

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# Semantically Grounded Visual Embeddings for Zero-Shot Learning

Shah Nawaz<sup>1\*†</sup>, Jacopo Cavazza<sup>1\*</sup>, Alessio Del Bue<sup>1,2</sup> <sup>1</sup>Pattern Analysis & Computer Vision (PAVIS) - Istituto Italiano di Tecnologia (IIT), <sup>2</sup>Visual Geometry & Modelling (VGM) - Istituto Italiano di Tecnologia (IIT)

{shah.nawaz,jacopo.cavazza,alessio.delbue}@iit.it

### Abstract

Zero-shot learning methods rely on fixed visual and semantic embeddings, extracted from independent vision and language models, both pre-trained for other large-scale tasks. This is a weakness of current zero-shot learning frameworks as such disjoint embeddings fail to adequately associate visual and textual information to their shared semantic content. Therefore, we propose to learn semantically grounded and enriched visual information by computing a joint image and text model with a two-stream network on a proxy task. To improve this alignment between image and textual representations, provided by attributes, we leverage ancillary captions to provide grounded semantic information. Our method, dubbed joint embeddings for zero-shot learning is evaluated on several benchmark datasets, improving the performance of existing state-of-the-art methods in both standard (+1.6% on aPY, +2.6% on FLO) and generalized (+2.1% on AWA2, +2.2% on CUB) zero-shot recognition.

Zero-shot learning (ZSL) aims at classifying images into new "unseen" categories at test time without having been provided any corresponding visual examples in the training stage. This is possible by taking advantage of *semantic* embeddings describing the visual categories by means of auxiliary text information i.e. attributes. In recent years, several approaches have been proposed to tackle ZSL: A metric/score can be learnt to quantify the compatibility between semantic embeddings and pre-computed visual features (visual embedding) [54]. Alternatively, visual features can be synthesized from semantic embeddings with adversarial generative training [34, 56, 60]. In these terms, ZSL and its extension named generalised ZSL (GZSL) [54], naturally appear as multimodal classification problems. However, visual and semantic embeddings adopted in ZSL approaches are usually extracted *independently* from ancillary



Figure 1. For these three bird species of CUB [51] dataset, the attributes "brown spotted wings" is not discriminative. In fact, a "Brown Thrasher" is characterized by the presence of white wingbars. An "Arcadian Flycatcher" features spots which are present over the whole side of the animal – not just over the wings. A "Song Sparrow" shows peculiar tips over the stripes. We posit that captions can provide ancillary semantic information which we propose to exploit altogether with attributes and visual data.

tasks. In fact, visual embedding are usually features trained for object/scene recognition on classical vision benchmarks such as ImageNet [21,46]. On the other hand, semantic embeddings in ZSL are either manually defined, by collecting attributes which describe a certain category, or extracted using natural language processing [11,21,33,46].

Independently extracted semantic embeddings convey only *general* information about a given visual category. For this reason, their sensitivity may be limited in describing *specific* semantic nuances that actually help in discriminating between another category whose instances look extremely similar (see Figure 1). In this paper, we posit that this problem is affecting zero-shot learning and, in order to mitigate it, we propose to exploit captions to inject semantic information into visual embeddings. This is achieved by means of a joint training, finalized at better enriching the embeddings using both visual and semantic cues, in order to ease a zero-shot learning recognition paradigm.

Generally, a jointly learned image and text representation is referred as *pretrain-then-transfer*, a defacto stan-

<sup>\*</sup>Equal contribution

<sup>†</sup> Current Affiliation: Deutsches Elektronen-Synchrotron (DESY) Email: shah.nawaz@desy.de

dard paradigm to process vision and language information [6,32] The success of *pretrain-then-transfer* learning is instrumental in developing methods which can capture connections between vision and language domains on so-called "proxy tasks" with large-scale image and text datasets i.e. MSCOCO [31], Flickr30K [39], Conceptional Captions (ConcCapt) [45], Visual Genome [27]. Therefore, in the recent years, an increased interest was registered in unifying models as to learn joint image and text representations [10, 32, 40, 48]. These models utilize cross-modal attention mechanisms to capture interaction between image regions and fine-grained text information. However, zero-shot learning methods cannot leverage these powerful models to improve visual or semantic information because the zero-shot learning problem does not account for finegrained text description for visual information. Therefore, we solve for this deficiency by presenting a framework to learn joint image and text representation to improve visual information for ZSL and GZSL tasks.

We propose a novel framework, called joint-embeddings for zero-shot learning (JE-ZSL), to learn a semanticallygrounded visual representation. We use a two-stream network, consisting of two layers of non-linearities on top of the image and captions, trained with a bi-directional loss function. We implement structure-preserving constraints to align the alternative visual and textual sources of information. We posit that, by means of our joint training, we can better capture the interaction between semantically-related image and text components, ultimately bridging the gap between visual and semantic representations and, in turn, improving the performance of a zero-shot learning mechanism. We evaluate our proposed semantically grounded embeddings on various well-know compatibility function and GAN based methods, improving standard and generalised zero-shot learning performance on coarse-grained and finegrained dataset. This empirically demonstrates the benefits of corroborating the representation adopted for visual data in zero-shot learning by means of semantically grounded information.

## 1. Related Work

We summarize previous work on both joint image and text representations (Section 1.1) along with standard/generalized zero-shot learning (Section 1.2).

#### **1.1. Joint Image and Text Representation**

In the recent years, several works have been proposed to learn joint image and text representation for various vision and languages tasks in a multimodal setting [4,6,12,17,36]. Usually, joint image and text representation strategies start with separate vision and language models (pre-trained for other large-scale tasks) and then perform a task-specific training [3, 12]. However, the last years have seen a paradigm shift in learning joint image and text representations on proxy tasks designed to bridge the semantic gap between visual and textual clues leveraging large-scale datasets [10, 29, 32, 40, 48]. These joint representations transfer to various downstream vision and language tasks including visual question answering, grounding referring expression, visual commonsense reasoning and image retrieval, without the necessity of extra-training. In this spirit, we proposed to learn semantically grounded enriched visual embedding using joint image and text representation on a proxy task. The key differences between our JE-ZSL approach and the other methods are twofold: (i) Interaction between image and text representations only occur at the last layer of the two-stream network; and (ii) JE-ZSL uses global visual and text representation to capture interaction between image and textual clues. These key differences enable our JE-ZSL model to handle downstream zero-shot learning tasks.

#### **1.2. Zero-shot Learning**

In ZSL, a compatibility function is used to measure the matching between visual and semantic embeddings. For that scope, either metric [24, 26, 42, 63], scoring [30, 47, 49] or ranking functions [9, 19] are adopted. However, the main drawback of those methods is the shallowness of the architectural design. As a result, an improvement in the recognition of unseen classes usually corresponds to a decreased performance in the seen classes and vice-versa [54], damaging the overall generalization capabilities. In this paper, we propose to tackle this issue by the usage of joint embeddings, which we conjecture that can help shallow models in better generalizing.

As an alternative strategy to compatibility functions, the pivotal paper [57] proposed to directly synthesize visual descriptors conditioned on semantic embeddings. A Generative Adversarial Network (GAN) is trained to synthesize visual features which are indistinguishable from the ones extracted from a ResNet-101 backbone. At the deployment stage, the very same model can be utilized to generate synthetic visual features for the unseen classes. Finally, a shallow softmax classifier is trained using both real features from seen classes and unseen features from the unseen ones. This baseline architectural design was recently advanced in various approaches, such as the adoption of variational inference to boost the generator [22,35,43,58,65]. Additional loss functions were investigated to better handle the problem, like cycle-consistency [14], triplet [7] or contrastive loss [1]. In this paper, we show that the semantic conditioning used to generate visual features can be furthermore boosted by the usage of joint embeddings: since they are endowed of semantically-grounded visual information, the "gap" between semantic and visual embeddings is already bridged at the input level already.



Figure 2. Overview of Joint Embeddings for Zero-Shot Learning (JE-ZSL). Our model consists of a visual and semantic embedding modules, V and S respectively, which are responsible for the joint training of the image embeddings starting from textual cues (here, sentences). We stack two fully-connected layers, using ReLU non-linearity, batch normalization (BN), L2 normalization and a Structure-Preserving Alignment Loss which is responsible of injecting semantic patterns into visual descriptors (see Section 2.2). Once the jointly trained visual embeddings x are optimized over a set of seen classes (in the figure, watercress, daffodil, primula, ...), they can be used to train the optimization objective of an arbitrary zero-shot learning framework. Consequently, even if we never saw an image of a "buttercup", we can still manage to classify it (decision boundary represented in red) while leveraging the side information which provides the attributes that the class "buttercup" is expected to own.

# 2. Joint Embeddings for Zero-Shot Learning: Methodology

Our proposed strategy JE-ZSL is to ground visual embeddings through the usage of textual information for the sake of enhancing zero-shot learning methods from the data representation point of view.

### 2.1. Overview of the Proposed Architecture

In Figure 2, we visualize the joint embeddings model that we propose. It is composed by three input streams, each of them corresponding to a different type of input data (images, sentences and class attributes). The modules V and S provide vectorial representation extracted from ResNet-101 [21] and Bidirectional Encoder Representations from Transformers (BERT) [11] models respectively to learn joint embeddings. On top of the output of V and S, JE-ZSL stacks two-fully connected layers, the first of which followed by a rectified linear unit (ReLU) and the second of which is followed by batch-norm operator followed by an L2 normalization. These two fully-connected layers are then responsible for the joint training of the visual embeddings x supported by the sentence embeddings

y, by means of a structure-preserving alignment loss [50] (see Section 2.2). Once visual embeddings are optimized, they can be used in an arbitrary zero-shot learning framework, together with attributes - or other class embeddings - without requiring any modification in the pipeline of the zero-shot learner, thus ensuring a broad applicability (additional details in Section 2.3).

#### 2.2. Structure-Preserving Alignment Loss

The JE-ZSL loss function is defined as follow. Given a visual embedding  $x_i$ , let  $Y_i^+$  and  $Y_i^-$  represent sets of positive and negative sentence embeddings respectively. The distance between a positive pair should be smaller than the distance between a negative pair with margin m:

$$d(x_i, y_j) + m < d(x_i, y_k) \quad \forall y_j \in Y_i^+, \forall y_k \in Y_i^-, (1)$$

where  $x_i$  and  $y_j$  represents positive pair while  $x_i$  and  $y_k$  is negative pair. Similarly, given a sentence  $y_{i'}$ :

$$d(x_{j'}, y_{i'}) + m < d(x_{k'}, y_{i'}) \quad \forall f_{j'} \in X_{i'}^+, \forall x_{k'} \in X_{i'}^-,$$
(2)

where  $X_{i'}^+$  and  $X_{i'}^-$  represents the sets of positive and negative images for  $y_{i'}$ .

In addition, the loss function includes neighborhood constraints to project semantically similar images and sentences close to each other in the embedding space. Let  $N(x_i)$  represent the neighborhood of  $x_i$  with images that contains the same concept. In our scenario, it contains images represented by the set of sentences associated with  $x_i$ Furthermore, a margin m is enforced between  $N(x_i)$  and any point outside of the neighborhood:

$$d(x_i, x_j) + m < d(x_i, x_k) \quad \forall x_j \in N(x_i), \forall x_k \notin N(x_i)$$
(3)

Similarly to Eq. 3, the constraint is enforced for sentences:

$$d(y_{i'}, y_{j'}) + m < d(y_{i'}, y_{k'}) \quad \forall y_{j'} \in N(y_{i'}), \forall y_{k'} \notin N(y_{i'})$$
(4)

Finally, constraints are converted to the training objective using hinge loss. The resulting loss function is given by:

$$\begin{split} L(X,Y) &= \sum_{i,j,k} \max\left[0, m + d\left(x_{i}, y_{j}\right) - d\left(x_{i}, y_{k}\right)\right] \\ &+ \lambda_{1} \sum_{i',j',k'} \max\left[0, m + d\left(x_{j'}, y_{i'}\right) - d\left(x_{k'}, y_{i'}\right)\right] \\ &+ \lambda_{2} \sum_{i,j,k} \max\left[0, m + d\left(x_{i}, x_{j}\right) - d\left(x_{i}, x_{k}\right)\right] \\ &+ \lambda_{3} \sum_{i',j',k'} \max\left[0, m + d\left(y_{i'}, y_{j'}\right) - d\left(y_{i'}, y_{k'}\right)\right]. \end{split}$$

The hyperparameter  $\lambda_1$  balances strengths of both ranking terms and the value is fixed to 2. Similarly,  $\lambda_2$  and  $\lambda_3$  controls the neighborhood constraint and values are set to 0.1 or 0.2 respectively [50]. The distance *d* is fixed to be the Euclidean distance. In addition, triplets are selected within the mini-batch only. To learn our joint embeddings, the network is trained in end-to-end fashion with largescale datasets MSCOCO [31], Flickr30K [39], Conceptional Captions (ConcCapt)) [45], using captioning annotations to leverage the fine-grained datasets Caltech-UCSD-Birds (CUB) [52] and Oxford Flowers (FLO) [38].

#### 2.3. Joint Embeddings for Zero-shot Learning

Once the joint embeddings x are extracted by means of the structure-preserving alignment loss, they are then fed to a second computational module which is responsible for the zero-shot learning task. In order to classify a pool of unseen classes  $c \in C_U$ , we take advantage of training data  $(x_i, c_i), i = 1, ..., N$  where the labels  $c_i$  belong to the seen classes  $(c_i \in C_S)$  for every *i*. In order to bridge the gap between seen and unseen classes, either of them need to be describes in terms of class-embeddings  $a = a_c$  such as manually defined attributes. The learning problem writes as minimizing the empirical risk defined through the following loss function

$$\sum_{i=1}^{N} \mathcal{L}_{\text{ZSL}}(x_i, a_{c_i}; \theta),$$
(5)

which depends upon the learnable parameters  $\theta$ . In turn, the parameters  $\theta$  can be optimized exploiting different strategies, using alternative expressions for  $\mathcal{L}_{\text{ZSL}}$  in eq (5). For methods based on a compatibility function,  $\mathcal{L}_{\text{ZSL}}$  is chosen to be either a distance [24, 26, 42, 63], scoring [30, 47, 49] or ranking function [9, 19]. The optimization attempts to match  $x_i$  and  $a_i$  when and only when they belongs to the same class. Additional regularization terms are used to make sure that the distance/scoring/ranking function  $\mathcal{L}_{\text{ZSL}}$ can generalize towards unseen classes. In fact, when a new test instance  $\tilde{x}$  occurs, inference is solved as

$$\arg\max F_{\rm ZSL}(\tilde{x},\theta),\tag{6}$$

being  $F_{\text{ZSL}}$  a 1-nearest neighbor classifier based on the learned compatibility function  $\mathcal{L}_{\text{ZSL}}$ . Alternatively,  $F_{\text{ZSL}}$ can be chosen to be a softmax classifier trained on real joint embeddings from seen classes  $C_S$  and synthesized features from the unseen classes  $C_U$ . For the feature synthesis, models such as generative adversarial networks [5,15,23,57,62] or variational autoencoders [18, 35, 44, 58] are adopted to generate visual descriptors conditioned on the class-specific embeddings  $a_c, c \in C_U$ .

### 3. Experiments

#### 3.1. Data Sources

**Datasets for Joint Embeddings.** We trained joint image and text model on three publicly available datasets including MSCOCO [31], Flickr30K [39] and Conceptional Captions (ConcCapt) [45]. MSCOCO dataset contains 123,287 images, and each image is annotated with five text descriptions or captions. Similarly, Flickr30K contains 31,783 images along with five captions for each image. While, ConcCapt dataset consists of ~3.3 million image and text pairs collected from alt-text enabled images on the web. At time of access, some links are broken which resulted in ~1.3 million image and text pairs.

**Datasets for Zero-shot Learning.** We evaluate JE-ZSL on two fine-grained and two coarse-grained zero-shot object recognition data sources: CUB [52], FLO [38], Animals with Attributes (AWA2) [55] and Attribute Pascal & Yahoo (aPY) [13] containing 200, 102, 717, 50 and 32 categories, respectively. For fair comparison, we followed the standard zero-shot splits proposed by Xian et al. [57].

#### 3.2. Backbone ZSL Methods

We evaluate the performance of JE-ZSL when adopting the following backbone zero-shot learning methods: We select one of the most efficient methods based on a compatibility function [25], while also including two state-of-theart approaches based on feature generation [35,56]. Further details on these methods are provided below, including how we combined them with our joint embeddings.

Coupled Dictionary Learning [25] (CDL) is composed of the following stages. First, given the semantic embeddings of the classes to be recognized, visual prototypes are learnt by preserving neighborhood relationship across the visual and semantic spaces. This is obtained through a structure alignment module and implemented in the form of visual and semantic dictionaries which are then learnt and coupled. We take advantage of our jointly training embeddings for the sake of learning these dictionary-based representations.

**f-CLSWGAN** [56] is implemented as a Wasserstein GAN which attempts to generate visual features, which are similar to the ones available for seen classes, while also performing semantic-to-visual translation. To better enhance the discrimination capability of the learnt feature, a classification loss is used. Instead of standard ResNet-101 descriptors, we take advantage of jointly training embedding as the real features, supervising the adversarial learning. Once features have been generated for the unseen classes, a softmax classifier is trained from real joint embeddings from the seen classes and synthesized features for the unseen ones. The trained softmax classifier is the model used for the final inference.

A recent generative approach, called **tf-VAEGAN** [35], improved upon [56] under multiple perspectives. First, the generation capability of a GAN are paired with a variational autoencoder and a cycle consistency loss is adopted, aiming at predicting attributes from synthesized features using a regression paradigm. We modify the original implementation by exploiting joint embeddings as either the real features from the seen classes and the input data for the approach regressing attributes. For the final inference stage, the concatenation of the generated visual features and regressed semantic embeddings are used to train the final softmax classifier. Joint embeddings are adopted also at this stage.

#### **3.3. Implementation Details**

We used listed hyperparameters to learn joint image and text embedding with a structure perserving alignment loss [50]. Similarly, we employed listed hyperparameters for zero-shot learning methods [25, 34, 56].

**Baseline Visual and Semantic Embeddings.** For visual embedding, we extract 2048-dim top layer pooling units of ResNet-101 [21]. For the semantic embeddings, we employ the class-level attributes for CUB (312 - d), AWA (85 - d) and aPY (64 - d). For FLO, fine-grained visual descriptions of image with size (1024 - d) are extracted from a character-based CNN-RNN [41].

Semantically Grounded Embeddings. The proposed embedding is extracted from joint image and text representation trained with structure preserving alignment loss. For fair comparison, we learn joint image and text representation with ResNet-101 embeddings.

**Evaluation Metrics.** To evaluate the performance on standard zero-shot learning, we utilized the recommended top-1 classification accuracy over seen classes (T1). Similarly, we adopted common metrics for generalized zero-shot learning as well [54]: Mean per-class accuracy over seen and unseen classes s and u, respectively, while reporting also their harmonic mean H.

# 4. Joint Embeddings for Zero-Shot Learning: Ablation Study and the Fine-Grained case

In this section, we provide an experimental analysis to understand and evaluate the impact of two distinct factors influencing the performance of JE-ZSL. First, we present an ablation study with to monitor if there is an effect of varying the dimensionality of our proposed joint embeddings while training them using the structure-preserving alignment loss (presented in Section 2.2) on various image-text datasets. Then, we will evaluate the impact of using fine-grained captions leveraging fine-grained datasets.

#### 4.1. Effect of the Embedding Size

In Table 1, we evaluate the impact of various embedding sizes applied to JE-ZSL. We take advantage of state-of-the-art methods (CDL [25], f-CLSWGAN [56] and tf-VAEGAN [35]) on coarse-grained benchmark datasets: AWA2 [28] and aPY [61]. These methods are fed with off-the-shelves conventional visual features, exploiting the standard ResNet-101 embeddings as provided by [56]. We then compare the very same zero-shot learning approaches as a baseline references to evaluate the impact in performance of having visual embeddings trained in a joint manner with text information, taking advantage of the architecture presented in Section 2.1. For the joint training, we take advantage of Flickr30K [39], MSCOCO [31] and ConcCapt [45] and, in all cases, we extract joint embedding of dimensions d = 256, 512, 1024 and 2048 to compare them with baseline visual embeddings.

**Discussion**. Table 1 shows the results in bold text font and color which are superior to the respective baseline values: It is evident that the joint training produces a widespread improvement on performance, being consistent with respect to different baseline approaches and multimodal datasets used. Notably, the idea of performing joint training for the embedding has the peculiar trait of ensuring a stable performance even at low embedding sizes. No matter which is the input embedding size, in fact, in a few cases, joint training with d = 256, d = 512, d = 1024

				AWA	2 [28]		aPY [61]					
			T1	u	s	H	T1	u	s	H		
	Baseline : C	CDL [25]	69.9	28.1	73.5	40.6	43.0	19.8	48.6	28.1		
JE-ZSL (ours)		d = 256	65.6	24.2	71.7	36.2	44.1	18.1	46.1	26.0		
	Flickr30K	d = 512	67.1	26.2	76.1	39.0	44.8	19.4	47.3	27.5		
	THERISOR	d=1024	69.2	27.1	74.6	39.8	45.6	20.1	<b>49.8</b>	28.6		
		d=2048	68.2	27.9	77.1	41.0	45.2	22.3	50.4	30.9		
JE-ZSL (ours)		d = 256	68.1	27.1	69.3	39.0	45.1	18.6	48.6	26.9		
	MSCOCO	d = 512	69.3	28.4	75.6	41.3	46.3	23.2	50.2	31.7		
		d=1024	69.5	29.0	75.6	41.9	45.3	20.4	49.9	28.7		
		d=2048	70.1	28.4	72.1	40.7	44.1	19.4	49.0	27.8		
		d = 256	68.3	26.1	71.7	38.3	42.1	20.2	46.7	28.2		
JE-ZSL (ours)	ConcCont	d = 512	69.9	28.1	73.5	40.7	43.3	21.8	49.2	30.2		
	Conceapt	d=1024	69.3	28.1	77.1	41.2	44.8	22.8	50.4	31.4		
		d=2048	71.0	29.3	77.1	42.5	45.9	24.5	51.1	33.1		
Baselin	68.2	57.9	61.4	59.6	40.5	25.8	59.5	36.0				
JE-ZSL (ours)		d = 256	67.0	55.4	58.2	56.8	38.6	28.1	18.0	22.0		
	Flickr30K	d = 512	65.6	59.0	62.1	60.5	39.3	24.5	50.6	33.0		
		d=1024	68.0	55.2	65.8	60.1	38.7	27.9	44.9	34.4		
		d=2048	65.0	55.8	67.6	61.1	37.6	23.4	63.2	34.1		
JE-ZSL (ours)	MSCOCO	d = 256	72.7	61.5	62.3	61.9	39.3	31.4	27.2	29.2		
		d = 512	70.7	63.4	64.0	63.8	41.7	26.5	48.5	34.3		
		d=1024	70.4	62.8	63.8	63.3	38.0	25.7	49.9	34.0		
		d=2048	70.3	60.8	69.1	64.7	38.4	28.2	50.6	36.2		
		d = 256	68.5	59.0	65.0	61.9	39.0	24.4	29.2	26.6		
JE-ZSL (ours)	ConcCont	d = 512	70.0	62.1	65.4	63.7	38.3	26.4	41.4	32.2		
	Conceapt	d=1024	68.4	60.0	65.8	62.6	40.1	26.4	57.3	36.2		
		d=2048	67.0	61.2	61.5	61.3	41.0	27.4	67.3	39.0		
Baseli	ine : tf-VAEG	72.2	59.8	75.1	66.6	40.8	30.8	54.6	39.3			
		d = 256	70.0	57.8	68.2	62.6	36.4	25.5	42.2	31.8		
JE-ZSL (ours)	Flickr30K	d = 512	67.4	57.3	73.8	64.5	38.9	26.8	50.9	35.0		
	THERISOR	d=1024	69.9	63.1	71.5	67.0	41.4	29.2	53.0	37.6		
		d=2048	69.9	63.1	61.5	67.0	40.2	29.9	47.1	36.8		
		d = 256	70.5	56.9	73.4	64.0	43.2	30.0	62.3	40.4		
IF-7SL (ours)	MSCOCO	d = 512	72.3	60.6	64.0	66.0	43.7	30.6	59.7	40.5		
JE-ZSL (OURS)	MISCOCO	d=1024	71.8	62.8	63.8	66.8	40.8	30.8	53.6	38.5		
		d = 2048	76.3	62.3	78.3	69.4	41.8	27.3	70.6	39.4		
JE-ZSL (ours)		d = 256	71.9	61.0	73.6	66.8	41.8	29.3	51.9	37.5		
	ConcCant	d = 512	69.4	58.8	71.2	64.4	39.6	29.8	52.5	38.0		
	Conceapt	d=1024	68.1	59.0	71.8	64.8	40.3	31.9	54.7	40.3		
		d = 2048	70.6	63.0	75.6	68.7	42.1	32.1	49.5	39.0		

Table 1. Joint Embeddings for zero-shot learning (**JE-ZSL**) - ablation study on the effect of the dimensionality. We mark in blue and bold any improvement scored by joint embeddings over the respective backbone zero-shot learning method used as baseline.

and d = 2048 is capable of improving upon the baseline methods in both standard zero-shot learning (T1 on aPY for both Flickr30K and MSCOCO with respect to CDL zero-shot learn method, T1 on AWA2 for MSCOCO with respect to f-CLSWGAN) and generalized zero-shot learning as well (*H* for MSCOCO and ConcCapt with respect to f-CLSWGAN and *H* for MSCOCO with respect to tfVAEGAN). This is especially remarkable if considering that baseline approaches exploit 2048-dimensional visual embeddings: Even if using lower dimensional embeddings, there are only very few cases in which severe drops are registered (such as for d = 256 on Flickr30K with respect to f-CLSWGAN) and performance are often better. Overall, methods such as CDL are improved in their standard and



Figure 3. Joint embeddings for zero-shot learning (JE-ZSL) to handle fine-grained datasets.

generalized zero-shot recognition by +2.9% (T1) and +5% (*H*), respectively using ConcCapt and d = 2048. Similarly, using the very same setup, we register analogous improvements of +0.5% for T1 and +3.0% for *H* (with respect to f-CLSWAGAN), while we enhance tf-VAEGAN by +4.1% (T1) on standard zero-shot learning and +2.8% (*H*) for generalized zero-shot learning using d = 2048 and MSCOCO. These results indicate that the joint training improve representation with visual and semantic cues which in turn eases zero-shot learning recognition task.

## 4.2. Fine-grained Joint Embeddings

Fig. 3 shows the comparison of the setup d = 2048 with Flickr30K [39], MSCOCO [31], ConcCapt [45] and finegrained datasets (CUB [51] and FLO [37]) for which captions annotating the images are available on seen classes. This helps us in understanding to which extent a joint training is capable of enhancing results, while disentangling it from the intrinsic difficulty of the dataset on which the ZSL evaluation is carried out. In fact, on CUB and FLO datasets, categories are fine-grained in nature (bird and flower species, respectively) and, therefore, much harder to classify - and we will evaluate joint embeddings also in this playground.

**Discussion**. Despite the fine-grained nature of the classes to recognize, dataset such as Flickr30K are enough to improve *s* and *H* metrics with respect to the baseline model CDL (on CUB). Joint embeddings trained using ConcCapt improve the metric *s* scored by CDL on FLO. Using MSCOCO, joint embeddings boost the performance of f-CLSWGAN (T1 on FLO) and tf-VAEGAN is improved as well (T1 on FLO) using either Flickr30K or MSCOCO. Thus, despite the challenging nature of fine-grained recognition task, still the idea of joint training provides subtle improvements.

These preliminary improvements are furthermore enhanced as soon as the fine-grained textual information is exploited on CUB and FLO: We improve CDL, f-CLSWGAN and tf-VAEGAN. More specifically, our proposed approach improves the T1 score of f-CLSWGAN on CUB (+7.9). On FLO dataset, CDL, f-CLSWGAN and tf-VAEGAN baseline methods are improved by +2.6%, +8.4% and +4.8%

in their zero-shot learning performance (T1). Similar improvements are registered in GZSL: the *H* value by CDL is improved by +4.6% on CUB and by +8.4% on FLO. The *H* value by f-CLSWGAN is improved by +3.1% on CUB and +8.4% on FLO. Finally, tf-VAEGAN achieves an improvement of *H* on both CUB (+2.2%) and FLO (+3.4%). The scored performance improvements on CUB and FLO datasets demonstrate that fine-grained textual information enhances visual embeddings when paired with our proposed joint training.

# 5. Comparison with the State-of-the-Art

In this section, we test our proposed joint embeddings against state-of-the-art zero-shot learning methods based on two mainstream paradigms: Compatibility functions and feature generation.

*Compatibility functions.* Among the state-of-the-art approaches, we included the idea of generating a latent representation on which visual and semantic information is projected (LATEM [53]) and the generation of synthetic classifiers for the unseen classes as a linear combination of classifiers of the seen ones (SynC [8]). We also compare against structured embedding obtained by combining different sources of side information (SJE [2]) and the deep visual and semantic embedding model known as DeVISE [16]. We include the kernelized approach to learn a compatibility function in max-margin sense (KerZSL [64]) and the coupled dictionary learning approach to learn and then match a set of visual and semantic atoms with which images and side information can be described (CDL [25]).

*Feature generating approaches.* We consider several methods to generate visual features for the unseen classes, known in terms of attributes only. Namely, the Wasserstein GAN trained with a classifier loss (f-CLSWGAN [57]) and its variant in which a cycle-consistency loss is added to predict attributes from synthesized features. (Cycle-WGAN [15]). We compare against the usage of confidence smoothing (COSMO [5]) which, combined with a feature generation approach, helps in balancing the recognition performance across seen and unseen classes. A similar effect is achieved by the adoption of dual learning to train feature

		AWA2 [28]			CUB [51]				FLO [37]				aPY [61]				
		T1	u	s	H	T1	u	s	H	T1	u	s	H	T1	u	s	H
Compatibility Func	LATEM [53]	55.1	7.3	71.7	13.3	49.6	15.2	57.3	24	40.4	6.6	47.6	11.5	36.8	5.7	65.6	10.4
	SynC [8]	49.3	8.9	87.3	16.2	53.0	11.5	70.9	19.8	-	_	_	_	23.9	7.4	66.3	13.3
	SJE [2]	65.6	11.3	74.6	19.6	53.9	23.5	59.2	33.6	53.4	13.9	47.6	21.5	31.7	1.3	71.4	2.6
	DeViSE [16]	54.2	13.4	68.7	22.4	52.0	23.8	53.0	32.8	45.9	9.9	44.2	16.2	37.0	3.5	78.4	6.7
	KerZSL [64]	70.5	18.9	82.7	30.8	51.7	21.6	52.8	30.6	-	_	_	_	45.3	10.5	76.2	18.5
	CDL [25]	69.9	28.1	73.5	40.6	54.5	23.5	55.2	32.9	59.6	32.0	64.8	42.8	43.0	19.8	48.6	28.1
	JE-ZSL (ours)	71.0	29.3	77.1	42.5	54.1	25.6	70.2	37.5	62.2	43.5	62.2	51.2	45.9	24.5	51.1	33.1
Feature Generation	f-CLSWGAN [56]	68.2	57.9	61.4	59.6	57.3	43.7	57.7	49.7	67.2	59.0	73.8	65.6	40.5	25.8	59.5	36.0
	Cycle-WGAN [14]	66.8	63.4	59.6	59.8	58.6	59.3	47.9	53.0	70.3	61.6	69.2	65.2	-	_	—	_
	f-VAEGAN-D2 [58]	71.1	57.6	70.6	63.5	61.0	48.4	60.1	53.6	67.7	56.8	74.9	64.6	-	_	_	—
	CADA-VAE [44]	68.2	55.8	75.0	63.9	57.3	51.6	53.5	52.4	-	_	_	_	-	_	_	—
	GDAN [23]	68.1	32.1	67.5	43.5	58.8	39.3	66.7	49.5	-	_	-	_	38.3	30.4	75.0	43.4
	COSMO [5]	-	52.8	80.0	63.6	-	44.4	57.8	50.2	-	59.6	81.4	68.8	-	_	_	—
	PGN [62]	71.2	48.0	83.6	61.0	68.3	48.5	57.2	52.5	81.4	63.6	77.8	70.0	-	_	_	—
	0-VAEGAN [20]	69.9	56.2	71.7	63.0	54.9	41.1	48.5	44.4	-	_	_	_	36.3	31.7	53.2	39.7
	tf-VAEGAN [35]	72.2	59.8	75.1	66.6	64.9	52.8	64.7	58.1	70.8	62.5	84.1	71.7	40.8	30.8	54.6	39.3
	JE-ZSL (ours)	70.6	63.0	75.6	68.7	70.2	60.0	60.7	60.3	75.6	67.6	84.7	75.1	42.1	32.1	49.5	39.0

Table 2. Comparison of our proposed **JE-ZSL** against state-of-the-art methods based on either compatibility functions or feature generation. Improvements scored by **JE-ZSL** over each of this class of methods is highlighted in blue and bolded.

generation and attribute prediction synchronously (GDAN [23]). We also compared against the idea of scheduling the generation of synthetic feature by means of episodic training (PGN [62]). In addition, we consider methods in which the feature generation stage is enhanced by the adoption of a variational autoencoder to pair a GAN: 0-VAE-GAN [18], f-VAEGAN-D2 [59], CADA-VAE [44], and tf-VAEGAN [35].

Table 2 shows the results of this analysis. We fix the dimensionality of the embedding to d = 2048 that were used to boost CDL when challenging other compatibility functions. When comparing against tf-VAEGAN, performance is enhanced with joint embedding while comparing against other feature generating approaches. The fine-grained setup (see Section 4.2) was used when handling both CUB and FLO datasets, while using Conceptual Captions for the joint training on AWA2 and aPY.

**Discussion**. With our proposed semantically grounded embeddings, the performance previously scored by compatibility functions is considerably improved. We observed performance boosts with respect to the metric *u* over all datasets (AWA2, CUB, FLO and aPY) along with *s* on FLO dataset. Additionally, on AWA2, the best performance on zero-shot learning (KerZSL) is improved by +0.5% and the *H* of the best scoring method on generalized zero-shot learning (CDL) is boosted by +1.9%. Similarly, the *H* metric is enhanced for generalized zero-shot learning on CUB (+5.9% with respect to SJE) on FLO (+8.4% with respect to CDL) and on aPY (+5% with respect to CDL). The best scored T1 classification accuracy of prior methods is improved on FLO (+2.6% with respect to CDL) and on aPY (+0.6% with respect to KerZSL). While considering feature generalized zero-shot learning feature generalized to KerZSL).

ating approaches, our proposed embedding boosts the H score on both AWA2 and CUB (+2.1% and +2.2%, respectively - both with respect to tf-VAEGAN). On aPY, we also improve f-CLSWGAN in the T1 score for ZSL (+1.6%).

## 6. Conclusions

This paper presents joint embeddings for zero-shot learning tasks. Unlike existing methods which utilize *fixed* visual and semantic information, we proposed to learn semantically grounded visual embedding by capturing interactions between images and text clues with a two-stream network on a proxy task leveraging large-scale unlabelled data sources. We evaluated our embeddings with various state-of-the-art methods on two fine-grained (CUB and FLO) and two coarse-grained (AWA2 and aPY) benchmark zero-shot learning datasets. Our evaluation showed that our proposed embeddings considerably improved standard and generalised zero-shot learning performance across various compatibility functions and GAN-based methods.

### References

- Transferable contrastive network for generalized zero-shot learning. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [2] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for finegrained image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 7, 8
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE*

international conference on computer vision, pages 2425–2433, 2015. 2

- [4] Omer Arshad, Ignazio Gallo, Shah Nawaz, and Alessandro Calefati. Aiding intra-text representations with visual context for multimodal named entity recognition. In 2019 15th IAPR International Conference on Document Analysis and Recognition (ICDAR). IEEE, 2019. 2
- [5] Yuval Atzmon and Gal Chechik. Adaptive confidence smoothing for generalized zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2019. 4, 7, 8
- [6] Tadas Baltrusaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2019. 2
- [7] Yannick Le Cacheux, Herve Le Borgne, and Michel Crucianu. Modeling inter and intra-class relations in the triplet loss for zero-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10333– 10342, 2019. 2
- [8] Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, and Fei Sha. Synthesized classifiers for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2016. 7, 8
- [9] Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, and Fei Sha. Classifier and exemplar synthesis for zero-shot learning. *The Springer International Journal of Computer Vision* (*IJCV*), 2019. 2, 4
- [10] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. ECCV, 2020.
   2
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1, 3
- [12] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. arXiv preprint arXiv:1707.05612, 2017. 2
- [13] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 1778–1785. IEEE, 2009. 4
- [14] Rafael Felix, Vijay BG Kumar, Ian Reid, and Gustavo Carneiro. Multi-modal cycle-consistent generalized zeroshot learning. In *The European Conference on Computer Vision (ECCV)*, 2018. 2, 8
- [15] Rafael Felix, Vijay BG Kumar, Ian Reid, and Gustavo Carneiro. Multi-modal cycle-consistent generalized zeroshot learning. In *European Conference on Computer Vision* (ECCV), 2018. 4, 7
- [16] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc'Aurelio Ranzato, and Tomas Mikolov. De-VISE: A deep visual-semantic embedding model. In Advances in Neural Information Processing Systems (NeurIPS), 2013. 7, 8
- [17] Ignazio Gallo, Alessandro Calefati, and Shah Nawaz. Multimodal classification fusion in real-world scenarios. In 2017

14th IAPR International Conference on Document Analysis and Recognition (ICDAR), volume 5, pages 36–41. IEEE, 2017. 2

- [18] Rui Gao, Xingsong Hou, Jie Qin, Jiaxin Chen, Li Liu, Fan Zhu, Zhao Zhang, and Ling Shao. Zero-vae-gan: Generating unseen features for generalized and transductive zero-shot learning. *IEEE Transactions on Image Processing*, 29:3665– 3680, 2020. 4, 8
- [19] Yuchen Guo, Guiguang Ding, Jungong Han, Xiaohan Ding, Sicheng Zhao, Zheng Wang, Chenggang Yan, and Qionghai Dai. Dual-view ranking with hardness assessment for zeroshot learning. In *The AAAI Conference on Artificial Intelligence*, volume 33, pages 8360–8367, 2019. 2, 4
- [20] Zongyan Han, Zhenyong Fu, and Jian Yang. Learning the redundancy-free features for generalized zero-shot object recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 8
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 3, 5
- [22] He Huang, Changhu Wang, Philip S Yu, and Chang-Dong Wang. Generative dual adversarial network for generalized zero-shot learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [23] He Huang, Changhu Wang, Philip S Yu, and Chang-Dong Wang. Generative dual adversarial network for generalized zero-shot learning. In *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2019. 4, 8
- [24] Huajie Jiang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Learning class prototypes via structure alignment for zero-shot recognition. In *European conference on computer* vision (ECCV), 2018. 2, 4
- [25] Huajie Jiang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Learning class prototypes via structure alignment for zero-shot recognition. In *European conference on computer* vision (ECCV), 2018. 5, 6, 7, 8
- [26] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *The IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2017. 2, 4
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 2
- [28] CH. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 5, 6, 8
- [29] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557, 2019. 2
- [30] Yan Li, Junge Zhang, Jianguo Zhang, and Kaiqi Huang. Discriminative learning of latent features for zero-shot recogni-

tion. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 4

- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 4, 5, 7
- [32] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In Advances in Neural Information Processing Systems, pages 13–23, 2019. 2
- [33] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In Advances in neural information processing systems, pages 3111–3119, 2013. 1
- [34] Sanath Narayan, Akshita Gupta, Fahad Shahbaz Khan, Cees GM Snoek, and Ling Shao. Latent embedding feedback and discriminative features for zero-shot classification. arXiv preprint arXiv:2003.07833, 2020. 1, 5
- [35] Sanath Narayan, Akshita Gupta, Fahad Shahbaz Khan, Cees GM Snoek, and Ling Shao. Latent embedding feedback and discriminative features for zero-shot classification. In *The European Conference on Computer Vision (ECCV)*, 2020. 2, 4, 5, 6, 8
- [36] Shah Nawaz, Muhammad Kamran Janjua, Ignazio Gallo, Arif Mahmood, Alessandro Calefati, and Faisal Shafait. Do cross modal systems leverage semantic relationships? In Proceedings of the IEEE International Conference on Computer Vision Workshops, pages 0–0, 2019. 2
- [37] M-E. Nilsback and A. Zisserman. A visual vocabulary for flower classification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006. 7, 8
- [38] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing, pages 722–729. IEEE, 2008. 4
- [39] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer* vision, pages 2641–2649, 2015. 2, 4, 5, 7
- [40] Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. arXiv preprint arXiv:2001.07966, 2020. 2
- [41] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–58, 2016. 5
- [42] Bernardino Romera-Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *The International Conference on Machine Learning (ICML)*, 2015. 2, 4
- [43] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero- and fewshot learning via aligned variational autoencoders. In *The*

*IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2

- [44] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero- and fewshot learning via aligned variational autoencoders. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2019. 4, 8
- [45] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 2556–2565, 2018. 2, 4, 5, 7
- [46] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 1
- [47] Jie Song, Chengchao Shen, Yezhou Yang, Yang Liu, and Mingli Song. Transductive unbiased embedding for zeroshot learning. In *The IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2018. 2, 4
- [48] Hao Tan and Mohit Bansal. Lxmert: Learning crossmodality encoder representations from transformers. arXiv preprint arXiv:1908.07490, 2019. 2
- [49] Vinay Kumar Verma and Piyush Rai. A simple exponential family framework for zero-shot learning. In *The European* conference on machine learning and knowledge discovery in databases (ECML-PKDD), 2017. 2, 4
- [50] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proceed*ings of the IEEE conference on computer vision and pattern recognition, pages 5005–5013, 2016. 3, 4, 5
- [51] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010. 1, 7, 8
- [52] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010. 4
- [53] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 7, 8
- [54] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning-a comprehensive evaluation of the good, the bad and the ugly. *The IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1, 2, 5
- [55] 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, 41(9):2251– 2265, 2018. 4
- [56] 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, pages 5542–5551, 2018. 1, 5, 6, 8

- [57] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 4, 7
- [58] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. F-VAEGAN-D2: A Feature Generating Framework for Any-Shot Learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 4, 8
- [59] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. f-VAEGAN-D2: A feature generating framework for any-shot learning. In *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2019. 8
- [60] Meng Ye and Yuhong Guo. Zero-shot classification with discriminative semantic representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7140–7148, 2017. 1
- [61] Y. Yu. aPascal-aYahoo Image Data Collection. Technical report, University of Illinois at Urbana-Champaign, 2009. 5, 6, 8
- [62] Yunlong Yu, Zhong Ji, Jungong Han, and Zhongfei Zhang. Episode-based prototype generating network for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 4, 8
- [63] Hongguang Zhang and Piotr Koniusz. Zero-shot kernel learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7670–7679, 2018. 2, 4
- [64] Hongguang Zhang and Piotr Koniusz. Zero-shot kernel learning. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018. 7, 8
- [65] Yizhe Zhu, Jianwen Xie, Bingchen Liu, and Ahmed Elgammal. Learning feature-to-feature translator by alternating back-propagation for generative zero-shot learning. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2