



From Classical to Generalized Zero-Shot Learning: a Simple Adaptation Process

Yannick Le Cacheux, Hervé Le Borgne and Michel Crucianu

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

October 29, 2018

From Classical to Generalized Zero-Shot Learning: a Simple Adaptation Process

Yannick Le Cacheux¹, Hervé Le Borgne¹, and Michel Crucianu²

¹ CEA LIST

{yannick.lecacheux,herve.le-borgne}@cea.fr

² CEDRIC – CNAM

michel.crucianu@cnam.fr

Abstract. Zero-shot learning (ZSL) is concerned with the recognition of previously *unseen* classes. It relies on additional semantic knowledge for which a mapping can be learned with training examples of *seen* classes. While classical ZSL considers the recognition performance on unseen classes only, generalized zero-shot learning (GZSL) aims at maximizing performance on both seen and unseen classes. In this paper, we propose a new process for training and evaluation in the GZSL setting; this process addresses the gap in performance between samples from unseen and seen classes by penalizing the latter, and enables to select hyper-parameters well-suited to the GZSL task. It can be applied to any existing ZSL approach and leads to a significant performance boost: the experimental evaluation shows that GZSL performance, averaged over eight state-of-the-art methods, is improved from 28.5 to 42.2 on CUB and from 28.2 to 57.1 on AWA2.

Keywords: Zero-Shot Learning · Multimodal Classification

1 Introduction

Zero-shot learning (ZSL) [14, 16, 17] aims to recognize classes for which no training example is available. This is often achieved by relying on additional semantic knowledge, consisting for example in vectors of attributes. During training, a relation between visual features and semantic attributes is learned from training examples belonging to the *seen* classes, for which both modalities (visual and semantic) are available. This model is then applied in the testing phase on examples from *unseen* classes, for which no visual instance was available during training. Predictions on these classes can thus be made on the basis of the inferred relation between visual and semantic features.

In classical ZSL, the test set only contains examples from the novel, unseen classes, and these classes alone can be predicted. Although this setting has enabled significant progress in methods linking visual content to semantic information in the last few years [27], it is hardly realistic. It seems much more reasonable to assume that objects which are to be classified can belong to either a seen class or an unseen class, since in real-life use-cases one could legitimately

want to recognize both former and novel classes. This setting is usually referred to as generalized zero-shot learning (GZSL).

However, recent work shows that a direct use of a ZSL model in a GZSL setting usually leads to unsatisfactory results. Indeed, in addition to the number of candidate classes being higher due to the presence of the seen classes among them, most samples from unseen classes are incorrectly classified as belonging to one of the seen classes [6]. Different methods have been proposed to measure this discrepancy, such as the area under the curve representing all the possible trade-offs between the accuracies on samples from seen classes versus samples from unseen classes [6], or their harmonic mean [27] to penalize models with strong imbalance between the two. While these proposals only measure the extent of the problem, we aim to explicitly address this issue in addition to quantifying its impact.

The main contribution of this paper is a new process for training and evaluating models in a GZSL setting. In accordance with recent studies, we show that the application of a ZSL model “out of the box” gives results that are far from optimal in the GZSL context. We demonstrate how two simple techniques – the calibration of similarities and the use of appropriately balanced regularization – can dramatically improve the performance of most models. The final score for the GZSL task can thus be increased up to a factor of two, with no change regarding the underlying hypotheses of the GZSL task or the data available at any given time, which means that our process is applicable to any ZSL model. We also provide new insights on the reasons why these two techniques are relevant and on the fundamental differences between samples from seen and unseen classes.

We extensively evaluate these techniques on several recent ZSL methods. For sanity-check, we independently reproduce results obtained in the literature before applying our process. We find that some models show a variability in performance with respect to their random initialization, so measures averaged over several runs should be preferred. We eventually find that, with fair comparison under unbiased conditions as enabled by our process, a regularized linear model can give results close to or even better than the state-of-the-art.

2 Related work

An early rigorous definition and evaluation of GZSL was put forward in [6]. The authors argue that this setting is more realistic than ZSL and highlight the gap between accuracies on seen and unseen classes when labels from *all* classes can be predicted (denoted respectively $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$, and formally defined in Sec. 3.1). They also introduce the idea of calibration to address this issue and suggest a new metric for GZSL, *Accuracy Under Seen-Unseen Curve (AUSUC)*, which measures the trade-off between the two accuracies but does not directly provide the expected performance in real use-cases.

An extensive evaluation of recent ZSL methods with a common protocol is provided in [27], both in ZSL and GZSL settings. The authors use a different




Fig. 1. Illustration of how the regularization parameter λ affects the accuracies on samples from seen and unseen classes $A_{U \rightarrow C_u}$ and $A_{S \rightarrow C_s}$ (see Sec. 3.1) as measured on CUB [23] (left) and AwA2 [25] (right). Optimal regularization is not the same in a ZSL setting, where performance is measured by $A_{U \rightarrow C_u}$ (red dotted line), and in a GZSL setting, where it is measured by the harmonic mean of $A_{U \rightarrow C}$ and $A_{S \rightarrow C}$ (black dotted line).

metric for GZSL, the harmonic mean between $A_{U \rightarrow C}$ and $A_{S \rightarrow C}$, which does not directly quantify the trade-off between accuracies but better estimates the practical performance of a given model. However, they do not explicitly address the gap between similarities evaluated on seen and unseen classes [6], which has a significant impact on the final performance as we show in Sec. 4.3.

Further GZSL results based on the harmonic mean metric are provided in [4, 13, 26]. All three methods rely on generators of artificial training examples from unseen classes. However, these methods assume that a semantic description of all unseen classes is available during training. This assumption is not necessarily met in practice and makes the inclusion of additional unseen classes more difficult.

Transductive ZSL methods [9, 11, 19] also assume that additional information, taking the form of unlabeled samples from unseen classes, is available during training. This can naturally lead to improved performance. In this article, we make none of these assumptions and consider that *no* information regarding unseen classes is available at training time.

3 Proposed approach

3.1 Problem statement

We denote by \mathcal{C}_s the set of classes *seen* during training and by \mathcal{C}_u the set of *unseen* classes. We define $\mathcal{C} = \mathcal{C}_s \cup \mathcal{C}_u$, with $\mathcal{C}_s \cap \mathcal{C}_u = \emptyset$. During the training phase, we consider N^{tr} training samples consisting of D -dimensional visual features $\mathbf{X}^{tr} = (\mathbf{x}_1^{tr}, \dots, \mathbf{x}_{N^{tr}}^{tr})^\top \in \mathbb{R}^{N^{tr} \times D}$ and corresponding labels $\mathbf{y}^{tr} = (y_1^{tr}, \dots, y_{N^{tr}}^{tr})^\top \in \mathcal{C}_s^{N^{tr}}$, as well as K -dimensional semantic *class prototypes* noted by $\mathbf{S}^{tr} = (\mathbf{s}_1^{tr}, \dots, \mathbf{s}_{|\mathcal{C}_s|}^{tr})^\top \in \mathbb{R}^{|\mathcal{C}_s| \times K}$. We seek to learn a function $f :$

$\mathbb{R}^D \times \mathbb{R}^K \rightarrow \mathbb{R}$ assigning a similarity score to each pair composed of a visual feature vector and a semantic representation so as to minimize the following regularized loss:

$$\frac{1}{N^{tr}} \sum_{n=1}^{N^{tr}} \sum_{c=1}^{|\mathcal{C}_s|} L(f(\mathbf{x}_n^{tr}; \mathbf{s}_c^{tr}), y_n^{tr}) + \lambda \Omega[f] \quad (1)$$

where L is the loss function and Ω the regularization term weighted by λ . During the testing phase, we consider N^{te} unlabeled visual samples $\mathbf{X}^{te} = (\mathbf{x}_1^{te}, \dots, \mathbf{x}_{N^{te}}^{te})^\top \in \mathbb{R}^{N^{te} \times D}$ and class prototypes for *candidate classes*. In a ZSL setting, the candidate classes \mathcal{C}^{te} are the unseen classes such that $\mathbf{S}^{te} \in \mathbb{R}^{|\mathcal{C}_u| \times K}$. In a GZSL setting, classes to be predicted can be in either \mathcal{C}_u or \mathcal{C}_s , such that $\mathcal{C}^{te} = \mathcal{C}$ and the class prototypes are $\mathbf{S}^{te} = (\mathbf{s}_1^{te}, \dots, \mathbf{s}_{|\mathcal{C}_s|}^{te}, \mathbf{s}_{|\mathcal{C}_s|+1}^{te}, \dots, \mathbf{s}_{|\mathcal{C}_s|+|\mathcal{C}_u|}^{te})^\top \in \mathbb{R}^{|\mathcal{C}| \times K}$. In both cases, given a function \hat{f} learned in the training phase, we want to estimate a prediction \hat{y} for a visual testing sample \mathbf{x} such that:

$$\hat{y} = \operatorname{argmax}_{c \in \mathcal{C}^{te}} \hat{f}(\mathbf{x}; \mathbf{s}_c^{te}) \quad (2)$$

In classical ZSL, performance is measured by the *accuracy of unseen classes among unseen classes*, noted $A_{\mathcal{U} \rightarrow \mathcal{C}_u}$, while in GZSL we are interested in the *accuracy of unseen classes among all classes* and the *accuracy of seen classes among all classes*, noted respectively $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$ as in [6]. $A_{\mathcal{S} \rightarrow \mathcal{C}_s}$ is similarly defined.

3.2 Calibration and GZSL split

As evidenced by [6], when a ZSL model is applied in a GZSL setting, $A_{\mathcal{S} \rightarrow \mathcal{C}}$ is usually significantly higher than $A_{\mathcal{U} \rightarrow \mathcal{C}}$. This is because most samples from unseen classes are incorrectly classified into one of the seen classes. To address this, a *calibration factor* γ is added in [6] to penalize seen classes. Eq. (2) then becomes:

$$\hat{y} = \operatorname{argmax}_{c \in \mathcal{C}^{te}} \left(\hat{f}(\mathbf{x}; \mathbf{s}_c^{te}) - \gamma \mathbf{1}[c \in \mathcal{C}_s] \right) \quad (3)$$

where $\mathbf{1}[\cdot]$ is an indicator function.

The *Accuracy Under Seen-Unseen Curve (AUSUC)* metric also proposed in [6] is defined as the area under the curve representing $A_{\mathcal{S} \rightarrow \mathcal{C}}$ versus $A_{\mathcal{U} \rightarrow \mathcal{C}}$ when γ varies from $-\infty$ to $+\infty$, which shows the trade-off between the two.

Instead of computing a metric involving all possible trade-offs between $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$, we look for a single specific value of γ , corresponding to the best compromise between the two as measured by the harmonic mean of $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$ [27]. We propose to determine the optimal value of γ with a cross-validation specific to GZSL. Usually in machine learning a dataset is divided at random into three parts: a training, a validation and a testing set. In classical ZSL, this splitting process is done with respect to the classes as opposed to the samples: a set of classes is used for training, a disjoint set for validation and a final mutually disjoint set for testing. In GZSL, a fraction (usually 20%) of the

samples from the validation and training sets are kept for testing time to be used as test samples from seen classes. We refer to this set as the *seen test set*. Note that here *seen* only indicates that these samples belong to seen classes, not that they have been used during training. To be able to cross-validate parameters for GZSL, we further keep an additional 20% of the remaining training set to be used as samples from seen classes when cross-validating parameters; we refer to this set as the *seen validation set*. Fig. 2 illustrates this partitioning.




Fig. 2. Illustration of the different splits. Each column is a class and each cell is a sample. In this example there are 20 different classes with 10 samples per class. Five classes are used for testing, five other for validation and the remaining ten for training. Among the samples from the validation and training classes, 20% are kept for testing (*seen test set*) and 20% more samples from training classes are kept for validation (*seen validation set*).

To determine the optimal value of γ we first train a model on the GZSL training set. We then use the (GZSL) validation set and the seen validation set to compute the GZSL metric (the harmonic mean) and keep the value γ^* that maximizes this metric. The ZSL model is subsequently re-trained on the training, validation and seen validation sets, then class similarities are computed for the test set. The value γ^* is subtracted from the similarities of seen classes and the resulting similarities are used to compute the final GZSL score.

3.3 Regularization for GZSL

The usual approach to optimize the regularized loss (Eq. (1)) in GZSL consists in using the value of λ determined on the ZSL task. We argue here that this is unlikely to be optimal and provide some insight to justify our position. Then, we propose a simple method to determine a better value of λ to improve performance in GZSL.

Figure 1 shows $A_{\mathcal{U} \rightarrow \mathcal{C}_u}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}_s}$ as a function of λ for a regularized linear model (ridge regression [3, 24]), measured on the first validation splits of the *proposed splits* of [27] on CUB [23] and AWA2 [25].

In each case, there is a value of λ that maximizes the ZSL score $A_{\mathcal{U} \rightarrow \mathcal{C}_u}$, indicated by the red dotted vertical line, that we note λ_{ZSL}^* . The overall tendency for $A_{\mathcal{S} \rightarrow \mathcal{C}_s}$ is to decrease as λ increases. This is not a concern for the ZSL task, since it only considers samples from unseen classes. However, for the GZSL




Fig. 3. (a) MSE of predicted attributes (averaged over attributes and samples) as a function of the regularization parameter λ ; (b) Illustration of the bias-variance decomposition.

task, we want the best trade-off between $A_{U \rightarrow C}$ and $A_{S \rightarrow C}$. Note that $A_{U \rightarrow C} \leq A_{U \rightarrow C_u}$ and $A_{S \rightarrow C} \leq A_{S \rightarrow C_s}$, with equality only if we are able to perfectly distinguish samples from seen and unseen classes. It follows that λ_{GZSL}^* , the value of λ that maximizes the GZSL score, is not necessarily the same as λ_{ZSL}^* : a small decrease from λ_{ZSL}^* can significantly increase $A_{S \rightarrow C_s}$ while only slightly penalizing $A_{U \rightarrow C_u}$. This has a similar impact on the maximum values obtainable by $A_{S \rightarrow C}$ and $A_{U \rightarrow C}$, and can ultimately improve the GZSL score. We quantify in Sec. 4.3 the gains attributed to the use of λ_{GZSL}^* .

The reason why λ affects $A_{U \rightarrow C_u}$ and $A_{S \rightarrow C_s}$ in this way can be explained with the bias-variance decomposition. For regression, we generally assume that we are given a dataset $\mathcal{D} = (\mathbf{X}, \mathbf{t})$, with samples (\mathbf{x}_n, t_n) independently drawn from a joint distribution $p(\mathbf{x}, t)$, such that $p(t|\mathbf{x}) = \mathcal{N}(t|h(\mathbf{x}), \sigma^2)$, where h is the true dependence. For a prediction function \hat{h} estimated from \mathcal{D} we can then write the expected loss on a new pair (\mathbf{x}, t) as:

$$\mathbb{E}_{\mathcal{D}, \mathbf{x}, t}[(t - \hat{h}(\mathbf{x}))^2] = \sigma^2 + \mathbb{E}_{\mathcal{D}, \mathbf{x}}[(h(\mathbf{x}) - \hat{h}(\mathbf{x}))^2] + \text{var}_{\mathcal{D}, \mathbf{x}}[\hat{h}(\mathbf{x})] \quad (4)$$

where the first term is the intrinsic noise of the dataset, the second is the (squared) bias of the predictor and the third is the variance in the estimation of the predictor. It can be shown [3, 24] that for ridge regression the bias increases and the variance decreases with the regularization parameter λ , as illustrated in Fig. 3(b).

In the case of ZSL, \mathbf{x} corresponds to visual samples and t to attribute(s) to be estimated from \mathbf{x} . The variance comes from both the differences between samples from the same class (intra-class variance) and from the differences between classes (inter-class variance). Intra-class variance is usually significantly smaller than inter-class variance in ZSL; therefore, most of the variance in Eq. (4) can be attributed to the choice of training classes C_s . For samples from unseen classes,

the bias-variance decomposition applies and there exists a λ corresponding to the best trade-off between the two. This is evidenced in Fig. 3(a), where the red curve shows the Mean Squared Error (MSE) in the predictions of attributes from unseen classes as a function of λ , for a regularized linear model on the first validation split of AwA2 [25].

For a sample from a seen class, the variance attributable to the choice of the training classes is much smaller since, by definition, the seen class is present in the training dataset. This allows to better estimate attributes from seen classes and most of the expected error therefore comes from the intrinsic noise and the bias. Thus, the expected error mostly increases with λ , as evidenced by the blue curve in Fig. 3(a). If we plausibly assume that the accuracy of predictions for samples from a given class depends on how well we estimate their attributes, this explains both why predictions are better for samples from seen classes than from unseen classes and why their behavior with respect to λ is different.

We then suggest the following procedure to select the optimal value of λ : we repeat the protocol described in Sec. 3.2 for selecting γ^* and we take the value of λ which gives the best result for the harmonic mean between $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$ on the validation set *after* having subtracted γ^* from the similarities of seen classes. The rest of the process is identical: we retrain the ZSL model on the training, validation and seen validation sets with the hyperparameter λ_{GZSL}^* that we just determined, we compute similarities for the test set, subtract γ^* from the similarities of seen classes and compute the resulting GZSL score.

4 Experimental evaluation

4.1 Methods

We independently reimplemented six methods frequently cited in the literature to evaluate them with our protocol: ALE [1], DeVISE [8], SJE [2], Sync [5], ESZSL [20] and SAE [12].

In addition, we also evaluate two simple linear models. $\text{Linear}_{\mathcal{V} \rightarrow \mathcal{S}}$ applies a linear mapping $\mathbf{W} \in \mathbb{R}^{K \times D}$ from the visual space \mathcal{V} to the semantic space \mathcal{S} to minimize standard MSE. With $\mathbf{T}^{tr} = (\mathbf{s}_{y_1^{tr}}^{tr}, \dots, \mathbf{s}_{y_N^{tr}}^{tr})^\top \in \mathbb{R}^{N^{tr} \times K}$ the matrix whose rows correspond to the class prototypes associated to each training sample based on its label, the loss function can be formulated as:

$$\frac{1}{N^{tr}} \|\mathbf{X}^{tr} \mathbf{W}^\top - \mathbf{T}^{tr}\|_F^2 + \lambda \|\mathbf{W}\|_F^2 \quad (5)$$

$\text{Linear}_{\mathcal{S} \rightarrow \mathcal{V}}$ is based on [21] where the authors argue that using the semantic space as the embedding space reduces the variance of the projected points and thus aggravates the hubness problem [18]. They suggest instead to project semantic class prototypes onto the visual space and to compute similarities in this space. Keeping $\mathbf{W} \in \mathbb{R}^{K \times D}$ as our linear mapping, we formulate the loss function as:

$$\frac{1}{N^{tr}} \|\mathbf{X}^{tr} - \mathbf{T}^{tr} \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2 \quad (6)$$

We can easily obtain closed-form solutions for the two models from the objective functions (5) and (6). For the Linear $\mathcal{V} \rightarrow \mathcal{S}$ model we have $\mathbf{W} = \mathbf{T}^{tr\top} \mathbf{X}^{tr} (\mathbf{X}^{tr\top} \mathbf{X}^{tr} + \lambda N^{tr} \mathbf{I}_D)^{-1}$ (7) and for the Linear $\mathcal{S} \rightarrow \mathcal{V}$ model $\mathbf{W} = (\mathbf{T}^{tr\top} \mathbf{T}^{tr} + \lambda N^{tr} \mathbf{I}_K)^{-1} \mathbf{T}^{tr\top} \mathbf{X}^{tr}$ (8).

4.2 Experimental setting

Datasets We perform our experiments on two standard datasets for ZSL: Caltech-UCSD-Birds 200-2011 (CUB) [23] and Animals with Attributes2³ (AwA2) [25]. CUB is a fine-grained dataset composed of 11788 pictures of birds from 200 species (*black footed albatross*, ..., *common yellowthroat*). It comes with 312-dimensional binary attributes for each picture, that are averaged by class to obtain semantic class prototypes. AwA2 is a coarse-grained dataset comprising 37322 pictures of 50 animal species (*antelope*, ..., *zebra*). For each class, 85-dimensional attributes are provided.

Splits The best performing ZSL methods usually rely on visual features obtained with deep neural networks pre-trained on ImageNet [7], such as GoogLeNet [22] or ResNet [10]. As evidenced by [27], this induces a huge bias for ZSL datasets whose classes are not disjoint from categories of ImageNet, as is the case with AwA2, since test classes cannot be considered truly unseen. We therefore adopt the approach of [27] and use their terms *Standard Split* (S.S.) for the split widely used in the literature and *Proposed Split* (P.S.) for the split they introduce. The training and validation splits are further divided for GZSL as described in Sec. 3.2.

Settings Attributes are normalized such that each class prototype has unit ℓ_2 norm. We use the 101-layered ResNet [10] pre-trained on ImageNet [7] as visual features extractor, keeping the $D = 2048$ activations of the last pooling units.

Metrics For ZSL, we evaluate the accuracy of samples from unseen classes among unseen classes $A_{\mathcal{U} \rightarrow \mathcal{C}_u}$. There are two possible ways to define accuracy: most of the literature uses *per sample* accuracy, defined as $100 \cdot \frac{1}{N^{te}} \sum_{n=1}^{N^{te}} \mathbb{1}[\hat{y}(\mathbf{x}_n^{te}) = y_n^{te}]$, while in [27] it is argued that *per class* accuracy, defined as $100 \cdot \frac{1}{|\mathcal{C}^{te}|} \sum_{c \in \mathcal{C}^{te}} \frac{1}{|\{n|y_n=c\}|} \sum_{y_n=c} \mathbb{1}[\hat{y}(\mathbf{x}_n^{te}) = y_n^{te}]$, better takes class imbalance into account.

We report per class accuracy for fair comparison with the extensive results of [27]. Nonetheless, to enable comparison with the rest of the literature, we also provide per sample accuracy results in Table 3. For GZSL we compute the harmonic mean between $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$, defined as $\frac{2 \cdot A_{\mathcal{U} \rightarrow \mathcal{C}} \cdot A_{\mathcal{S} \rightarrow \mathcal{C}}}{A_{\mathcal{U} \rightarrow \mathcal{C}} + A_{\mathcal{S} \rightarrow \mathcal{C}}}$.

Accuracy is again assumed to be per class unless otherwise stated.

³ AwA2 was recently proposed in [25] as a replacement for the Animals with Attributes (AwA) dataset [15] whose images are not publicly available.

Table 1. ZSL score: per-class accuracy $A_{U \rightarrow C_u}$, as reported in [27] and independently reproduced. S.S.: Standard Split, P.S.: Proposed Split [27]. Averaged over 5 runs.

Method	CUB [23]				AwA2 [25]			
	Reported in [27]		Reproduced		Reported in [27]		Reproduced	
	S.S.	P.S.	S.S.	P.S.	S.S.	P.S.	S.S.	P.S.
Linear $\nu \rightarrow s$	<i>n/a</i>	<i>n/a</i>	41.0 \pm 0.0	41.8 \pm 0.0	<i>n/a</i>	<i>n/a</i>	68.2 \pm 0.0	49.7 \pm 0.0
Linear $s \rightarrow \nu$	<i>n/a</i>	<i>n/a</i>	56.0 \pm 0.0	53.5 \pm 0.0	<i>n/a</i>	<i>n/a</i>	85.5 \pm 0.0	68.9 \pm 0.0
ALE [1]	53.2	54.9	54.8 \pm 0.8	54.0 \pm 1.2	80.3	62.5	80.3 \pm 2.2	62.9 \pm 2.3
DeViSE [8]	53.2	52.0	52.5 \pm 0.9	52.6 \pm 1.3	68.6	59.7	76.6 \pm 1.6	62.1 \pm 1.6
SJE [2]	55.3	53.9	53.8 \pm 2.3	49.2 \pm 1.4	69.5	61.9	80.4 \pm 2.9	62.2 \pm 1.2
ESZSL [20]	55.1	53.9	34.9 \pm 0.0	34.9 \pm 0.0	75.6	58.6	70.5 \pm 0.0	50.8 \pm 0.0
Sync [5]	54.1	55.6	56.4 \pm 0.9	54.8 \pm 0.6	71.2	46.6	65.6 \pm 0.8	58.1 \pm 0.8
SAE [12]	33.4	33.3	56.2 \pm 0.0	53.3 \pm 0.0	80.7	54.1	81.1 \pm 0.0	62.8 \pm 0.0

4.3 Results

Table 2. GZSL score (harmonic mean of $A_{U \rightarrow C}$ and $A_{S \rightarrow C}$, per class accuracy) with and without calibration and GZSL regularization. On Proposed Split [27], averaged over 5 runs.

Method	CUB [23]			AwA2 [25]		
	Reported in [27]	Ours		Reported in [27]	Ours	
with calibration	-	-	✓ ✓ ✓	-	-	✓ ✓ ✓
with λ_{GZSL}^*	-	-	- ✓ ✓	-	-	- ✓ ✓
Linear $\nu \rightarrow s$	<i>n/a</i>	18.2 34.3 35.5		<i>n/a</i>	8.3 47.3 48.1	
Linear $s \rightarrow \nu$	<i>n/a</i>	32.5 41.9 43.5		<i>n/a</i>	44.3 62.7 64.0	
ALE [1]	34.4	35.6 45.1 46.2		23.9	26.9 55.8 55.8	
DeViSE [8]	32.8	35.1 43.6 43.4		27.8	17.4 54.6 54.6	
SJE [2]	33.6	29.7 41.2 44.2		14.4	28.9 58.2 59.0	
ESZSL [20]	21.0	17.9 33.7 33.9		11.0	39.9 53.6 53.7	
Sync [5]	19.8	33.2 46.2 47.6		18.0	30.6 61.0 61.0	
SAE [12]	13.6	25.7 43.1 43.1		2.2	29.5 60.2 60.2	
Average	25.9	28.5 41.1 42.2		16.2	28.2 56.7 57.1	

We first evaluate the performances of the different methods in a classical ZSL setting. Table 1 shows the average per class accuracy measured on testing sets of the Standard Splits (S.S.) and the Proposed Splits (P.S.) [27] of CUB [23] and AwA2 [25]. We report the average score and the standard deviation over 5 runs with different random initializations. We also report the results from [27]. We see that some methods such as SJE [2] have high variability with respect to the initialization; for such methods, it is good practice to report average results since a single test run may not be representative of the true performance of the model. On the other hand, methods with closed-form or deterministic solutions

Table 3. ZSL and GZSL scores with 10-crop features, evaluated with per class (p.c.) and per sample (p.s.) accuracies. With calibration and λ_{GZSL}^* . On P.S. [27], averaged over 5 runs.

Method	CUB [23]				AwA2 [25]			
	ZSL		GZSL		ZSL		GZSL	
	Acc. p.c.	Acc. p.s.	H. p.c.	H. p.s.	Acc. p.c.	Acc. p.s.	H. p.c.	H. p.s.
Linear $\nu \rightarrow \mathcal{S}$	45.6	45.6	39.8	39.8	51.0	43.6	49.0	45.6
Linear $\mathcal{S} \rightarrow \nu$	57.1	57.2	47.7	48.0	70.4	69.3	65.1	68.7
ALE [1]	57.4	57.5	49.2	49.3	63.0	61.1	56.9	55.5
DeViSE [8]	52.9	52.9	42.4	42.5	63.1	62.2	55.0	50.6
SJE [2]	51.9	52.1	46.7	46.9	63.8	61.6	59.4	57.6
ESZSL [20]	39.0	38.8	38.7	38.6	52.6	51.9	54.4	57.9
Sync [5]	57.5	57.6	48.9	49.1	59.3	56.1	62.6	63.2
SAE [12]	56.1	56.2	46.3	46.6	63.5	65.4	62.3	63.6

such as the Linear $\nu \rightarrow \mathcal{S}$, Linear $\mathcal{S} \rightarrow \nu$, ESZSL [20] or SAE [12] are not dependent on the initialization and thus have a standard deviation of 0.

Most of the reproduced scores are consistent with [27], with two notable exceptions: first, a significant increase in performance is observed with SAE [12] and can be explained by the fact that similarities are computed in the visual space, with results close to those of the Linear $\mathcal{S} \rightarrow \nu$ model (results are close to those of Linear $\nu \rightarrow \mathcal{S}$ when similarities are computed in the semantic space). Second, the score for ESZSL [20] is significantly lower than reported in [27]. We found that the use of non-normalized attributes enables to reach performances comparable with [27], but we could not reproduce the reported results for ESZSL [20] with normalized attributes. For the sake of consistency, we chose to report results obtained with normalized attributes.

Table 2 shows results for GZSL. We measure the harmonic mean between per class accuracies $A_{\mathcal{U} \rightarrow \mathcal{C}}$ and $A_{\mathcal{S} \rightarrow \mathcal{C}}$ on the testing set of the Proposed Split [27]. We evaluate three settings: a ZSL model applied directly in a GZSL setting, i.e. with no calibration and a regularization specific to ZSL (λ_{ZSL}^*) as opposed to GZSL (λ_{GZSL}^*); a ZSL model with calibration and ZSL regularization λ_{ZSL}^* ; and a ZSL model with calibration and regularization λ_{GZSL}^* specific to the GZSL problem. We report the average score over 5 runs; standard deviations are available in the supplementary material. We also report the results from [27], which correspond to the setting with no calibration and no λ_{GZSL}^* . We can see that the calibration process significantly improves GZSL performance: in our experiments, the average score for all models improves from 28.5 with no calibration to 41.1 with calibration on CUB, and from 28.2 to 56.7 on AwA2. It is worth noting that the lowest score with calibration is close to or higher than the highest score without. The use of a regularization parameter specific to the GZSL task can lead to an additional improvement in performance. In some cases, the optimal λ is the same for the ZSL task and the GZSL task on the validation set, leading to no additional improvement over the score with calibration. However,

every time they are different, λ_{GZSL}^* is smaller than λ_{ZSL}^* , as expected from the results in Sec. 3.3. The only exception is with DeVISE [8] on CUB: a λ_{GZSL}^* higher than λ_{ZSL}^* was selected during cross-validation, probably due to random noise, resulting in a slightly lower final GZSL score.

Table 3 shows results with improved visual features; each original 256×256 image is cropped into ten 224×224 images: one in each corner and one in the center for both the original image and its horizontal symmetry. The ResNet features of the resulting images are averaged to obtain a 2048-dimensional vector. We report results for ZSL ($\mathcal{A}_{C \rightarrow C_u}$, abbreviated *Acc.*) and GZSL (using the harmonic mean metric, abbreviated *H.*) on the testing set of the Proposed Split [27]. In order to facilitate fair comparison with the rest of the literature, both per class (*p.c.*) and per sample (*p.s.*) metrics are reported. Results with 10-cropped visual features are almost always better than the results with standard visual features in Table 2. The per sample metrics are on average not very different from the per class metrics. This is not surprising since classes in both CUB and AWA2 are fairly balanced.

5 Conclusion

We proposed a simple process for applying ZSL methods in a GZSL setting. This process is based on the empirical observation that ZSL models perform differently on samples from seen and unseen classes. We provided insights about why this should be expected and suggested steps to overcome these problems. Through extensive experiments, we showed that this process enables significant improvements in performance for many existing ZSL methods. Finally, we provided results under optimal conditions for these methods with different metrics to support fair comparison with the rest of the state-of-the-art.

References

1. Akata, Z., Perronnin, F., Harchaoui, Z., Schmid, C.: Label-embedding for image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(7), 1425–1438 (2016)
2. Akata, Z., Reed, S., Walter, D., Lee, H., Schiele, B.: Evaluation of output embeddings for fine-grained image classification. In: *Proc. CVPR 2015*. pp. 2927–2936. IEEE (2015)
3. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer (2006)
4. Bucher, M., Herbin, S., Jurie, F.: Generating visual representations for zero-shot classification. In: *ICCV Workshops: TASK-CV*. IEEE (2017)
5. Changpinyo, S., Chao, W.L., Gong, B., Sha, F.: Synthesized classifiers for zero-shot learning. In: *Proc. CVPR 2016*. pp. 5327–5336. IEEE (2016)
6. Chao, W.L., Changpinyo, S., Gong, B., Sha, F.: An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In: *Proc. ECCV 2016*. pp. 52–68. Springer (2016)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Proc. CVPR 2009*. pp. 248–255. IEEE (2009)

8. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Mikolov, T., et al.: Devise: A deep visual-semantic embedding model. In: Proc. NIPS 2013. pp. 2121–2129 (2013)
9. Fu, Y., Hospedales, T.M., Xiang, T., Gong, S.: Transductive multi-view zero-shot learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **37**(11), 2332–2345 (2015)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proc. CVPR 2016. pp. 770–778. IEEE (2016)
11. Kodirov, E., Xiang, T., Fu, Z., Gong, S.: Unsupervised domain adaptation for zero-shot learning. In: Proc. CVPR 2015. pp. 2452–2460. IEEE (2015)
12. Kodirov, E., Xiang, T., Gong, S.: Semantic autoencoder for zero-shot learning. In: Proc. CVPR 2017. pp. 4447–4456. IEEE (2017)
13. Kumar Verma, V., Arora, G., Mishra, A., Rai, P.: Generalized zero-shot learning via synthesized examples. In: Proc. CVPR 2010. pp. 4281–4289. IEEE (2018)
14. Lampert, C.H., Nickisch, H., Harmeling, S.: Learning to detect unseen object classes by between-class attribute transfer. In: Proc. CVPR 2009. pp. 951–958. IEEE (2009)
15. Lampert, C.H., Nickisch, H., Harmeling, S.: Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(3), 453–465 (2014)
16. Larochelle, H., Erhan, D., Bengio, Y.: Zero-data learning of new tasks. In: AAAI. vol. 1, p. 3 (2008)
17. Palatucci, M., Pomerleau, D., Hinton, G.E., Mitchell, T.M.: Zero-shot learning with semantic output codes. In: Proc. NIPS 2009. pp. 1410–1418 (2009)
18. Radovanović, M., Nanopoulos, A., Ivanović, M.: Hubs in space: Popular nearest neighbors in high-dimensional data. *Journal of Machine Learning Research* **11**(Sep), 2487–2531 (2010)
19. Rohrbach, M., Ebert, S., Schiele, B.: Transfer learning in a transductive setting. In: Proc. NIPS 2013. pp. 46–54 (2013)
20. Romera-Paredes, B., Torr, P.: An embarrassingly simple approach to zero-shot learning. In: Proc. ICML 2015. pp. 2152–2161 (2015)
21. Shigeto, Y., Suzuki, I., Hara, K., Shimbo, M., Matsumoto, Y.: Ridge regression, hubness, and zero-shot learning. In: Proc. ECML PKDD 2015. pp. 135–151. Springer (2015)
22. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proc. CVPR 2015. pp. 1–9. IEEE (2015)
23. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 dataset (2011)
24. van Wieringen, W.N.: Lecture notes on ridge regression. arXiv preprint arXiv:1509.09169 (2015)
25. Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: Zero-shot learning – a comprehensive evaluation of the good, the bad and the ugly. arXiv preprint arXiv:1707.00600 (2017)
26. Xian, Y., Lorenz, T., Schiele, B., Akata, Z.: Feature generating networks for zero-shot learning. In: Proc. CVPR 2018. IEEE (2018)
27. Xian, Y., Schiele, B., Akata, Z.: Zero-shot learning – the good, the bad and the ugly. In: Proc. CVPR 2017. pp. 3077–3086. IEEE (2017)