

Generalised Zero-Shot Learning with Domain Classification in a Joint Semantic and Visual Space

Rafael Felix^{1,2}, Ben Harwood^{1,3}, Michele Sasdelli^{1,2} and Gustavo Carneiro^{1,2}

Australian Centre for Robotic Vision¹

The University of Adelaide²

Monash University³

{rafael.felixalves,michele.sasdelli,gustavo.carneiro}@adelaide.edu.au, ben.harwood@monash.edu

Abstract—Generalised zero-shot learning (GZSL) is a classification problem where the learning stage relies on a set of seen visual classes and the inference stage aims to identify both the seen visual classes and a new set of unseen visual classes. Critically, both the learning and inference stages can leverage a semantic representation that is available for the seen and unseen classes. Most state-of-the-art GZSL approaches rely on a mapping between latent visual and semantic spaces without considering if a particular sample belongs to the set of seen or unseen classes. In this paper, we propose a novel GZSL method that learns a joint latent representation that combines both visual and semantic information. This mitigates the need for learning a mapping between the two spaces. Our method also introduces a domain classification that estimates whether a sample belongs to a seen or an unseen class. Our classifier then combines a class discriminator with this domain classifier with the goal of reducing the natural bias that GZSL approaches have toward the seen classes. Experiments show that our method achieves state-of-the-art results in terms of harmonic mean, the area under the seen and unseen curve and unseen classification accuracy on public GZSL benchmark data sets. Our code will be available upon acceptance of this paper.

I. INTRODUCTION

Humans have a powerful ability to learn about new visual objects without actually seeing them. This process generally involves the use of language to describe how a new visual object would look like. The textual description then allows for a new class of object to be formed in a person’s mind. Our understanding of exactly how the human brain functions for this task is limited, but it is clear that humans make some sort of association between visual objects and semantic textual descriptions. Conceptually, objects with similar descriptions can naturally be viewed as being near to each other in some latent space, representing visual and semantic information. The research topic is known as generalised zero-shot learning (GZSL) aims to mimic this recognition ability of humans. In general, GZSL approaches employ an auxiliary set of semantic information that describes a set of visual classes. This additional information, such as tags or descriptions, can be utilised to overcome missing visual information in some of the classes [1].

Traditional GZSL approaches aim to recognise the visual classes available during the training process (i.e. the **seen**, source or known classes), and also classes that are not available during training (i.e. **unseen**, target or novel classes). Due to this constraint, GZSL approaches are intrinsically divided into two main tasks: (1) the training of a model that learns a

transformation from the visual to the semantic space, using the visual samples and semantic information from seen classes; and (2) the transformation of a new test image by the model above into the semantic space, followed by a search of the closest semantic sample representing a seen or unseen class. In recent years, GZSL researchers have become increasingly interested in pairwise functions for disentangling these domains [2], and deep generative models [3], [4] for learning to transform between the visual and semantic representations. In general, GZSL methods do not try to estimate if a test sample belongs to the set of seen or unseen classes – this issue inevitably biases GZSL approaches toward seen classes. Only recently this issue has been acknowledged with a method that automatically combines the classification of Zero-Shot Learning (ZSL) for unseen classes with the classification of seen classes, by automatically weighting (using the test sample) the contribution of each classifier [5]. Although that approach is in the right direction, it has the issue of relying on the training of multiple classifiers. Another issue with the methods above is that they do not consider a latent space jointly optimised for the visual and semantic representation, which we believe is a crucial part of the inference process performed by humans that should be imitated by GZSL methods. In Fig. 1, we illustrate the idea explored in this paper for GZSL. The visual and semantic samples are represented in a joint latent space. This space is used to learn a classifier of visual classes and a domain classifier for seen and unseen domains.

In this paper, we aim to explore two observations about the latent space for the domain classification. The first observation is that samples from unseen classes that are visually similar to one of the seen classes tend to be projected relatively close to other seen classes distributions, instead of outside of the distribution of seen classes, as proposed by Socher et. al [6]. Our second observation is that samples from unseen classes that are visually different from any of the seen classes, tend to be projected outside the distribution of seen classes [6]. Atzmon and Chechik [5] propose a general framework that combines domain expert classifiers, such as DAP [7] for unseen classes, and LAGO for the seen classes [5]. However, this method relies on the disjoint training of both experts models, and the assumption that unseen samples are projected outside the distribution of seen classes [6]. Hence, this method can be considered to be in general sub-optimal. We propose a general framework for learning and combining the visual and domain classifiers using the latent space. More specifically, we first introduce a general framework for latent space learning

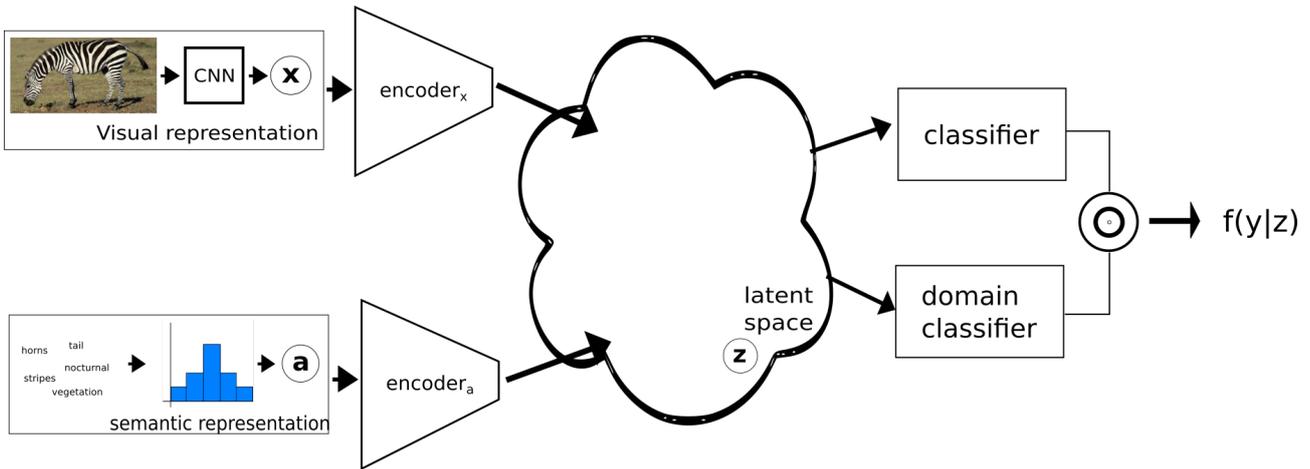


Fig. 1: Depiction of the method proposed in this paper – our approach learns the latent space for the visual and semantic modalities. We train two classifiers using samples from this latent space: one to classify all the seen and unseen visual classes, and another to classify between the seen and unseen domains. The final classification combines the results of these two classifiers.

from cycle-WGAN [3] and CADA-VAE [4]. Then, we propose a novel method for the seen and unseen domain classification from this latent space. Finally, we introduce a way to combine the visual and domain classifiers. The empirical results show that our proposed framework outperforms previous approaches in terms of unseen accuracy and harmonic mean (**H-mean**) on several GZSL benchmark data sets, such as CUB [8], SUN [1], AWA1 [7], [1] and AWA2 [7], [1]. In terms of unseen accuracy, our method shows improvements of 4.5%, 5.6%, 2.5%, 1.5% for CUB, SUN, AWA1, and AWA2, respectively. Moreover, our method shows substantial improvements in terms of area under the curve of seen and unseen accuracy (**AUSUC**) [9]. For AUSUC we improved from 0.3698, 0.5238, 0.5216 to 0.3743, 0.5247, 0.5219, on CUB, AWA1 and AWA2, respectively.

II. RELATED WORK

In this section, we discuss relevant literature that motivates and contextualises our work.

A. Traditional Zero-Shot Learning

Zero-shot learning (ZSL) is similar to GZSL, with a crucial difference: during inference, only the visual samples from the unseen classes are considered [10], [1]. This difference makes ZSL a special case of GZSL. Therefore, critical problems present in GZSL are not considered in this approach, such as the natural bias of the visual classifier toward the seen classes. Unfortunately, this setup not only reduces the applicability of ZSL methods but also makes it unrealistic for real-world applications [11], [3]. Also, ZSL fails to handle jointly the seen and unseen data [9], [12]. Due to the simplicity and unrealistic assumptions of ZSL, the whole field moved toward the GZSL problem, which is introduced in the next section.

B. Generalised Zero-Shot Learning

In GZSL, the algorithm is trained using visual samples from the seen classes, but the inference involves the analysis of samples from the seen and unseen classes. The main issue faced by GZSL methods is the bias toward the seen

classes naturally present during inference, so a great deal of research has focused on mitigating this problem [3], [11]. Particularly important examples of this type of research are anomaly detection [6], domain balancing [9] and generative data augmentation for GZSL [3], [4], [11]. Despite the advances in GZSL with the approaches mentioned above, we note that little attention has been devoted to addressing the seen/unseen domain classification in GZSL based on a latent space that is jointly learned to represent the visual and semantic representations. Moreover, we argue that the multi-modal nature of this joint latent space carries interesting properties to perform domain classification. In this paper, we show that classifying the seen and unseen domains plays an important role in improving domain balancing in GZSL.

C. Data Augmentation for Zero-Shot Learning

A particularly successful GZSL method is based on data augmentation, where artificial visual samples of the unseen classes are generated from the semantic representation to train the visual classifier [3], [4], [11], [13]. This approach has produced the current state-of-the-art results in GZSL benchmark data sets. Overall, these studies focus on how to learn generative models conditioned on the semantic information that is used to augment the data set for the unseen classes. Among the main approaches, we observe the use of Generative Adversarial Networks (GAN) [3], [11] and Variational Autoencoders (VAE) [4], [13]. In this paper, we formalise these approaches as a framework for generative probabilistic latent space learning. Additionally, we show that these latent spaces have interesting properties that allow our approach to classifying samples into the seen or unseen domains for GZSL.

D. Domain Classification

Recent research has tackled the problem of GZSL as a novelty detection problem [6]. This approach assumes that unseen classes are projected out of the distribution of seen classes. Therefore, these unseen classes samples can be handled as an outlier of the seen classes distribution [6]. However, this

approach fails to notice that samples from unseen classes can be projected relatively close to one of the seen classes. Atzmon and Chechik [5] aims to tackle this novelty detection issue by providing a framework that handles domain classification for GZSL. The gist of that approach consists of a gating method that performs domain adaptation to combine an unseen class classifier (e.g., DAP [10], DeVISE [2]), and CMT [6]), and a seen class classifier [5]. Even though this method achieves remarkable performance in GZSL, it still relies on a sub-optimal disjoint training of multiple classifiers. In this paper, we mitigate these two issues by combining a seen/unseen class discriminator with a domain classifier that uses samples from a latent space that is trained to represent both the visual and semantic spaces.

III. METHOD

In this section, we introduce the problem formulation and our proposed approach.

A. Generalised Zero-Shot Learning

In order to formulate the method of learning a classifier that can recognise visual samples from unseen visual classes, we define a visual data set $\mathcal{D} = \{(\mathbf{x}, y)_i\}_{i=1}^N$, where $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^K$ denotes the visual representation, and $y \in \mathcal{Y} = \{1, \dots, C\}$ denotes the visual class. Recent research shows that such visual representation, \mathbf{x} , can be acquired from networks specialised in feature extraction. These are widely available in the literature, such as pre-trained deep residual nets [14].

In GZSL, the set of classes \mathcal{Y} is split into two domains: seen domain $\mathcal{Y}^S = \{1, \dots, |S|\}$, and the unseen domain $\mathcal{Y}^U = \{|S| + 1, \dots, |S| + |U|\}$. Hence, the total number of classes is $C = |S| + |U|$, with $\mathcal{Y} = \mathcal{Y}^S \cup \mathcal{Y}^U$, $\mathcal{Y}^S \cap \mathcal{Y}^U = \emptyset$. During training, we can only access visual samples from \mathcal{Y}^S , but during testing, samples can come from any class in \mathcal{Y} . This lack of visual samples from unseen classes during training is compensated with a semantic data set that includes semantic information for the seen and unseen classes. Therefore, we introduce the semantic data set $\mathcal{R} = \{(\mathbf{a}, y)_j\}_{j \in \mathcal{Y}}$, which associates visual classes with semantic samples, where $\mathbf{a} \in \mathcal{A} \subseteq \mathbb{R}^L$ represents a semantic feature (e.g., set of continuous features such as *word2vec* [1], or *BoW*). Note that the semantic data set only has a single element per class.

In comparison with the supervised learning paradigm, the problem of GZSL has a distinct setup. The data set \mathcal{D} is divided into mutually exclusive training and testing visual subsets \mathcal{D}^{Tr} and \mathcal{D}^{Te} , respectively. The \mathcal{D}^{Tr} contains a subset of the visual samples belonging to the seen classes, and \mathcal{D}^{Te} contains the visual samples from the seen classes that are held out from training and all samples from the unseen classes. The training data set is composed of the semantic data set \mathcal{R} and the training visual subset \mathcal{D}^{Tr} , while the testing data set relies only on the testing visual subset \mathcal{D}^{Te} .

B. Data Augmentation Framework

In this section, we first introduce the components for the latent space learning applied to GZSL models, then we describe CADA-VAE and cycle-WGAN. Finally, we introduce the domain classification for these latent space.

In recent years, we note an increasing number of models that use data augmentation for GZSL models [3], [4], [11], [15], [16], [17]. Overall, these methods aim to learn a generative model that produces artificial samples from unseen visual classes conditioned on their semantic representation. These artificial samples lie in a latent space. In this paper, we aim to demonstrate that our proposed domain classification can be adapted to GZSL models that rely on data augmentation, such as CADA-VAE [4] and cycle-WGAN [3]. Although these two models consist of different training approaches, we observe that their components can be generally described as a framework for latent space learning. Below, we introduce three components of such models: the encoder (or generator), the decoder (or regressor), and the discriminator.

The encoder transforms samples from an input space (i.e., visual or semantic) into a latent space. We represent the encoder with

$$\mathbf{z}_x = \text{Encoder}_x(\mathbf{x}) \quad (1)$$

for the visual space and similarly for the semantic space with $\mathbf{z}_a = \text{Encoder}_a(\mathbf{a})$, where the vector $\mathbf{z}_{\{x,z\}} \in \mathbb{R}^Z$ lies in the latent space. The decoder transforms from the latent space into one of the input modalities. We represent the decoder with

$$\tilde{\mathbf{x}} = \text{Decoder}_x(\mathbf{z}), \quad (2)$$

and similarly for the semantic space with $\tilde{\mathbf{a}} = \text{Decoder}_a(\mathbf{z})$. The latent space discriminator, used to determine whether a sample \mathbf{z} belongs to the latent space given the input \mathbf{x} , is represented by

$$p(\mathbf{z} | \mathbf{x}) = \text{Discriminator}(\mathbf{z}; \mathbf{x}). \quad (3)$$

We consider the simplified models above to describe CADA-VAE [4] and cycle-WGAN [3] as the latent space learning models.

CADA-VAE: This model is a special type of variational autoencoder (VAE) for GZSL [4]. In this approach, the VAE aims to learn the latent space with cross alignment and distribution alignment losses, as depicted in Fig. 2. The overall loss by Schonfeld et al. [4] can be described with

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{VAE} + \gamma \left(\sum_i^L \sum_{j \neq i}^L \|\mathbf{x}^{(j)} - \tilde{\mathbf{x}}^{(i)}\| \right) \\ & + \delta \left(\|\mu^{(j)} - \mu^{(i)}\|_2^2 + \|\Sigma_{(j)}^{\frac{1}{2}} - \Sigma_{(i)}^{\frac{1}{2}}\|_{Frobenius}^2 \right), \end{aligned} \quad (4)$$

where the first term represents the VAE loss [4], the second term denotes the reconstruction error between L modalities – that is, during training, the encoder projects input samples in the latent space (e.g. Encoder_x for \mathbf{x}), then the decoder of a different modality is used (e.g. Decoder_a from \mathbf{z}_x – see Fig. 2), which constrains the visual and semantic projections to be in the same region of the latent space represented by the mean μ and variance Σ of the samples produced by the encoder [4].

cycle-WGAN: Fig. 3 depicts the model cycle-WGAN [3]. This model is optimised as a Generative Adversarial Network (GAN), regularised by a cycle consistent term, described with

$$\mathcal{L} = \mathcal{L}_{WGAN} + \gamma \left(\|\mathbf{a} - \tilde{\mathbf{a}}\|_2^2 \right), \quad (5)$$

where the first term, \mathcal{L}_{WGAN} , represents a *Wasserstein* Generative Adversarial Loss (WGAN [3]), and the second term

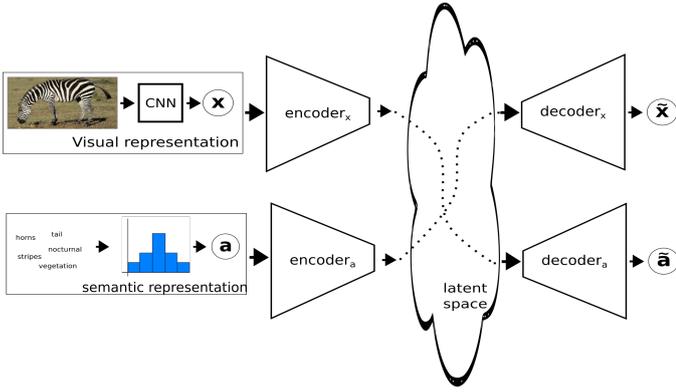


Fig. 2: Depiction of the method CADA-VAE [4]. In this method encoders for the visual and semantic representation project samples into a shared latent space.

denotes the reconstruction loss (cycle) for the semantic representation. Thus, the generative projection of a given semantic representation into the latent space is encouraged to be back projected near the original semantic representation.

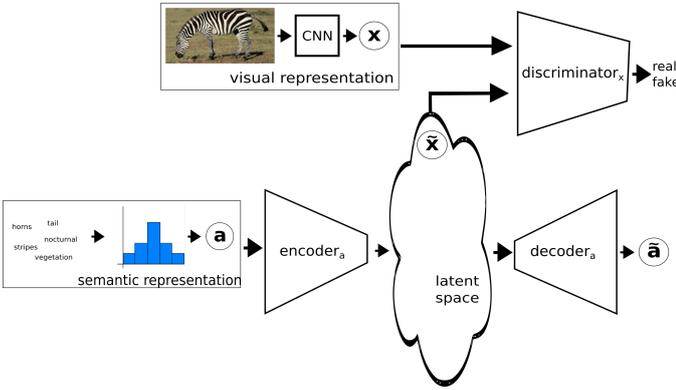


Fig. 3: Depiction of the cycle-WGAN method [3]. This method encodes the semantic space into a latent visual space. The decoder produces semantic vectors that are used to regularise the learning process.

C. Domain Classification

From the previous section, we note that the latent space is an embedding space for visual and semantic samples. Therefore, we can use this latent space to learn a discriminative model given by

$$f(y | \mathbf{x}) = \int_v \int_{\mathbf{z}} p(\mathbf{z} | \mathbf{x}) f(y, v | \mathbf{z}) dv d\mathbf{z}, \quad (6)$$

where the function $f(\cdot)$ represents the GZSL classifier and can be described in terms of domains, $v \in \{s, u\}$ ($s = \text{seen}$ and $u = \text{unseen}$), with

$$f(y | \mathbf{x}) = \sum_{v \in \{s, u\}} p(y | \mathbf{z}_x) f(v | \mathbf{z}_x, y), \quad (7)$$

where we assume from (6) that $p(\mathbf{z} | \mathbf{x})$ is a delta function at $\mathbf{z}_x = \text{Encoder}_x(\mathbf{x})$. The term $p(y | \mathbf{z}_x)$ in (7) is represented

by a simple deep learning classifier with softmax activation. We define the function $f(\cdot)$ in (7) by

$$f(v | \mathbf{z}_x, y) = \begin{cases} p(v | \mathbf{z}_x, y), & \text{if } v, y \text{ are in same domain} \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where "same domain" means the domain of seen or the unseen classes, and $p(v | \mathbf{z}_x, y)$ is denoted by a deep learning classifier with softmax activation. The function in (8) represents our proposed domain classifier (DC). During the DC training, for training samples of the seen domain, we optimise $p(v = s | \mathbf{z}, y)$ with samples drawn from the latent space. These samples are acquired from visual and semantic representations projected in the latent space. For the unseen domain, $p(v = u | \mathbf{z}, y)$, we use the semantic projections in the latent space.

IV. EXPERIMENTS

In this section, we present the benchmark datasets, as well as the evaluation criteria for our experimental setup. We then show the results of our method and compare them with the current state-of-the-art. Finally, we provide ablation studies to explore our method.

A. Data Sets

We assess our method on four publicly available benchmark GZSL data sets: CUB-200-2011 [8]; SUN [1]; AWA1 [7], [1], and AWA2 [7], [1]. To guarantee that our experiments are reproducible, we use the GZSL experimental setup described by Xian et al. [1]. As the CUB data set is generally regarded as fine-grained, there is an intrinsic expectation that the novel unseen classes tend to have their class modes close to the seen classes. Thus, such dense visual representation space is a challenging problem for GZSL approaches. We also explore the use of coarse data sets, such as AWA1, AWA2, and SUN. Given the diversity of classes for such coarse data sets, there is an intrinsic expectation that novel classes will be projected far away from the samples of seen classes in the latent space, making the domain classification a trivial task. However, we argue that this statement does not always hold, particularly for classes that are visually similar (e.g. zebra/horse, whale/dolphin, leopard/bobcat), as depicted in Fig. 4. Table I contains some basic information about the data sets in terms of the number of seen and unseen classes and the number of training and testing images.

We represent the visual space by extracting image features from the activation of the 2048-dimensional top pooling layer of ResNet-101 [14]. For the semantic representation of the data set CUB-200-2011 [1], we use the 1024-dimensional vector produced by CNN-RNN [19]. These semantic samples represent a written description of each image using 10 sentences per image. To define a unique semantic sample per-class, we average the semantic samples of all images belonging to each class [1]. We use manually annotated semantic samples containing 102 and 85 dimensions respectively, for the data sets SUN [1], AWA1 [1], and AWA2 [1]. To prevent a violation of the ZSL constraints, where the test classes should not be accessed during training, all the features were extracted according to training splits proposed in [1].

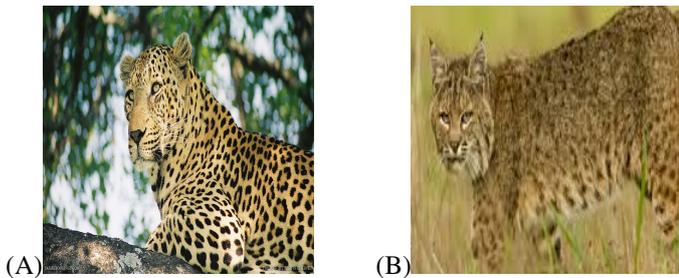


Fig. 4: Example of two classes that are visually similar from the benchmark dataset AWA1 [1]. (A) the sample leopard belongs to the seen classes, and (B) the sample bobcat belongs to the unseen classes. We speculate that samples from these two classes will lie close to each other in the latent space even though they come from different domains, challenging the view that samples from new unseen classes will lie far from samples of the seen classes in the latent space.

TABLE I: The benchmarks for GZSL: CUB[8], SUN [18], AWA1[1], and AWA2 [1]. Column (1) shows the number of seen classes, denoted by $|\mathcal{Y}^S|$, split into the number of training and validation classes (train+val), (2) presents the number of unseen classes $|\mathcal{Y}^U|$, (3) displays the number of samples available for training $|\mathcal{D}^{Tr}|$ and (4) shows number of testing samples that belong to the unseen classes $|\mathcal{D}_U^{Te}|$ and number of testing samples that belong to the seen classes $|\mathcal{D}_S^{Te}|$ from [3], [11]

Name	$ \mathcal{Y}^S $ (train+val)	$ \mathcal{Y}^U $	$ \mathcal{D}^{Tr} $	$ \mathcal{D}_U^{Te} + \mathcal{D}_S^{Te} $
CUB	150 (100+50)	50	7057	1764+2967
SUN	745 (580+65)	72	14340	2580+1440
AWA1 ¹	40 (27+13)	10	19832	4958+5685
AWA2	40 (27+13)	10	23527	5882+7913

B. Evaluation Protocol

Xian et. al [1] formalised the current evaluation protocol for GZSL. We first compute the average per-class top-1 accuracy measured independently for each class, then we calculate the overall mean. We calculate the mean-class accuracy for each domain separately, i.e., the seen (\mathcal{Y}^S) and the unseen (\mathcal{Y}^U) classes. Then, we also compute the harmonic mean (H-mean) of the seen and unseen domains accuracy [1]. Furthermore, we show results by measuring the area under the seen and unseen curve (AUSUC) [9] by varying the domain expertise [9]. This domain expertise consists of a hyper-parameter to perform the trade-off between the performance in the seen and unseen classes [9].

C. Implementation Details

In this section, we describe the architecture and training procedures for learning the proposed latent space. As described in Sec. III, we extend the following two models for our experimental setup: CADA-VAE [4] and cycle-WGAN [3]. The model CADA-VAE contains the following models that are parameterised as neural networks: $Encoder_x(\cdot)$, $Encoder_a(\cdot)$ in (1), $Decoder_x(\cdot)$, and $Decoder_a(\cdot)$ in (2). The training of CADA-VAE aims to produce a latent space that satisfies (4). In terms of the model architecture and hyper-parameters (e.g.

the number of epochs, batch size, the number of layers, learning rate, and, weight decay), we followed the specifications provided by [4]. The encoder for visual representation is parameterised with 1560 hidden neurons, and the encoder for the semantic representation is parameterised with 1450 hidden neurons. The decoders for the visual and semantic representation are parameterised with 1560, 660 hidden neurons, respectively. For both modalities, the encoders project samples into the latent space, which is represented with 64-dimension vectors in the latent space. The model is optimised with Adam for 100 epochs [20]. We use an adaptive scheduling rate for the hyper-parameters γ, δ , by (0.044, 0.0026), with respective epochs (21–75, 0–90) [4]. We also extended cycle-WGAN [3], as explained in Sec. III. The model cycle-WGAN contains the following functions that are parameterised as neural networks: $Encoder_a(\cdot)$ in (1), $Decoder_a(\cdot)$ in (2), and $Discriminator(\cdot)$ in (3). We followed the hyper-parameters choice (e.g. number of epochs, batch size, number of layers, learning rate, and weight decay, learning rate decay) defined in [3]. The encoder is parameterised with a single hidden layer containing 4096 nodes with LeakyReLU activation [21], and the output layer, with 2048 nodes, has a ReLU activation [22]. The decoder is parameterised with a linear layer, and the discriminator is a network with a single hidden layer with 4096 nodes. The network has a LeakyReLU activation, and the output layer has no activation.

The domain classifier (DC)² is implemented as a neural network with binary output, representing the seen and unseen domains. The model is trained with Adam optimiser [20] to recognise the domains. The output probability of the domain classifier tends not to be well calibrated [23], [5]. Therefore, we calibrate the model output using the validation set [23], [1]. Then, the domain classification is performed as described in (7) [9].

D. Results

In this section, we present the results for our proposed approach. The first question aimed to be answered in this paper consists of whether the proposed latent space contains relevant information that enables our approach to learn the domain classifier for GZSL. Thus, we provide numerical evidence that our method outperforms both baselines (i.e., CADA-VAE and cycle-WGAN) and previous GZSL. In Table III, we show the results in terms of unseen class accuracy \mathcal{Y}^U , seen class accuracy \mathcal{Y}^S and harmonic mean H , as described in Sec. IV-B. These results are given for the data sets CUB, SUN, AWA1 and AWA2. We compare our approach with 12 leading GZSL methods, which are divided into three groups: semantic (SJE [24], ALE [25], LATEM [26], ESZSL [27], SYNC [12], DEVISE [2]), latent space learning (SAE [15], f-CLSWGAN [11], cycle-WGAN [3] and CADA-VAE [4]) and domain classification (CMT[6] and DAZSL [5]). The semantic group contains methods that only use the seen class visual and semantic samples to learn a transformation function from the visual to the semantic space, and classification is based on nearest neighbour classification in that semantic space. The latent space learning group relies on visual samples from seen classes and semantic samples from seen and unseen classes

²The code will be available upon acceptance, and we intend to add the link to the Github repository here.

during training, and are detailed in Sec. III. The domain classification group relies on methods that weight the classification of seen and unseen classes. We discuss the numeral results in Table III in Section V.

E. Ablation Studies

In Table II we report the area under the curve of seen and unseen accuracy (AUSUC) [9] for the benchmark data sets CUB, SUN, AWA1, and AWA2. We compare the results of the original CADA-VAE [4] and cycle-WGAN [3] with and without the DC. Similar to harmonic mean, the AUSUC is an evaluation metric that measures the trade-off between the seen and unseen domains.

TABLE II: Area under the curve of seen and unseen accuracy (AUSUC). The highlighted values per column represent the best results in each data set. The notation * represents the results that we reproduced.

Classifier	CUB	SUN	AWA1	AWA2
EZSL	0.3020	0.1280	0.3980	—
DAZSL [5]	0.3570	0.2390	0.5320	—
f-CLSWGAN [11]	0.3550	0.2200	0.4610	—
cycle-WGAN [3]*	0.4180	0.2321	0.4730	—
CADA-VAE [4]*	0.3698	0.2362	0.5238	0.5216
cycle-WGAN + DC	0.4262	0.2321	0.4744	—
CADA + DC	0.3743	0.2364	0.5247	0.5219

V. DISCUSSIONS

In this section, we discuss the main contributions presented by our approach. We performed our experiments by combining previous GZSL approaches (such as CADA-VAE [4]) and cycle-WGAN [3]) with our Domain Classification in order to enhance the balancing of the seen and unseen domains for GZSL.

Firstly, in Table III we provide quantitative information that shows that our method outperforms existing methods in terms of unseen accuracy, \mathcal{Y}^U . This demonstrates that by learning to classify the domain for each sample, our method improves the classification of the unseen classes. Specifically, for CUB, SUN, AWA1 and AWA2 data sets, the baseline unseen classification results of 48.4%, 45.1%, 55.0%, and 55.5% have become 52.9%, 50.7%, 57.5%, and 57.0%. This improvement was achieved given a minor trade-off with the seen classes.

Secondly, despite the trade-off mentioned above, our approach is still able to achieve minor improvements in terms of **H-mean**. Table III shows an improvement of 0.2%, 0.1%, 0.1% and 0.4%, when compared to the baseline CADA-VAE. Although these results can be considered minor, we argue that our model does not directly optimise the H-mean. Thus, this improvement indicates that our approach has a more balanced performance than previous models.

We note similar behaviour for the cycle-WGAN model [3], where the proposed method achieves improvement for **H-mean** from 52.2% to 52.7% for CUB, from 39.2% to 40.3% for SUN, and from 59.7% to 60.0% for AWA1. However, such improvement is achieved due to the positive trade-off towards the seen domain. We argue that this difference, when compared to CADA-VAE, is due to the inherent differences in the latent

space learning of each of the approaches. In fact, the approach CADA-VAE is directly optimised by a variational autoencoder, where the control on the latent space is guided by a divergence measure for the visual and semantic representation jointly. On the other hand, the cycle-WGAN model is directly optimised by an adversarial loss from a generative adversarial network conditioned mainly on the semantic representation.

In terms of AUSUC, the proposed approach achieves improvements for both cycle-WGAN [3] and CADA-VAE [4]. For CADA-VAE, the domain classification yielded improvements from 0.3698, 0.2362, 0.5238, 0.5216 to 0.3743, 0.2364, 0.5247, 0.5219, for CUB, SUN, AWA1 and AWA2, respectively. Likewise, for cycle-WGAN [3], the DC provided improvements from 0.4180, 0.4730 to 0.4268, 0.4744 for CUB and AWA1, respectively.

VI. CONCLUSION AND FUTURE WORK

In this paper, we introduce a principled method to classify the seen and unseen domains in GZSL. In particular, we presented our domain classifier that learns directly from the latent space of visual and semantic information. We have demonstrated that our proposed approach can be combined with previous latent space learning models, such as CADA-VAE and cycle-WGAN. Our approach yielded improvements for each one of those models by automatically balancing the seen and unseen domains in benchmark experiments on four available data sets: CUB, SUN, AWA1, and AWA2.

Our experimental results show that our proposed approach has achieved state-of-the-art H-mean results for CUB, AWA1 and AWA2, and unseen accuracy for CUB, SUN, AWA1, and AWA2. In particular, our results are substantially better than the state of the art on CUB and SUN, which contain a large number of classes. On AWA1, AWA2, which are smaller data sets, our results are marginally better. Furthermore, our model produces substantial improvements in terms of AUSUC results for CUB, AWA1 and marginally better on AWA2.

As stated previously, our domain classification learns to discriminate between samples from the seen and unseen domains. We observe that the improvement of CADA-VAE and cycle-WGAN are different. The CADA-VAE model tends to improve in terms of the unseen domain when the DC is applied. Whereas cycle-WGAN tends to improve in terms of the seen domain. On one hand, we note that the training strategy for both models follows different guidelines, VAE and GAN. On the other hand, our model does not impose direct constraints in order to optimise GZSL metrics, such as accuracy or H-mean. In fact, we believe that these aspects are the main factors for the contrasting outcomes for CADA-VAE and cycle-WGAN models. With that in mind, we believe that the differences between these two data augmentation approaches should be studied in future generalised zero-shot learning research.

In the future, we intend to further study the reasons behind the performance difference observed between the data sets. Moreover, we also plan to develop a more extensive framework that can incorporate domain classification for approaches that do not rely on latent space learning.

TABLE III: GZSL results using per-class average top-1 accuracy on the test sets of unseen classes \mathcal{Y}^U , seen classes \mathcal{Y}^S , and H-mean result H ; and ZSL results on the unseen classes exclusively – all results shown in percentage. The results from previously proposed methods in the field were extracted from [1]. The highlighted values represent the best ones in each column. The methods below the double horizontal line represent the ones that use the semantic vectors from unseen classes during training. The notation * represents the results that we reproduced, and results represented with – were not available in the literature, or hyper-parameters were not given.

Classifier	CUB			SUN			AWA1			AWA2		
	\mathcal{Y}^S	\mathcal{Y}^U	H									
Semantic approach												
SJE [24]	59.2	23.5	33.6	30.5	14.7	19.8	74.6	11.3	19.6	73.9	8.0	14.4
ALE [25]	62.8	23.7	34.4	33.1	21.8	26.3	76.1	16.8	27.5	81.8	14.0	23.9
LATEM [26]	57.3	15.2	24.0	28.8	14.7	19.5	71.7	7.3	13.3	77.3	11.5	20.0
ESZSL [27]	63.8	12.6	21.0	27.9	11.0	15.8	75.6	6.6	12.1	77.8	5.9	11.0
SYNC [12]	70.9	11.5	19.8	43.3	7.9	13.4	87.3	8.9	16.2	90.5	10.0	18.0
DEWISE [2]	53.0	23.8	32.8	27.4	16.9	20.9	68.7	13.4	22.4	74.7	17.1	27.8
Generative approach												
SAE [15]	18.0	8.8	11.8	54.0	7.8	13.6	77.1	1.8	3.5	82.2	1.1	2.2
f-CLSWGAN [11]	57.7	43.7	49.7	36.6	42.6	39.4	61.4	57.9	59.6	68.9	52.1	59.4
cycle-WGAN [3]	60.3	46.0	52.2	33.1	48.3	39.2	63.5	56.4	59.7	–	–	–
CADA-VAE [4]	53.5	51.6	52.4	35.7	47.2	40.6	72.8	57.3	64.1	75.0	55.8	63.9
CADA-VAE [4]*	57.2	48.4	52.4	36.8	45.1	40.6	76.6	55.0	64.1	75.3	55.5	63.9
Domain Classification												
CMT [6]	49.8	7.2	12.6	21.8	8.1	11.8	87.6	0.9	1.8	90.0	0.5	1.0
DAZSL [5]	56.9	47.6	51.8	37.2	45.6	41.4	76.9	54.7	63.9	–	–	–
cycle-WGAN + DC (ours)	61.9	45.9	52.7	39.3	41.3	40.3	68.6	53.4	60.0	–	–	–
CADA-VAE + DC (ours)	52.4	52.9	52.6	34.0	50.7	40.7	72.6	57.5	64.2	74.9	57.0	64.3

ACKNOWLEDGMENT

Supported by Australian Research Council through grants *DP180103232*, *CE140100016* and *FL130100102*. R.F. would like to acknowledge the donation of a TitanXp by Nvidia. G.C. acknowledges the support by the Alexander von Humboldt-Stiftung for the renewed research stay sponsorship.

REFERENCES

- [1] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata, “Zero-shot learning - A comprehensive evaluation of the good, the bad and the ugly,” *CoRR*, vol. abs/1707.00600, 2017.
- [2] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov *et al.*, “DeVise: A deep visual-semantic embedding model,” in *Advances in neural information processing systems*, 2013, pp. 2121–2129.
- [3] R. Felix, B. V. Kumar, I. Reid, and G. Carneiro, “Multi-modal cycle-consistent generalized zero-shot learning,” in *European Conference on Computer Vision*. Springer, 2018, pp. 21–37.
- [4] E. Schonfeld, S. Ebrahimi, S. Sinha, T. Darrell, and Z. Akata, “Generalized zero- and few-shot learning via aligned variational autoencoders,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8247–8255.
- [5] Y. Atzmon and G. Chechik, “Adaptive confidence smoothing for generalized zero-shot learning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 11 671–11 680.
- [6] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng, “Zero-shot learning through cross-modal transfer,” in *Advances in Neural Information Processing Systems*, 2013, pp. 935–943.
- [7] C. H. Lampert, H. Nickisch, and S. Harmeling, “Learning to detect unseen object classes by between-class attribute transfer,” in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, June 2009, pp. 951–958.
- [8] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona, “Caltech-ucsd birds 200,” 2010.
- [9] W.-L. Chao, S. Changpinyo, B. Gong, and F. Sha, “An empirical study and analysis of generalized zero-shot learning for object recognition in the wild,” in *European Conference on Computer Vision*. Springer, 2016, pp. 52–68.
- [10] C. H. Lampert, H. Nickisch, and S. Harmeling, “Attribute-based classification for zero-shot visual object categorization,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 3, pp. 453–465, 2014.
- [11] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata, “Feature generating networks for zero-shot learning,” *arXiv*, 2017.
- [12] S. Changpinyo, W.-L. Chao, B. Gong, and F. Sha, “Synthesized classifiers for zero-shot learning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5327–5336.
- [13] V. K. Verma, G. Arora, A. Mishra, and P. Rai, “Generalized zero-shot learning via synthesized examples,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [15] T. X. Elyor Kodirov and S. Gong, “Semantic Autoencoder for Zero-shot Learning,” *IEEE CVPR 2017*, 2017.
- [16] A. Mishra, S. Krishna Reddy, A. Mittal, and H. A. Murthy, “A generative model for zero shot learning using conditional variational autoencoders,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 2188–2196.
- [17] Y. Zhu, M. Elhoseiny, B. Liu, X. Peng, and A. Elgammal, “A generative adversarial approach for zero-shot learning from noisy texts,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1004–1013.
- [18] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, “Sun database: Large-scale scene recognition from abbey to zoo,” in *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*. IEEE, 2010, pp. 3485–3492.
- [19] S. Reed, Z. Akata, H. Lee, and B. Schiele, “Learning deep representations of fine-grained visual descriptions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 49–58.
- [20] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv*, 2014.

- [21] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. icml*, vol. 30, no. 1, 2013, p. 3.
- [22] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 807–814.
- [23] Y. Gal, J. Hron, and A. Kendall, "Concrete dropout," in *Advances in Neural Information Processing Systems*, 2017, pp. 3581–3590.
- [24] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele, "Evaluation of output embeddings for fine-grained image classification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 2927–2936.
- [25] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, "Label-embedding for image classification," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 7, pp. 1425–1438, 2016.
- [26] Y. Xian, Z. Akata, G. Sharma, Q. Nguyen, M. Hein, and B. Schiele, "Latent embeddings for zero-shot classification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 69–77.
- [27] B. Romera-Paredes and P. Torr, "An embarrassingly simple approach to zero-shot learning," in *International Conference on Machine Learning*, 2015, pp. 2152–2161.