# Towards Recognizing Unseen Categories in Unseen Domains

 $\begin{array}{c} {\rm Massimiliano\ Mancini^{1,2}[0000-0001-8595-9955],}\\ {\rm Zeynep\ Akata^{2}[0000-0002-1432-7747],\ Elisa\ Ricci^{3,4}[0000-0002-0228-1147],\ and \\ {\rm Barbara\ Caputo}^{5,6}[0000-0001-7169-0158] \end{array}$ 

 <sup>1</sup> Sapienza University of Rome, <sup>2</sup> University of Tübingen, <sup>3</sup> University of Trento, <sup>4</sup> Fondazione Bruno Kessler,
 <sup>5</sup> Politecnico di Torino, <sup>6</sup> Italian Instituite of Technology mancini@diag.uniroma1.it

Abstract. Current deep visual recognition systems suffer from severe performance degradation when they encounter new images from classes and scenarios unseen during training. Hence, the core challenge of Zero-Shot Learning (ZSL) is to cope with the *semantic-shift* whereas the main challenge of Domain Adaptation and Domain Generalization (DG) is the domain-shift. While historically ZSL and DG tasks are tackled in isolation, this work develops with the ambitious goal of solving them jointly, i.e. by recognizing unseen visual concepts in unseen domains. We present CuMix (Curriculum Mixup for recognizing unseen categories in unseen domains), a holistic algorithm to tackle ZSL, DG and ZSL+DG. The key idea of CuMix is to simulate the test-time domain and semantic shift using images and features from unseen domains and categories generated by mixing up the multiple source domains and categories available during training. Moreover, a curriculum-based mixing policy is devised to generate increasingly complex training samples. Results on standard ZSL and DG datasets and on ZSL+DG using the DomainNet benchmark demonstrate the effectiveness of our approach.

Keywords: Zero-Shot Learning, Domain Generalization

# 1 Introduction

Despite their astonishing success in several applications [12, 34], deep visual models perform poorly for the classes and scenarios that are unseen during training. Most existing approaches are based on the assumptions that (a) training and test data come from the same underlying distribution, i.e. domain shift, and (b) the set of classes seen during training constitute the only classes that will be seen at test time, i.e. semantic shift. These assumptions rarely hold in practice and, in addition to depicting different semantic categories, training and test images may differ significantly in terms of visual appearance in the real world.

To address these limitations, research efforts have been devoted to designing deep architectures able to cope with varying visual appearance [7] and with novel



**Fig. 1.** Our ZSL+DG problem. During training we have images of multiple categories (*e.g. elephant,horse*) and domains (*e.g. photo, cartoon*). At test time, we want to recognize unseen categories (*e.g. dog, giraffe*), as in ZSL, in unseen domains (*e.g. paintings*), as in DG, exploiting side information describing seen and unseen categories.

semantic concepts [47]. In particular, the domain-shift problem [15] has been addressed by proposing domain adaptation (DA) models [7] that assume the availability target domain data during training. To circumvent this assumption, a recent trend has been to move to more complex scenarios where the adaptation problem must be either tackled online [14, 24], with the help of target domain descriptions [23], auxiliary data [32] or multiple source domains [25, 26, 36]. For instance, domain generalization (DG) methods [19, 21, 5] aim to learn domainagnostic prediction models and to generalize to any unseen target domain.

Regarding semantic knowledge, multiple works have designed approaches for extending deep architectures to handle new categories and new tasks. For instance, continual learning methods [18] attempt to sequentially learn new tasks while retaining previous knowledge, tackling the catastrophic forgetting issue. Similarly, in open-world recognition [4] the goal is to detect unseen categories and successfully incorporate them into the model. Another research thread is Zero-Shot Learning (ZSL) [47], where the goal is to recognize objects unseen during training given external information about the novel classes provided in forms of semantic attributes [17], visual descriptions [2] or word embeddings [27].

Despite these significant efforts, an open research question is whether we can tackle the two problems jointly. Indeed, due to the large variability of visual concepts in the real world, in terms of both semantics and acquisition conditions, it is impossible to construct a training set capturing such variability. This calls for a holistic approach addressing them together. Consider for instance the case depicted in Fig. 1. A system trained to recognize elephants and horses from realistic images and cartoons might be able to recognize the same categories in another visual domain, like art paintings (Fig. 1, bottom) or it might be able to describe other quadrupeds in the same training visual domains (Fig. 1, top).

On the other hand, how to deal with the case where new animals are shown in a new visual domain is not clear.

To our knowledge, our work is the first attempt to answer this question, proposing a method that is able to recognize unseen semantic categories in unseen domains. In particular, our goal is to jointly tackle ZSL and DG (see Fig.1). ZSL algorithms usually receive as input a set of images with their associated semantic descriptions, and learn the relationship between an image and its semantic attributes. Likewise, DG approaches are trained on multiple source domains and at test time are asked to classify images, assigning labels within the same set of source categories but in an unseen target domain. Here we want to address the scenario where, during training, we are given a set of images of multiple domains and semantic categories and our goal is to build a model able to recognize images of unseen concepts, as in ZSL, in unseen domains, as in DG.

To achieve this, we need to address challenges usually not present when these two classical tasks, i.e. ZSL and DG, are considered in isolation. For instance, while in DG we can rely on the fact that the multiple source domains permit to disentangle semantic and domain-specific information, in ZSL+DG we have no guarantee that the disentanglement will hold for the unseen semantic categories at test time. Moreover, while in ZSL it is reasonable to assume that the learned mapping between images and semantic attributes will generalize also to test images of the unseen concepts, in ZSL+DG we have no guarantee that this will happen for images of unseen domains.

To overcome these issues, during training we simulate both the semantic and the domain shift we will encounter at test time. Since explicitly generating images of unseen domains and concepts is an ill-posed problem, we sidestep this issue and we synthesize unseen domains and concepts by interpolating existing ones. To do so, we revisit the *mixup* [53] algorithm as a tool to obtain partially unseen categories and domains. Indeed, by randomly mixing samples of different categories we obtain new samples which do not belong to a single one of the available categories during training. Similarly, by mixing samples of different domains, we obtain new samples which do not belong to a single source domain available during training.

Under this perspective, mixing samples of both different domains and classes allows to obtain samples that cannot be categorized in a single class and domain of the one available during training, thus they are *novel* both for the semantic and their visual representation. Since higher levels of abstraction contain more taskrelated information, we perform *mixup* at both image and feature level, showing experimentally the need for this choice. Moreover, we introduce a curriculumbased mixing strategy to generate increasingly complex training samples. We show that our CuMix (**Cu**rriculum **Mix**up for recognizing unseen categories in unseen domains) model obtains state-of-the-art performances in both ZSL and DG in standard benchmarks and it can be effectively applied to the combination of the two tasks, recognizing unseen categories in unseen domains.<sup>7</sup>

<sup>&</sup>lt;sup>7</sup> The code is available at https://github.com/mancinimassimiliano/CuMix

To summarize, our contributions are as follows. (i) We introduce the ZSL+DG scenario, a first step towards recognizing unseen categories in unseen domains. (ii) Being the first holistic method able to address ZSL, DG, and the two tasks together, our method is based on simulating new domains and categories during training by mixing the available training domains and classes both at image and feature level. The mixing strategy becomes increasingly more challenging during training, in a curriculum fashion. (iii) Through our extensive evaluations and analysis, we show the effectiveness of our approach in all three settings: namely ZSL, DG and ZSL+DG.

# 2 Related Works

**Domain Generalization (DG).** Over the past years the research community has put considerable efforts into developing methods to contrast the domain shift. Opposite to domain adaptation [7], where it is assumed that target data are available in the training phase, the key idea behind DG is to learn a domain agnostic model to be applied to any unseen target domain.

Previous DG methods can be broadly grouped into four main categories. The first category comprises methods which attempt to learn domain-invariant feature representations [28] by considering specific alignment losses, such as maximum mean discrepancy (MMD), adversarial loss [22] or self-supervised losses [5]. The second category of methods [19, 15] develop from the idea of creating deep architectures where both domain-agnostic and domain-specific parameters are learned on source domains. After training, only the domain-agnostic part is retained and used for processing target data. The third category devises specific optimization strategies or training procedures in order to enhance the generalization ability of the source model to unseen target data. For instance, in [20] a meta-learning approach is proposed for DG. Differently, in [21] an episodic training procedure is presented to learn models robust to the domain shift. The latter category comprises methods which introduce data and feature augmentation strategies to synthesize novel samples and improve the generalization capability of the learned model [39, 43, 42]. These strategies are mostly based on adversarial training [39, 43].

Our work is related to the latter category since we also generate synthetic samples with the purpose of learning more robust target models. However, differently from previous methods, we specifically employ mixup to perturb feature representations. Recently, works have considered mixup in the context of domain adaptation [51] to *e.g.* reinforce the judgments of a domain discrimination. However, we employ mixup from a different perspective *i.e.* simulating semantic and domain shift we will encounter at test time. To this extent, we are not aware of previous methods using mixup for DG and ZSL.

**Zero-Shot Learning (ZSL).** Traditional ZSL approaches attempt to learn a projection function mapping images/visual features to a semantic embedding space where classification is performed. This idea is achieved by directly predicting image attributes e.g. [17] or by learning a linear mapping through

<sup>4</sup> M. Mancini et al.

margin-based objective functions [1, 2]. Other approaches explored the use of non-linear multi-modal embeddings [45], intermediate projection spaces [54, 55] or similarity-based interpolation of base classifiers [6]. Recently, various methods tackled ZSL from a generative point of view considering Generative Adversarial Networks [48], Variational Autoencoders (VAE) [38] or both of them [50]. While none of these approaches explicitly tackled the domain shift, i.e. visual appearance changes among different domains/datasets, various methods proposed to use domain adaptation technique, e.g. to refine the semantic embedding space, aligning semantic and projected visual features [38] or, in transductive scenarios, to cope with the inherent domain shift existing among the appearance of attributes in different categories [16, 9, 10]. For instance, in [38] a distance among visual and semantic embedding projected in the VAE latent space is minimized. In [16] the problem is addressed through a regularised sparse coding framework, while in [9] a multi-view hypergraph label propagation framework is introduced.

Recently, works have considered also coupling ZSL and DA in a transductive setting. For instance, in [56] a semantic guided discrepancy measure is employed to cope with the asymmetric label space among source and target domains. In the context of image retrieval, multiple works addressed the sketch-based image retrieval problem [52, 8], even across multiple domains. In [40] the authors proposed a method to perform cross-domain image retrieval by training domain-specific experts. While these approaches integrated DA and ZSL, none of them considered the more complex scenario of DG, where no target data are available.

# 3 Method

In this section, we first formalize the Zero-Shot Learning under Domain Generalization (ZSL+DG). We then describe our approach, CuMix , which, by performing curriculum learning through mixup, simulates the domain- and semantic-shift the network will encounter at test time, and can be holistically applied to ZSL, DG and ZSL+DG.

## 3.1 Problem Formulation

In the ZSL+DG problem, the goal is to recognize unseen categories (as in ZSL) in unseen domains (as in DG). Formally, let  $\mathcal{X}$  denote the input space (e.g. the image space),  $\mathcal{Y}$  the set of possible classes and  $\mathcal{D}$  the set of possible domains. During training, we are given a set  $\mathcal{S} = \{(x_i, y_i, d_i)\}_{i=1}^n$  where  $x_i \in \mathcal{X}, y_i \in \mathcal{Y}^s$  and  $d_i \in \mathcal{D}^s$ . Note that  $\mathcal{Y}^s \subset \mathcal{Y}$  and  $\mathcal{D}^s \subset \mathcal{D}$  and, as in standard DG, we have multiple source domains (*i.e.*  $\mathcal{D}^s = \bigcup_{j=1}^m d_j$ , with m > 1) with different distributions *i.e.*  $p_{\mathcal{X}}(x|d_i) \neq p_{\mathcal{X}}(x|d_j), \forall i \neq j$ .

Given  $\mathcal{S}$  our goal is to learn a function h mapping an image x of domains  $\mathcal{D}^u \subset \mathcal{D}$  to its corresponding label in a set of classes  $\mathcal{Y}^u \subset \mathcal{Y}$ . Note that in standard ZSL, while the set of train and test domains are shared, *i.e.*  $\mathcal{D}^s \equiv \mathcal{D}^u$ , the label sets are disjoint *i.e.*  $\mathcal{Y}^s \cap \mathcal{Y}^u \equiv \emptyset$ , thus  $\mathcal{Y}^u$  is a set of *unseen* classes. On the other hand, in DG we have a shared output space, *i.e.*  $\mathcal{Y}^s \equiv \mathcal{Y}^u$ , but

a disjoint set of domains between training and test *i.e.*  $\mathcal{D}^s \cap \mathcal{D}^u \equiv \emptyset$ , thus  $\mathcal{D}^u$  is a set of *unseen* domains. Since the goal of our work is to recognize unseen classes in unseen domains, we unify the settings of DG and ZSL, considering both semantic- and domain-shift at test time *i.e.*  $\mathcal{Y}^s \cap \mathcal{Y}^u \equiv \emptyset$  and  $\mathcal{D}^s \cap \mathcal{D}^u \equiv \emptyset$ .

In the following we divide the function h into three parts: f, mapping images into a feature space  $\mathcal{Z}$ , *i.e.*  $f : \mathcal{X} \to \mathcal{Z}$ , g going from  $\mathcal{Z}$  to a semantic embedding space  $\mathcal{E}$ , *i.e.*  $g : \mathcal{Z} \to \mathcal{E}$ , and an embedding function  $\omega : \mathcal{Y}^t \to \mathcal{E}$  where  $\mathcal{Y}^t \equiv \mathcal{Y}^s$ during training and  $\mathcal{Y}^t \equiv \mathcal{Y}^u$  at test time. Note that  $\omega$  is a learned classifier for DG while it is a fixed semantic embedding function in ZSL, mapping classes into their vectorized representation extracted from external sources. Given an image x, the final class prediction is obtained as follows:

$$y^* = \operatorname{argmax}_{y} \omega(y)^{\mathsf{T}} g(f(x)). \tag{1}$$

In this formulation, f can be any learnable feature extractor (*e.g.* a deep neural network), while g any ZSL predictor (*e.g.* a semantic projection layer, as in [46] or a compatibility function among visual features and labels, as in [1, 2]). The first solution to address the ZSL+DG problem could be training a classifier using the aggregation of data from all source domains. In particular, for each sample we could minimize a loss function of the form:

$$\mathcal{L}_{\text{AGG}}(x_i, y_i) = \sum_{y \in \mathcal{Y}^s} \ell(\omega(y)^{\mathsf{T}} g(f(x_i)), y_i)$$
(2)

with  $\ell$  an arbitrary loss function, *e.g.* the cross-entropy loss. In the following, we show how we can use the input to Eq. (2) to effectively recognize unseen categories in unseen domains.

### 3.2 Simulating Unseen Domains and Concepts through Mixup

The fundamental problem of ZSL+DG is that, during training, we have neither access to visual data associated to categories in  $\mathcal{Y}^u$  nor to data of the unseen domains  $\mathcal{D}^u$ . One way to overcome this issue in ZSL is to generate samples of unseen classes by learning a generative function conditioned on the semantic embeddings in  $\mathcal{W} = \{\omega(y) | y \in \mathcal{Y}^s\}$  [48, 50]. However, since no description is available for the unseen target domain(s) in  $\mathcal{D}^u$ , this strategy is not feasible in ZSL+DG. On the other hand, previous works on DG proposed to synthesize images of unseen domains through adversarial strategies of data augmentation [43, 39]. However, these strategies are not applied to ZSL since they cannot easily be extended to generate data for unseen semantic categories  $\mathcal{Y}^u$ .

To circumvent this issue, we introduce a strategy to simulate, during training, novel domains and semantic concepts by interpolating from the ones available in  $\mathcal{D}^s$  and  $\mathcal{Y}^s$ . Simulating novel domains and classes allows to train the network to cope with both semantic- and domain-shift, the same situation our model will face at test time. Since explicitly generating inputs of novel domains and categories is a complex task, in this work we propose to achieve this goal, by *mixing* images and features of different classes and domains, revisiting the popular *mixup* [53] strategy.



**Fig. 2.** Our CuMix Framework. Given an image (bottom, *horse*, *photo*), we randomly sample one image from the same (middle, *photo*) and one from another (top, *cartoon*) domain. The samples are mixed through  $\phi$  (white blocks) both at image and feature level, with their features and labels projected into the embedding space  $\mathcal{E}$  (by g and  $\omega$  respectively) and there compared to compute our final objective. Note that  $\phi$  varies during training (top part), changing the mixing ratios in and across domains.

In practice, given two elements  $a_i$  and  $a_j$  of the same space (e.g.  $a_i, a_j \in \mathcal{X}$ ), mixup [53] defines a mixing function  $\varphi$  as follows:

$$\varphi(a_i, a_j) = \lambda \cdot a_i + (1 - \lambda) \cdot a_j \tag{3}$$

with  $\lambda$  sampled from a beta distribution, *i.e.*  $\lambda \sim \text{Beta}(\beta, \beta)$ , with  $\beta$  an hyperparameter. Given two samples  $(x_i, y_i)$  and  $(x_j, y_j)$  randomly drawn from a training set  $\mathcal{T}$ , a new loss term is defined as:

$$\mathcal{L}_{\text{MIXUP}}((x_i, y_i), (x_j, y_j)) = \mathcal{L}_{\text{AGG}}(\varphi(x_i, x_j), \varphi(\bar{y}_i, \bar{y}_j))$$
(4)

where  $\bar{y}_i \in \Re^{|\mathcal{Y}^s|}$  is the one-hot vectorized representation of label  $y_i$ . Note that, when mixing two samples and label vectors with  $\varphi$ , a single  $\lambda$  is drawn and applied within  $\varphi$  in both image and label spaces. The loss defined in Eq.(4) forces the network to disentangle the various semantic components (*i.e.*  $y_i$  and  $y_j$ ) contained in the mixed inputs (*i.e.*  $x_i$  and  $x_j$ ) plus the ratio  $\lambda$  used to mix them. This auxiliar task acts as a strong regularizer that helps the network to *e.g.* being more robust against adversarial examples [53]. Note however that the function  $\varphi$  creates input and targets which do not represent a single semantic concept in  $\mathcal{T}$  but contains characteristics taken from multiple samples and categories, synthesising a *new* semantic concept from the interpolation of existing ones.

For recognizing unseen concepts in unseen domains at test time, we revisit  $\varphi$  to obtain both cross-domain and cross-semantic mixes during training, simulating both semantic- and domain-shift. While simulating the semantic-shift is a by-product of the original *mixup* formulation, here we explicitly revisit  $\varphi$  in order to perform cross-domain mixups. In particular, instead of considering a pair

of samples from our training set, we consider a triplet  $(x_i, y_i, d_i), (x_j, y_j, d_j)$  and  $(x_k, y_k, d_k)$ . Given  $(x_i, y_i, d_i)$ , the other two elements of the triplet are randomly sampled from  $\mathcal{S}$ , with the only constraint that  $d_i = d_k, i \neq k$  and  $d_j \neq d_i$ . In this way, the triplet contains two samples of the same domain  $(i.e. d_i)$  and a third of a different one  $(i.e. d_j)$ . Then, our mixing function  $\phi$  is defined as follows:

$$\phi(a_i, a_j, a_k) = \lambda a_i + (1 - \lambda)(\gamma a_j + (1 - \gamma)a_k) \tag{5}$$

with  $\gamma$  sampled from a Bernoulli distribution  $\gamma \sim \mathcal{B}(\alpha)$  and a representing either the input x or the vectorized version of the label y, *i.e.*  $\bar{y}$ . Note that we introduced a term  $\gamma$  which allows to perform either intra-domain (with  $\gamma = 0$ ) or cross-domain (with  $\gamma = 1$ ) mixes.

To learn a feature extractor f and a semantic projection layer g robust to domain- and semantic-shift, we propose to use  $\phi$  to simulate both samples and features of novel domains and classes during training. Namely, we simulate the semantic- and domain-shift at two levels, i.e. image and class levels. Given a sample  $(x_i, y_i, d_i) \in S$  we define the following loss:

$$\mathcal{L}_{\text{M-IMG}}(x_i, y_i, d_i) = \mathcal{L}_{\text{AGG}}(\phi(x_i, x_j, x_k), \phi(\bar{y}_i, \bar{y}_j, \bar{y}_k)).$$
(6)

where  $(x_i, y_i, d_i), (x_j, y_j, d_j), (x_k, y_k, d_k)$  are randomly sampled from  $\mathcal{S}$ , with  $d_i = d_k$  and  $d_j \neq d_k$ . The loss term in Eq. (6) enforces the feature extractor to effectively process inputs of mixed domains/semantics obtained through  $\phi$ . Additionally, to also act at classification level, we design another loss which forces the semantic consistency of mixed features in  $\mathcal{E}$ . This loss term is defined as:

$$\mathcal{L}_{\mathrm{M-F}}(x_i, y_i, d_i) = \sum_{y \in \mathcal{Y}^s} \ell \Big( \omega(y)^{\mathsf{T}} g\big( \phi(f(x_i), f(x_j), f(x_k)) \big), \phi(\bar{y}_i, \bar{y}_j, \bar{y}_k) \Big)$$
(7)

where, as before,  $(x_j, y_j, d_j), (x_k, y_k, d_k) \sim S$ , with  $d_i = d_k, i \neq k$  and  $d_j \neq d_k$ and  $\ell$  is a generic loss function *e.g.* the cross-entropy loss. This second loss term forces the classifier  $\omega$  and the semantic projection layer g to be robust to features with mixed domains and semantics.

While we can simply use a fixed mixing function  $\phi$ , as defined in Eq. (5), for Eq. (6) and Eq. (7), we found that it is more beneficial to devise a dynamic  $\phi$ which changes its behaviour during training, in a curriculum fashion. Intuitively, minimizing the two objectives defined in Eq.(6) and Eq.(7) requires our model to correctly disentangle the various semantic components used to form the mixed samples. While this is a complex task even for intra-domain mixes (*i.e.* when only the semantic is mixed), mixing samples across domains makes the task even harder, requiring to isolate also domain specific factors. To effectively tackle this task, we choose to act on the mixing function  $\phi$ . In particular, we want our  $\phi$  to create mixed samples with progressively increased degree of mixing both with respect to content and domain, in a curriculum-based fashion.

During training we regulate both  $\alpha$  (weighting the probability of crossdomain mixes) and  $\beta$  (modifying the probability distribution of the mix ratio  $\lambda$ ), changing the probability distribution of the mixing ratio  $\lambda$  and of the crossdomain mix  $\gamma$ . In particular, given a warm-up step of N epochs and being s the current epoch we set  $\beta = \min(\frac{s}{N}\beta_{max},\beta_{max}))$ , with  $\beta_{max}$  as hyperparameter, while  $\alpha = \max(0, \min(\frac{s-N}{N}, 1))$ . As a consequence, the learning process is made of three phases, with a smooth transition among them. We start by solving the plain classification task on a single domain (*i.e.*  $s < N, \alpha = 0, \beta = \frac{s}{N}\beta_{max}$ ,). In the subsequent step ( $N \leq s < 2N$ ) samples of the same domains are mixed randomly, with possibly different semantics (*i.e.*  $\alpha = \frac{s-N}{N}, \beta = \beta_{max}$ ). In the third phase ( $s \geq 2N$ ), we mix up samples of different domains (*i.e.*  $\alpha = 1$ ), simulating the domain shift the predictor will face at test time. Figure 2, shows a representation of how  $\phi$  varies during training (top, white block).

**Final objective.** The full training procedure, is represented in Figure 2. Given a training sample  $(x_i, y_i, d_i)$ , we randomly draw other two samples,  $(x_j, y_j, d_j)$ and  $(x_k, y_k, d_k)$ , with  $d_i = d_k, i \neq k$  and  $d_j \neq d_i$ , feed them to  $\phi$  and obtain the first mixed input. We then feed  $x_i, x_j, x_k$  and the mixed sample through f, to extract their respective features. At this point we use features extracted from other two randomly drawn samples (in the figure, and just for simplicity,  $x_j$  and  $x_k$  with same mixing ratios  $\lambda$  and  $\gamma$ ), to obtain the feature level mixed features needed to build the objective in Eq.(7). Finally, the features of  $x_i$  and the two mixed variants at image and feature level, are fed to the semantic projection layer g, which maps them to the embedding space  $\mathcal{E}$ . At the same time, the labels in  $\mathcal{Y}^s$  are projected in  $\mathcal{E}$  through  $\omega$ . Finally, the objectives defined in Eq.(2),Eq.(6) and Eq.(7) functions are then computed in the semantic embedding space. Our final objective is:

$$\mathcal{L}_{\text{CuMIX}}(\mathcal{S}) = |\mathcal{S}|^{-1} \sum_{(x_i, y_i, d_i) \in \mathcal{S}} \mathcal{L}_{\text{AGG}}(x_i, y_i) + \eta_{\text{I}} \mathcal{L}_{\text{M-IMG}}(x_i, y_i, d_i) + \eta_{\text{F}} \mathcal{L}_{\text{M-F}}(x_i, y_i, d_i)$$
(8)

with  $\eta_{\rm I}$  and  $\eta_{\rm F}$  hyperparameters weighting the importance of the two terms. As  $\ell(x, y)$  in both  $\mathcal{L}_{\rm AGG}$ ,  $\mathcal{L}_{\rm M-IMG}$  and  $\mathcal{L}_{\rm M-F}$ , we use the standard cross-entropy loss, even if any ZSL objective can be applied. Finally, we highlight that the optimization is performed batch-wise, thus also the sampling of the triplet considers the current batch and not the full training set S. Moreover, while in Figure 2 we show for simplicity that the same samples are drawn for  $\mathcal{L}_{\rm M-IMG}$  and  $\mathcal{L}_{\rm M-F}$ , in practice, given a sample, the random sampling procedure of the other two members of the triplet is held-out twice, one at the image level and one at the feature level. Similarly, the sampling of the mixing ratios  $\lambda$  and cross domain factor  $\gamma$  of  $\phi$  is held-out sample-wise and twice, one at image level and one at feature level. As in Eq. (3),  $\lambda$  and  $\gamma$  are kept fixed across mixed inputs/features and their respective targets in the label space.

**Discussion.** We present similarities between our framework with DG and ZSL methods. In particular, presenting the classifier with noisy features extracted by a non-domain specialist network, has a similar goal as the episodic strategy for DG described in [21]. On the other hand, here we sidestep the need to train domain experts by directly presenting as input to our classifier features of novel domains that we obtain by interpolating the available sources samples. Our



Fig. 3. ZSL results on CUB, SUN, AWA and FLO datasets with ResNet-101 features.

method is also linked to *mixup* approaches developed in DA [51]. Differently from them, we use *mixup* to simulate unseen domains rather than to progressively align the source to the given target data.

Our method is also related to ZSL frameworks based on feature generation [48, 50]. While the quality of our synthesized samples is lower since we do not exploit attributes for conditional generation, we have a lower computational cost. In fact, during training we simulate the test-time semantic shift without generating samples of unseen classes. Moreover, we do not require additional training phases on the generated samples or the availability of unseen class attributes to be available beforehand.

# 4 Experimental results

#### 4.1 Datasets and implementation details

We assess CuMix in three scenarios: ZSL, DG and the proposed ZSL+DG setting. **ZSL**. We conduct experiments on four standard benchmarks: Caltech-UCSD-Birds 200-2011 (CUB) [44], SUN attribute (SUN) [31], Animals with Attributes (AWA) [17] and Oxford Flowers (FLO) [29]. CUB contains 11,788 images of 200 bird species, with 312 attributes, SUN 14,430 images of 717 scenes annotated with 102 attributes, and AWA 30,475 images of 50 animal categories with 85 attributes. Finally, FLO is a fine-grained dataset of flowers, containing 8,189 images of 102 categories. As semantic representation, we use the visual descriptions of [35], following [48, 46]. For each dataset, we use the train, validation and test split provided by [47]. In all the settings we employ features extracted from the second-last layer of a ResNet-101 [13] pretrained on ImageNet as image representation. For CuMix , we consider f as the identity function and as g a simple fully connected layer, perform our version of mixup directly at the feature-level while applying our alignment loss in the embedding space. All hyperparameters have been set following [47].

**DG**. We perform experiments on the PACS dataset [19] with 9,991 images of 7 semantic classes in 4 different visual domains, *art paintings, cartoons, photos* and *sketches*. For this experiment we use the standard train and test split defined in [19], with the same validation protocol. We use as base architecture a ResNet-18

[13] pretrained on ImageNet. For our model, we consider f to be the ResNet-18 while g to be the identity function. We use the same training hyperparameters and protocol of [21].

**ZSL+DG**. Since no previous work addressed the problem of ZSL+DG, there is no benchmark on this task. As a valuable benchmark, we choose Domain-Net [33], a recently introduced dataset for multi-source domain adaptation [33] with a large variety of domains, visual concepts and possible descriptions. It contains approximately 600'000 images from 345 categories and 6 domains, *clipart*, *infograph*, *painting*, *quickdraw*, *real* and *sketch*.

To convert this dataset from a DA to a ZSL scenario, we need to define an unseen set of classes. Since our method uses a network pretrained on ImageNet [37], the set of unseen classes can not contain any of the classes present in ImageNet following the good practices in [49]. We build our validation + test set with 100 classes that contain at least 40 images per domain and that has no overlap with ImageNet. We reserve 45 of these classes for the unseen test set, matching the number used in [40], and the remaining 55 classes for the unseen validation set. The remaining 245 classes are used as seen classes during training.

We set the hyperparameters of each method by training on all the images of the seen classes on a *subset* of the source domains and validating on all the images of the validation set from the held-out source domain. After the hyperparameters are set, we retrain the model on the training set, i.e. 245 classes, and validation set, i.e. 55 classes, of a total number of 300 classes. Finally, we report the final results on the 45 unseen classes. As semantic representation we use word2vec embeddings [27] extracted from the Google News corpus and L2-normalized, following [40]. For all the baselines and our method, we employ as base architecture a ResNet-50 [13] pretrained on ImageNet, using the same number of epochs and SGD with momentum as optimizer, with the same hyperparameters of [40].

#### 4.2 Results

**ZSL.** In the ZSL scenario, we choose as baselines standard inductive methods plus more recent approaches. In particular we report the results of ALE [1], SJE [2], SYNC [6], GFZSL [41] and SPNet [46]. ALE [1] and SJE [2] are linear compatibility methods using a ranking loss and the structural SVM loss respectively. SYNC [6] learns a mapping from the feature space and the semantic embedding space by means of phantom classes and a weighted graph. GFZSL [41] employs a generative framework where each class-conditional distribution is modeled as a multivariate Gaussian. Finally, SPNet [46] learns a semantic projection function from the feature space through the image embedding space by minimizing the standard cross-entropy loss.

Our results grouped by datasets are reported in Figure 3. Our model achieves performance either superior or comparable to the state-of-the-art in all benchmarks but AWA. We believe that in AWA learning a better alignment between visual features and attributes may not be as effective as improving the quality of the visual features. Especially, although the names of the test classes do not appear in the training set of ImageNet, for AWA being a non-fine-grained dataset,

	AGG	DANN	MLDG	CrossGrad	MetaReg	JiGen	Epi-FCR	CuMix
Target		[11]	[20]	[39]	[3]	[5]	[21]	
Photo	94.9	94.0	94.3	94.0	94.3	96.0	93.9	95.1
Art	76.1	81.3	79.5	78.7	79.5	79.4	82.1	82.3
Cartoon	73.8	73.8	77.3	73.3	75.4	75.3	77.0	76.5
Sketch	69.4	74.3	71.5	65.1	72.2	71.4	73.0	72.6
Average	78.5	80.8	80.7	80.7	77.8	80.4	81.5	81.6

Table 1. Domain Generalization accuracies on PACS with ResNet-18.

the information content of the test classes is likely represented by the ImageNet training classes. Moreover, for non-fine-grained datasets, finding labeled training data may not be as challenging as it is in fine-grained datasets. Hence, we argue that zero-shot learning is of higher practical interest in fine-grained settings. Indeed our proposed model is effective in fine-grained scenarios (*i.e.* CUB, SUN, FLO) where it consistently outperforms the state-of-the-art approaches.

These results show that our model based on *mixup* achieves competitive performances on ZSL by simulating the semantic shift the classifier will experience at test time. To this extent, our approach is the first to show that mixup can be a powerful regularization strategy for the challenging ZSL setting.

**DG.** The second series of experiments consider the standard DG scenario. Here we test our model on the PACS dataset using a ResNet-18 architecture. As baselines for DG we consider the standard model trained on all source domains together (AGG), the adversarial strategies in [11] (DANN) and [39], the meta learning-based strategy MLDG [20] and MetaReg [3]. Moreover we consider the episodic strategy presented in [21] (Epi-FCR).

As shown in Table 1, our model achieves competitive results comparable to the state-of-the-art episodic strategy Epi-FCR [21]. Remarkable is the gain obtained with respect to the adversarial augmentation strategy CrossGrad [39]. Indeed, synthesizing novel domains for domain generalization is an ill-posed problem, since the concept of unseen domain is hard to capture. However, with CuMix we are able to simulate inputs/features of novel domains by simply interpolating the information available in the samples of our sources. Despite containing information available in the original sources, our approach allows to produce a model more robust to domain shift.

Another interesting comparison is against the self-supervised approach JiGen [5]. Similarly to [5] we employ an additional task to achieve higher generalization abilities to unseen domains. While in [5] the JigSaw puzzles [30] are used as a secondary self-supervised task, here we employ the mixed samples/features in the same manner. The improvement in the performances of our method highlights that recognizing the semantic of mixed samples acts as a more powerful secondary task to improve robustness to unseen domains.

Finally, it is worth noting that CuMix performs a form of episodic training, similar to Epi-FCR [21]. However, while Epi-FCR considers multiple domainspecific architectures (to simulate the domain experts needed to build the episodes), we require a single domain agnostic architecture. We build our episodes by mak-

$\mathcal{L}_{\mathrm{AGG}}$	$\mathcal{L}_{\mathrm{M-IMG}}$	$\mathcal{L}_{\mathrm{M-F}}$	Curriculum	Art	Cartoon	Photo	Sketch	Avg.
~				76.1	73.8	94.9	69.4	78.5
~	1			78.4	72.7	94.7	59.5	76.3
1		1		81.8	76.5	94.9	71.2	81.1
1	1	1		82.7	75.4	95.4	71.5	81.2
1	1	1	1	82.3	76.5	95.1	72.6	81.6

Table 2. Ablation on PACS dataset with ResNet-18 as backbone.

ing the *mixup* among images/features of different domains increasingly more drastic. Despite not requiring any domain experts, CuMix achieves comparable performances to Epi-FCR, showing the efficacy of our strategy to simulate

unseen domain shifts.

Ablation study. In this section, we ablate the various components of our method. We performed the ablation on the PACS benchmark for DG, since this allows us to show how different choices act on the generalization to unseen domains. In particular, we ablate the following implementation choices: 1) mixing samples at the image level, feature level or both 2) impact of our curriculum-based strategy for mixing features and samples.

As shown in Table 2, mixing samples at feature level produces a clear gain on the results with respect to the baseline, while mixing samples only at image level can even harm the performance. This happens particularly in the *sketch* domain, where mixing samples at feature level produces a gain of 2% while at image level we observe a drop of 10% with respect to the baseline. This could be explained by mixing samples at image level producing inputs that are too noisy for the network and not representative of the actual shift experienced at test time. Mixing samples at feature level instead, after multiple layers of abstractions, allows to better synthesize the information contained in the different samples, leading to more reliable features for the classifier. Using both of them allows to obtain higher results in almost all domains.

Finally, we analyze the impact of the curriculum-based strategy for mixing samples and features. As the table shows, adding the curriculum strategy allows to boost the performances for the most difficult cases (i.e. sketches) producing a further accuracy boost. Moreover, applying this strategy allows to stabilize the training procedure, as demonstrated experimentally.

**ZSL+DG.** On the proposed ZSL+DG setting we use the DomainNet dataset, training on five out of six domains and reporting the average per-class accuracy on the held-out one. We report the results for all possible target domains but one, *i.e.* real photos, since our backbone has been pretrained on ImageNet, thus the photo domain is not an unseen one. Since no previous methods addressed the ZSL+DG problem, in this work we consider simple baselines derived from the literature of both ZSL and DG. The first baseline is a standard ZSL model without any DG algorithm (i.e. the standard AGG): as ZSL method we consider SPNet [46]. The second baseline is a DG approach coupled with a ZSL algorithm. To this extent we select the state-of-the-art Epi-FCR as the DG ap-

Table 3. ZSL+DG scenario on the DomainNet dataset with ResNet-50 as backbone.

Method	Clipart	Infograph	Painting	Quickdraw	Sketch	Avg.
SPNet	26.0	16.9	23.8	8.2	21.8	19.4
mixup + SPNet	27.2	16.9	24.7	8.5	21.3	19.7
Epi-FCR+SPNet	26.4	16.7	24.6	9.2	23.2	20.0
CuMix	27.6	17.8	25.5	9.9	22.6	20.7

proach, coupling it with SPNet. As a reference, we also evaluate the performance of standard *mixup* coupled with SPNet.

As shown in Table 3, our method achieves competitive performances in ZSL+DG setting when compared to a state-of-the-art approach for DG (Epi-FCR) coupled with a state-of-the-art one for ZSL (SPNet), outperforming this baseline in almost all settings but *sketch* and, in average by almost 1%. Particularly interesting are the results on the *infograph* and *quickdraw* domains. These two domains are the ones where the shift is more evident as highlighted by the lower results of the baseline. In these scenarios, our model consistently outperforms the competitors, with a remarkable gain of more than 1.5% in average accuracy per class with respect to the ZSL only baseline. We want to highlight also that DomainNet is a challenging dataset, where almost all standard DA approaches are ineffective or can even lead to negative transfer [33]. Our method however is able to overcome the unseen domain shift at test time, improving the performance of the baselines in all scenarios. Our model consistently outperforms SPNet coupled with the standard *mixup* strategy in every scenario. This demonstrates the efficacy of the choices in CuMix for revisiting *mixup* in order to recognize unseen categories in unseen domains.

# 5 Conclusions

In this work, we propose the novel ZSL+DG scenario. In this setting, during training, we are given a set of images of multiple domains and semantic categories and our goal is to build a model able to recognize unseen concepts, as in ZSL, in unseen domains, as in DG. To solve this problem we design CuMix, the first algorithm which can be holistically and effectively applied to DG, ZSL and ZSL+DG. CuMix is based on simulating inputs and features of new domains and categories during training by mixing the available source domains and classes, both at image and feature level. Experiments on public benchmarks show the effectiveness of CuMix, achieving state-of-the-art performances in almost all settings in all tasks. Future works will investigate the use of alternative data-augmentation schemes in the ZSL+DG setting.

Acknowledgments We thank the ELLIS Ph.D. student program and the ERC grants 637076 - RoboExNovo (B.C.) and 853489 - DEXIM (Z.A.). This work has been partially funded by the DFG under Germany's Excellence Strategy – EXC number 2064/1 – Project number 390727645.

# References

- Akata, Z., Perronnin, F., Harchaoui, Z., Schmid, C.: Label-embedding for attributebased classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 819–826 (2013)
- Akata, Z., Reed, S., Walter, D., Lee, H., Schiele, B.: Evaluation of output embeddings for fine-grained image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2927–2936 (2015)
- Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: Advances in Neural Information Processing Systems. pp. 998–1008 (2018)
- 4. Bendale, A., Boult, T.: Towards open world recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1893–1902 (2015)
- Carlucci, F.M., D'Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2229–2238 (2019)
- Changpinyo, S., Chao, W.L., Gong, B., Sha, F.: Synthesized classifiers for zero-shot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5327–5336 (2016)
- Csurka, G.: A Comprehensive Survey on Domain Adaptation for Visual Applications, pp. 1–35. Springer International Publishing (2017)
- Dutta, A., Akata, Z.: Semantically tied paired cycle consistency for zero-shot sketch-based image retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5089–5098 (2019)
- 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)
- Gan, C., Yang, T., Gong, B.: Learning attributes equals multi-source domain generalization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 87–97 (2016)
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. The Journal of Machine Learning Research 17(1), 2096–2030 (2016)
- Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hoffman, J., Darrell, T., Saenko, K.: Continuous manifold based adaptation for evolving visual domains. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 867–874 (2014)
- Khosla, A., Zhou, T., Malisiewicz, T., Efros, A.A., Torralba, A.: Undoing the damage of dataset bias. In: European Conference on Computer Vision. pp. 158–171 (2012)
- Kodirov, E., Xiang, T., Fu, Z., Gong, S.: Unsupervised domain adaptation for zeroshot learning. In: Proceedings of the IEEE international conference on computer vision. pp. 2452–2460 (2015)
- Lampert, C.H., Nickisch, H., Harmeling, S.: Attribute-based classification for zeroshot visual object categorization. IEEE transactions on pattern analysis and machine intelligence 36(3), 453–465 (2013)

- 16 M. Mancini et al.
- Lange, M.D., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G.G., Tuytelaars, T.: Continual learning: A comparative study on how to defy forgetting in classification tasks. arXiv:1909.08383 (2019)
- Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: Proceedings of the IEEE international conference on computer vision. pp. 5542–5550 (2017)
- Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Learning to generalize: Metalearning for domain generalization. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
- Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1446–1455 (2019)
- Li, H., Jialin Pan, S., Wang, S., Kot, A.C.: Domain generalization with adversarial feature learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5400–5409 (2018)
- Mancini, M., Bulo, S.R., Caputo, B., Ricci, E.: Adagraph: Unifying predictive and continuous domain adaptation through graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6568–6577 (2019)
- Mancini, M., Karaoguz, H., Ricci, E., Jensfelt, P., Caputo, B.: Kitting in the wild through online domain adaptation. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1103–1109 (2018)
- Mancini, M., Porzi, L., Bulo, S.R., Caputo, B., Ricci, E.: Inferring latent domains for unsupervised deep domain adaptation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
- Mancini, M., Porzi, L., Rota Bulò, S., Caputo, B., Ricci, E.: Boosting domain adaptation by discovering latent domains. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3771–3780 (2018)
- Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. 1st International Conference on Learning Representations Workshop Track Proceedings (2013)
- Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: International Conference on Machine Learning. pp. 10– 18 (2013)
- Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing. pp. 722–729. IEEE (2008)
- Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conference on Computer Vision. pp. 69–84. Springer (2016)
- Patterson, G., Hays, J.: Sun attribute database: Discovering, annotating, and recognizing scene attributes. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 2751–2758. IEEE (2012)
- Peng, K.C., Wu, Z., Ernst, J.: Zero-shot deep domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 764–781 (2018)
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1406–1415 (2019)
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)

17

- Reed, S., Akata, Z., Lee, H., Schiele, B.: Learning deep representations of finegrained visual descriptions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 49–58 (2016)
- Roy, S., Siarohin, A., Sangineto, E., Bulo, S.R., Sebe, N., Ricci, E.: Unsupervised domain adaptation using feature-whitening and consensus loss. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9471–9480 (2019)
- 37. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision **115**(3), 211–252 (2015)
- Schonfeld, E., Ebrahimi, S., Sinha, S., Darrell, T., Akata, Z.: Generalized zero-and few-shot learning via aligned variational autoencoders. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8247–8255 (2019)
- Shankar, S., Piratla, V., Chakrabarti, S., Chaudhuri, S., Jyothi, P., Sarawagi, S.: Generalizing across domains via cross-gradient training. International Conference on Learning Representations (2018)
- Thong, W., Mettes, P., Snoek, C.G.: Open cross-domain visual search. arXiv preprint arXiv:1911.08621 (2019)
- Verma, V.K., Rai, P.: A simple exponential family framework for zero-shot learning. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 792–808. Springer (2017)
- Volpi, R., Murino, V.: Addressing model vulnerability to distributional shifts over image transformation sets. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 7980–7989 (2019)
- Volpi, R., Namkoong, H., Sener, O., Duchi, J.C., Murino, V., Savarese, S.: Generalizing to unseen domains via adversarial data augmentation. In: Advances in Neural Information Processing Systems. pp. 5334–5344 (2018)
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-ucsd birds 200 (2010)
- 45. Xian, Y., Akata, Z., Sharma, G., Nguyen, Q., Hein, M., Schiele, B.: Latent embeddings for zero-shot classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 69–77 (2016)
- 46. Xian, Y., Choudhury, S., He, Y., Schiele, B., Akata, Z.: Semantic projection network for zero-and few-label semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8256–8265 (2019)
- Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: 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)
- Xian, Y., Lorenz, T., Schiele, B., Akata, Z.: Feature generating networks for zeroshot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5542–5551 (2018)
- 49. Xian, Y., Schiele, B., Akata, Z.: Zero-shot learning-the good, the bad and the ugly. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4582–4591 (2017)
- Xian, Y., Sharma, S., Schiele, B., Akata, Z.: f-vaegan-d2: A feature generating framework for any-shot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10275–10284 (2019)
- Xu, M., Zhang, J., Ni, B., Li, T., Wang, C., Tian, Q., Zhang, W.: Adversarial domain adaptation with domain mixup. In: The Thirty-Fourth AAAI Conference on Artificial Intelligence. pp. 6502–6509. AAAI Press (2020)

- 18 M. Mancini et al.
- Yelamarthi, S.K., Reddy, S.K., Mishra, A., Mittal, A.: A zero-shot framework for sketch based image retrieval. In: European Conference on Computer Vision. pp. 316–333. Springer (2018)
- 53. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (2018)
- Zhang, Z., Saligrama, V.: Zero-shot learning via semantic similarity embedding. In: Proceedings of the IEEE international conference on computer vision. pp. 4166– 4174 (2015)
- 55. Zhang, Z., Saligrama, V.: Zero-shot learning via joint latent similarity embedding. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6034–6042 (2016)
- Zhuo, J., Wang, S., Cui, S., Huang, Q.: Unsupervised open domain recognition by semantic discrepancy minimization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 750–759 (2019)