# Zero-Shot Image Classification Using Coupled Dictionary Embedding

Mohammad Rostami[1,*], Soheil Kolouri[2,*1], Zak Murez[3], Yuri Owechko[4],
Eric Eaton[5], Kuyngnam Kim[6]
mrostami@isi.edu;soheil.kolouri@vanderbilt.edu;zak@murez.com
yowechko@hrl.com,eeaton@seas.upenn.edu;ai.guru@sk.com

[1]*USC Information Sciences Institute, CA, USA*
[2]*Vanderbilt University, Nashville, TN, USA*
[3] *Wayve, London, UK*
*HRL Laboratories, Malibu, CA, USA,*
[5] *University of Pennsylvania, Philadelphia, PA, USA*
[6] *SK Telecom, Seoul, South Korea*

---

[1]Equal Contribution

**Abstract**

Zero-shot learning (ZSL) is a framework to classify images that belong to unseen visual classes using their semantic descriptions about the unseen classes. We develop a new ZSL algorithm based on coupled dictionary learning. The core idea is to enforce the visual features and the semantic attributes of an image to share the same sparse representation in an intermediate embedding space, modeled as the shared input space of two sparsifying dictionaries. In the ZSL training stage, we use images from a number of seen classes for which we have access to both the visual and the semantic attributes to train two coupled dictionaries that can represent both the visual and the semantic feature vectors of an image using a single sparse vector. In the ZSL testing stage and in the absence of labeled data, images from unseen classes are mapped into the attribute space by finding the joint-sparse representations using solely the visual dictionary via solving a LASSO problem. The image is then classified in the attribute space given semantic descriptions of unseen classes. We also provide attribute-aware and transductive formulations to tackle the "domain-shift" and the "hubness" challenges for ZSL, respectively. Experiments on four primary datasets using VGG19 and GoogleNet visual features, are provided. Our performances using VGG19 features are 91.0%, 48.4%, and 89.3% on the SUN , the CUB, and the AwA1 datasets, respectively. Our performances on the SUN, the CUB, and the AwA2 datasets are 57.0%,49.7%, and 71.7%, respectively, when GoogleNet features are used. Comparison with existing methods demonstrates that our method is effective and compares favorably against the sate-of-the-art. In particular, our algorithm leads to decent performance on the all four datasets. [1].

*Keywords:* Zero shot learning, coupled dictionary learning, semantic attribute embedding, domain-shift, hubness

## 1. Introduction

Advances in deep learning have led to a remarkable performance improvement in a wide range of classification and categorization tasks. This

---

[1]Early partial results of this paper is presented at 2018 AAAI [22]

success is primarily due to the fact that deep neural networks automate the process of feature engineering using a blind end-to-end supervised training procedure [37, 28]. However, the cost for this success is the need for huge annotated datasets to implement supervised training. Emergence of crowdsourcing data annotation platforms such as Amazon Mechanical Turk has made data annotation easier [50], but manual data annotation is still not feasible in many cases, including:

1. Data annotation for tine-grained multi-class classification (e.g., thousands of classes for animal categorization) is a challenging task because it requires training annotators to be able to provide accurate labels. Additionally, the process is more time-consuming and usually requires more annotators because the labels are noisier.

2. In many domains, including medical domains, sharing data with annotators is not easy due to privacy regulations that limits sharing data. As a result, it is highly challenging to hire annotators that have sufficient clearances to process data.

3. When the domain is a specialized domain, e.g., synthetic aperture radar images, only people with years of prior training are able to annotate data. As a result, qualified data annotators are limited and expensive to hire.

4. The persistent and dynamic emergence of new classes (e.g., new products on shopping websites) makes data annotation a continual time-consuming procedure. Additionally, retraining the model to incorporate new classes can be computationally expensive.

5. In some application, there are classes with highly infrequent members (e.g., rare event classification). Preparing training instances for rare event annotation is challenging.

Consequently, training deep learning model is not a feasible solution when the above and other potentially similar challenges are present. To circumvent these challenges, it is desirable to improve existing systems by enabling them to benefit from knowledge transfer. For example, it is desirable to learn using a few training samples [59, 63, 62] and even learning *unseen* classes with no accessible training samples [40, 60, 39, 27, 75, 32, 34]. Additionally, learning from past experiences accumulatively, i.e., continual or lifelong learning, can help to avoid learning redundant information [7, 74, 55, 54, 31, 53]. Such abilities help to classify new emerging classes

more efficiently, to relax the need for persistent data annotation and model retraining, and to benefit from past learned experiences [56].

A primary approach to reduce dependence on large annotated datasets is to benefit from secondary domains of information [48]. Domain adaptation [41, 44, 57, 49, 61] and zero-shot learning (ZSL) [40, 60, 39, 27, 75, 32, 34] are two primary learning settings to benefit from auxiliary domains to relax data annotation. Most works within domain adaptation use the strong assumption that the two domains are homogeneous and also share the same set of classes. It is assumed that we have fully annotated data in one of the domains and in the second domain only unannotated data is accessible. The goal would be to train a classifier for the unannotated domain by traninferring knowledge from the annotated domain [52, 47]. In contrast, zero-shot learning considers learning unseen classes in a single domain, usually visual data, by coupling it with an auxiliary domain, usually natural language information, using a number of seen classes for which we have bi-view annotations for both domains. We focus on ZSL in this work. ZSL is inspired by the ability of humans to recognize new visual classes using their semantic descriptions. Humans remarkably are extremely good at learning enormous numbers of classes from little data using descriptions in natural language. Consider the problem of classifying animal images. It is estimated that as many as one million different species of animals have been identified, with as many as ten thousand new species being discovered annually. This classification problem is a case when ZSL can be extremely helpful. Most people probably have not seen an image of a 'tardigrade', nor heard of this species. If you have not heard of this animal specie, we can intuitively demonstrate how ZSL can be possible for this class. Consider the following sentence from Wikipedia: "Tardigrades" (also known as water bears or moss piglets) are water-dwelling, eight-legged, segmented micro animals." Given this textual description, most humans can easily identify the creature in Figure 1 (a) (left) as a Tardigrade, even though they may have never seen one before. Humans can easily perform this ZSL task by: 1) identifying the semantic features that describe the class *Tardigrade* as 'bear-like', 'piglet-like', 'water-dwelling', 'eight-legged', 'segmented', and 'microscopic animal', 2) parsing the image into its visual attributes (see Figure 1 (a)), and 3) matching the parsed visual features to the parsed textual information. In other words, humans can transfer knowledge from the domain of natural language to solve problems in the vision domain.

4

ZSL has been implemented in computer vision based on the above intuition. To this end, we can parse textual features into a vector of either predetermined binary attributes (e.g., water-dwelling) or continues features using, e.g., using *word2vec* [36]. We can also use pretrained deep convolutional neural networks (CNNs) to extract visually rich features from natural images to parse the visual information. ZSL algorithms generally learn a mapping between the visual features and semantic attributes using a shared intermediate embedding space [40, 60, 39, 27, 75, 32, 34]. After model training, his intermediate space can be used to transfer knowledge across the two domains. ZSL has been found to be practically beneficial on applications, including face verification [24], video understanding [67], and object recognition [26]. We can categorize the ZSL algorithms into two primary subgroups. A group of ZSL methods model the cross-domain mapping as a linear function [14, 46, 1, 2]. Since a simple hypothesis space is used to learn the cross-domain mapping, learning the cross-domain mapping is computationally efficient. However, nonlinear relations between the domains may not be encoded well. In contrast, more recent methods use deep neural networks to model the mapping [68, 70, 30, 29]. Although deep neural networks usually lead to state-of-the-art performance for diverse set of applications, training a neural network for ZSL will require more data instances of seen classes which goes against the very goal of ZSL. ZSL methods that can learn the cross-domain mapping as a nonlinear function and at the same time do not have significant complexity are desirable. In this paper, we follow a middle-ground between the above two subgroups to develop a ZSL algorithm which has a competitive performance despite having a less computational model training complexity.

We develop a new ZSL algorithm based on coupled dictionary learning (CDL) [72] to relate the visual features and the semantic attributes. CDL in essence is a (semi-)linear model but it can encode nonlinearities. This ability stems from the fact that the hypothesis space of a dictionary is the union of linear subspaces, i.e., a nonlinear space. Our specific contributions include:
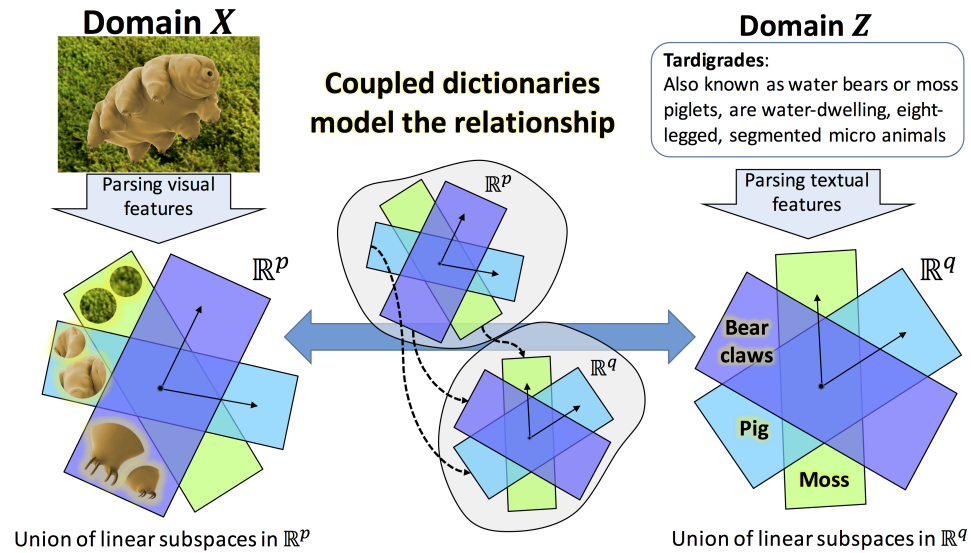
1. We formulate ZSL as a coupled dictionary learning problem and demonstrate by solving a dictionary learning problem, ZSL can be performed.
2. We provide an efficient algorithm to solve the resulting joint dictionary learning optimization problem.

3. We also address the challenges of hubness [12] and domain-shift [15] in ZSL using by augmenting our base optimization problem with suitable regularization terms.
4. We provide theoretical analysis which establishes PAC-learnability of our proposed algorithm.
5. We perform experiments on primary benchmark datasets and demonstrate that our method is effective and compares favorably with respect to the state-of-the-art.
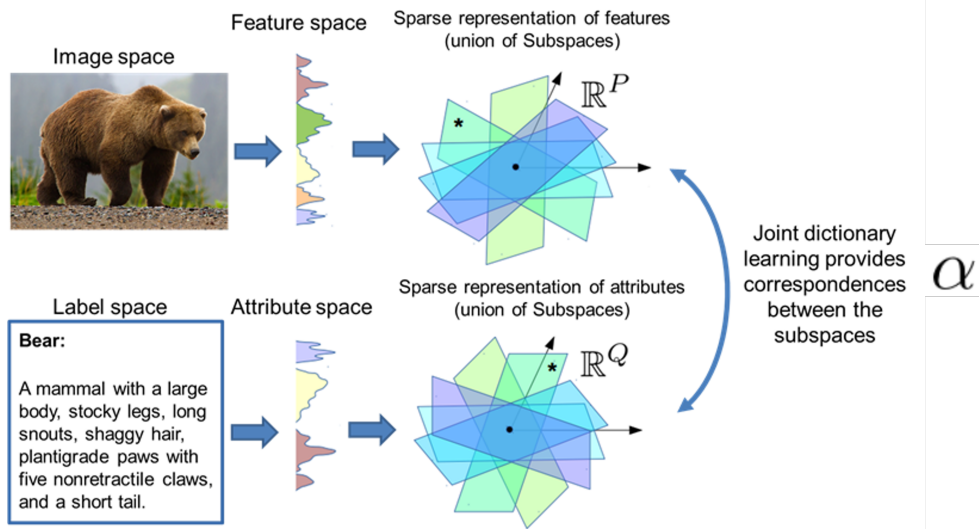
The remaining of the paper is as follows. In section 2, we will explain the formulation we use for ZSL and our high-level algorithmic idea. Section 3 summarizes our algorithm to tackle ZSL. Section 4 provides a theoretical analysis for our algorithm. We have provides experimental validation in section 5. The paper finally is concluded in section 6 with a brief discussion.

## 2. Problem Formulation and Technical Rationale

We follow Palatucci et. al. to formulate ZSL as a two stage estimation problem [40]. Consider a visual feature metric space $\mathcal{F}$ of dimension $p$, a semantic metric space $\mathcal{A}$ with dimension of $q$ as well as a class label set $\mathcal{Y}$ with dimension $K$ which ranges over a finite alphabet of size $K$ (images can potentially have multiple memberships in the classes). As an example $\mathcal{F} = \mathbb{R}^p$ for the visual features extracted from a deep CNN and $\mathcal{A} = \{0,1\}^q$ when a binary code of length $q$ is used to identify the presence/absence of various characteristics in an object [27]. We are given a labeled dataset $\mathcal{D} = \{((\mathbf{x_i}, \mathbf{z_i}), \mathbf{y_i})\}_{\mathbf{i=1}}^{\mathbf{N}}$ of features of seen images and their corresponding semantic attributes, where $\forall i : \mathbf{x_i} \in \mathcal{F}, \mathbf{z_i} \in \mathcal{A}$, and $\mathbf{y_i} \in \mathcal{Y}$. We are also given the unlabeled attributes of unseen classes $\mathcal{D}' = \{\mathbf{z'_j}, \mathbf{y'_j}\}_{\mathbf{j=1}}^{\mathbf{M}}$ (i.e., we have access to textual information for a wide variety of objects but not have access to the corresponding visual information). Following the standard assumption in ZSL, we assume that the set of the seen and the unseen classes are disjoint. The challenge is how to learn a model on the labeled set and transfer the learned knowledge to the unlabeled set. We also assume that the same semantic attributes could not describe two different classes of objects, i.e., by knowing semantic attribute of an image one can classify that image. The goal is to learn from the labeled dataset to classify images of unseen classes. For further clarification, consider an instance of ZSL in which features extracted from images of horses and tigers are included in

6

**Domain $X$**

**Coupled dictionaries model the relationship**

**Domain $Z$**

**Tardigrades**: Also known as water bears or moss piglets, are water-dwelling, eight-legged, segmented micro animals

Parsing visual features

$\mathbb{R}^p$

Parsing textual features

$\mathbb{R}^p$

$\mathbb{R}^q$

$\mathbb{R}^q$

**Bear claws**

**Pig**

**Moss**

Union of linear subspaces in $\mathbb{R}^p$

Union of linear subspaces in $\mathbb{R}^q$

(a) Training

Feature space

Sparse representation of features (union of Subspaces)

Image space

$\mathbb{R}^P$

Joint dictionary learning provides correspondences between the subspaces

$\alpha$

Label space

Attribute space

Sparse representation of attributes (union of Subspaces)

**Bear:**

A mammal with a large body, stocky legs, long snouts, shaggy hair, plantigrade paws with five nonretractile claws, and a short tail.

$\mathbb{R}^Q$

(b) Testing

Figure 1: High-level overview of our approach: (a) we consider that the visual features (left) and the attribute features (right) can be represented sparsely using unions of subspaces that are modeled using two dictionaries. We train the two dictionaries such that these two space are matched (middle), leading to coupling the visual and the attribute features in the shared space. (b) During testing, we solve for the sparse representation of the input images and then find its corresponding semantic description, y first solving for the joint sparse vector using the visual dictionary and then searching for the closest semantic description.

7

seen visual features $X = [\mathbf{x_1}, ..., \mathbf{x_N}]$, where $\mathbf{x_i} \in \mathcal{F}$, but $X$ does not contain features of zebra images. On the other hand, the semantic attributes contain information of all the seen images $Z = [\mathbf{z_1}, ..., \mathbf{z_N}]$ for $\mathbf{z_i} \in \mathcal{A}$ and the unseen images $Z' = [\mathbf{z'_1}, ..., \mathbf{z'_M}]$ for $\mathbf{z'_j} \in \mathcal{A}$ including the zebras. The goal is by learning the relationship between the image features and the attributes "horse-like" and "has stripes" from the seen images, we assign an unseen zebra image to its corresponding attribute.

Within this paradigm, ZSL can be performed by a two stage estimation. First the visual features can be mapped into the semantic space and then the label is estimated in the semantic space. More formally, we want to learn the mapping $\phi : \mathcal{F} \rightarrow \mathcal{A}$ which relates the visual space and the attribute space. We also assume that $\psi : \mathcal{A} \rightarrow \mathcal{Y}$ is the mapping between the semantic space and the label space. The mapping $\psi$ can be as simple as nearest neighbor, i.e., we assign labels according to the closest semantic attribute in the semantic attribute space. Having learned this mapping, one can recover the corresponding attribute vector for an unseen image using the image features and then classify the image using a second mapping $\boldsymbol{y} = (\psi \circ \phi)(\mathbf{x})$, where '$\circ$' represents function composition. The goal is to introduce a type of bias to learn both mappings using the labeled dataset. Having learned both mappings, ZSL is feasible in the testing stage. Because if the mapping $\phi(\cdot)$ can map an unseen image close enough to its true semantic features, then intuitively the mapping $\psi(\cdot)$ can still recover the corresponding class label. Following our example, if the function $\phi(\cdot)$ can recover that an unseen image of a zebra is "horse-like" and "has stripes", then it is likely that the mapping $\psi(\cdot)$ can classify the unseen image. Our core idea is to benefit from coupled dictionary learning [72] to model these mappings.

## 2.1. Proposed Idea

The idea of using coupled dictionaries to map data from a given metric space to a second related metric space was first proposed by Yang et al. [72] for single image super-resolution problem [45]. Their pioneer idea is to assume that the high-resolution and the corresponding low-resolution patches of image can be represented with a unique joint sparse vector in two low- and high-resolution dictionaries. The core idea is that in the absence of a high-resolution image and given its low-resolution version, the corresponding joint sparse representation can be computed using sparse signal recovery. The sparse vector is then can be used to generate the

8

high-resolution image patches using the low-resolution image. They also propose an efficient algorithm to learn the low- and the high-resolution dictionaries using a training set, consisting of both low- and high-resolution version of natural images. Our goal is to follow the same approach but replacing the low- and the high-resolution metric spaces with the visual and the semantic spaces, respectively.

As a big picture to understand our approach, Figure 1 captures the gist of our idea. In Figure 1 (a), visual features are extracted via CNNs (left sub-figure). For example, the last fully-connected layer of a trained CNN can be removed and the rest of the deep net can be used as a feature extractor given an input image. These features have been demonstrated to be highly descriptive and lead to the state-of-the-art performance for many computer vision and image processing tasks. To perform ZSL, we need textual description of the classes, too. Textual description of many classes are cheap to obtain, e.g., Wikipedia. The semantic attributes then can be provided via textual feature extractors like word2vec or potentially via human annotations (right sub-figure). It is assumed that both the visual features and the semantic attributes can be represented sparsely using the visual and the attribute dictionaries that is modeled as a shared union of linear subspaces (left and right sub-figures). The idea here is that the sparse representation vectors for both feature vectors are equal. Thus, one can map an image to its textual description in this space using the joint sparse vector (middle sub-figure).

The intuition from a co-view perspective is that both the visual and the attribute features provide information about the same class or entity, and so each can augment learning of the other to improve performance [20, 51]. Each underlying class is common to both views, and so we can find task embeddings that are consistent for both the visual features and their corresponding attribute features. The main challenge is to learn these dictionaries for the visual and the attribute spaces. Having learned these two dictionaries, zero-shot classification can be performed by mapping images of unseen classes into the attribute space, where classification can be simply done via nearest neighbor or more advanced clustering approaches (Figure 1 (b)). Given the coupled nature of the learned dictionaries, an image can be mapped to its semantic attributes by first finding the sparse representation with respect to the visual dictionary. Our algorithm is equipped with a novel entropy minimization regularizer [17], which facilitates a better solution for the ZSL problem. The entropy regularization helps to tackle

9

the challenge of domain-shift in zero-shot learning. Next the semantic attribute dictionary can be used to recover the attribute vector from the joint sparse representation which can then be used for classification (Figure 1 (b)). We also show that a transductive approach applied to our attribute-aware JD-ZSL formulation improves the perofrmance via mitigating the challenge of hubness in high dimensions. Our experiments demonstrate that our algorithm leads to competitive performance on four standard ZSL benchmark datasets.

### 2.2. Technical Rationale

For the rest of our discussion we assume that $\mathcal{F} = \mathbb{R}^p$, $\mathcal{A} = \mathbb{R}^q$, and $\mathcal{Y} \subset \mathbb{R}^K$. Most ZSL algorithms focus on learning $\phi(\cdot)$ because even using a simple method like nearest neighbor classification for $\psi(\cdot)$ yields descent ZSL performance. The simplest ZSL approach is to assume that the mapping $\phi : \mathbb{R}^p \to \mathbb{R}^q$ is linear, $\phi(\mathbf{x}) = W^T \mathbf{x}$ where $W \in \mathbb{R}^{p \times q}$, and then minimize the regression error $\frac{1}{N} \sum_i \|W^T \mathbf{x}_i - \mathbf{z}_i\|_2^2$ to learn $W$. Although closed form solution exists for $W$, the solution contains the inverse of the covariance matrix of $X$, $(\frac{1}{N} \sum_i (x_i x_i^T))^{-1}$, which requires a large number of data points for accurate estimation. To overcome this problem, various regularizations are considered for $W$. Decomposition of $W$ as $W = P \Lambda Q$, where $P \in \mathbb{R}^{p \times l}$, $\Lambda \in \mathbb{R}^{l \times l}$, $Q \in \mathbb{R}^{l \times q}$, and $l < min(p,q)$ can be helpful. Intuitively, $P$ is a right linear operator that projects $\mathbf{x}$'s into a shared low-dimensional subspace, $Q$ is a left linear operator that projects $\mathbf{z}$ into the same shared subspace, and $\Lambda$ provides a bi-linear similarity measure in the shared subspace. The regression problem can then be transformed into maximizing $\frac{1}{N} \sum_i \mathbf{x}_i^T P \Lambda Q \mathbf{z}_i$, which is a weighted correlation between the embedded $\mathbf{x}$'s and $\mathbf{z}$'s. This is the essence of many ZSL techniques including Romera-Paredes et al. [46]. This technique can be extended to nonlinear mappings using kernel methods. However, the choice of kernels remains an open challenge.

The mapping $\phi : \mathbb{R}^p \to \mathbb{R}^q$ can also be chosen to be highly nonlinear, e.g., deep neural networks. Let a deep net be denoted by $\phi(.; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ represents the synaptic weights and biases. ZSL can then be addressed by minimizing $\frac{1}{N} \sum_i \|\phi(\mathbf{x}_i; \boldsymbol{\theta}) - \mathbf{z}_i\|_2^2$ with respect to $\boldsymbol{\theta}$. Alternatively, one can nonlinearly embed $\mathbf{x}$'s and $\mathbf{z}$'s in a shared metric space via deep nets, $p(\mathbf{x}; \boldsymbol{\theta}_x) : \mathbb{R}^p \to \mathbb{R}^l$ and $q(\mathbf{z}; \theta_z) : \mathbb{R}^q \to \mathbb{R}^l$, and maximize their similarity measure in the embeding space, $\frac{1}{N} \sum_i p(\mathbf{x}_i; \boldsymbol{\theta}_x)^T q(\mathbf{z}_i; \boldsymbol{\theta}_z)$. This approach might improve performance for particular data sets, but in turn would

10

require more training samples. Note however, this might not be plausible for ZSL because the very reason and motivation to perform ZSL is to learn from as few labeled data points as possible.

Comparing the above approaches, nonlinear methods are computationally expensive and require a large training dataset. In contrast, linear ZSL algorithms are efficient, but their performances are lower. As a compromise, we can model nonlinearities in data distributions as a union of linear subspaces using coupled dictionaries [9]. The relationship between the metric spaces is also reflected in the learned dictionaries. This allows a nonlinear scheme with a computational complexity comparable to linear techniques.

## 3. Zero-Shot Learning using Coupled Dictionary Learning

In the standard dictionary learning framework, a sparsifying dictionary is learned using a given training sample set $X = [\mathbf{x}_1, ..., \mathbf{x}_N]$ for a particular class of signals. Unlike standard dictionary learning, coupled dictionary learning has been proposed to couple related features from two metric spaces to learn the mapping function between these spaces. Following the same framework, the gist of our approach is to learn the mapping $\phi : \mathbb{R}^p \rightarrow \mathbb{R}^q$ through two dictionaries, $D_x \in \mathbb{R}^{p \times r}$ and $D_z \in \mathbb{R}^{q \times r}$ for $X$ and $[Z, Z']$ sets, respectively, where $r > max(p, q)$. The goal is to find a shared sparse representation $\mathbf{a}_i$ for $\mathbf{x}_i$ and $\mathbf{z}_i$, such that $\mathbf{x}_i = D_x \mathbf{a}_i$ and $\mathbf{z}_i = D_z \mathbf{a}_i$. The shared sparse representation couples the semantic and visual feature spaces. We first explain the training procedure for the two dictionaries, and then the way we use these dictionaries to estimate $\phi(\cdot)$.

*3.1. Training Phase*

The standard dictionary learning is based upon minimizing the empirical average estimation error $\frac{1}{N}\|X - D_x A\|_F^2$ on a given training set $X$, where an additive $\ell_1$ regularization penalty term on $A$ enforces sparsity:

$$D_x^*, A^* = \underset{D_x, A}{\mathrm{argmin}} \left\{ \frac{1}{N}\|X - D_x A\|_F^2 + \lambda \|A\|_1 \right\}$$
$$\text{s.t. } \left\|D_x^{[i]}\right\|_2^2 \leq 1 \ . \tag{1}$$

Here $\lambda$ is the regularization parameter that controls sparsity level and $D_x^{[i]}$ is the $i^{th}$ column of $D_x$. The columns of the dictionary are normalized to

obtain a unique dictionary. Alternatively, following the Lagrange multiplier technique, the Frobenius norm of $D_x$ could be used as a regularizer in place of the constraint. The above problem is not a convex optimization problem, but is convex in each variable alone; it is biconvex with respect to the variables $D_x$ and $A$. As a result, most dictionary learning algorithms use alternation on variables $D_x$ and $A$ to solve (1) which leads to iterations on two separate optimizations. Each optimization problem is performed solely on one of the variables, assuming the other variable to be constant. Upon a suitable initialization, Eq. (1) reduces to a number of parallel sparse recovery when the dictionary is fixed, i.e., LASSO problems which can be solved efficiently. Then, for a fixed $A$, Eq. (1) reduces to a standard quadratically constrained quadratic program (QCQP) problem which can be solved efficiently with iterative methods such as conjugate gradient descent algorithms even for high-dimensional (large $p$) and huge problems (large $r$). This alternative procedure on the variables is repeated until some convergence criterion is met.

In our coupled dictionary learning framework, we aim to learn coupled dictionaries $D_x$ and $D_z$ such that they share the sparse coefficients $A$ to represent the seen visual features $X$ and their corresponding attributes $Z$, respectively. Intuitively this means that visual features for an object have corresponding semantic features. On the other hand, $D_z$ also needs to sparsify the semantic attributes of other (unseen) classes, $Z'$, to enable performing ZSL. Hence, we propose the following optimization problem to learn both dictionaries:

$$
\begin{aligned}
D_x^*, A^*, D_z^*, B^* = & \underset{D_x, A, D_z, B}{\operatorname{argmin}} \left\{ \frac{1}{Np} \left( \|X - D_x A\|_F^2 + \frac{p\lambda}{r} \|A\|_1 \right) \right. \\
& \left. + \frac{1}{Nq} \|Z - D_z A\|_F^2 + \frac{1}{Mq} \left( \|Z' - D_z B\|_F^2 + \frac{q\lambda}{r} \|B\|_1 \right) \right\} \\
& \text{s.t.:} \left\| D_x^{[i]} \right\|_2^2 \leq 1, \ \left\| D_z^{[i]} \right\|_2^2 \leq 1 .
\end{aligned}
\tag{2}
$$

The above formulation combines the dictionary learning problems for $X$ and $Z$ by coupling them via the joint sparse code matrix $A$, and also enforces $D_z$ to be a sparsifying dictionary for $Z'$ with the sparse codes $B$. Similar to Eq. (1), the optimization in Eq. (2) is biconvex in $(D_x, D_z)$ and $(A, B)$. Hence, we use an alternative scheme to update $D_x$ and $D_z$ for a local solution.

12

---
**Algorithm 1** Coupled Dictionary Learning ($\{X, Z, Z'\}, \lambda, r, itr$)
---
1: $D_x \leftarrow \text{RandomMatrix}_{p,r}$, $D_z \leftarrow \text{RandomMatrix}_{q,r}$
2: $D_x \leftarrow \text{update}(D_x, \{X\}, \lambda, \beta)$                  Eq. 3
3: $D_z \leftarrow \text{update}(D_z, \{Z, Z', A\}, \lambda, \beta)$        Eq. 4
---

First we add the constraints on dictionary atoms in (2) as a penalty term:

$$\min_{A, D_x} ||X - D_x A||_2^2 + \lambda ||A||_1 + \beta ||D_x||_2^2 \tag{3}$$

We can solve Eq. (3) by alternately solving LASSO for $A$ and taking gradient steps with respect to $D_x$. Next we solve the following problem:

$$\min_{B, D_z} ||Z - D_z A||_2^2 + ||Z' - D_z B||_2^2 + \lambda ||B||_1 + \beta ||D_z||_2^2 \tag{4}$$

by alternately solving the LASSO for $B$ and taking gradient steps with respect to $D_z$, (while holding $A$ fixed as the solution found in Eq. (3).

Algorithm 1 summarizes the coupled dictionary learning procedure. The learned dictionaries then can be used to perform ZSL in the testing phase (see Figure 1 (b)).

### 3.2. Prediction of Unseen Attributes

In the testing phase, we are only given the extracted features from unseen images and the goal is to predict their corresponding semantic attributes. We propose two different methods to predict the semantic attributes of the unseen images based on the learned dictionaries in the training phase, namely attribute-agnostic prediction and attribute-aware prediction methods.

### 3.2.1. Attribute-Agnostic Prediction

The attribute-agnostic (AAg) method is the naive way of predicting the semantic attributes from an unseen image $\mathbf{x}'_i$. In the attribute-agnostic formulation, we first find the sparse representation $\alpha_i$ of the unseen image $\mathbf{x}'_i$ by solving the following LASSO problem,

$$\alpha_i = \text{argmin}_{\mathbf{a}} \left\{ \frac{1}{p} ||\mathbf{x}_i - D_x \mathbf{a}||_2^2 + \frac{\lambda}{r} ||\mathbf{a}||_1 \right\} . \tag{5}$$

13

and its corresponding attribute is estimated by $\hat{z}_i = D_z \boldsymbol{\alpha}_i$. We call this formulation attribute-agnostic because the sparse coefficients are found without any information from the attribute space. We use AAg as a baseline to demonstrate the effectiveness of the attribute-aware prediction.

### 3.2.2. Attribute-Aware Prediction

In the attribute-aware (AAw) formulation, we would like to find the sparse representation $\boldsymbol{\alpha}_i$ to not only approximate the input visual feature, $\mathbf{x}'_i \approx D_x \boldsymbol{\alpha}_i$, but also provide an attribute prediction, $\hat{z}_i = D_z \boldsymbol{\alpha}_i$, that is well resolved in the attribute space. This mean that ideally we want to have $\hat{z}_i = \mathbf{z}'_m$, for some $m \in \{1, ..., M\}$. To achieve this, we define the soft assignment of $\hat{z}_i$ to $\mathbf{z}'_m$, denoted by $p_m$, using the Student's t-distribution as a kernel to measure similarity between $\hat{z}_i = D_z \boldsymbol{\alpha}_i$ and $\mathbf{z}'_m$,

$$
p_m(\boldsymbol{\alpha}_i) = \frac{\left(1 + \frac{\|D_z \boldsymbol{\alpha}_i - \mathbf{z}'_m\|_2^2}{\rho}\right)^{-\frac{\rho+1}{2}}}{\sum_k \left(1 + \frac{\|D_z \boldsymbol{\alpha}_i - \mathbf{z}'_k\|_2^2}{\rho}\right)^{-\frac{\rho+1}{2}}} ,
\tag{6}
$$

where $\rho$ is the kernel parameter. We chose t-distribution as it is less sensitive to the choice of kernel parameter, $\rho$.

Ideally, $p_m(\boldsymbol{\alpha}_i) = 1$ for some $m \in \{1, ..., M\}$ and $p_j(\boldsymbol{\alpha}_i) = 0$ for $j \neq m$. In other words, the ideal soft-assignment $\mathbf{p} = [p_1, p_2, ..., p_M]$ would be one-sparse and hence would have minimum entropy. This motivates our attribute-aware formulation, which penalizes Eq. 5 with the entropy of $\mathbf{p}$.

$$
\boldsymbol{\alpha}_i = \operatorname{argmin}_{\mathbf{a}} \left\{ \underbrace{\frac{1}{p} \|\mathbf{x}'_i - D_x \mathbf{a}\|_2^2 - \gamma \sum_m p_m(\mathbf{a}) \log(p_m(\mathbf{a}))}_{g(\mathbf{a})} + \frac{\lambda}{r} \|\mathbf{a}\|_1 \right\} ,
\tag{7}
$$

where $\gamma$ is the regularization parameter for entropy of the soft-assignment probability vector $\mathbf{p}$, $H_{\mathbf{p}}(\boldsymbol{\alpha})$. The entropy minimization has been successfully used in several works [17] either as a sparsifying regularization or to boost the confidence of classifiers. Such regularization, however, turns the optimization in Eq. (7) into a nonconvex problem. However, since $g(\mathbf{a})$ is differentiable and the $\ell_1$ norm is continuous, we can apply proximal gradient descent [42] or ADMM [3] to solve it. In practice, we found

14

---

**Algorithm 2** Zero-shot Prediction ($\mathbf{x}_i$ $\lambda$)

---

 1: Attribute-Agnostic prediction:
 2: $\boldsymbol{\alpha}_i \leftarrow \text{argmin}_\mathbf{a} \frac{1}{p}\|\mathbf{x}_i - D_x\mathbf{a}\|_2^2 + \frac{\lambda}{r}\|\mathbf{a}\|_1$
 3: $\hat{\mathbf{z}}_i = D_z\boldsymbol{\alpha}_i$
 4: $\mathbf{z}'_m = \text{argmin}_{\mathbf{z}' \in Z'} \|\mathbf{z}' - \hat{\mathbf{z}}_i\|_2$
 5: Attribute-Aware prediction:
 6: $\boldsymbol{\alpha}_i \leftarrow \text{argmin}_\mathbf{a} \frac{1}{p}\|\mathbf{x}'_i - D_x\mathbf{a}\|_2^2 - \gamma H_{\boldsymbol{p}}(\boldsymbol{\alpha}) + \frac{\lambda}{r}\|\mathbf{a}\|_1$
 7: $\hat{\mathbf{z}}_i = D_z\boldsymbol{\alpha}_i$
 8: $\mathbf{z}'_m = \text{argmin}_{\mathbf{z}' \in Z'} \|\mathbf{z}' - \hat{\mathbf{z}}_i\|_2$
 9: Transducer Prediction
10: Solve 7 to predict $\boldsymbol{\alpha}_i$ for all unseen samples.
11: Use label propagation to spread the labels.

---

that gradient descent applied directly on Eq. (7) works fine since Eq. (7) is differentiable almost everywhere. Note, however, a good initialization is needed for achieving an accurate solution, due to the non-convex nature of the objective function. Therefore we initialize $\boldsymbol{\alpha}$ from the solution of the attribute-agnostic formulation which is an approximate solution. Finally the corresponding attributes are estimated as $\hat{z}_i = D_z\boldsymbol{\alpha}_i$, for $i = 1, ..., l$.

### 3.3. From Predicted Attributes to Labels

To predict the image labels, one needs to assign the predicted attributes to the $M$ attributes of the unseen classes $Z'$. We performed this task in two approaches, namely the inductive approach and the transductive approach.

#### 3.3.1. Inductive Approach

In the inductive approach, the inference can be performed using a nearest neighbor (NN) approach in which the label of each individual $\hat{\mathbf{z}}_i$ is assigned to be the label of its nearest neighbor $\mathbf{z}'_m$:

$$\mathbf{z}'_m = \text{argmin}_{\mathbf{z}' \in Z'} \left\{ \|\mathbf{z}' - \hat{\mathbf{z}}_i\|_2 \right\} , \tag{8}$$

and the corresponding label of $\mathbf{z}'_m$ is assigned to $\hat{\mathbf{z}}_i$. Note, however, the structure of $\hat{\mathbf{z}}_i$'s is not taken into account if we use Eq. (8). Looking at the t-SNE embedding visualization of $\hat{\mathbf{z}}_i$'s and $\mathbf{z}'_m$'s in Figure 2 (b) (details are explained later), it can be seen that nearest neighbor will not provide an optimal label assignment.

### 3.3.2. Transductive Learning

In the transductive attribute-aware (TAA) method, the attributes for all test images (i.e., unseen) are first predicted to form $\hat{Z} = [\hat{\mathbf{z}}_1, ..., \hat{\mathbf{z}}_L]$. Next, a graph is formed on $[Z', \hat{Z}]$, where the labels for $Z'$ are known and the task is to infer the labels of $\hat{Z}$. Intuitively, we want the data points that are close together to have similar labels. This problem can be formulated as a graph-based semi-supervised label propagation.

We follow the work of Zhou et al. [78] and spread the labels of $Z'$ to $\hat{Z}$. We form a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where the set of nodes $\mathcal{V} = \{\mathbf{v}\}_1^{M+L} = [\mathbf{z}_1', ..., \mathbf{z}_M', \hat{\mathbf{z}}_1, ..., \hat{\mathbf{z}}_L]$, and $\mathcal{E}$ is the set of edges whose weights reflect the similarities between the attributes. Note that the first $M$ nodes are labeled and our task is to use these labels to predict the labels of the rest of the nodes. We use a Gaussian kernel to measure the edge weights between the connected nodes, $W_{mn} = exp\{-\|\mathbf{v}_m - \mathbf{v}_n\|^2/2\sigma^2\}$, where $\sigma$ is the kernel parameter and $W_{ii} = 0$. To construct the graph $\mathcal{G}$ one can utilize efficient k-NN graph construction methods, where the assumption is that a neighbor of a neighbor is more likely to be a neighbor. Let $F \in \mathbb{R}^{M \times (M+L)}$ corresponds to a classification of the nodes, where $F_{mn}$ is the probability of $\mathbf{v}_n$ belonging to the $m$'th class. Let $Y \in \mathbb{R}^{M \times (M+L)} = [\mathbf{I}_{M \times M}, \mathbf{0}_{M \times L}]$ represent the initial labels, where $\mathbf{I}$ denotes an identity matrix and $\mathbf{0}$ denotes a zeros matrix. From a Graph-Signal Processing point of view, $F$ is a signal defined on the graph $\mathcal{G}$, and one requires this signal to be smooth. Zhou et al. [78] proposed to obtain a smooth signal on graph $\mathcal{G}$ that fits the initial known labels,

$$\text{argmin}_F \left\{ \frac{1}{2} \left( \sum_{m,n} W_{mn} \| \frac{F_m}{\sqrt{D_{mm}}} - \frac{F_n}{\sqrt{D_{nn}}} \|^2 + \mu \sum_m \|F_m - Y_m\|^2 \right) \right\}, \quad (9)$$

where $D \in \mathbb{R}^{(M+L) \times (M+L)}$ is the diagonal degree matrix of graph $\mathcal{G}$, $D_{mm} = \sum_n W_{mn}$, and $\mu$ is the fitness regularization. Note that the first term in Eq. (9) enforces the smoothness of signal $F$ and the second term enforces the fitness of $F$ to the initial labels. Then Eq. (9) would have the following solution:

$$F = \frac{\mu}{1+\mu} \left( \mathbf{I} - \frac{1}{1+\mu} (D^{-\frac{1}{2}} W D^{-\frac{1}{2}}) \right)^{-1} Y. \quad (10)$$

Algorithm 2 summarizes the zero-shot label prediction procedure.

16

### 4. Theoretical Discussion

In this section, we establish PAC-learnability of the proposed algorithm. We provide a PAC style generalization error bound for the proposed ZSL algorithm. The goal is to establish conditions under which, our ZSL algorithm can identify instances from unseen classes. We use the framework developed by Palatucci et. al. [40], to derive this bound. The core idea is that if we are able to recover the semantic attributes of a given image with high accuracy, then the correct label can be recovered with high probability as well. Note that three probability events are involved in the probability event of predicting an unseen class label correctly, denoted by $P_t$:

1. Given a certain confidence parameter $\delta$ and error parameter $\epsilon$, a dictionary can be learned with $M_{\epsilon,\delta}$ samples. We denote this event by $\mathcal{D}_\epsilon$. Hence $P(\mathcal{D}_\epsilon) = 1 - \delta$ and $\mathbb{E}(\|\boldsymbol{x} - D\boldsymbol{a}\|_2^2) \leq \epsilon$, where $\mathbb{E}(\cdot)$ denotes statistical expectation.

2. Given the event $\mathcal{D}_\epsilon$ (learned dictionaries), the semantic attribute can be estimated with high probability. We denote this event by $\mathcal{S}_\epsilon|\mathcal{D}_\epsilon$.

3. Given the event $\mathcal{S}_\epsilon|\mathcal{D}_\epsilon$, the true label can be predicted. We denote this event by $\mathcal{T}|\mathcal{S}_\epsilon$ and so $P(\mathcal{T}|\mathcal{S}_\epsilon) = 1 - \zeta$.

Therefore, the event $P_t$ can be expressed as the following probability decoupling by multiplying the above probabilities:

$$P_t = P(\mathcal{D}_\epsilon)P(\mathcal{S}_\epsilon|\mathcal{D}_\epsilon)P(\mathcal{T}|\mathcal{S}_\epsilon) \ . \tag{11}$$

Our goal is given the desired values for confidence parameters $\zeta$ and $\delta$ for the two ZSL stages, i.e., $P(\mathcal{D}_\epsilon) = 1 - \delta$ and $P(\mathcal{T}|\mathcal{S}_\epsilon) = 1 - \zeta$, we compute the necessary $\epsilon$ for that level of prediction confidence as well as $P(\mathcal{S}_\epsilon|\mathcal{D}_\epsilon)$. We also need to compute the number of required training samples to secure the desired errors. Given $P(\mathcal{T}|\mathcal{S}_\epsilon) = 1 - \zeta$, we compute $\epsilon$ and the conditional probability $P(\mathcal{S}_\epsilon|\mathcal{D}_\epsilon)$.

To establish the error bound, we need to compute the maximum error in predicting the semantic attributes of a given image, for which we still can predict the correct label with high probability. Intuitively, this error depends on geometry of $\mathcal{A}$ and probability distribution of semantic attributes of the classes in this space, $\mathcal{P}$. For example, if semantic attributes of two classes are very close, then error tolerance for those classes will be less than two classes with distant attributes. To model this intuition, we focus our analysis on nearest neighbor label recovery. Let $\hat{\boldsymbol{z}}$ denote the predicted attribute for a given image by our algorithm. Let $d(\hat{\boldsymbol{z}}, \boldsymbol{z}') : \mathbb{R}^q \times \mathbb{R}^q \to \mathbb{R}$

denote the distance between this point and another point in the semantic space. We denote the distribution function for this distance as $R_{\hat{z}}(t) = P(d(\hat{z}, z') \leq t)$. Let $T_{\hat{z}}$ denote the distance to the nearest neighbor of $\hat{z}$ and $W_{\hat{z}}(t) = P(T_{\hat{z}} \leq t)$ denotes its probability distribution. The latter distribution has been computed by Ciaccia and Patella [8] as:

$$W_{\hat{z}}(t) = 1 - \left(1 - R_{\hat{z}}(t)\right)^{n} , \tag{12}$$

where $n$ is the number of points drawn from the distribution $\mathcal{P}$. Note that the function $R_{\hat{z}}(t)$ is an empirical distribution which depends on the distribution of semantic feature space, $\mathcal{P}$, and basically is the fraction of sampled points from $\mathcal{P}$ that are less than some distance $t$ away from $\hat{z}$.

Following the general PAC learning framework, given a desired probability (confidence) $\zeta$, we want the distance $T_{\hat{z}}$ to be less than the distance of the predicted attribute $\hat{z}$ from the true semantic description of the true class that it belongs to, i.e., or $W_{\hat{z}}(\tau_{\hat{z}}) \leq \zeta$. Now note that since $W_{\hat{z}}(\cdot)$ is a cumulative distribution (never decreasing), $W_{\hat{z}}^{-1}(\cdot)$ is well-defined as $W_{\hat{z}}^{-1}(\zeta) = \text{argmax}_{\tau_{\hat{z}}}[W_{\hat{z}}(\tau_{\hat{z}}) \leq \zeta]$. If $\tau_{\hat{z}} \leq W_{\hat{z}}^{-1}(\zeta)$, then the correct label can be recovered with probability of $1 - \zeta$. Hence, prior to label prediction (which itself is done for a given confidence parameter $\delta$), the semantic attributes must be predicted with the true error at most $\epsilon_{max} = W_{\hat{z}}^{-1}(\zeta)$ and we need to ensure that semantic attribute prediction achieves this error bound, that is $\mathbb{E}_{\mathbf{z}}\left(\|\mathbf{z} - D_{z}\mathbf{a}^*\|_2^2\right) \leq W_{\hat{z}}^{-1}(\zeta)$. To ensure this to happen, we rely on the following theorem on PAC-learnability of the dictionary learning (1) derived by Gribonval et. al [16]:

**Theorem 1 [16]**: Consider dictionary learning problem in (1), and the confidence parameter $\delta$ $(P(\mathcal{D}_\epsilon) = 1 - \delta)$ and the error parameter $\epsilon_{max} = W_{\hat{z}}^{-1}(\zeta)$ in standard PAC-learning setting. Then the number of required samples to learn the dictionary $M_{W_{\hat{z}}^{-1}, \delta}$ satisfies the following relation:

$$W_{\hat{z}}^{-1}(\zeta) \geq 3\sqrt{\frac{\beta \log(M_{W_{\hat{z}}^{-1}, \delta})}{M_{W_{\hat{z}}^{-1}, \delta}}} + \sqrt{\frac{\beta + \log(2/\delta)/8}{M_{W_{\hat{z}}^{-1}, \delta}}} \tag{13}$$

$$\beta = \frac{pr}{8}\max\{1, \log(6\sqrt{8}L)\} ,$$

where $L$ is a contestant that depends on the loss function which measures the data fidelity. Given all parameters, Eq. (12) can be solved for $M_{W_{\hat{z}}^{-1}, \delta}$.

So, according to Theorem 1 if we use at least $M_{W_{\hat{z}}^{-1}, \delta}$ sample images to learn the coupled dictionaries, we can achieve the required error rate $\epsilon_{max} = W_{\hat{z}}^{-1}(\zeta)$. Now we need to determine what is the probability of recovering the true label in ZSL regime or $P(\mathcal{S}_\epsilon | \mathcal{D}_\epsilon)$. Note that the core step for predicting the semantic attributes in our scheme is to compute the joint sparse representation for an unseen image. Also note that Eq. 1 can be interpreted as result of a maximum a posteriori (MAP) inference. This means that from a probabilistic perspective, $\boldsymbol{\alpha}$'s are drawn from a Laplacian distribution and the dictionary $D$ is a Gaussian matrix with elements drawn i.i.d: $d_{ij} \sim \mathcal{N}(\mathbf{0}, \epsilon)$. This means that given a drawn dataset, we learn MAP estimate of the Gaussian matrix $[D_x, D_z]^\top$ and then use the Gaussian matrix $D_z$ to estimate $\mathbf{a}$ in ZSL regime. To compute the probability of recovering $\mathbf{a}$ in this setting, we rely on the following theorem:

**Theorem 2 (Theorem 3.1 in [38])**: Consider the linear system $\mathbf{x}_i = D_x \mathbf{a}_i + \boldsymbol{n}_i$ with a sparse solution, i.e., $\|\mathbf{a}_i\|_0 = k$, where $D_x \in \mathbb{R}^{p \times r}$ is a random Gaussian matrix and $\|\boldsymbol{n}_i\|_2 \leq \epsilon$). Then the unique solution of this system can be recovered by solving eq. (5) with probability of $(1 - e^{p\tilde{\xi}})$ as far as $k \leq c'p \log(\frac{r}{p})$, where $c'$ and $\tilde{\xi}$ are two constant parameters.

Theorem 2 suggests that we can use eq. (5) to recover the sparse representation and subsequently unseen attributes with high probability $P(\mathcal{S}_\epsilon | \mathcal{D}_\epsilon) = (1 - e^{p\tilde{\xi}})$. This theorem also suggests that for our approach to work, existence of a good sparsifying dictionary as well as rich attribute data is essential. Therefore, given desired error parameters $1 - \zeta$ and $1 - \delta$ for the two stages of ZSL algorithm and an error parameter $\epsilon$, the probability event of predicting an unseen class label correctly can be computed as:

$$P_t = (1 - \delta)(1 - e^{p\tilde{\xi}})(1 - \zeta) \ , \tag{14}$$

which concludes our proof on PAC-learnability of the algorithm.

*4.1. Computational Complexity*

A major criterion for choosing a particular approach from a set of competitors for solving a specific problem is computational complexity. Since in ZSL, model training is supposedly performed once, training computational complexity is a secondary criterion. As a result, it is more important to analyze the testing computational complexity. As explained, during testing, our baseline approach first solves for the joint sparse vector by solving the LASSO problem (5) and then the label is predicted from the recovered

attribute using nearest neighbor. Given that $D_x \in \mathbb{R}^{p \times r}$ and $r \geq p$, the computational complexity for solving the LASSO problem would be $O(r^3)$ (5) [13]. Upon sparse vector recovery, we can estimate the attribute by the matrix multiplication $\hat{z}_i = D_z \alpha_i$ which has the computational complexity $O(qr)$. Finally. the nearest neighbor computational complexity would be $O(qM)$ and hence the overall computational complexity for recovering the label for an unseen class would be $O(r^3 + qr + qM)$.

To judge the efficiency of our algorithm, we can compare it against a heuristic deep learning method. For simplicity, we relax the problem and assume that due to existence of enough labeled data from the seen classes, we can train a deep networks $\mathcal{G}(\cdot) : \mathbb{R}^p \to \mathbb{R}^q$ that maps a given $x$ to the corresponding attribute $z$ at its output. As a result, the computational complexity of using a deep net for ZSL during testing would be computational complexity of forward pass. If we denote the number of nodes in the network by $m_1, \ldots, m_n$, then computational complexity of forward pass would be $O(pqm_1 \ldots m_n)$ and assuming that, nearest neighbor is used at the output, then overall complexity would be $O(pqm_1 \ldots m_n + qM)$. It is quite clear that this computational complexity can grow fast when the network is deep which is typical of the current practical network with several layers with many nodes. This demonstrates that using coupled dictionary learning is more practical if computational complexity is a major concern.

## 5. Experiments

We carried out experiments on four benchmark ZSL datasets and empirically evaluated the resulting performances against existing ZSL algorithms.

**Datasets:** We conducted our experiments on four benchmark datasets namely: the Animals with Attributes (AwA1) [27], (AwA2), [69] the SUN attribute [43], and the Caltech-UCSD-Birds 200-2011 (CUB) bird [65] datasets.

The AwA1 dataset is a coarse-grained dataset containing 30475 images of 50 types of animals with 85 corresponding attributes for these classes. Semantic attributes for this dataset are obtained via human annotations. The images for the AWA1 dataset are not publicly available; therefore we use the publicly available features of dimension 4096, extracted from a VGG19 convolutional neural network, which was pretrained on the ImageNet dataset. Following the conventional usage of this dataset, 40 classes are used as the source classes to learn the model and the remaining

10 classes are used as the target (unseen) classes to test the performance of zero-shot classification. The major disadvantage of AwA1 dataset is that only extracted features are available for this dataset. The AwA2 dataset is developed to compensate for this weakness by providing the original images. The AWA2 dataset has a similar structure with the same 50 animal classes and 85 attributes, but with 37322 images. Because the original images are available, one can use alternative deepnet structures for feature extraction.

The SUN dataset is a fine-grained dataset and contains 717 classes of different scene categories with 20 images per category (14340 images total). Each image is annotated with 102 attributes that describe the corresponding scene. There are two general approaches to split this dataset into training and testing sets. Following [75], 707 classes are used to learn the dictionaries and the remaining 10 classes are used for testing. Following the second approach [27], we used 645 classes to learn the dictionaries and 72 classes are used for testing. Both splits are informative because together help analyzing the effect of the training set size on the performance.

The CUB200 dataset is a fine-grained dataset containing 200 classes of different species of birds with 11788 images with 312 attributes and boundary segmentation for each image. The attributes are obtained via human annotation. The dataset is divided into four almost equal folds, where three folds are used to learn the model and the fourth fold is used for testing.

For each dataset, except for AwA1 (where images are not available), we use features extracted by the final layer prior to classification of VGG19 [58], Inception [64], ResNet [18], and DenseNet [19]. For AwA1, AwA2, and CUB200-2011 the networks were trained on ImageNet [23]. For SUN, they were trained on Places [77].

We used flat hit@K classification accuracy, to measure the performance. This means that a test image is said to be classified correctly if it is classified among the top $K$ predicted labels. We report hit@1 rate to measure ZSL image classification performance and hit@3 and hit@5 for image retrieval performance.

**Results:** Each experiment is performed ten times and the mean is reported in Table 1. For the sake of an ablative study, we have included results for the AAg formulation using nearest neighbor, the AAw using nearest neighbor, and AAw using the transductive approach, denoted as transductive attribute-aware (TAAw) formulation. In other words, we can

| Feature \ Method | AAg (5) | AAw (6) | TAAw | AAg(hit@3) | AAw(hit@3) | TAAw(hit@3) | AAg(hit@5) | AAw(hit@5) | TAAw(hit@5) |
|---|---|---|---|---|---|---|---|---|---|
| AwA1 Dataset | | | | | | | | | |
| VGG19 | 77.30 | 79.48 | 89.35 | 96.05 | 96.54 | 97.52 | 98.56 | 98.67 | 98.51 |
| AwA2 Dataset | | | | | | | | | |
| VGG19 | 41.68 | 45.54 | 69.93 | 74.80 | 78.62 | 88.77 | 91.36 | 92.56 | 93.34 |
| Inception | 39.05 | 47.61 | 71.72 | 82.15 | 84.58 | 97.12 | 90.64 | 92.08 | 97.66 |
| ResNet50 | 43.55 | 47.81 | 81.99 | 80.09 | 83.32 | 94.76 | 92.92 | 93.91 | 95.37 |
| DenseNet161 | 40.72 | 43.47 | 78.14 | 77.08 | 80.17 | 98.09 | 94.63 | 95.89 | 98.44 |
| CUB | | | | | | | | | |
| VGG19 | 35.29 | 40.62 | 48.41 | 60.52 | 67.67 | 67.75 | 72.14 | 74.44 | 78.57 |
| Inception | 35.32 | 40.31 | 49.65 | 51.17 | 55.52 | 63.78 | 67.05 | 71.37 | 75.33 |
| ResNet50 | 24.81 | 29.79 | 44.19 | 48.22 | 56.52 | 67.03 | 58.93 | 66.69 | 75.60 |
| DenseNet161 | 28.91 | 33.55 | 51.03 | 51.51 | 59.57 | 73.13 | 61.06 | 68.25 | 79.63 |
| SUN Dataset (645/72 Split) | | | | | | | | | |
| VGG19 | 42.36 | 45.69 | 48.40 | 57.50 | 61.48 | 67.50 | 71.94 | 75.76 | 82.01 |
| Inception | 55.66 | 56.02 | 57.03 | 80.10 | 80.65 | 81.18 | 87.06 | 87.22 | 87.72 |
| ResNet50 | 44.60 | 45.49 | 53.09 | 70.13 | 70.76 | 75.06 | 79.53 | 79.81 | 81.79 |
| DenseNet161 | 42.76 | 43.48 | 51.22 | 68.24 | 68.76 | 74.65 | 77.71 | 78.40 | 81.35 |
| SUN Dataset (707/10 Split) | | | | | | | | | |
| VGG19 | 85.50 | 89.25 | 91.00 | 93.95 | 96.50 | 98.05 | 97.15 | 98.05 | 98.50 |
| Inception | 83.30 | 83.80 | 84.95 | 96.80 | 96.80 | 96.95 | 98.85 | 98.85 | 98.80 |
| ResNet50 | 76.10 | 83.60 | 84.60 | 93.20 | 97.35 | 97.10 | 96.70 | 96.95 | 99.05 |
| DenseNet161 | 74.65 | 75.10 | 86.65 | 93.05 | 93.05 | 97.30 | 96.60 | 96.70 | 99.25 |

Table 1: Zero-shot classification and image retrieval results for the proposed algorithm.

study the effect of the absence of each of the entropy regularization and the transductive prediction on the performance to demonstrate their positive effect. As can be seen, while the AAw formulation significantly improves the AAg formulation, adding the transductive approach (i.e., label propagation on predicted attributes) to the AAw formulation further boosts the classification accuracy, as also shown in Figure 2. These results also support the logic behind our approach that: 1) the attribute aware optimization always boosts the performance, and 2) the transductive prediction of labels leads to a secondary boost in performance of our method. Finally, for completeness hit@3 and hit@5 rates measure image retrieval performance.

Figure 2 demonstrates the 2D t-SNE embedding for predicted attributes and actual class attributes of the AWA1 dataset. It can be seen that our algorithm can cluster the dataset in the attribute space. The actual attributes are depicted by the colored circles with black edges. The first column of Figure 2 demonstrates the attribute prediction for AAg and AAw formulations. It can be seen that the entropy regularization in AAw formulation improves the clustering quality, decreases data overlap, and reduces the domain-shift problem. The nearest neighbor label assignment is shown in the second
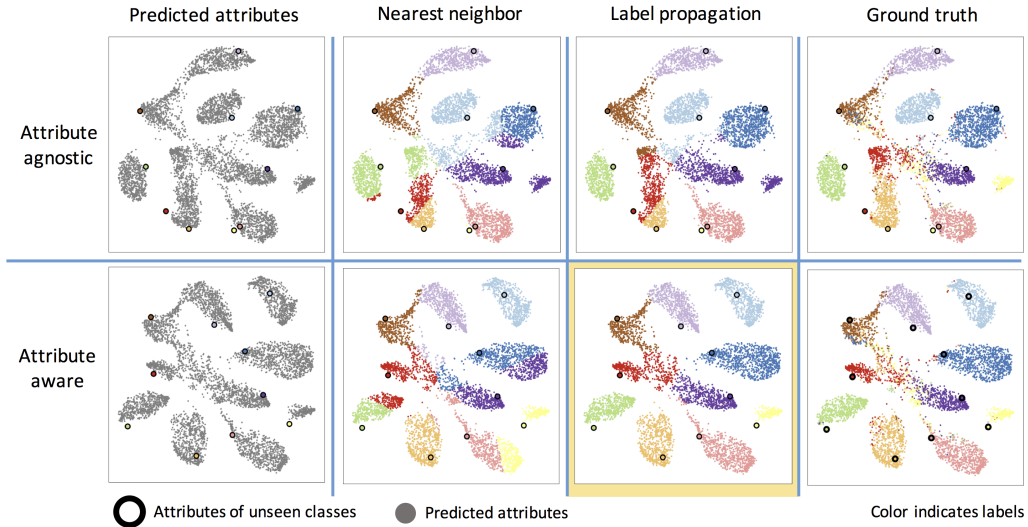
Figure 2: Attributes predicted from the input visual features for the unseen classes of images for AWA1 dataset using our attribute-agnostic and attribute-aware formulations respectively in top and bottom rows. The nearest neighbor and label propagation assignment of the labels together with the ground truth labels are visualized. It can be seen that the attribute-aware formulation together with the label propagation scheme overcomes the hubness and domain-shift problems, enclosed in yellow margins. Best seen in color.

column, which demonstrates the domain-shift and hubness problems with NN label assignment in the attribute space. The third column of Figure 2 shows the transductive approach in which a label propagation is performed on the graph of the predicted attributes. Note that the label propagation addresses the domain-shift and hubness problem and when used with the AAw formulation improves the accuracy significantly.

Performance comparison results using VGG19 and GoogleNet extracted features are summarized in Table 2 and and Table 3. Note that in Table 3 we used AwA2 dataset in order to extract the ResNet and GoogleNet features. As pointed out by Xian et al. [69] the variety of used image features (e.g., various DNNs and various combinations of these features) as well as the variation of used attributes (e.g., word2vec, human annotation), and different data splits make direct comparison with the ZSL methods in the literature very challenging. In Table 2 and Table 3 we provide comparison of our JDZSL performance to the recent methods in the literature. All compared methods use the same visual features and the same attributes (i.e.,

| Method | | SUN | CUB | AwA1 |
|---|---|---|---|---|
| Romera-Paredes and Torr [46] | | 82.10 | - | 75.32 |
| Zhang and Saligrama [75][†] | | 82.5 | 30.41 | 76.33 |
| Zhang and Saligrama [76][†] | | 83.83 | 42.11 | 80.46 |
| Bucher, Herbin, and Jurie [4][†] | | 84.41 | 43.29 | 77.32 |
| Xu et. al. [71][†] | | 84.5 | 53.6 | 83.5 |
| Long et. al. [33] [†] | | 80.5 | - | 82.12 |
| Lu et. al. [34] [†] | | 84.67 | 44.67 | 65.89 |
| Ye and Guo [73][†] | | 85.40 | 57.14 | 87.22 |
| Ding, Shao, and Fu [10][†] | | 86.0 | 45.2 | 82.8 |
| Wang and Chen [66][†] | | - | 42.7 | 79.8 |
| Kodirov, Xiang, and Gong [21][†] | | 91.0 | 61.4 | 84.7 |
| Ding et. al. [11][†] | | 88.2 | 48.50 | 84.74 |
| Ours | AAg (5) | 85.5 | 35.29 | 77.30 |
| Ours | AAw (6) | 89.3 | 40.62 | 79.48 |
| Ours | Transductive AAw (TAAw) | 91.00 | 48.41 | 89.35 |

Table 2: Zero-shot classification results for four benchmark datasets. All methods use VGG19 features trained on the ImageNet dataset and the original continuous (or binned) attributes provided by the datasets. Here, † indicates that the results are extracted directly from the corresponding paper, ‡ indicates that the results are reimplemented with VGG19 features, and − indicates that the results are not reported.

the continuous or binned) provided in the dataset to make the comparison fair. Table 2 and Table 3 conclude the following remarks:

1. Comparing Table 2 with Table 3 reveals that the visual features affect the ZSL performance significantly. This is a natural observation as ZSL depends on how discriminative the features are across different classes.

2. We observe that our method achieves state-of-the-art or close to the state-of-the-art performance for both zero-shot scene and object recognition tasks. Quite importantly, while some of the other methods perform better on a specific dataset, our algorithm leads to competitive performance on all the four benchmark datasets.

3. We observe that despite not using a deep neural network for modelling the cross-domain mapping function, we are able to achieve a competitive performance. This observation concludes that our method can potentially work better than some of the recent deep learning-based algorithms when instances from the seen classes are

| Method | SUN | CUB | AwA2 |
|---|---|---|---|
| Romera-Paredes and Torr [46][†] | 18.7 | 44.0 | 64.5 |
| Norouzi et. al. [39][†] | 51.9 | 36.2 | 63.3 |
| Mensink et. al. [35][†] | 47.9 | 40.8 | 61.8 |
| Akata et. al. [2][†] | 56.1 | 50.1 | 66.7 |
| Lampert et. al. [25][†] | 44.5 | 39.1 | 60.5 |
| Changpinyo et. al. [6] [†] | 62.7 | 54.5 | 72.9 |
| Bucher, Herbin, and Jurie [4][†] | - | 43.3 | 77.3 |
| Xian et. al. [68][†] | - | 45.5 | 71.9 |
| Bucher et. al. [5][†] | 56.4 | 60.1 | 55.3 |
| Zhang and Saligrama [75][†] | - | 30.4 | 76.3 |
| Long et al. [32][†] | - | 58.40 | 79.30 |
| Ours AAg (5) | 55.7 | 35.3 | 39.1 |
| Ours AAw (6) | 56.0 | 40.3 | 47.6 |
| Ours Transductive AAw (TAAw) | 57.0 | 49.7 | 71.7 |

Table 3: Zero-shot classification results for three benchmark datasets. All methods use Inception features trained on the ImageNet dataset and the original continuous (or binned) attributes provided by the datasets. Here $-$ indicates that the results are not reported.

limited.

4. Considering progressive improvement of our results when using AAw and TAAw solutions, we conclude that secondary mechanisms to address hubness [12] and domain-shift [15] are necessary to improve ZSL algorithms.

## 6. Conclusions and Discussions

In this paper, we developed a new zero-shot learning (ZSL) algorithm by recasting the ZSL problem as a coupled dictionary learning problem. In our formulation, the relationship between visual features and semantic attributes of data points are captured via a shared sparse representation vector in the two dictionary domains. We can use this formulation because representing signals that share some level of commonality in a union of subspaces is feasible. In the ZSL setting, since we primarily focus on classifying classes within one domain, they share a good level of commonalities. As a result of learning the two dictionaries, the shared sparse domain acts as a shared embedding space that is used to map images to their semantic descriptions. We established theoretical results for PAC-learnability of our

method. Our analysis supports that training these two dictionaries given a sufficient number of samples is feasible.

In addition to the baseline algorithm based on CDC, we also demonstrated that an entropy regularization scheme can help with the domain-shift. We face domain-shift on ZSL because the dictionaries are trained primarily based on the seen classes. As a result, a recovered sparse representation vector is biased towards the representations for seen classes. Entropy regularization is helpful to tackle domain-shift because it biases the recover sparse vector to be close to the sparse representations of attributes of an unseen class. As a result, domain-shift challenge is mitigated. Our results also demonstrate that a transductive approach towards assigning labels to the predicted attributes can boost the performance considerably and lead to state-of-the-art zero-shot classification by mitigating the hubness challenge. Hubness challenge occurs as a result of curse of dimensionality which makes Euclidean distance a non-perfect measure of similarity because it is only a point-wise measure of similarity. Our trandusctive approach considers structure of the data points to compute similarity to mitigate hubness in high dimensions. Our empirical ablative results demonstrate that both approaches are effective. We also compared our method with the state-of-the-art approaches in the literature and demonstrated its competitiveness on four primary ZSL benchmark datasets. An advantage of our method is that it preforms decently all the four datasets, despite the diversity between these datasets.

Our method is not an end-to-end training method and requires preprocessed suitable visual and semantic features. This limits the applicability of our method in situations that suitable feature extraction methods are not accessible. Note, however, in many common domains, pretrained models are able to generate discriminative features. The upside of our method is that dictionary learning is less data-greedy compared to the end-to-end methods based on deep learning both in terms of training and also during model execution. Hence, compared to the ZSL methods that use deep neural networks, dictionary learning is more effective when the number of annotated training data is small. An unexplored aspect for future work is extension to generalized ZSL setting. In a generalized ZSL setting, we encounter samples from both seen and unseen classes during testing. As a result, domain-shift will become a more challenging obstacle because the model is biased to recover sparse respresentations of the seen classes.

## References

[1] Akata, Z., Perronnin, F., Harchaoui, Z., and Schmid, C. (2016). Label-embedding for image classification. *IEEE Trans. on Pattern anal. and Machine Intel.*, 38(7):1425–1438.

[2] Akata, Z., Reed, S., Walter, D., Lee, H., and Schiele, B. (2015). Evaluation of output embeddings for fine-grained image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2927–2936. IEEE.

[3] Boyd, S., Parikh, N., and Chu, E. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc.

[4] Bucher, M., Herbin, S., and Jurie, F. (2016). Improving semantic embedding consistency by metric learning for zero-shot classiffication. In *European Conference on Computer Vision*, pages 730–746. Springer.

[5] Bucher, M., Herbin, S., and Jurie, F. (2017). Generating visual representations for zero-shot classification. In *Proc. of the IEEE Int. Conf. on Computer Vision*, pages 2666–2673.

[6] Changpinyo, S., Chao, W., Gong, B., and Sha, F. (2016). Synthesized classifiers for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5327–5336.

[7] Chen, Z. and Liu, B. (2018). Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207.

[8] Ciaccia, P. and Patella, M. (2000). Pac nearest neighbor queries: Approximate and controlled search in high-dimensional and metric spaces. In *Data Engineering, 2000. Proceedings. 16th International Conference on*, pages 244–255. IEEE.

[9] Das, D. and Lee, C. (2019). Zero-shot image recognition using relational matching, adaptation and calibration. *IEEE International Joint Conference on Neural Networks*.

[10] Ding, Z., S., M., and Fu, Y. (2017). Low-rank embedded ensemble semantic dictionary for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2050–2058.

[11] Ding, Z., Shao, M., and Fu, Y. (2018). Generative zero-shot learning via low-rank embedded semantic dictionary. *IEEE transactions on pattern analysis and machine intelligence*.

[12] Dinu, G., Lazaridou, A., and Baroni, M. (2014). Improving zero-shot learning by mitigating the hubness problem. *arXiv preprint arXiv:1412.6568*.

[13] Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *The Annals of statistics*, 32(2):407–499.

[14] Frome, A., Corrado, G. S., Shlens, J., Bengio, S., Dean, J., Mikolov, T., et al. (2013). Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129.

[15] Fu, Y., Hospedales, T. M., Xiang, T., and Gong, S. (2015). Transductive multi-view zero-shot learning. *IEEE transactions on pattern analysis and machine intelligence*, 37(11):2332–2345.

[16] Gribonval, R., Jenatton, R., Bach, F., Kleinsteuber, M., and Seibert, M. (2015). Sample complexity of dictionary learning and other matrix factorizations. *IEEE Transactions on Information Theory*, 61(6):3469–3486.

[17] Guo, L. and Garland, M. (2006). The use of entropy minimization for the solution of blind source separation problems in image analysis. *Pattern Recognition*, 39(6):1066–1073.

[18] He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Int. Conf. on Computer Vision*, pages 770–778.

[19] Huang, G., Liu, Z., Weinberger, K. Q., and van der Maaten, L. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4700–4708.

[20] Isele, D., Rostami, M., and Eaton, E. (2016). Using task features for zero-shot knowledge transfer in lifelong learning. In *Proc. of International Joint Conference on Artificial Intelligence*, pages 1620–1626.

[21] Kodirov, E., Xiang, T., and Gong, S. (2017). Semantic autoencoder for zero-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3174–3183.

[22] Kolouri, S., Rostami, M., Owechko, Y., and Kim, K. (2018). Joint dictionaries for zero-shot learning. In *AAAI Conf. on AI*, pages 3431–3439.

[23] Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.

[24] Kumar, N., Berg, A. C., Belhumeur, P. N., and Nayar, S. K. (2009). Attribute and simile classifiers for face verification. In *2009 IEEE 12th international conference on computer vision*, pages 365–372. IEEE.

[25] Lampert, C. H., Nickisch, H., and Harmeling, S. (2009). Learning to detect unseen object classes by between-class attribute transfer. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 951–958.

[26] Lampert, C. H., Nickisch, H., and Harmeling, S. (2013). Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465.

[27] Lampert, C. H., Nickisch, H., and Harmeling, S. (2014). Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):453–465.

[28] LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436–444.

[29] Li, J., Jing, M., Lu, K., Ding, Z., Zhu, L., and Huang, Z. (2019). Leveraging the invariant side of generative zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7402–7411.

[30] Liu, S., Long, M., Wang, J., and Jordan, M. I. (2018). Generalized zero-shot learning with deep calibration network. In *Advances in Neural Information Processing Systems*, pages 2005–2015.

[31] Lomonaco, V., Pellegrini, L., Cossu, A., Carta, A., Graffieti, G., Hayes, T. L., De Lange, M., Masana, M., Pomponi, J., van de Ven, G. M., et al. (2021). Avalanche: an end-to-end library for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3600–3610.

[32] Long, T., Xu, X., Shen, F., Liu, L., Xie, N., and Yang, Y. (2018a). Zero-shot learning via discriminative representation extraction. *Pattern Recognition Letters*, 109:27–34.

[33] Long, Y., Liu, L., Shen, F., Shao, L., and Li, X. (2018b). Zero-shot learning using synthesised unseen visual data with diffusion regularisation. *IEEE Trans. on Pattern Anal. and Machine Intel.*, 40(10):2498–2512.

[34] Lu, J., Li, J., Yan, Z., Mei, F., and Zhang, C. (2018). Attribute-based synthetic network (abs-net): Learning more from pseudo feature representations. *Pattern Recognition*, 80:129–142.

[35] Mensink, T., Gavves, E., and Snoek, C. G. M. (2014). Costa: Co-occurrence statistics for zero-shot classification. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 2441–2448.

[36] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.

[37] Morgenstern, Y., Rostami, M., and Purves, D. (2014). Properties of artificial networks evolved to contend with natural spectra. *Proceedings of the National Academy of Sciences*, 111(Supplement 3):10868–10872.

[38] Negahban, S., Yu, B., Wainwright, M., and Ravikumar, P. (2009). A unified framework for high-dimensional analysis of $m$-estimators with decomposable regularizers. In *Advances in neural information processing systems*, pages 1348–1356.

[39] Norouzi, M., Mikolov, T., Bengio, S., Singer, Y., Shlens, J., Frome, A., Corrado, G. S., and Dean, J. (2014). Zero-shot learning by convex combination of semantic embeddings. *International Conference on Learning Representations*.

[40] Palatucci, M., Pomerleau, D., Hinton, G. E., and Mitchell, T. M. (2009). Zero-shot learning with semantic output codes. In *Advances in neural information processing systems*, pages 1410–1418.

[41] Pan, S., Tsang, I. W., Kwok, J. T., and Yang, Q. (2011). Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 22(2):199–210.

[42] Parikh, N., Boyd, S., et al. (2014). Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239.

[43] Patterson, G., Xu, C., Su, H., and Hays, J. (2014). The sun attribute database: Beyond categories for deeper scene understanding. *International Journal of Computer Vision*, 108(1-2):59–81.

[44] Pinheiro, P. O. (2018). Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8004–8013.

[45] Rehman, A., Rostami, M., Wang, Z., Brunet, D., and Vrscay, E. R. (2012). Ssim-inspired image restoration using sparse representation. *EURASIP Journal on Advances in Signal Processing*, 2012(1):1–12.

[46] Romera-Paredes, B. and Torr, P. (2015). An embarrassingly simple approach to zero-shot learning. In *Int. Conf. on Machine Learning*, pages 2152–2161.

[47] Rostami, M. (2021a). Lifelong domain adaptation via consolidated internal distribution. In *Advances in neural information processing systems*.

[48] Rostami, M. (2021b). *Transfer Learning Through Embedding Spaces*. CRC Press.

[49] Rostami, M. and Galstyan, A. (2020). Sequential unsupervised domain adaptation through prototypical distributions. *arXiv e-prints*, pages arXiv–2007.

[50] Rostami, M., Huber, D., and Lu, T.-C. (2018a). A crowdsourcing triage algorithm for geopolitical event forecasting. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pages 377–381. ACM.

[51] Rostami, M., Isele, D., and Eaton, E. (2020a). Using task descriptions in lifelong machine learning for improved performance and zero-shot transfer. *Journal of Artificial Intelligence Research*, 67:673–704.

[52] Rostami, M., Kolouri, S., Eaton, E., and Kim, K. (2019a). Sar image classification using few-shot cross-domain transfer learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0.

[53] Rostami, M., Kolouri, S., Kim, K., and Eaton, E. (2018b). Multi-agent distributed lifelong learning for collective knowledge acquisition. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 712–720.

[54] Rostami, M., Kolouri, S., Pilly, P., and McClelland, J. (2020b). Generative continual concept learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5545–5552.

[55] Rostami, M., Kolouri, S., and Pilly, P. K. (2019b). Complementary learning for overcoming catastrophic forgetting using experience replay. In *Proceedings of the IJCAI Conference*.

[56] Rostami, M., Spinoulas, L., Hussein, M., Mathai, J., and Abd-Almageed, W. (2021). Detection and continual learning of novel face presentation attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14851–14860.

[57] Saito, K., Watanabe, K., Ushiku, Y., and Harada, T. (2018). Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3723–3732.

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

[59] Snell, J., Swersky, K., and Zemel, R. (2017). Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems*, 30:4077–4087.

[60] Socher, R., Ganjoo, M., Manning, C. D., and Ng, A. (2013). Zero-shot learning through cross-modal transfer. In *Advances in neural information processing systems*, pages 935–943.

[61] Stan, S. and Rostami, M. (2021). Unsupervised model adaptation for continual semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2593–2601.

[62] Sun, Q., Liu, Y., Chua, T.-S., and Schiele, B. (2019). Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 403–412.

[63] Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P. H., and Hospedales, T. M. (2018). Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1199–1208.

[64] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *IEEE Conf. on Computer Vision and Pattern Recog.*, pages 1–9.

[65] Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. (2011). The caltech-ucsd birds-200-2011 dataset. Technical report, CalTech.

[66] Wang, Q. and Chen, K. (2017). Zero-shot visual recognition via bidirectional latent embedding. *Int. Journal of Comp. Vision*, 124(3):356–383.

[67] Wu, Z., Fu, Y., Jiang, Y.-G., and Sigal, L. (2016). Harnessing object and scene semantics for large-scale video understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3112–3121.

[68] Xian, Y., Akata, Z., Sharma, G., Nguyen, Q., Hein, M., and Schiele, B. (2016). Latent embeddings for zero-shot classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 69–77.

[69] Xian, Y., Lampert, C. H., Schiele, B., and Akata, Z. (2018a). Zero-shot learning-a comprehensive evaluation of the good, the bad and the ugly. *IEEE Trans. on Pattern anal. and Machine Intel.*

[70] Xian, Y., Lorenz, T., Schiele, B., and Akata, Z. (2018b). Feature generating networks for zero-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5542–5551.

[71] Xu, X., Shen, F., Yang, Y., Zhang, D., Shen, H. T., and Song, J. (2017). Matrix tri-factorization with manifold regularizations for zero-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3798–3807.

[72] Yang, J., Wright, J., Huang, T. S., and Ma, Y. (2010). Image super-resolution via sparse representation. *IEEE Trans. on Image Proc.*, 19(11):2861–2873.

[73] Ye, M. and Guo, Y. (2017). Zero-shot classification with discriminative semantic representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7140–7148.

[74] Zenke, F., Poole, B., and Ganguli, S. (2017). Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR.

[75] Zhang, Z. and Saligrama, V. (2015). Zero-shot learning via semantic similarity embedding. In *Int. Conf. on Computer Vision*, pages 4166–4174.

[76] Zhang, Z. and Saligrama, V. (2016). Zero-shot learning via joint latent similarity embedding. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 6034–6042.

[77] Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. (2018). Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464.

[78] Zhou, D., Bousquet, O., Lal, T. N., Weston, J., and Schölkopf, B. (2003). Learning with local and global consistency. In *Advances in neural information processing systems*, volume 16, pages 321–328.