A Generative Framework for Zero Shot Learning with Adversarial Domain Adaptation

Varun Khare, Divyat Mahajan, Homanga Bharadhwaj, Vinay Verma, Piyush Rai
Department of Computer Science and Engineering
IIT Kanpur, India
{varun, divyatm, homangab, vkverma, piyush }@iitk.ac.in

Abstract

In this paper, we present a domain adaptation based generative framework for zero shot learning. We explicitly target the problem of domain shift between the seen and unseen class distribution in Zero-Shot Learning (ZSL) and seek to minimize it by developing a generative model and training it via adversarial domain adaptation. Our approach is based on end-to-end learning of the class distributions of seen classes and unseen classes. To enable the model to learn the class distributions of unseen classes, we parameterize these class distributions in terms of the class attribute information (which is available for both seen and unseen classes). This provides a very simple way to learn the class distribution of any unseen class, given only its class attribute information, and no labeled training data. Training this model with adversarial domain adaptation provides robustness against the distribution mismatch between the data from seen and unseen classes. Through a comprehensive set of experiments, we show that our model yields superior accuracies as compared to various state-of-the-art zero shot learning models, on a variety of benchmark datasets.

1 Introduction

In the conventional image classification tasks, examples from all classes are available during training of the model. This assumption rarely holds in real-world problems. There exist a plethora of real-world concepts that do not have the corresponding ubiquity of representative images. In addition, it is common knowledge that humans do not require prior visual evidence of a category to recognize an example from that category. Given that a child sees a picture of a horse and reads a description about zebra's appearance, he/she would more likely than not be able to easily recognize a zebra when an image is shown. The Zero Shot Learning (ZSL) problem [Socher *et al.*, 2013; Xian *et al.*, 2018] in machine learning is motivated by similar considerations and seeks to exploit the existence of a labeled training set of 'seen' classes and the knowledge about how each 'unseen' class relates semantically to the seen classes.

The success of ZSL lies in learning an effective semantic representation (e.g. attributes / textual features) for successful transfer of knowledge from the seen to the unseen classes. While, in Sec. 3, we provide a detailed overview of the prior work on ZSL, generative ZSL methods [Xian *et al.*, 2015; Wang *et al.*, 2017b; Verma and Rai, 2017; Verma *et al.*, 2018] in particular, have become popular in recent works, by virtue of their ability to *generate* labeled examples for the unseen classes. However, a key requirement in such methods is the reliable estimation of the class distribution of seen and unseen classes. Moreover, a generative model alone is not sufficient as there may be a domain shift between the original data and the synthesized data.

The presence of acute domain-shift between the seen and unseen classes hinders the performance of ZSL models [Kodirov *et al.*, 2015a]. We note that by explicitly enforcing domain adaptation to tackle problem settings where the train and test distributions are significantly far apart, the model's performance can be greatly improved. While domain adaptation applies to machine learning problems in general, it can be of particular significance in ZSL, since the ground-truth supervised data is very limited. An earlier approach [Kodirov *et al.*, 2015a] used the idea of *joint* sparse coding for minimizing domain shift between the seen and unseen class data, however, since then there have been developments in adversarial domain adaptation that enable robust detection and resolution of domain shift [Tzeng *et al.*, 2017; Hoffman *et al.*, 2017]. Adversarial learning and adaptation methods have also found applicability in a wide range of fields from robotics navigation [Bharadhwaj *et al.*, 2018b] to recommender systems [Bharadhwaj *et al.*, 2018a; Wang *et al.*, 2017a]. However, several adversarial adaptation techniques like ADDA[Tzeng *et al.*, 2017] require explicit source-target pairs of data points. Such a luxury is not present in Zero-Shot transductive setting where the test data is unlabeled. Similarly, unsupervised domain matching methods like CycleGAN[Zhu *et al.*,] use cyclic consistency to find the data point most similar to the source sample and then minimize the gap between these two. Such a formulation destroys the inherent clusters and often distributes a source cluster into several clusters in the target domain. This has serious repercussions over the classification accuracy in the transformed domain.

In order to develop an adversarial domain adaptation

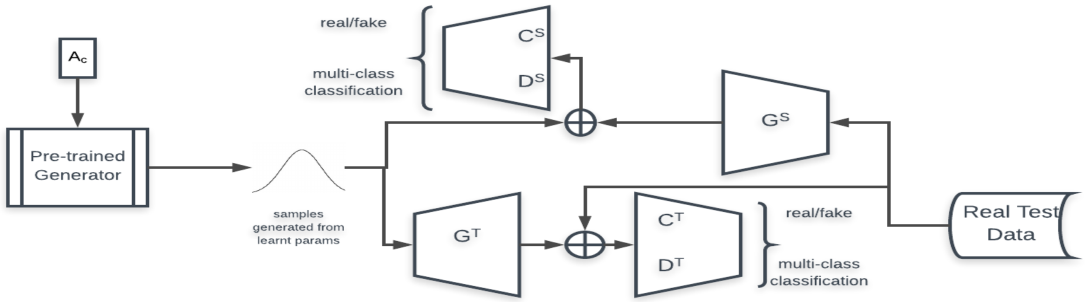

Figure 1: The overall architecture of the proposed approach. All the notations are consistent with that described in Section 2. $\mathbf{A_c}$ denotes the class attributes for all classes i.e. $\{\mathbf{a_c}\}_{c=1}^{S+U}$.

method for ZSL, we require a generative framework that models the data distribution and can act as the generator in the adversarial optimization scheme. A simple, yet principled, way to construct generative models for ZSL is to learn the class distributions for the seen and unseen classes [Verma and Rai, 2017; Wang *et al.*, 2017b]. While this is straightforward for seen classes (for which we have access to labeled data), it can't be done for the unseen classes. In a recent work, [Verma and Rai, 2017] used exponential family to model the distributions of the class conditionals in terms of learnable parameters (which in turn are parameterized on the class attribute vectors). This is an effective model; however, their approach does not extend to non-exponential family distributions. Moreover, they used offline learning techniques to learn the parameters of seen classes, and rely on linear and non-linear kernel-based regression to estimate the class parameters, given the class attributes. The model also faces numerical stability issues and requires explicit thresholding and huge regularization parameters. This makes it necessary to have very precise hyperparameter tuning. Our model, on the other hand, exploits the advantages of neural nets and end-to-end training to provide stability during the learning phase and remains less susceptible to hyperparameter variations.

Motivated by these desiderata, in this work, we develop an Adversarial Domain Adaptation framework for ZSL that leverages a ZSL model to improve upon the classification. Our model is able to transform the synthesized samples into the test domain while maintaining the data clusters associations. We first learn a generative model for the class conditional distribution of the seen and unseen classes by utilizing labeled data from the seen classes and the class attribute vectors of both the seen and the unseen classes. Then, by domain adaptation, we explicitly bring closer the generated distribution and the true distribution of the class conditionals. To the best of our knowledge, there is no adversarial framework for semi-supervised domain matching where explicit pairs of data points are not given but an external agent associates noisy labels to the samples. We employ a scheme of cyclically consistent adversarial domain adaptation [Hoffman *et al.*, 2017] to minimize domain shift without assuming any particular parametric form of the source and target distributions. The generative framework can in principle assume any form, some of the popular ones being GAN and VAE based models. However, they lack explainability, unlike probabilistic models. To this end, we consider a mixture of class conditional distributions as the generating framework.

2 The Proposed Approach

Our approach consists of a pre-training phase followed by adversarial domain adaptation (ADA). We first describe the generative model and then elaborate on the ADA setting. A detailed illustration of our method is shown in Figure 1.

2.1 The Generative Model

We model the data from class c by a class conditional probability $p(\mathbf{x}|c, \mathbf{\Theta})$ where $\mathbf{\Theta}$ denotes the global parameters of the model. We do not have any restriction on the the type of distribution chosen. Let us denote the total number of classes whose examples are seen during training by S, and the classes, none of whose examples are seen during training by U. Then, for an observation $\mathbf{x}$ from either a seen or unseen class c, where $c \in [1, S+U]$, we have $y_n = c$, and, assume the input to be generated as $x_n \sim p(\mathbf{x}|c, \mathbf{\Theta})$

Under this framework, given test example $\mathbf{x}_+$, the predicted class $\hat{y}_+$ can be given by computing the most-probable class as follows $\hat{y}_+ = \text{argmax}_c\, p(c|\mathbf{x}_+, \mathbf{\Theta})$ and using Baye's Rule we have,

$$p(c|\mathbf{x}_+, \mathbf{\Theta}) = \frac{p(\mathbf{x}_+|c, \mathbf{\Theta})p(c|\mathbf{\Theta})}{\sum_{c\in[1,S]} p(\mathbf{x}_+|c, \mathbf{\Theta})p(c|\mathbf{\Theta})} \tag{1}$$

Thus

$$\hat{y}_+ = \underset{c}{\text{argmax}}\; p(\mathbf{x}_+|c, \mathbf{\Theta})p(c|\mathbf{\Theta}) \tag{2}$$

From our preliminary empirical evaluation, we observed that the estimation of class probabilities didn't have a significant impact on the model accuracy and hence for simplicity, we chose to treat them as equal for all unknown classes. Thus the prediction rule is simply

$$\hat{y}_+ = \underset{\mathbf{c}}{\text{argmax}}\; p(\mathbf{x}_+|c, \mathbf{\Theta}) \tag{3}$$

If labeled training data for all the classes are available, then standard inference techniques like Maximum Likelihood Estimation (MLE), Maximum-a-Posteriori (MAP) Estimation, or fully Bayesian Inference can be used to determine the class conditional distributions. However, since the unseen classes do not have labeled training examples, we need a way to "extrapolate" the seen class distribution parameters to unseen class distribution parameters. This will be done via the class attribute vectors as we describe next.

First, asuming $\mathbf{X}$ to denote the inputs and $\mathbf{C}$ ($c_k \in S \cup U$) to denote the associated output class labels, we take a generative approach and maximize their joint distribution $\mathbb{P}[\mathbf{X}^{\mathbf{S}\cup\mathbf{U}}, \mathbf{C}^{\mathbf{S}\cup\mathbf{U}}|\mathbf{\Theta}]$. Assuming i.i.d. observations, we have

$$\mathbb{P}[\mathbf{X}^{\mathbf{S}\cup\mathbf{U}}, \mathbf{C}^{\mathbf{S}\cup\mathbf{U}}|\mathbf{\Theta}] = \mathbb{P}[\mathbf{X}^{S}, \mathbf{C}^{\mathbf{S}}|\mathbf{\Theta}]\mathbb{P}[\mathbf{X}^{\mathbf{U}}, \mathbf{C}^{\mathbf{U}}|\mathbf{\Theta}] \quad (4)$$

Since, $\mathbf{X}^{U}, \mathbf{C}^{\mathbf{U}}$ are unavailable for parameter estimation, we learn $\mathbf{\Theta}$ to maximize $\mathbb{P}[\mathbf{X}^{\mathbf{S}}, \mathbf{C}^{\mathbf{S}}|\mathbf{\Theta}]$

$$\mathbb{P}[\mathbf{X}^{\mathbf{S}}, \mathbf{C}^{\mathbf{S}}|\mathbf{\Theta}] = \prod_{\mathbf{x},c\sim S} p(\mathbf{x}, c|\mathbf{\Theta}) \quad (5)$$

$$\Rightarrow log(\mathbb{P}[\mathbf{X}^{\mathbf{S}}, \mathbf{C}^{\mathbf{S}}|\mathbf{\Theta}]) = \sum_{\mathbf{x},c\sim S} log(p(\mathbf{x}, c|\mathbf{\Theta}))$$

$$= \sum_{\mathbf{x},c\sim S} log(p(\mathbf{x}|c, \mathbf{\Theta})) + log(p(c|\mathbf{\Theta})) \quad (6)$$

Since we are not modelling the class probability distribution $p(c|\mathbf{\Theta})$, the objective becomes

$$\underset{\mathbf{\Theta}}{\operatorname{argmax}} \underset{\mathbf{x},c\sim S}{\mathbb{E}}[log(p(\mathbf{x}|c, \mathbf{\Theta}))] \quad (7)$$

Mapping Class Attributes to Class Parameters

Since each class is described in terms of attribute vectors $\mathbf{a}_c$, we condition the class distribution on their respective attribute vector $\mathbf{a}_c$. Let these class specific parameters be ζ_c which can be uniquely determined from the class attribute vector $\mathbf{a_c}$ and global parameters $\mathbf{\Theta}$ by a functional mapping f. This mapping for most purposes will be a complicated relationship and using a linear mapping (e.g., as done in [Verma and Rai, 2017]) here would severely affect the generation quality of the network. Hence, we model this function $f : \{\mathbf{a}_c\} \rightarrow \{\zeta_c\}$ using neural networks having parameters $\mathbf{\Theta}$ which bring with them extensive expressiveness and hierarchical relationships among attribute features. This leads to a stable training procedure and enables us to perform joint learning of $f_{\mathbf{\Theta}}$ and class parameters $\{\zeta_c\}$ unlike a two-stage procedure employed by the model of [Verma and Rai, 2017].

To elaborate further, we model the class parameters as

$$\zeta_{\mathbf{c}} = f_{\Theta}(\mathbf{a_c}) \quad (8)$$

For simplicity of exposition, we take $p(\mathbf{x}|c, \mathbf{\Theta})$ to be a Gaussian distribution with parameters $\zeta_c = \{\mu_c, \Sigma_c\}$ where $c \in S$. We model $\mu_{\mathbf{c}}$ and $\mathbf{\Sigma}_c^{-1}$ as non-linear functions of the attribute vector $\mathbf{a_c}$ with learnable parameters $\mathbf{\Theta} = \{\theta^{\mu}, \theta^{\mathbf{\Sigma}}\}$ in the following manner,

$$\mu_{\mathbf{c}} = f_{\theta^{\mu}}(\mathbf{a_c}),\ \mathbf{\Sigma_c}^{-1} = diag(f_{\theta^{\mathbf{\Sigma}}}(\mathbf{a_c})),\ \mathbf{x} \sim N(\mu_{\mathbf{c}}, \mathbf{\Sigma_c})$$

To ensure the condition of the covariance matrix (Σ_c) being a positive semi-definite matrix we model the inverse covariance to a diagonal matrix with positive diagonal entries. Thus $f_{\theta^{\Sigma}}$ outputs a vector in $\mathbb{R}^{d}_{>0}$ where d is the dimension of mean vector (equivalently the dimension of semantic space). The overall objective function becomes:

$$\underset{\theta^{\mu},\theta^{\mathbf{\Sigma}}}{\operatorname{argmax}} \underset{(\mathbf{x},c)\sim S}{\mathbb{E}} \left[log(\mathbf{\Sigma_c}^{-1}) - (\mathbf{x} - \mu_{\mathbf{c}})^{T} \mathbf{\Sigma_c}^{-1} (\mathbf{x} - \mu_{\mathbf{c}})\right] \quad (9)$$

Also we take $\mathbf{x_n}$ as the features extracted from dataset images by resnet-101[He *et al.*, 2016] pre-trained on the Imagenet dataset [Russakovsky *et al.*, 2015]. We again emphasize the fact that choosing Gaussian distribution is only for expositional purposes and one can also try other non-exponential family distributions as a part of inductive bias. The model does not restrict the choice of class conditional distribution. For the rest of the paper, $\{\mathbf{x}_n\}_c$ denotes the entire test data comprising of samples from all the classes. Similarly, $\{y_n\}_c$ denotes the samples generated from the generative model (as defined here) for all the classes.

2.2 Adversarial Domain Adaptation

The procedure described in the previous section only leverages the data from seen classes to predict the parameters of the class conditional distributions of all the classes. In our overall architecture, we denote the process of learning the ZSL model parameters $\mathbf{\Theta}$ as 'pre-training.' Based on the generative framework learned during pre-training, we can sample the class-conditional distribution for unseen classes to generate the unseen samples. We then minimize the domain gap between the generated distribution of the unseen classes and the actual distribution.

In this section, we denote the source domain as S and the target domain as T. Through adversarial adaptation, we aim to bring the target distribution of $\{y_n\}_c$ closer to the source distribution of $\{x_n\}_c$; hence we learn a function $G^{T}(\{y_n\}_c)$ that maps class conditionals from the generated distribution $\{y_n\}_c$ to the real test distribution$\{x_n\}_c$ for all unseen classes $\{c\}_{c=1}^{U}$. Hence, $G^{T} : S \rightarrow T$ is a mapping from source S to target T. Similarly, we define another function $G^{S} : T \rightarrow S$ that maps the class conditionals from the real test distribution to the same latent space as the class conditionals from the generated distribution. D^{T} and D^{S} are the corresponding discriminators.

Our design is inspired by CycleGAN [Zhu *et al.*,] and we make a number of modifications to its base architecture for supporting zero shot learning.

Label Augmentation

Inspired by conditional GAN[Reed *et al.*, 2016], we augment the input to the generators G^{T} and G^{S} with the resepctive class labels, which facilitates preservation of relationships between the synthesized data and their correct class labels. For G^{S} the input data (test data) is unlabeled, hence we use the predictions from our pre-trained ZSL model as the guiding labels. Note that these labels are noisy labels and both the generators and discriminators should be capable of handling data corruption during the training phase. This is yet again a problem with the conventional GAN architectures.

Classifiers

We handled label recovery by adding two classifiers (C^{T}, C^{S}) in parallel with the discriminators. The parameters of the classifiers are trained jointly with the corresponding discriminators. Recently, parallel to our work, [Thekumparampil *et al.*, 2018] gave theoretical support to the use of external classifier with conditional GAN architecture to counter noisy data labels. The classifiers provide an additional benefit of enforcing clustering for the generator. We provide the justification for clustering in the next section.

Optimization Function

Let the loss defined in CycleGAN which consists of cyclic consistency loss ($\mathcal{L}_{cyc}$) and the adversarial loss ($\mathcal{L}_{adv}$) for domains T and S be $\mathcal{L}$

$$\mathcal{L} = \mathcal{L}_{cyc} + \mathcal{L}_{adv}^{T} + \mathcal{L}_{adv}^{S} \tag{10}$$

For our case, L_1 norm worked the best for cyclic consistency loss, while Wasserstein loss [Arjovsky *et al.*, 2017] was found suitable for $\mathcal{L}_{adv}$. We add the classification loss ($\mathcal{L}_{clf}$) of the real data (not generated by G) to the discriminator loss during adversarial training. We ensured that the classification loss is not added for data transformed by generators in accordance with mismatch loss addition for only real images [Reed *et al.*, 2016]. We used the complete cross entropy loss which handles both correct and mismatched pairs of label-image. This enforces a stronger clustering than the mismatched label loss term in [Reed *et al.*, 2016]. With χ, ξ as tunable hyper parameters, the overall loss function then becomes

$$\mathcal{L} = \mathcal{L}_{adv}^{T} + \mathcal{L}_{adv}^{S} + \chi\mathcal{L}_{cyc} + \xi\mathcal{L}_{clf}^{T} + \xi\mathcal{L}_{clf}^{S} \tag{11}$$

3 Related Work

Due to its ability to overcome the drawbacks of conventional classification problems, ZSL has attained tremendous recent interest for a wide range of AI problems, including those in computer vision. Earlier works [Lampert *et al.*, 2014a; Lampert *et al.*, 2014b] on ZSL were based on directly or indirectly mapping the instances of specific examples to their class-attributes. The learned mapping was then used during inference; this mapping works by first projecting the unseen data to the class-attribute space and then using the nearest neighbor search to classify the unseen image. Another approach for ZSL focuses on learning the map of bi-linear compatibility between the visual space and the semantic space. ALE [Akata *et al.*, 2013], DEVISE [Frome *et al.*, 2013], SJE [Akata *et al.*, 2015], ESZSL [Romera and Torr, 2015], and SAE [Kodirov *et al.*, 2017] are based on the approach of measuring the bi-linear compatibility.

Generative models [Mishra *et al.*, 2017; Chen *et al.*, 2018; Xian *et al.*, 2015; Verma and Rai, 2017; Guo *et al.*, 2017; Bucher *et al.*, 2017; Wang *et al.*, 2017b] have shown promising results for both ZSL and GZSL setups. Another advantage of the generative approach is that by using synthesized samples, we can convert the ZSL problem to the conventional supervised learning problem that can handle the biases towards the seen classes. The [Verma and Rai, 2017] used a simple generative model based on the exponential family framework while [Guo *et al.*, 2017] synthesize the classifier. While recent generative approaches for the ZSL are deep generative models based on the VAE [Kingma and Welling, 2014] and GAN [Goodfellow *et al.*, 2014]. The approach [Verma *et al.*, 2018; Bucher *et al.*, 2017; Mishra *et al.*, 2017] is based on the VAE architecture while [Xian *et al.*, 2015; Chen *et al.*, 2018; Lu *et al.*, 2017] used the adversarial sample generation based on the class conditioned attribute.

In ZSL, the train and test classes are disjoint and hence there is a high probability of domain shift for the unseen classes. This is another challenge in the ZSL setup and needs to be handled. Previously, very few works have handled the domain shift problem and worked on both the transductive as well as inductive settings. [Verma and Rai, 2017] adapted to the new domain by simple Gaussian mixture model updates. [Song *et al.*, 2018] used the unbiased embedding in the transductive setting. [Kodirov *et al.*, 2015b; Ye and Guo,] proposed unsupervised domain adaption for the ZSL. [Zhang and Saligrama, 2016] used the structural SVM formulation for domain adaption.

In this paper, we propose the design of a deep generative model that achieves a significant improvement in the ZSL setting. There are stark differences between previously proposed deep generative models for ZSL, which were based on VAE/GAN and our proposed approach. The VAE based architecture minimizes the ELBO [Kingma and Welling, 2014] by using a scheme of approximate optimization, making is less robust in handling domain shift. The GAN based generative approach is difficult to train, requiring a lot of seen class examples during training to facilitate learning of reasonable class conditionals for the unseen classes. Moreover, they need the attribute vectors of unseen classes at the beginning of the procedure while our model can handle on the fly addition of new classes. To this end, we propose a simple CNN based architecture that has the ability to learn any parametric distribution with exact optimization, and unlike the GAN based approach, has stable training. This makes it especially suitable for augmenting explicit domain shift minimization by adapting the distribution of the unseen classes by adversarial transfer. Our approach is general enough to be of use in both the inductive and the transductive settings.

4 Experiments and Results

To demonstrate the effectiveness of the proposed approach we performed extensive experimentation on the three standard datasets for ZSL, namely AWA2 [Xian *et al.*, 2018], CUB-200 [Welinder *et al.*, 2010] and SUN [Xiao *et al.*, 2010]. In all the experiments, we follow the newly proposed train test split suggested by [Xian *et al.*, 2018]. Since we are use the pre-trained resnet-101 model, therefore, we first sought to make sure that any class that belongs to the test classes is not present in the ImageNet [Russakovsky *et al.*, 2015] training samples. This was already rectified in the split proposed by [Xian *et al.*, 2018] for ZSL. The complete description of the bench mark datasets are as follows:

Dataset	Attribute/Dim	#Image	Seen/Unseen Class
AWA2	A/85	37322	40/10
CUB	A/312	11788	150/50
SUN	A/102	14340	645/72

Table 1: Datasets used in our experiments, and their statistics

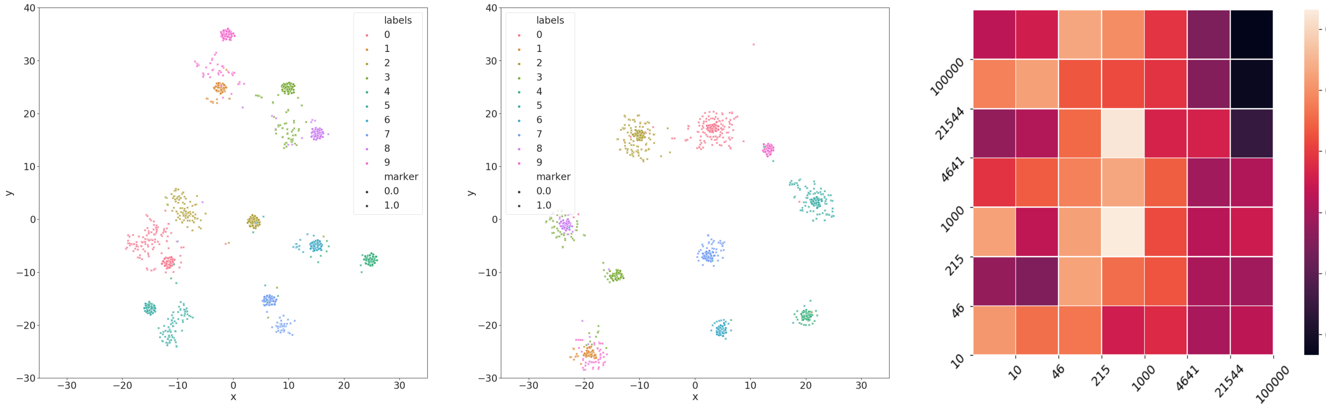

(a) (b) (c)

Figure 2: **(a)** shows the t-SNE plot for the output of the generative model as compared to the test data. Crosses represent the test data while dots represent the generated data. The domain shift is clearly visible in this plot. **(b)** shows the t-SNE plot after domain shift minimization with our model. The scale for the axes in **(a)** and **(b)** is kept constant for comparison. The model is able to allot the clusters properly, considering the corruption of labels. **(c)** shows the stability of the generative model wrt regularization coefficients on AWA2 dataset. The xlabels and ylabels are the weight decay in Adam optimizer for learning the NN parameters predicting the Mean and Sigma of the class-conditional distributions respectively. The shaded grid values represent the top-1 accuracy obtained for the given configuration of hyperparameters.

Animal with Attribute (AWA2): The dataset has 50 classes of animals. Each class also has a human annotated 85-dimensional attribute vector associated with it. Here 40 class are used for the training data and the rest 10 classes are used for the test samples.

CUB-200: This is a fine-grained dataset, containing 200 classes of birds. It has 11788 data points and 312-dimensional human annotated attribute vectors for each class. For the ZSL setup, 150 classes are used for the training while the rest 50 classes are left for testing. There is little inherent domain shift in the data but lack of sufficient data poses a problem for typical ZSL models.

SUN Seen Recognition: There are a total of 14340 images between 717 classes. Hence, every class has nearly 20 samples. Each class is associated with a 102-dimensional human annotated attribute vector. For the ZSL setup, 645 classes are used for training and the rest 72 classes are used for testing.

4.1 Implementation Details

Generative Framework

In this section, we describe the architecture that yields our reported results in the ZSL setting wherein there are no images from seen classes in test samples. For SUN dataset, both the networks used for modeling mean and co-variance have linear (1800 and 2048 nodes), batch normalization and Relu layers. Additionally, the co-variance is restricted to be in the range of 0.5 to 1.5 for numerical stability via sigmoid activation. Both the networks are trained with ADAM optimizer [Kingma and Ba, 2014] using 0.001 and 0.1 as regularizer coefficients for means and covariance respectively.

For AWA, the generator networks have an architecture similar to SUN but consist of an additional dropout layer with probability 0.1. The parameters of the means are regularized with the coefficient 10^3 while the parameters of covariance are regularized with coefficient 10^4.

For CUB, the generator networks have three linear layers (1200,1800,2048 nodes), 2 Relu, 2 batch normalization and 2 dropout layers. Their regularization coefficients are 0.01 and 0.1 for mean and covariance respectively. All the above networks are trained with a learning rate of 10^{-5}. The training of these networks was additionally regularized via early stopping

Adversarial Domain Adaption

As described above, ADA model comprises of two discriminators ($D^{S,T}$), two classifiers ($C^{S,T}$) and two domain adapting generators ($G^{S,T}$).

The discriminators $D^{S,T}$ are 5 layered neural networks comprising of 2 Linear layers of 1600 nodes and 1 node respectively, a single leaky Relu layer with a negative slope of 0.2 and batch normalization. The final output (output $\in (0,1]$) is generated by passing the output of discriminator through a sigmoid activation. The classifiers are single layered networks with the number of nodes equal to the number of classes. $log(softmax(\mathbf{x}))$ is used as activation function for the classifiers $C^{S,T}$. The generators $G^{S,T}$ consist of three linear layers (1200,1200 and 2048 nodes), dropout layers, batch normalization and leaky Relu.

The overall objective is minimized using ADAM ([Kingma and Ba, 2014]) optimizer with a learning rate of 0.00001. A manual seed of 100 has been used for all the ADA experiments.

4.2 Zero-Shot Learning (ZSL)

We report per class accuracy as is the convention in standard ZSL. It is a better metric to report the accuracy of the model as compared to the overall (across classes) accuracy when the classes are unbalanced. We use the newly proposed splits for dividing the train and test examples. We use the corresponding attribute vector provided against each dataset. Please refer to table-1 for details on the dataset.

The results of the ZSL setting are shown in the table-2. We can see from the table the proposed approach shows significant improvement over the previous state-of-art approach. On the SUN [Xiao *et al.*, 2010] and AWA2 [Lampert *et al.*, 2009] we have our top-1 accuracy as 63.3% and 70.4% respectively, which is better than its close competitor [Verma and Rai, 2017] where they report the top-1 accuracy as 63.1% and 67.0% respectively. Also, their top-1 accuracy on the fine-grained CUB dataset is reduced to 49.2%, as compared to our model's top-1 accuracy of 53.5% on the CUB dataset.

Also, the models ESZSL[Romera-Paredes and Torr, 2015] and SYNC[Changpinyo *et al.*, 2016] that perform slightly better than our approach on CUB dataset, they are significantly behind in terms of accuracy on AWA2 and SUN dataset as compared to our model. The model SYNC[Changpinyo *et al.*, 2016] achieves 46.6% on AWA2 and 56.3% on SUN dataset compared to our model's top-1 accuracy of 70.4% and 63.3% respectively. Hence, our model performs well consistently on all three benchmark datasets.

Additionally, our approach is more stable to hyperparameter variations as compared to the other competing generative approaches like GFZSL[Verma and Rai, 2017]. This is clearly visible in *figure 2, (c)* where the accuracy remains stable across the grid. We train our model in an online fashion, which leads to simultaneous updates of attribute mappings while estimating the class distribution parameters from the current batch of data.

We further improve the accuracy by augmenting our model with Adversarial Domain Adaptation.

	SUN	CUB	AWA2
Method	**PS**	**PS**	**PS**
DAP[Lampert *et al.*, 2014a]	39.9	40.0	46.1
IAP [Lampert *et al.*, 2014b]	19.4	24.0	35.9
CONSE [Norouzi *et al.*, 2013]	38.8	34.3	44.5
CMT [Socher *et al.*, 2013]	39.9	34.6	37.9
SSE [Zhang and Saligrama, 2016]	51.5	43.9	61.1
LATEM [Xian *et al.*, 2016]	55.3	49.3	55.8
DEVISE [Frome *et al.*, 2013]	56.5	52.0	59.7
SJE [Akata *et al.*, 2015]	53.7	53.9	61.9
ESZSL [Romera and Torr, 2015]	54.5	53.9	58.6
SYNC[Changpinyo *et al.*, 2016]	56.3	**55.6**	46.6
SAE [Kodirov *et al.*, 2017]	40.3	33.3	54.1
DEM [Zhang *et al.*, 2017]	61.9	51.7	67.1
GFZSL[Verma and Rai, 2017]	63.1	49.2	67.0
CVAE-ZSL[Mishra *et al.*, 2017]	61.7	52.1	65.8
W/O ADA (Ours)	**63.3**	53.5	**70.4**

Table 2: Zero Shot Learning Accuracy on the SUN, CUB, and AWA2 dataset. Here PS is the proposed split recently adopted in the ZSL community.

4.3 Domain Adaption

In ZSL, since $S \cap U = \phi$, there is a high probability that the seen and unseen data do not come from the same underlying domain. Hence, the existence of the domain shift between the seen and unseen class distributions is a reasonable assumption. This implies that the estimated parameters for the unseen classes, based on the training data of the seen class are likely to deviate from their optimal values. To this end, we propose an Adversarial Domain Adaptation (ADA) method to explicitly handle the domain shift problem. The proposed domain adaption method is based on ADDA and CycleGAN, that learns a common latent space for the two domains by enforcing cyclic consistency and yields stable adversarial training.

Method	**SUN**	**CUB**	**AWA2**
ALE [Akata *et al.*, 2013]	57.1	54.6	71.3
GFZSL [Verma and Rai, 2017]	64.2	50.5	**78.6**
With ADA (Ours)	**65.4**	**56.7**	77.2

Table 3: Transductive Zero-Shot Learning results on the SUN, CUB, and AWA2 dataset. Transductive setting for our model corresponds to ADA.

In Table-3 we show the results of the proposed ADA method and compare against the previous transductive setting approaches. The result of ALE [Akata *et al.*, 2013] and GFZSL [Verma and Rai, 2017] are taken from the Figure 8 of [Xian *et al.*, 2018]. Here we observe that using the domain adaption method boosts the generative model's performance. In the case of the AWA2 dataset without domain adaption, the top-1 accuracy was 70.3% while with the domain adaption it rises to 77.2%. A similar pattern is observed for the CUB and SUN dataset also. The increment is lesser in CUB and SUN due to small inherent domain shift in the data. Moreover, our domain adaptation is able to minimize domain shift (apparent in figure 2 (a),(b)) in accordance with the clusters allotted by ZSL model. We can see that the model associates wrong clusters for only two classes owing to the low prediction accuracy of the generative model for these classes which, itself is due to strong overlap in test clusters of these classes. Thus, a reduction in label corruption will definitely improve the domain matching. This justifies the utility of the adversarial framework.

5 Conclusion

In this paper, we propose a scheme for explicitly minimizing domain shift between the distributions of the seen and unseen classes in zero-shot learning that leverages a generative framework. We adopt an end-to-end approach for generative modeling that requires minimal hyper-parameter tuning and captures non-linear dynamics better as compared to previous state-of-the-art approaches. The proposed approach first learns the class conditional distributions for both the seen and unseen classes by leveraging the data from only the seen classes through an end-to-end training scheme. Following this, we explicitly minimize the domain shift by using a cyclic consistency based adversarial scheme that tunes the generative model previously learned. We show through detailed experimentation, that our proposed generative model, although much simpler than GAN/VAE based frameworks, outperform existing models in the ZSL setting by significant margins. In addition, we show that our scheme of explicitly minimizing domain shift significantly improves performance, as compared to the standard transductive setting methods adopted by previous approaches.

References

[Akata *et al.*, 2013] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for attribute-based classification. In *CVPR*, pages 819–826, 2013.

[Akata *et al.*, 2015] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *Proceedings of the IEEE Conference on CVPR*, pages 2927–2936, 2015.

[Arjovsky *et al.*, 2017] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.

[Bharadhwaj *et al.*, 2018a] Homanga Bharadhwaj, Homin Park, and Brian Y Lim. Recgan: recurrent generative adversarial networks for recommendation systems. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pages 372–376. ACM, 2018.

[Bharadhwaj *et al.*, 2018b] Homanga Bharadhwaj, Zihan Wang, Yoshua Bengio, and Liam Paull. A data-efficient framework for training and sim-to-real transfer of navigation policies. *arXiv preprint arXiv:1810.04871*, 2018.

[Bucher *et al.*, 2017] Maxime Bucher, Stéphane Herbin, and Frédéric Jurie. Generating visual representations for zero-shot classification. *CoRR*, abs/1708.06975, 2017.

[Changpinyo *et al.*, 2016] Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, and Fei Sha. Synthesized classifiers for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5327–5336, 2016.

[Chen *et al.*, 2018] Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, and Shih-Fu Chang. Zero-shot visual recognition using semantics-preserving adversarial embedding networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[Frome *et al.*, 2013] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NIPS*, pages 2121–2129, 2013.

[Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[Guo *et al.*, 2017] Yuchen Guo, Guiguang Ding, Jungong Han, and Yue Gao. Synthesizing samples for zero-shot learning. In *IJCAI*, 2017.

[He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.

[Hoffman *et al.*, 2017] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017.

[Kingma and Ba, 2014] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[Kingma and Welling, 2014] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014.

[Kodirov *et al.*, 2015a] Elyor Kodirov, Tao Xiang, Zhenyong Fu, and Shaogang Gong. Unsupervised domain adaptation for zero-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2452–2460, 2015.

[Kodirov *et al.*, 2015b] Elyor Kodirov, Tao Xiang, Zhenyong Fu, and Shaogang Gong. Unsupervised domain adaptation for zero-shot learning. In *ICCV*, pages 2452–2460, 2015.

[Kodirov *et al.*, 2017] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. *arXiv preprint arXiv:1704.08345*, 2017.

[Lampert *et al.*, 2009] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 951–958. IEEE, 2009.

[Lampert *et al.*, 2014a] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on PAMI*, 36(3):453–465, 2014.

[Lampert *et al.*, 2014b] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on PAMI*, 36(3):453–465, 2014.

[Lu *et al.*, 2017] Jiang Lu, Jin Li, Ziang Yan, and Changshui Zhang. Zero-shot learning by generating pseudo feature representations. *arXiv preprint arXiv:1703.06389*, 2017.

[Mishra *et al.*, 2017] Ashish Mishra, M Reddy, Anurag Mittal, and Hema A Murthy. A generative model for zero shot learning using conditional variational autoencoders. *arXiv preprint arXiv:1709.00663*, 2017.

[Norouzi *et al.*, 2013] Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. *arXiv preprint arXiv:1312.5650*, 2013.

[Reed *et al.*, 2016] Scott E. Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *CoRR*, abs/1605.05396, 2016.

[Romera and Torr, 2015] Bernardino Romera, Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *International Conference on Machine Learning*, pages 2152–2161, 2015.

[Romera-Paredes and Torr, 2015] Bernardino Romera-Paredes and Philip HS Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, pages 2152–2161, 2015.

[Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, pages 211–252, 2015.

[Socher *et al.*, 2013] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, pages 935–943, 2013.

[Song *et al.*, 2018] Jie Song, Chengchao Shen, Yezhou Yang, Yang Liu, and Mingli Song. Transductive unbiased embedding for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1024–1033, 2018.

[Thekumparampil *et al.*, 2018] Kiran K Thekumparampil, Ashish Khetan, Zinan Lin, and Sewoong Oh. Robustness of conditional gans to noisy labels. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 10271–10282. Curran Associates, Inc., 2018.

[Tzeng *et al.*, 2017] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017.

[Verma and Rai, 2017] Vinay Kumar Verma and Piyush Rai. A simple exponential family framework for zero-shot learning. In *ECML-PKDD*, pages 792–808. Springer, 2017.

[Verma *et al.*, 2018] Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, and Piyush Rai. Generalized zero-shot learning via synthesized examples. *CVPR*, 2018.

[Wang *et al.*, 2017a] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 515–524. ACM, 2017.

[Wang *et al.*, 2017b] Wenlin Wang, Yunchen Pu, Vinay Kumar Verma, Kai Fan, Yizhe Zhang, Changyou Chen, Piyush Rai, and Lawrence Carin. Zero-shot learning via class-conditioned deep generative models. *arXiv preprint arXiv:1711.05820*, 2017.

[Welinder *et al.*, 2010] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010.

[Xian *et al.*, 2015] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[Xian *et al.*, 2016] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 69–77, 2016.

[Xian *et al.*, 2018] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning-a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 2018.

[Xiao *et al.*, 2010] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR, 2010*, pages 3485–3492. IEEE, 2010.

[Ye and Guo,] Meng Ye and Yuhong Guo. Zero-shot classification with discriminative semantic representation learning.

[Zhang and Saligrama, 2016] Ziming Zhang and Venkatesh Saligrama. Learning joint feature adaptation for zero-shot recognition. *arXiv preprint arXiv:1611.07593*, 2016.

[Zhang *et al.*, 2017] Li Zhang, Tao Xiang, Shaogang Gong, et al. Learning a deep embedding model for zero-shot learning. 2017.

[Zhu *et al.*,] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks.