t$. Note that the labels of the selected k instances are updated as $c(\mathbf{x}^t)$ while the labels of the remaining instances in $\mathcal{U}$ stay unchanged.

After moving the k most confident instances from $\mathcal{U}$ into $\mathcal{L}$, we update $\mathbf{P}$ by changing some predicted labels in $\mathcal{U}$ (*i.e.*, $\mathbf{X}^u \mathbf{P}$) while using $\mathcal{L}$ as weak supervision. In particular, we tend to keep the predicted labels of confident instances unchanged while changing the predicted labels of some unconfident instances selected by a group-lasso regularizer [36] $\|\mathbf{X}^u \mathbf{P} - \mathbf{Y}^u\|_{2,1}$, in which the rows with non-zero entries correspond to the selected instances. We change the predicted labels of selected instances based on the similarities among unconfident instances and the similarities among unseen categories, which will be explained as follows. 1) For the similarities among unconfident instances, we employ a standard Laplacian regularizer based on the smoothness assumption that the predicted label vectors of two unconfident instances should be close to each other when their visual features are similar. 2) For the similarities among unseen categories, we construct a transition matrix $\hat{\mathbf{S}}$ to characterize the probabilities that one category label is changed to another, in which $\hat{S}_{i,j}$ is the cosine similarity $\frac{\mathbf{a}_i^t \mathbf{a}_j^t}{\|\mathbf{a}_i^t\| \|\mathbf{a}_j^t\|}$, similar to that in Section IV-A1.

With the transition matrix $\hat{\mathbf{S}}$, we employ a coherent regularizer $\text{tr}(\mathbf{Y}^u \hat{\mathbf{S}} \mathbf{P}' \mathbf{X}^u)$, which enforces the predicted labels $\mathbf{P}' \mathbf{X}^u$ to be coherent with the expected transited labels $\mathbf{Y}^u \hat{\mathbf{S}}$ based on the transition probabilities. Note that we set the diagonal elements of $\hat{\mathbf{S}}$ as zeros to encourage the labels to be changed. To this end, the formulation to update $\mathbf{P}$ in each iteration can be written as

$$\min_{\mathbf{P}} \frac{1}{2} \|\mathbf{X}^l \mathbf{P} - \mathbf{Y}^l\|_F^2 + \frac{\gamma_1}{2} \|\mathbf{X}^u \mathbf{P} - \mathbf{Y}^u\|_{2,1} - \gamma_2 \text{tr}(\mathbf{Y}^u \hat{\mathbf{S}} \mathbf{P}' \mathbf{X}^u) + \frac{\gamma_3}{2} \text{tr}(\mathbf{P}' \mathbf{X}^u \mathbf{H}^u \mathbf{X}^u \mathbf{P}), \quad (11)$$

where γ_1 , γ_2 , and γ_3 are trade-off parameters, which can be obtained using cross-validation, and $\mathbf{H}^u$ is the Laplacian matrix constructed based on $\mathbf{X}^u$ following [21].

The problem in (11) is not easy to solve due to the regularizer $\|\mathbf{X}^u \mathbf{P} - \mathbf{Y}^u\|_{2,1}$. According to [37], solving (11) is equivalent to solving the following problem iteratively:

$$\min_{\mathbf{P}} \frac{1}{2} \|\mathbf{X}^l \mathbf{P} - \mathbf{Y}^l\|_F^2 + \frac{\gamma_1}{2} \text{tr}((\mathbf{P}' \mathbf{X}^u - \mathbf{Y}^u) \mathbf{D} (\mathbf{X}^u \mathbf{P} - \mathbf{Y}^u)) - \gamma_2 \text{tr}(\mathbf{Y}^u \hat{\mathbf{S}} \mathbf{P}' \mathbf{X}^u) + \frac{\gamma_3}{2} \text{tr}(\mathbf{P}' \mathbf{X}^u \mathbf{H}^u \mathbf{X}^u \mathbf{P}), \quad (12)$$

where $\mathbf{D}$ is a diagonal matrix with the i -th diagonal element being $\frac{1}{2\|\mathbf{q}_i\|_2}$, in which $\mathbf{q}_i$ is the i -th row of $\mathbf{Q}$ with $\mathbf{Q} = \mathbf{X}^u \mathbf{P} - \mathbf{Y}^u$. In each iteration of solving $\mathbf{P}$, we calculate $\mathbf{D}$ based on previous $\mathbf{P}$ and then update $\mathbf{P}$ by solving (12). As claimed in [37], the updated $\mathbf{P}$ in each iteration will decrease the objective of (11).

Algorithm 2 The algorithm of progressive label refinement

- 1: **Input:** $\mathbf{X}^u, \mathbf{Y}^u, \hat{\mathbf{S}}, k, \gamma_1, \gamma_2, \gamma_3$
 - 2: Initialize $\mathcal{L}$ as $\emptyset$ and $\mathcal{U}$ as the entire test set. Initialize $\mathbf{P}$ based on the learnt $\mathbf{W}^c$'s from Algorithm 1.
 - 3: **repeat**
 - 4: Move k most confident instances selected by $\mathbf{P}$ from $\mathcal{U}$ to $\mathcal{L}$. Update $\mathbf{X}^l, \mathbf{X}^u, \mathbf{Y}^l, \mathbf{Y}^u, \mathbf{H}^u$ accordingly.
 - 5: **repeat**
 - 6: Update $\mathbf{D}$ based on its definition below (12).
 - 7: Update $\mathbf{P}$ by using (13).
 - 8: **until** The objective of (11) converges.
 - 9: **until** The unconfident set $\mathcal{U}$ is empty.
 - 10: **Output:** $\mathbf{Y}^l, \mathbf{P}$.
-

By setting the derivative of (12) w.r.t. $\mathbf{P}$ to zeros, we can obtain the close-form solution to $\mathbf{P}$ as

$$\mathbf{P} = (\mathbf{X}^l \mathbf{X}^{l'} + \gamma_1 \mathbf{X}^u \mathbf{D} \mathbf{X}^{u'} + \gamma_3 \mathbf{X}^u \mathbf{H}^u \mathbf{X}^{u'} + \nu \mathbf{I})^{-1} (\mathbf{X}^l \mathbf{Y}^l + \gamma_1 \mathbf{X}^u \mathbf{D} \mathbf{Y}^u + \gamma_2 \mathbf{X}^u \mathbf{Y}^u \hat{\mathbf{S}}), \quad (13)$$

in which ν is a very small number (*i.e.*, $\nu \rightarrow 0$). We update $\mathbf{P}$ by solving (12) iteratively until the objective of (11) converges. According to [37], the output $\mathbf{P}$ is the global optimum solution to (11) because $\frac{1}{2} \|\mathbf{X}^{l'} \mathbf{P} - \mathbf{Y}^l\|_F^2 - \gamma_2 \text{tr}(\mathbf{Y}^u \hat{\mathbf{S}} \mathbf{P}' \mathbf{X}^u) + \frac{\gamma_3}{2} \text{tr}(\mathbf{P}' \mathbf{X}^u \mathbf{H}^u \mathbf{X}^{u'} \mathbf{P})$ is convex w.r.t. $\mathbf{P}$. The whole algorithm of the proposed progressive label refinement is summarized in Algorithm 2. We refer to our AEZSL with Label Refinement as AEZSL_LR.

Discussion: Since $\mathbf{p}^c = \mathbf{W}^c \mathbf{a}_c^t$ (see Section IV-A1) and $\mathbf{a}_c^t$'s are fixed, updating $\mathbf{p}^c$'s implies updating $\mathbf{W}^c$. In other words, by updating $\mathbf{p}^c$ based on the similarities among test instances and among unseen categories, we actually adapt the mapping matrices $\mathbf{W}^c$'s to the test set implicitly.

B. Deep Adaptive Embedding Zero-Shot Learning

One of the important applications of zero-shot learning is large-scale classification since it is difficult to obtain sufficient well-labeled training data exhaustively for all the categories. However, our AEZSL method is not very efficient in this scenario, because we need to learn one visual-semantic mapping for each unseen category (*e.g.*, over 20,000 test categories in the ImageNet 2011 21K dataset). This concern motivates us to design a model which is more suitable for large-scale ZSL. In order to mitigate the burden induced by a large set of unseen categories, we design a deep adaptive embedding model named DAEZSL, which only needs to be trained once but can generalize to arbitrary number of unseen categories.

Similar as in Section IV-A1, we assume the visual-semantic mappings should be category-specific and related to the semantic representation of each category. To avoid learning one mapping for each unseen category as in Section IV-A1, one possible approach is learning a projection $f(\cdot)$ from semantic representation $\mathbf{a}_c$ to visual-semantic mapping $\mathbf{W}_c$, *i.e.*, $f(\mathbf{a}_c) = \mathbf{W}_c$, so that given a new unseen category with semantic representation $\mathbf{a}_{\tilde{c}}$, we can easily obtain its visual-semantic mapping $\mathbf{W}_{\tilde{c}}$ by using $\mathbf{W}_{\tilde{c}} = f(\mathbf{a}_{\tilde{c}})$.

However, the size of $\mathbf{W}_c$ is usually very large (*i.e.*, $d \times a$), so we adopt an alternative approach for the sake of computational

efficiency, that is, learning the projection $g(\cdot)$ from semantic representation $\mathbf{a}_c$ to feature mask $\mathbf{m}_c$, *i.e.*, $g(\mathbf{a}_c) = \mathbf{m}_c$. The feature mask $\mathbf{m}_c$ is applied on the visual feature via elementwise product, and thus has the same dimension d as visual feature. Since the size of $\mathbf{m}_c$ (*i.e.*, d) is much smaller than that of $\mathbf{W}_c$ (*i.e.*, $d \times a$), it is much more efficient to learn the projection $g(\cdot)$ instead of $f(\cdot)$. In fact, applying the feature masks of different categories can be treated as implicitly adapting the visual-semantic mapping to different categories, which will be explained as follows. Assume we have a common visual-semantic mapping $\mathbf{W}$, given the i -th training instance with visual feature $\mathbf{x}_i^s$, its decision values after applying the feature mask $\mathbf{m}_c^s$ is

$$(\mathbf{x}_i^{s'} \circ \mathbf{m}_c^{s'}) \mathbf{W} \mathbf{A}^s = \mathbf{x}_i^{s'} (\bar{\mathbf{M}}_c^s \circ \mathbf{W}) \mathbf{A}^s, \quad (14)$$

in which $\bar{\mathbf{M}}_c^s$ is horizontally stacked a copies of $\mathbf{m}_c^s$. Then, we define the implicit visual-semantic mapping $\bar{\mathbf{W}}_c^s$ as

$$\bar{\mathbf{W}}_c^s = \bar{\mathbf{M}}_c^s \circ \mathbf{W}, \quad (15)$$

from which we can see that learning feature mask $\mathbf{m}_c^s$ is equivalent to adapting $\mathbf{W}$ to $\bar{\mathbf{W}}_c^s$ by using the mask $\bar{\mathbf{M}}_c^s$. Note that we simplify the task of adapting $\mathbf{W}$ by using the mask $\bar{\mathbf{M}}_c^s$ with all columns being the same $\mathbf{m}_c^s$, so that the number of variables (*i.e.*, size of $\mathbf{m}_c^s$) generated by $g(\cdot)$ is greatly reduced compared with the size of visual-semantic mapping.

In practice, we only need to learn $\mathbf{W}$ and $g(\cdot)$ without explicitly producing $\bar{\mathbf{W}}_c^s$. In the remainder of this section, we use implicit $\bar{\mathbf{W}}_c^s$ merely for the purpose of better explanation. Specifically, we expect implicit $\bar{\mathbf{W}}_c^s$ to satisfy two properties: generalizability and specificity, analogous to AEZSL in (5) (category-specific $\mathbf{W}^c$'s for specificity and co-regularizer for generalizability).

Generalizability: On one hand, we expect any $\bar{\mathbf{W}}_c^s$ can correctly classify the training instances from any category, which can be achieved by minimizing the square loss $\sum_{c=1}^{C^s} \|\mathbf{X}^{s'} \bar{\mathbf{W}}_c^s \mathbf{A}^s - \mathbf{Y}^s\|_F^2$ with $\mathbf{X}^s$ and $\mathbf{Y}^s$ being the same as defined in (5). After defining the one-hot label vector of $\mathbf{x}_i^s$ as $\mathbf{y}_i^s$ with the $c(i)$ -th entry being 1 and $\mathbf{M}^s \in \mathcal{R}^{d \times C^s}$ as the aggregated feature masks over all C^s categories with the c -th column being $\mathbf{m}_c^s$, we can have

$$\begin{aligned} & \sum_{c=1}^{C^s} \|\mathbf{X}^{s'} \bar{\mathbf{W}}_c^s \mathbf{A}^s - \mathbf{Y}^s\|_F^2, \\ &= \sum_{c=1}^{C^s} \sum_{i=1}^{n^s} \|\mathbf{x}_i^{s'} \bar{\mathbf{W}}_c^s \mathbf{A}^s - \mathbf{y}_i^s\|^2 \\ &= \sum_{i=1}^{n^s} \sum_{c=1}^{C^s} \|\mathbf{x}_i^{s'} (\bar{\mathbf{M}}_c^s \circ \mathbf{W}) \mathbf{A}^s - \mathbf{y}_i^s\|^2 \\ &= \sum_{i=1}^{n^s} \sum_{c=1}^{C^s} \|(\mathbf{x}_i^{s'} \circ \mathbf{m}_c^{s'}) \mathbf{W} \mathbf{A}^s - \mathbf{y}_i^s\|^2 \\ &= \sum_{i=1}^{n^s} \|(\bar{\mathbf{X}}_i^{s'} \circ \mathbf{M}^{s'}) \mathbf{W} \mathbf{A}^s - \bar{\mathbf{Y}}_i^s\|_F^2, \end{aligned} \quad (16)$$

in which $\bar{\mathbf{X}}_i^s \in \mathcal{R}^{d \times C^s}$ is horizontally stacked C^s copies of $\mathbf{x}_i^s$ and $\bar{\mathbf{Y}}_i^s \in \mathcal{R}^{C^s \times C^s}$ is vertically stacked C^s copies of $\mathbf{y}_i^s$. Hence, in our implementation, given the i -th training instance,

Fig. 1: Deep adaptive embedding model for zero-shot learning. In the top flow, the feature of each input image is duplicated to C copies. In the bottom flow, the semantic representations of all C categories pass through multi-layer perceptrons (MLP) and generate the feature masks for C categories, which are applied on the duplicated features via elementwise product. Then, the masked features are fed into a fully connected (fc) layer (i.e., visual-semantic mapping), followed by dot product with semantic representations, and finally output the decision value matrix, based on which we minimize the training loss in the training stage and predict test instances in the testing stage.

we duplicate $\mathbf{x}_i^s$ and $\mathbf{y}_i^s$ to C^s copies, and apply the aggregated feature masks $\mathbf{M}^s$ over all C^s categories on the duplicated visual features $\bar{\mathbf{X}}_i^s$. Then, the decision value matrix of the i -th training instance $(\bar{\mathbf{X}}_i^s \circ \mathbf{M}^s) \mathbf{W} \mathbf{A}^s$ can be easily obtained.

Specificity: On the other hand, we expect that $\bar{\mathbf{W}}_c^s$ can better fit the classification task corresponding to the c -th category. For the i -th training instance, we use $\mathbf{J}^i = (\bar{\mathbf{X}}_i^s \circ \mathbf{M}^s) \mathbf{W} \mathbf{A}^s$ to denote its decision value matrix, in which J_{c_1, c_2}^i is the decision value of c_2 -th category obtained by using $\bar{\mathbf{W}}_{c_1}^s$. In the following, we only consider the decision values of its ground-truth category $c(i)$, i.e., the $c(i)$ -th column of $\mathbf{J}^i$. In this case, similar as in (14), the decision value of $c(i)$ -th seen category obtained by using $\bar{\mathbf{W}}_c^s$ is

$$J_{\tilde{c}, c(i)}^i = (\mathbf{x}_i^{s'} \circ \mathbf{m}_{\tilde{c}}^{s'}) \mathbf{W} \mathbf{a}_{c(i)}^s = \mathbf{x}_i^{s'} \bar{\mathbf{W}}_{\tilde{c}}^s \mathbf{a}_{c(i)}^s. \quad (17)$$

Then, we expect $J_{c(i), c(i)}^i$ obtained by $\bar{\mathbf{W}}_{c(i)}^s$ should be larger than $J_{\tilde{c}, c(i)}^i$ obtained by $\bar{\mathbf{W}}_{\tilde{c}}^s$ for $\tilde{c} \neq c(i)$. With this aim, we employ the hinge loss $R^i = \sum_{\tilde{c} \neq c(i)} \max(0, J_{\tilde{c}, c(i)}^i - J_{c(i), c(i)}^i + \rho)$ to push $J_{c(i), c(i)}^i$ to be larger than $J_{\tilde{c}, c(i)}^i$ by margin ρ for $\tilde{c} \neq c(i)$. In our experiments, we empirically set ρ as 0.5 considering that $\mathbf{J}^i$ is regressed to binary label matrix.

By taking both generalizability and specificity into consideration, the loss function of our deep adaptive embedding model is designed as

$$\min_{g(\cdot), \mathbf{W}} \sum_{i=1}^{n^s} (\|\mathbf{J}^i - \bar{\mathbf{Y}}_i^s\|_F^2 + R^i). \quad (18)$$

Based on (18), we aim to have implicit $\bar{\mathbf{W}}_c^s$'s which can generalize to other categories by minimizing $\|\mathbf{J}^i - \bar{\mathbf{Y}}_i^s\|_F^2$ and simultaneously better fit the classification task corresponding to its own category by minimizing R^i . Note that semantic representations of unseen categories and unlabeled test instances are not utilized in the training stage.

Our deep adaptive embedding model is illustrated in Fig. 1, from which we can see that $g(\cdot)$ is modeled by multi-layer perceptrons (MLP) and $\mathbf{W}$ is modeled by fully connected (fc)

layer. Specifically, during the training process, we input the training images and the semantic representations of all C^s seen categories, i.e., $\mathbf{A}^{s'} \in \mathcal{R}^{C^s \times a}$. In the top flow, d -dim feature of the i -th training image is duplicated to C^s copies, i.e., $\bar{\mathbf{X}}_i^{s'} \in \mathcal{R}^{C^s \times d}$. In the bottom flow, $\mathbf{A}^{s'}$ passes through MLP and generates the feature masks $\mathbf{M}^{s'} \in \mathcal{R}^{C^s \times d}$ for C^s categories. After applying the generated feature masks on the duplicated visual features via elementwise product, the masked features $\bar{\mathbf{X}}_i^{s'} \circ \mathbf{M}^{s'}$ are fed into fc layer (i.e., the common visual-semantic mapping matrix $\mathbf{W}$), and output $(\bar{\mathbf{X}}_i^{s'} \circ \mathbf{M}^{s'}) \mathbf{W}$. Finally, we perform dot product on $(\bar{\mathbf{X}}_i^{s'} \circ \mathbf{M}^{s'}) \mathbf{W}$ and $\mathbf{A}^s$, leading to the decision value matrix $\mathbf{J}^i$, based on which we employ the loss function in (18). Note that we train an end-to-end system to minimize the loss function in (18), during which the parameters of CNN, MLP, and fc layer in Fig. 1 are jointly optimized. More details about the network architecture and training process can be found in Section V-B.

In the prediction stage, given an unseen category with semantic representation $\mathbf{a}_c^t$, we can easily generate its feature mask $\mathbf{m}_c^t$ via the projection function $g(\cdot)$, which implies an adaptive visual-semantic mapping $\bar{\mathbf{W}}_c^t$. Intuitively, if $\mathbf{a}_c^t$ is similar to $\mathbf{a}_c^s$ of the c -th seen category, their generated feature masks $\mathbf{m}_c^t$ and $\mathbf{m}_c^s$ should be similar. Furthermore, their visual-semantic mapping matrices $\bar{\mathbf{W}}_c^t$ and $\bar{\mathbf{W}}_c^s$ should be close to each other. Therefore, the visual-semantic mapping of a given unseen category is expected to be in correlation with those of similar seen categories, analogous to learning $\mathbf{W}^c$ based on the similarity matrix $\mathbf{S}^c$ in (5).

We further elaborate on the prediction procedure based on Fig. 1. In particular, given the i -th test image $\mathbf{x}_i^t$ and the set of unseen categories with size C^t , we use this test image and the semantic representations of all C^t unseen categories $\mathbf{A}^{t'} \in \mathcal{R}^{C^t \times a}$ as input. The test image passes through the top flow and generates C^t duplicated copies of visual features, i.e., $\bar{\mathbf{X}}_i^{t'} \in \mathcal{R}^{C^t \times d}$, while $\mathbf{A}^{t'}$ passes through the bottom flow and generates the feature masks $\mathbf{M}^{t'} \in \mathcal{R}^{C^t \times d}$ for all unseen

categories. After performing $(\bar{\mathbf{X}}_i^{t'} \circ \mathbf{M}^{t'})\mathbf{W}\mathbf{A}^t$, we can obtain a decision value matrix $\mathbf{J}^i \in \mathcal{R}^{C^t \times C^t}$ for the i -th test image. Finally, we get the diagonal of $\mathbf{J}^i$ as decision value vector and classify this test image as the category corresponding to the highest decision value. Note that only the diagonal of $\mathbf{J}^i$ is used for prediction because we assume that for the c -th unseen category, the decision value $J_{c,c}^i = \mathbf{x}_i^{t'} \bar{\mathbf{W}}_c^t \mathbf{a}_c^t$ (see (17)) obtained based on $\bar{\mathbf{W}}_c^t$ can best fit the classification task for this category.

Relation to ESZSL: When fixing the feature masks for all categories, *i.e.*, $\mathbf{M}^s$, as all-one matrix without learning MLP in Fig. 1, our DAEZSL model approximately reduces to ESZSL, in which the visual-semantic mappings of all categories are the same.

Relation to AEZSL: Our DAEZSL model shares the similar idea with our AEZSL method, that is, visual-semantic mapping should be category-specific and related to the semantic representation of each category. Besides, we design the training loss in (18) considering the generalizability and specificity of visual-semantic mappings, in analogy to learning category-specific $\mathbf{W}^c$'s with co-regularizer in (5). Moreover, in our DAEZSL model, the visual-semantic mapping of a given unseen category is expected to be close to those of seen categories with similar semantic representations, resembling the first regularizer based on the similarity matrix $\mathbf{S}^c$ in (5).

V. EXPERIMENTS

In this section, we conduct experiments for image classification on three small-scale datasets (*e.g.*, CUB, SUN, and Dogs) and one large-scale dataset (*e.g.*, ImageNet). On the small-scale datasets, we compare our AEZSL and AEZSL_LR methods with their special cases as well as standard/semi-supervised ZSL baseline methods. Note that our DAEZSL method is designed for large-scale ZSL and the small-scale training set can easily cause overfitting, especially for small number of seen categories, so we omit the results of DAEZSL on the small-scale datasets. On the large-scale dataset, due to the inefficiency of AEZSL for a large number of unseen categories, we only evaluate our DAEZSL method, which is specifically designed for large-scale ZSL, and compare with recently reported state-of-the-art results.

A. Zero-Shot Learning on Small-Scale Datasets

Experimental Settings: We conduct experiments on the following three popular benchmark datasets which are commonly used for zero-shot learning tasks:

- CUB [38]: Caltech-UCSD Bird (CUB) has in total 11,788 images distributed in 200 bird categories. Following the ZSL setting in [39], we use the standard train-test split with 150 (*resp.*, 50) categories as seen (*resp.*, unseen) categories. The CUB dataset contains a 312-dim binary human specified attribute vector for each image, so we average the attribute vectors of the images within each category and use it as the semantic representation of that category.
- SUN [40]: Scene UNDERstanding (SUN) dataset has 20 images in each scene category. Following the ZSL setting

TABLE I: Accuracies (%) of different baseline methods and our methods on three benchmark datasets. The best results are highlighted in boldface.

Dataset	CUB	SUN	Dogs	Avg
ESZSL [24]	49.74	82.50	40.27	57.50
LatEm [22]	45.50	83.50	37.41	55.47
Zhang and Saligrama [23]	46.11	83.83	43.42	57.79
Bucher <i>et al.</i> [14]	43.29	84.41	36.18	54.63
AMP [9]	43.12	82.50	38.02	54.55
COSTA [11]	44.19	76.00	33.98	51.39
SSE [13]	40.64	82.50	38.63	53.92
SJE [5]	51.70	84.00	37.11	57.60
DAP/IAP [1]	41.40	72.00	37.11	50.17
Changpinyo <i>et al.</i> [28]	54.50	83.50	42.01	60.50
ConSE [12]	37.33	73.00	23.49	44.61
RKT [43]	53.73	81.00	42.11	58.95
EXEM [44]	58.54	86.00	42.90	62.48
Li <i>et al.</i> [21]	43.94	87.50	43.50	58.31
Kodirov <i>et al.</i> [18]	57.42	86.00	48.59	64.00
Zhang and Saligrama [45]	56.90	89.50	48.53	64.98
Shojaee and Baghshah [19]	58.80	86.16	47.49	64.15
Li and Guo [10]	57.14	88.50	47.84	64.49
SMS [25]	57.14	83.50	47.95	62.86
Xu <i>et al.</i> [15]	55.74	85.50	43.09	61.44
AEZSL_sim	57.55	87.50	43.98	63.01
AEZSL	59.73	88.00	44.62	64.12
AEZSL_LR (one-step)	60.08	89.00	48.90	65.99
AEZSL_LR	64.44	92.50	50.62	69.18

in [41], we use the provided list of 10 categories as unseen categories and the rest of 707 categories as seen categories. Similar to CUB, the averaged 102-dim attribute vector is used as the semantic representation for each category.

- Dogs [42]: ZSL was firstly performed on the Stanford Dogs dataset in [5], which uses 19,501 images from 113 breeds of dogs. We follow the provided train-test split in [5], *i.e.*, 85 (*resp.*, 28) categories as seen (*resp.*, unseen) categories. Since there is no manually annotated attribute for the Dogs dataset, we combine two types of output embeddings learned from online corpus (*i.e.*, 3,850-dim Bag-of-Words embedding and 163-dim WordNet-derived similarity embedding) as the semantic representation for each category, which has demonstrated superior performance in [5].

For CUB and SUN datasets, we extract the 4,096-dim output of the 6-th layer of the pretrained VGG [46] as the visual feature for each image, following myriad of previous ZSL works such as [23], [45], [13], [14]. For the Dogs dataset, we use the 1,024-dim output of the top layer of the pretrained GoogleNet [47] following [22], [5].

Baselines: We compare our AEZSL and AEZSL_LR methods with two sets of ZSL baselines: standard ZSL methods and semi-supervised/transductive ZSL methods. For the first set, we include the following methods, [22], [23], [14], [9], [11], [13], [5], [1], [28], [12], [43], [44], which do not utilize unlabeled test instances in the training stage. For the second set, we compare with the following transductive or semi-supervised ZSL methods [21], [18], [45], [19], [10], [25], [15], which utilize the unlabeled test instances and semantic representations of unseen categories in the training phase. Note that our AEZSL belongs to standard ZSL methods while

AEZSL_LR belongs to semi-supervised ZSL methods.

Besides two sets of baselines mentioned above, we consider a special case of our AEZSL method, *i.e.*, without using the co-regularizer $\sum_{c < \tilde{c}} \|\mathbf{W}^c - \mathbf{W}^{\tilde{c}}\|_F^2$ ($\lambda_3 = 0$), which is referred to as AEZSL_sim in Table I. We also report the results of one special cases of our AEZSL_LR method, namely, AEZSL_LR (one-step), which is a non-progressive approach by treating the entire test set as the unconfident set $\mathcal{U}$ and using the method in (11) without the regularizer $\|\mathbf{X}^{l'}\mathbf{P} - \mathbf{Y}^l\|_F^2$.

We use multi-class accuracy for performance evaluation. For the baselines, if the experimental settings (*i.e.*, train-test split, visual feature, semantic representation, evaluation metric, *etc*) mentioned in their papers are exactly the same as ours, we directly copy their reported results. Otherwise, we run their methods using our experimental setting for fair comparison.

Parameters: Our AEZSL method has three trade-off parameters λ_1 , λ_2 , and λ_3 in (5). Besides, our label refinement strategy has three additional trade-off parameters γ_1 , γ_2 , and γ_3 (see (11)). We use cross-validation strategy to determine the above trade-off parameters. Specifically, following [19], we choose the first C^c categories based on the default category indices from C^s seen categories as validation categories, in which C^c satisfies $\frac{C^c}{C^s} = \frac{C^t}{C^s + C^t}$. In our experiments, we first learn models based on the seen categories excluding the validation categories with λ_1 , λ_2 , and λ_3 set within the range $[10^{-3}, 10^{-2}, \dots, 10^3]$, and then test the learnt models on the validation categories. After determining the optimal trade-off parameters through the grid search, we learn the final model based on all seen categories. Note that we also have a hyper parameter k , which is the number of selected confident test instances in each iteration. We empirically fix k as 100 for CUB and Dogs datasets, and 10 for the SUN dataset given that there are only 200 test instances in the SUN dataset.

Based on our experimental observation, our methods are relatively robust when the parameters are set within certain range. By taking λ_3 and k as examples, we vary λ_3 (*resp.*, k) in the range of $[10^{-1}, \dots, 10^3]$ (*resp.*, $[80, 90, \dots, 120]$) and evaluate our AEZSL_LR method on the Dogs dataset. The performance variance is illustrated in the middle and right subfigure in Fig. 5, from which we can see that our method is insensitive to λ_3 and k within certain range. We have similar observations for the other parameters on the other datasets.

Experimental Results: The experimental results for our AEZSL and AEZSL_LR methods as well as all the baselines are reported in Table I. It can be seen that ESZSL achieves competitive results compared with other standard ZSL baselines, which demonstrates the effectiveness of ESZSL despite its simplicity. By comparing AEZSL_sim with AEZSL, we observe that AEZSL achieves better results after employing the co-regularizer, so it is beneficial to encourage the mappings from different categories to share some common parts. We also observe that our AEZSL method not only outperforms ESZSL and AEZSL_sim, but also outperforms the standard ZSL baselines [22], [23], [14], [9], [11], [13], [5], [1], [28], [12], [43], [44], which indicates the advantage of learning a visual-semantic mapping for each unseen category.

From Table I, we also observe that the transductive or semi-supervised baselines [21], [18], [45], [19], [10], [25],

Fig. 2: Illustration of two instance images from the category “flea market” and their corresponding mapped values for the attributes “cloth” and “cluttered space” obtained by using ESZSL and our AEZSL method.

Fig. 3: The first (*resp.*, second) row contains the instance images from different categories with the attribute “cloth” (*resp.*, “cluttered space”).

[15] generally achieve better results than those standard ZSL baselines, indicating that it is helpful to utilize the unlabeled test instances in the training stage. Another observation is that our AEZSL_LR method outperforms AEZSL, which demonstrates the effectiveness of the progressive label refinement. Moreover, AEZSL_LR also performs better than AEZSL_LR (one-step). This is because the selected confident set with more accurate predicted labels can provide weak supervision. Finally, our AEZSL_LR method outperforms all the baselines and achieves the state-of-the-art results on three datasets, which again shows the advantage of adapting the visual-semantic mapping to each unseen category.

Qualitative Analysis of the Attributes Learnt by AEZSL: In order to show the advantage of learning category-specific visual-semantic mappings intuitively, we take the SUN dataset as an example to investigate why our AEZSL can correctly classify the test instances which are misclassified by ESZSL. As illustrated in Figure 2, the two images are correctly

Fig. 4: The instance images from four nearest neighboring seen categories of the unseen category “flea market” with the attributes “cloth” and “cluttered space”.

Fig. 5: The left subfigure shows the accuracy variation of predicted test labels *w.r.t.* the number of iterations in label refinement on the Dogs dataset. The middle (*resp.*, right) subfigure shows the performance variance of our AEZSL_LR method *w.r.t.* parameter k (*resp.*, λ_3) on the Dogs dataset, in which the dash line indicates the parameter we use in Table I.

classified as “flea market” by our AEZSL method but are misclassified as “shoe shop” by ESZSL. Based on the semantic representations of “flea market” and “shoe shop”, we find that among the five attributes with maximum values in the semantic representation of “flea market”, only two attributes (*i.e.*, “cloth” and “cluttered space”) have larger values than those of “shoe shop”. We simply treat these two attributes as the representative attributes to distinguish “flea market” from “shoe shop”. Figure 2 shows the mapped values of the above two confusing images corresponding to two representative attributes obtained by using ESZSL and our AEZSL method, which are calculated by using $\mathbf{X}^s\mathbf{W}$ and $\mathbf{X}^s\mathbf{W}^c$ ($\mathbf{W}^c$ is the category-specific mapping matrix for “flea market”) respectively. It can be seen that our AEZSL method obtains larger mapped values for the two representative attributes than ESZSL, which contributes to the correct classification of two confusing images as “flea market”.

Based on the fact that AEZSL obtains larger mapped values corresponding to the two representative attributes, we conjecture that the learnt mapping using our method can better capture the semantic meanings of “cloth” and “cluttered space”. To verify this point, we first show some instance images from different categories with the attribute “cloth” (*resp.*, “cluttered space”) in the first (*resp.*, second) row in Figure 3. Together with the instance images in Figure 2, we can observe that for different categories, the visual appearances of “cloth” and “cluttered space” are considerably different as discussed in Section I. Thus, a general mapping matrix learnt by ESZSL cannot tell the subtle discrepancy in the semantic meanings of the same attribute between different categories.

In contrast, our AEZSL method learns the category-specific mapping matrix for “flea market” based on the similarities between this unseen category and each seen category, in which the major transfer comes from more similar seen categories. In Figure 4, we show the instance images from four nearest neighboring seen categories of “flea market” with the attributes “cloth” and “cluttered space”, from which it can be seen that in terms of the visual appearances of “cloth” and “cluttered space”, these neighboring categories resemble “flea market” much better than other categories such as those in Figure 3. Therefore, AEZSL can learn a better fitting visual-semantic mapping for “flea market”. We have similar observations for the other unseen categories and on the other datasets.

Performance Variation *w.r.t.* Iteration in AEZSL_LR:

Since our proposed label refinement method is a progressive approach, we are interested in the variation of the accuracy of the predicted test labels (*i.e.*, the union of $\mathbf{Y}^l$ and $\mathbf{Y}^u$) *w.r.t.* the number of iterations. By taking the Dogs dataset as an example, we plot the label accuracy *w.r.t.* the number of iterations in the left subfigure in Figure 5, from which we can observe that the accuracy of predicted test labels increases from 44.62% to 50.62% steadily within 50 iterations. Recall that in each iteration, we move the top k confident instances from the unconfident set with their refined predicted labels into the confident set. Thus, we can infer that in most iterations, the accuracy of the predicted labels of the selected k confident instances is improved using the updated visual classifiers, which verifies the effectiveness of progressive label refinement. We have similar observations on the other datasets.

Generalized ZSL: Most of existing ZSL methods assume

TABLE II: Accuracies (%) of different baseline methods and our DAEZSL method on the ImageNet dataset. The best results are highlighted in boldface.

Setting	Method	Flat Hit@K					Hierarchical precision@K			
		1	2	5	10	20	2	5	10	20
2-hop test	DeVISE [6]	6.0	10.0	18.1	26.4	36.4	15.2	19.2	21.7	23.3
	ConSE [12]	9.4	15.1	24.7	32.7	41.8	21.4	24.7	26.9	28.4
	Changpinyo <i>et al.</i> [28]	10.5	16.7	28.6	40.1	52.0	25.1	27.7	30.3	32.1
	EXEM [44]	12.5	19.5	32.3	43.7	55.2	26.9	29.1	31.2	33.3
	ESZSL	8.6	13.8	24.1	35.0	47.4	21.8	24.7	27.4	30.2
	DAEZSL	13.9	21.3	33.8	45.2	57.1	28.4	29.8	32.5	34.8
3-hop test	DeVISE [6]	1.7	2.9	5.3	8.2	12.5	3.7	19.1	21.4	23.6
	ConSE [12]	2.7	4.4	7.8	11.5	16.1	5.3	20.2	22.4	24.7
	Changpinyo <i>et al.</i> [28]	2.9	4.9	9.2	14.2	20.9	8.0	23.7	26.4	28.6
	EXEM [44]	3.6	5.9	10.7	16.1	23.1	8.2	25.3	27.8	30.1
	ESZSL	2.4	4.1	7.6	11.8	17.5	5.8	20.9	23.1	25.2
	DAEZSL	4.3	6.6	11.9	17.6	24.5	8.0	26.7	28.9	31.8
“all” test	DeVISE [6]	0.8	1.4	2.5	3.9	6.0	1.7	7.2	8.5	9.6
	ConSE [12]	1.4	2.2	3.9	5.8	8.3	2.5	7.8	9.2	10.4
	Changpinyo <i>et al.</i> [28]	1.5	2.4	4.5	7.1	10.9	3.6	9.6	11.0	12.5
	EXEM [44]	1.8	2.9	5.3	8.2	12.2	3.7	10.4	12.1	13.5
	ESZSL	1.2	2.0	3.8	5.9	9.1	2.8	9.3	10.7	11.9
	DAEZSL	2.0	3.2	5.9	8.7	13.6	3.5	11.8	12.9	14.7

that in the testing stage, the test instances only come from the unseen categories, which is actually an unrealistic setting because the instances from seen categories may also be encountered in the testing stage. So it is more useful to predict a test instance from either seen categories or unseen categories instead of assuming that the test instances are only from unseen categories. However, when using the mixture of test instances from both seen and unseen categories for testing, the performance will be significantly degraded due to the bias of prediction scores between seen category label space and unseen category label space, as demonstrated in [7], [48]. This more challenging test setting is referred to as generalized zero-shot learning (GZSL) in [48]. To validate the effectiveness of our AEZSL method under the more realistic GZSL setting, we additionally conduct experiments by mixing the test instances from both seen and unseen categories as the test set on the CUB dataset, which is the overlapped dataset with those used in [48]. In particular, we follow the setting in [48], that is, we move 20% of the instances from each of 150 seen categories to the test set, and thus the test set which originally has 50 unseen categories is expanded to 200 test categories including both seen categories and unseen categories.

When applying our AEZSL method to generalized ZSL setting, we adopt the same strategy as in (5) and the only difference is that we learn $(C^s + C^t)$ instead of C^t category-specific visual-semantic mappings. Particularly, we learn category-specific mappings for both unseen and seen categories by assigning different weights on different classification tasks of different seen categories based on the similarity between each category and all the seen categories. For fair comparison with state-of-the-art results under the generalized ZSL setting, we further employ the existing calibrated stacking strategy proposed in [48] on the results obtained by our AEZSL method. The idea of calibrated stacking strategy is simple yet very effective, that is, to reduce the prediction scores in the seen category label space by a threshold, which is learnt based on a performance metric called Area Under SeenUnseen

TABLE III: Accuracies (%) of different methods under the generalized ZSL setting on the CUB dataset. The best results are highlighted in boldface.

Method	ESZSL	AEZSL	AEZSL(CS)	Chao <i>et al.</i> [48]
Accuracy(%)	26.35	30.76	42.08	35.60

accuracy Curve (AUSUC) on the validation set, according to the observation that the prediction scores in the seen category label space are often higher than those in the unseen category label space. Thus, after employing the calibrated stacking strategy, we expect to obtain unbiased prediction scores.

In Table III, we report the results obtained by ESZSL and our AEZSL method as well as our AEZSL after employing Calibrated Stacking (CS) strategy, which is referred to as AEZSL(CS). We also include the state-of-the-art result reported in [48] for comparison. From Table III, we observe that our AEZSL method outperforms ESZSL, which shows that our AEZSL method is also effective under the generalized ZSL setting. Moreover, after employing the calibrated stacking strategy, the accuracy obtained by our AEZSL method is greatly improved, which demonstrates that our AEZSL method can be perfectly integrated with the existing calibrated stacking strategy. Finally, we observe that AEZSL(CS) achieves significantly better result 42.08% than the state-of-the-art result 35.60% reported in [48], which again indicates the effectiveness of our method under the generalized ZSL setting.

B. Zero-Shot Learning on Large-Scale Dataset

In this section, we apply our deep adaptive embedding model DAEZSL on the ImageNet dataset and compare with the state-of-the-art reported results.

Experimental Settings: We strictly follow the experimental settings in [6], [28]. Specifically, we use the 1000 categories in ImageNet ILSVRC 2012 1K [49] as seen categories, and perform evaluation on three test scenarios, which are chosen from ImageNet 2011 21K dataset and built based on ImageNet

label hierarchy with increasing difficulty. The three test sets are listed as follows.

- 2-hop test set: 1,509 unseen categories within two tree hops of the 1000 seen categories according to the ImageNet label hierarchy¹, which are semantically and visually similar with 1000 seen categories.
- 3-hop test set: 7,678 unseen categories within three tree hops of 1000 seen categories, which is constructed in a similar way to 2-hop test set.
- “all” test set: all the 20,345 unseen categories in the ImageNet 2011 21K dataset which do not belong to the ILSVRC 2012 1K dataset.

Note that 2-hop (*resp.*, 3-hop) test set is a subset of 3-hop (*resp.*, “all”) test set, and all the test sets have no overlap with the training set (*e.g.*, ImageNet ILSVRC 2012 1K).

Semantic Representations: For both seen categories and unseen categories, we use the 500-dim word vector for each category provided in [28], which is obtained based on a skip-gram language model [50] trained on the latest Wikipedia corpus with more than 3 billion words. Note that one ImageNet category may have more than one word according to its synset, we simply average the word vectors of all the words appearing in its synset as the word vector for that category.

Evaluation Metrics: Evaluating ZSL methods on the large-scale ImageNet dataset is a nontrivial task considering the large number of unseen categories and the semantic overlap of different categories in the ImageNet label hierarchy. So we use different evaluation metrics compared with multi-class accuracy used on three small-scale datasets (*i.e.*, CUB, SUN, and Dog) in Section V-A. Following [6], we use two metrics Flat hit@K and Hierarchical precision@K for performance evaluation. To be exact, Flat hit@K is defined as the percentage of test instances for which the ground-truth category is in its top K predictions. When $K = 1$, Flat hit@K is identical with the multi-class accuracy. Compared with Flat hit@K, Hierarchical precision@K takes the hierarchical structure of categories into consideration. Given each test image, we generate a feasible set of K nearest categories of its ground-truth category in the ImageNet label hierarchy and calculate the overlap ratio between the feasible set and the top K predictions of this test image, that is, the precision of top K predictions. When generating a feasible set of K nearest categories for a ground-truth category, we enlarge the searching radius around the ground-truth category in the ImageNet label hierarchy iteratively and add the searched categories belonging to the test set to the feasible set. This procedure is repeated until the size of feasible set exceeds K . For more details, please see the Appendix of [6] or the Supplementary of [28]. Note that when $K = 1$, Flat hit@K is equal to Hierarchical precision@K, so we omit Hierarchical precision@1 in Table II to avoid redundancy.

Network Structure: In terms of CNN and MLP in Fig. 1, we use GoogleNet as the CNN structure in our experiments, which is initialized with released model [47] pretrained on ImageNet dataset. The dimension of output from CNN is 1,024. The MLP we use has one hidden layer with its

size empirically set as the average of feature dimension and attribute dimension, *i.e.*, $\lfloor \frac{d+a}{2} \rfloor$. Besides, we add a dropout layer following the hidden layer in MLP with 50% ratio of dropped output. Additionally, we add a sigmoid layer following MLP to normalize the feature mask in the range of (0, 1). The network is implemented using TensorFlow and we use AdaGrad optimizer for training, with batchsize as 128 and learning rate as 0.001.

Experimental Results: To the best of our knowledge, there are few ZSL papers [6], [12], [28], [44] reporting their performances on the ImageNet dataset in 2-hop, 3-hop, and “all” test settings. We compare our DAESL method with the reported results of DeVISE [6], ConSE [12], EXEM [44], and [28] in Table II. We also include ESZSL as a baseline, in which we train the DAEZSL network using all-one feature masks without learning MLP. From Table II, we can observe that DAEZSL achieves far better results than ESZSL, which demonstrates the advantage of category-specific feature mask. Our DAEZSL also outperforms all the baseline methods in most cases (25 out of 27), indicating that it is beneficial to learn deep adaptive embedding model which can implicitly adapt visual-semantic mapping to different categories by using category-specific feature mask.

VI. CONCLUSION

In this paper, we have proposed our AEZSL method, which aims to learn a category-specific visual-semantic mapping for each unseen category based on the similarities between unseen categories and seen categories, followed by progressive label refinement. Moreover, we additionally propose a deep adaptive embedding model named DAEZSL for large-scale ZSL, which only needs to be trained once on the seen categories but can be applied to arbitrary number of unseen categories. Comprehensive experimental results demonstrate the effectiveness of our proposed methods.

REFERENCES

- [1] 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*, vol. 36, no. 3, pp. 453–465, 2014.
- [2] Y. Guo, G. Ding, J. Han, and Y. Gao, “Zero-shot learning with transferred samples,” *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3277–3290, 2017.
- [3] C. Luo, Z. Li, K. Huang, J. Feng, and M. Wang, “Zero-shot learning via attribute regression and class prototype rectification,” *IEEE Transactions on Image Processing*, vol. 27, no. 2, pp. 637–648, 2018.
- [4] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth, “Describing objects by their attributes,” in *Proceedings of the 22nd IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL, Jun. 2009, pp. 1778–1785.
- [5] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele, “Evaluation of output embeddings for fine-grained image classification,” in *Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition*, Boston, MA, Jun. 2015, pp. 2927–2936.
- [6] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov *et al.*, “DeViSE: A deep visual-semantic embedding model,” in *Proceedings of the 27th Annual Conference on Neural Information Processing Systems*, Lake Tahoe, Nevada, United States, Dec. 2013, pp. 2121–2129.
- [7] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng, “Zero-shot learning through cross-modal transfer,” in *Proceedings of the 27th Annual Conference on Neural Information Processing Systems*, Lake Tahoe, Nevada, United States, 2013, pp. 935–943.

¹http://www.image-net.org/api/xml/structure_released.xml

- [8] F. X. Yu, L. Cao, R. S. Feris, J. R. Smith, and S.-F. Chang, "Designing category-level attributes for discriminative visual recognition," in *Proceedings of the 26th IEEE Conference on Computer Vision and Pattern Recognition*, Portland, OR, Jun. 2013, pp. 771–778.
- [9] Z. Fu, T. Xiang, E. Kodirov, and S. Gong, "Zero-shot object recognition by semantic manifold distance," in *Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition*, Boston, MA, Jun. 2015, pp. 2635–2644.
- [10] X. Li and Y. Guo, "Max-margin zero-shot learning for multi-class classification," in *The 18th International Conference on Artificial Intelligence and Statistics*, San Diego, California, May 2015, pp. 626–634.
- [11] T. Mensink, E. Gavves, and C. G. Snoek, "COSTA: co-occurrence statistics for zero-shot classification," in *Proceedings of the 27th IEEE Conference on Computer Vision and Pattern Recognition*, Columbus, OH, Jun. 2014, pp. 2441–2448.
- [12] 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 *International Conference on Learning Representations*, Banff, Canada, Apr. 2014.
- [13] Z. Zhang and V. Saligrama, "Zero-shot learning via semantic similarity embedding," in *Proceedings of the 15th International Conference on Computer Vision*, Santiago, Chile, Dec. 2015, pp. 4166–4174.
- [14] M. Bucher, S. Herbin, and F. Jurie, "Improving semantic embedding consistency by metric learning for zero-shot classification," in *Proceedings of the 14th European Conference on Computer Vision*, Amsterdam, The Netherlands, Oct. 2016, pp. 730–746.
- [15] X. Xu, T. Hospedales, and S. Gong, "Transductive zero-shot action recognition by word-vector embedding," *International Journal of Computer Vision*, vol. 123, no. 3, pp. 309–333, 2017.
- [16] D. Jayaraman, F. Sha, and K. Grauman, "Decorrelating semantic visual attributes by resisting the urge to share," in *Proceedings of the 27th IEEE Conference on Computer Vision and Pattern Recognition*, Columbus, OH, Jun. 2014, pp. 1629–1636.
- [17] E. Kodirov, T. Xiang, and S. Gong, "Semantic autoencoder for zero-shot learning," *Proceedings of the 30th IEEE Conference on Computer Vision and Pattern Recognition*, Jul. 2017.
- [18] E. Kodirov, T. Xiang, Z. Fu, and S. Gong, "Unsupervised domain adaptation for zero-shot learning," in *Proceedings of the 15th International Conference on Computer Vision*, Santiago, Chile, Dec. 2015, pp. 2452–2460.
- [19] S. M. Shojaei and M. S. Baghshah, "Semi-supervised zero-shot learning by a clustering-based approach," *arXiv preprint arXiv:1605.09016*, 2016.
- [20] Y. Fu, T. M. Hospedales, T. Xiang, Z. Fu, and S. Gong, "Transductive multi-view embedding for zero-shot recognition and annotation," in *Proceedings of the 13th European Conference on Computer Vision*, Zurich, Switzerland, Sep. 2014, pp. 584–599.
- [21] X. Li, Y. Guo, and D. Schuurmans, "Semi-supervised zero-shot classification with label representation learning," in *Proceedings of the 15th International Conference on Computer Vision*, Santiago, Chile, Dec. 2015, pp. 4211–4219.
- [22] Y. Xian, Z. Akata, G. Sharma, Q. Nguyen, M. Hein, and B. Schiele, "Latent embeddings for zero-shot classification," in *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, Jun. 2016, pp. 69–77.
- [23] Z. Zhang and V. Saligrama, "Zero-shot learning via joint latent similarity embedding," in *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, Jun. 2016, pp. 6034–6042.
- [24] B. Romera-Paredes and P. H. Torr, "An embarrassingly simple approach to zero-shot learning," in *Proceedings of the 32nd IEEE International Conference on Machine Learning*, Lille, France, Jul. 2015, pp. 2152–2161.
- [25] Y. Guo, G. Ding, X. Jin, and J. Wang, "Transductive zero-shot recognition via shared model space learning," in *The 30th AAAI Conference on Artificial Intelligence*, Phoenix, Arizona, United States, Feb. 2016, pp. 3434–3500.
- [26] Y. Long, L. Liu, and L. Shao, "Attribute embedding with visual-semantic ambiguity removal for zero-shot learning," in *Proceedings of the British Machine Vision Conference*, York, UK, Sep. 2016.
- [27] Z. Al-Halah, M. Tapaswi, and R. Stiefelhofen, "Recovering the missing link: Predicting class-attribute associations for unsupervised zero-shot learning," in *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, Jun. 2016, pp. 5975–5984.
- [28] S. Changpinyo, W.-L. Chao, B. Gong, and F. Sha, "Synthesized classifiers for zero-shot learning," in *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, Jun. 2016, pp. 5327–5336.
- [29] M. Rohrbach, S. Ebert, and B. Schiele, "Transfer learning in a transductive setting," in *Proceedings of the 27th Annual Conference on Neural Information Processing Systems*, Lake Tahoe, Nevada, United States, Dec. 2013, pp. 46–54.
- [30] Y. Yang and T. M. Hospedales, "A unified perspective on multi-domain and multi-task learning," May 2015.
- [31] J. Lei Ba, K. Swersky, S. Fidler *et al.*, "Predicting deep zero-shot convolutional neural networks using textual descriptions," in *Proceedings of the 15th International Conference on Computer Vision*, Santiago, Chile, Dec. 2015, pp. 4247–4255.
- [32] L. Zhang, T. Xiang, and S. Gong, "Learning a deep embedding model for zero-shot learning," *arXiv preprint arXiv:1611.05088*, 2016.
- [33] R. Qiao, L. Liu, C. Shen, and A. van den Hengel, "Less is more: zero-shot learning from online textual documents with noise suppression," in *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, Vegas, NV, Jun. 2016, pp. 2249–2257.
- [34] V. Simoncini, "Computational methods for linear matrix equations," *SIAM Review*, vol. 58, no. 3, pp. 377–441, 2016.
- [35] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proceedings of the 11th annual conference on Computational learning theory*, Madison, WI, Jul. 1998, pp. 92–100.
- [36] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [37] F. Nie, H. Huang, X. Cai, and C. H. Ding, "Efficient and robust feature selection via joint 2_1 -norms minimization," in *Proceedings of the 24th Annual Conference on Neural Information Processing Systems*, Vancouver, Canada, Dec. 2010, pp. 1813–1821.
- [38] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The caltech-ucsd birds-200-2011 dataset," California Institute of Technology, Tech. Rep. CNS-TR-2011-001, 2011.
- [39] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, "Label embedding for attribute-based classification," in *Proceedings of the 26th IEEE Conference on Computer Vision and Pattern Recognition*, Portland, OR, Jun. 2013.
- [40] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *Proceedings of the 23th IEEE Conference on Computer Vision and Pattern Recognition*, San Francisco, CA, Jun. 2010, pp. 3485–3492.
- [41] D. Jayaraman and K. Grauman, "Zero-shot recognition with unreliable attributes," in *Proceedings of the 28th Annual Conference on Neural Information Processing Systems*, Montreal, Canada, Dec. 2014, pp. 3464–3472.
- [42] A. Khosla, N. Jayadevaprakash, B. Yao, and L. Fei-Fei, "Novel dataset for fine-grained image categorization," in *Proceedings of the 24th IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, Jun. 2011.
- [43] D. Wang, Y. Li, Y. Lin, and Y. Zhuang, "Relational knowledge transfer for zero-shot learning," in *The 30th AAAI Conference on Artificial Intelligence*, Phoenix, Arizona, United States, Feb. 2016, pp. 2145–2151.
- [44] S. Changpinyo, W.-L. Chao, and F. Sha, "Predicting visual exemplars of unseen classes for zero-shot learning," *ICCV*, Oct. 2017.
- [45] Z. Zhang and V. Saligrama, "Zero-shot recognition via structured prediction," in *Proceedings of the 14th European Conference on Computer Vision*, Amsterdam, The Netherlands, Oct. 2016, pp. 533–548.
- [46] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [47] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *CVPR*, 2015.
- [48] W.-L. Chao, S. Changpinyo, B. Gong, and F. Sha, "An empirical study and analysis of generalized zero-shot learning for object recognition in the wild," in *Proceedings of the 14th European Conference on Computer Vision*, Amsterdam, The Netherlands, Oct. 2016, pp. 52–68.
- [49] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [50] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proceedings of the 27th Annual Conference on Neural Information Processing Systems*, Lake Tahoe, Nevada, United States, Dec. 2013, pp. 3111–3119.