

Semantics Disentangling for Generalized Zero-Shot Learning

Zhi Chen¹ Yadan Luo¹ Ruihong Qiu¹ Sen Wang¹ Zi Huang¹ Jingjing Li² Zheng Zhang³

¹The University of Queensland, Australia

²University of Electronic Science and Technology of China ³Harbin Institute of Technology, Shenzhen

{zhi.chen,y.luo,r.qiu,sen.wang}@uq.edu.au

huang@itee.uq.edu.au, lijin117@yeah.net, darrenzz219@gmail.com

Abstract

Generalized zero-shot learning (GZSL) aims to classify samples under the assumption that some classes are not observable during training. To bridge the gap between the seen and unseen classes, most GZSL methods attempt to associate the visual features of seen classes with attributes or to generate unseen samples directly. Nevertheless, the visual features used in the prior approaches do not necessarily encode semantically related information that the shared attributes refer to, which degrades the model generalization to unseen classes. To address this issue, in this paper, we propose a novel semantics disentangling framework for the generalized zero-shot learning task (SDGZSL), where the visual features of unseen classes are firstly estimated by a conditional VAE and then factorized into semantic-consistent and semantic-unrelated latent vectors. In particular, a total correlation penalty is applied to guarantee the independence between the two factorized representations, and the semantic consistency of which is measured by the derived relation network. Extensive experiments conducted on four GZSL benchmark datasets have evidenced that the semantic-consistent features disentangled by the proposed SDGZSL are more generalizable in tasks of canonical and generalized zero-shot learning. Our source code is available at <https://github.com/uqzhichen/SDGZSL>.

1. Introduction

Human beings have a remarkable ability to learn new notions based on prior experience without seeing them in advance. For example, given the clues that zebras appear like horses yet with black-and-white stripes, one can quickly recognize a zebra if he/she has seen horses before. Nevertheless, unlike humans, supervised machine learning algorithms can only classify samples belonging to the classes that have already appeared during the training phase, and they are not able to handle samples from previously unseen




Figure 1: An illustration of the visual features (red boxes) that are not associated with the annotated attributes. Learning from such visual features that are semantically unrelated may jeopardize the model generalization to unseen classes.

categories. This challenge motivates the study of generalizing models to the unseen classes by transferring knowledge from intermediate semantics (e.g., attributes), which typically refers to zero-shot learning (ZSL). intro

Particularly, the core idea of ZSL [19, 34, 1, 17] lies in learning to map features between the semantic space and visual space, thereby closing the gap between the seen and unseen classes. While effective, conventional ZSL techniques are built upon the assumption that the test set only contains samples from the unseen classes, which can be easily violated in practice. Hence, it is more reasonable to consider a new protocol called generalized zero-shot learning (GZSL), where seen and unseen images are both to be identified.

Existing GZSL techniques can be roughly grouped into two types: embedding-based [9, 34, 22, 21, 13] and generative-based [38, 35, 28, 24, 15] approaches. The former group learns a projection or an embedding function to associate the visual features of seen classes with the respective semantic vectors, while the latter one learns a visual generator for the unseen classes based on the seen points and semantic representations of both classes. However, most GZSL approaches directly leverage the visual features extracted from the pre-trained deep models, such

as ResNet101 [11] pre-trained on ImageNet, which are not tailored for ZSL tasks. In [31], it is observed that not all the dimensions of the extracted visual features are semantically related to the pre-defined attributes, which triggers the bias on learning semantic-visual alignment and causes negative transfer to unseen classes. Given an example from the AWA dataset shown in Figure 1, despite the features of animals’ ears are visually salient for discriminating image samples, it is ignored in the manually annotated attributes. When generalizing to unseen classes such as cats, it is easy for them to be misclassified as tigers because the visual features corresponding to the concepts “*Big, Strong, Muscle*” are not highlighted. From this case, we believe that GZSL will benefit from using the visual features that can consistently align with the respective semantic attributes. We define this type of visual features as the semantic-consistent features, which are agnostic to both seen and unseen classes. In contrast, those visual features that are irrelevant to manually annotated attributes are defined as semantic-unrelated.

To unravel semantic-consistent and semantic-unrelated features from the original visual spaces, we present a novel framework, namely Semantics Disentangling for Generalized Zero-Shot Learning (SDGZSL), as shown in Figure 2. Specifically, we disentangle the underlying information of the extracted visual features into two disjoint latent vectors h_s and h_n . They are learned in an encoder-decoder architecture with a relation module and a total correlation penalty. The encoder network projects the original visual features to h_s and h_n . To make h_s consistent with the semantic embeddings, the relation module calculates a compatibility score between h_s and semantic information to guide the learning of h_s . We further apply the total correlation penalty to enforce the independence between h_s and h_n . Afterward, we reconstruct the original visual features \bar{x} from the two latent representations. This reconstruction objective ensures the two latent representations to cover both semantic-consistent and semantic-unrelated information. The disentangling modules are incorporated into a conditional variational autoencoder and trained in an end-to-end manner. The proposed framework is evaluated on various GZSL benchmarks and achieves better performance compared to the state-of-the-art methods. The main contributions of this work are summarized as follows:

- We propose a novel feature disentangling framework, namely Semantics Disentangling for Generalized Zero-Shot Learning (SDGZSL), to disentangle the underlying information of visual features into two latent representations that are semantic-consistent and semantic-unrelated, respectively. Exploiting the semantic-consistent representations can substantially increase the performance in GZSL comparing to directly using entangled visual features that are extracted from the pre-trained CNN models.

- To facilitate the feature disentanglement of the semantic-consistent and semantic-unrelated representations, by introducing a total correlation penalty in our framework we arrive at a more accurate characterization of the semantically annotated features.
- Extensive experiments conducted on four benchmark datasets evidence that the proposed method performs better than the state-of-the-art methods.

2. Related Work

Recent state-of-the-art approaches for GZSL using generative models have achieved promising performance. Generative models can synthesize an unlimited number of visual features from side information of the unseen classes, *e.g.*, manually annotated attributes. With these synthesized features, ZSL problems become a relatively straightforward supervised classification task. The two most commonly used generative models are generative adversarial networks (GANs) [10] and variational autoencoders (VAEs) [14]. Often, both models are jointly used to form generative architectures for ZSL tasks. f-CLSWGAN [35] leverages Wasserstein GANs (WGAN) [3] to synthesize vivid visual features. CADA-VAE [28] leverages two aligned variational autoencoders to learn the shared latent representations between different modalities. SE-ZSL [18] adopts an autoencoder followed by an attribute regressor to train a model with three alignments: visual-to-attribute, attribute-to-visual and visual-to-attribute. E-PGN [37] integrates the meta-learning approach into ZSL by formulating both the visual prototype generation and the class semantic inference into an adversarial framework. TF-VAEGAN [24] proposes a feedback module in a VAE-GAN model to modulate the latent representation of the generator. However, the CNN visual features contain semantic-unrelated information, *e.g.*, background noise and unannotated characteristics, which may jeopardise the semantic-visual alignment, we propose to factorize out the semantic-unrelated features and leverage the remaining semantic-consistent features as the generation target.

In most literature, disentanglement refers to independence among features of one representation [4, 7, 12, 5]. Total correlation [33] is a measurement of independence for multiple random variables. In information theory, total correlation is one of many generalizations of mutual information to random variables. It has been a key component in recent methods towards disentanglement. FactorVAE [16] is proposed to disentangle features by making the distribution of representations to be factorial and thus achieves independence across dimensions. Specifically, dimension-wise independence is achieved by applying the total correlation penalty on an original representation and its random permuted representation across dimensions. In contrast, we aim to enforce independence between

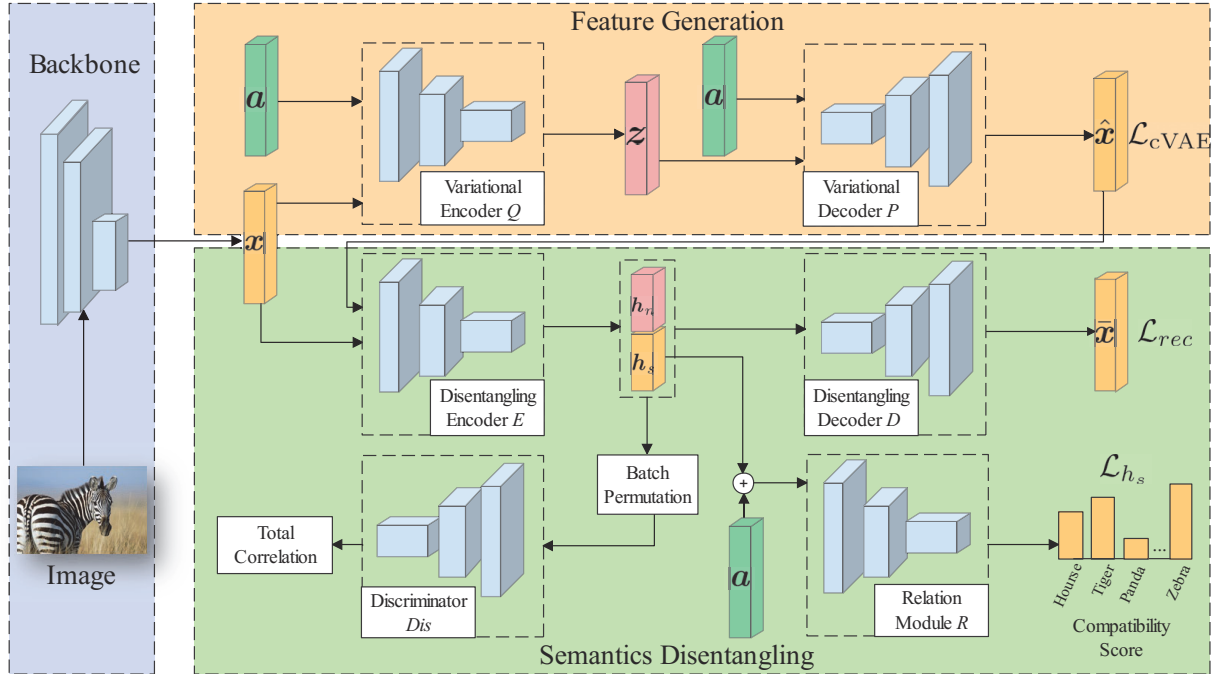


Figure 2: An illustration of our proposed SDGZSL, which is comprised of: (i) a conditional VAE model for visual feature generation (in the orange box); (ii) disentangling modules to learn two factorized latent representations (in the green box). Particularly, the inference network Q is learned to infer a latent variable z from the visual feature x and semantic embedding a of seen samples, where z is then applied to recover unseen visual features. The encoder E is subsequently trained to factorize the semantic-consistent h_s and semantic-unrelated h_n representations, and the independence between them is guaranteed by a total correlation penalty. By concatenating h_s with the random semantic vectors a , the derived relation module R distinguishes the matched pairs from the mismatched ones, thus forcing h_s to be semantically correlated.

two representations rather than keeping each dimension independent to each other. To our best knowledge, DLFZRL [31] is the only work proposed to consider feature disentangling for ZSL. DLFZRL proposes a hierarchical disentangling approach to learn the discriminative latent features. This approach is designed in a two-step fashion, starting with feature selection and then learning to generate the selected features. The original visual features are factorized into three latent representations, including semantic, non-semantic, and non-discriminative features. However, we argue that the non-semantic visual features do not align with the semantic embeddings. Hence, it is hard to transfer the non-semantic visual features from seen to unseen classes. In our approach, we filter out such non-semantic visual features. Moreover, we combine the generative model and the disentangling modules so that our approach can be trained in an end-to-end manner.

3. Proposed Approach

This section firstly introduces the problem formulation and notations, then depicts the proposed approach of factorizing semantic-consistent representations for GZSL.

3.1. Preliminaries

Let $\{\mathcal{X}^s, \mathcal{Y}^s\}$ be the dataset with S seen classes, which contains N^s training samples $\mathcal{X}^s = \{\mathbf{x}_{(i)}^s\}_{i=1}^{N^s}$ and the corresponding class labels $\mathcal{Y}^s = \{y_{(i)}^s\}_{i=1}^{N^s}$. The class labels span from 1 to S , $y^s \in \mathbf{L}^s = \{1, \dots, S\}$. Given another dataset $\{\mathcal{X}^u, \mathcal{Y}^u\}$, in which the classes are related to the seen dataset (e.g., all classes in both datasets correspond to animals). The dataset has U unseen classes and consists of N^u data instances $\mathcal{X}^u = \{\mathbf{x}_{(j)}^u\}_{j=1}^{N^u}$ with the corresponding labels $\mathcal{Y}^u = \{y_{(j)}^u\}_{j=1}^{N^u}$. The class labels thus range from $S+1$ to $S+U$, $y^u \in \mathbf{L}^u = \{S+1, \dots, S+U\}$. The seen and unseen classes are mutually exclusive, i.e., $\mathbf{L}^s \cap \mathbf{L}^u = \emptyset$. For the semantic information $\mathcal{A} = \{\mathbf{a}_{(k)}\}_{k=1}^{S+U}$, each class from both seen and unseen is associated with a class-level semantic vector, which can be embeddings or attributes. We denote \mathcal{A}^s and \mathcal{A}^u as the semantic vectors of seen and unseen classes.

3.2. Semantics Disentanglement

We start with defining three concepts in GZSL concerning the visual features and the semantic information.

Semantic-consistent Representations. We define semantic-consistent representations \mathbf{h}_s to represent the characteristics of images annotated with attributes. The visual features of images are extracted by a deep model, *e.g.*, ResNet101, pretrained on ImageNet. These visual features are not task-specific for ZSL datasets, in which the classes are usually related (*e.g.*, they all correspond to birds). Thus, the extracted visual features may unexpectedly involve redundant information that could compromise the semantic-visual relationship learning. The elegant solution to ZSL is attribute-based learning, which consists in introducing the intermediate semantic space \mathcal{A} . Such a space enables parameter sharing between classes. Ideally, if the visual features only contain information that corresponds to the annotated attributes \mathbf{h}_s , the visual-semantic relation can be appropriately learned from seen classes and further transferred to unseen classes which could be beneficial to the GZSL task.

Semantic-unrelated Representations. In analogy to \mathbf{h}_s , we define semantic-unrelated visual representations \mathbf{h}_n to represent the information contained in visual features which may help classification but not correspond to the annotated attributes, *e.g.*, in Figure 1 the ear shape is an intuitively discriminative characteristic that can help distinguish animals but does not appear in the annotated attributes.

Independence between \mathbf{h}_s and \mathbf{h}_n . In GZSL, the core mission is to transfer the semantic-visual relationship from seen classes to unseen classes. However, as the visual concepts of \mathbf{h}_n are not consistent with the manually annotated semantic information, it is difficult to align the relationship from semantic space to visual space and further transfer to unseen classes. As in the definitions of \mathbf{h}_s and \mathbf{h}_n , in the setting of GZSL, the visual features can be disentangled into semantic-consistent \mathbf{h}_s and semantic-unrelated \mathbf{h}_n representations. To achieve this, we enforce the independence between \mathbf{h}_s and \mathbf{h}_n . Independence can be measured by the mutual information, and the total correlation is one of the generalizations of the mutual information. Thus, we aim to minimize the entanglement of the visual information in the two representations by leveraging the total correlation to measure how \mathbf{h}_s and \mathbf{h}_n are independent.

3.2.1 Visual Feature Preservation by Reconstruction

The disentangling modules are developed with an encoder-decoder architecture, coupled with a total correlation penalty and a relation module. An encoder E_ψ parameterized with ψ is adopted to map a visual feature vector \mathbf{x} to a latent representation \mathbf{h} , *i.e.*, $E_\psi: \mathbb{R}^d \rightarrow \mathbb{R}^{l+m}$, where l and m are the dimensions of the semantic-consistent and semantic-unrelated representations. Then, we have:

$$E_\psi(\mathbf{x}) = \mathbf{h} = [\mathbf{h}_s, \mathbf{h}_n]. \quad (1)$$

In order to preserve the visual features in \mathbf{h} , a decoder D_ω is learned to transform \mathbf{h} to the original visual feature \mathbf{x} ,

i.e., $D_\omega: \mathbb{R}^{l+m} \rightarrow \mathbb{R}^d$. The reconstruction objective can be formulated as:

$$\mathcal{L}_{rec} = \sum_{\mathbf{x} \in \mathcal{X}^s} \|\mathbf{x} - D_\omega(\mathbf{h}_s, \mathbf{h}_n)\|^2, \quad (2)$$

where we calculate the mean squared error between the original visual features \mathbf{x} and the reconstructed visual features $\tilde{\mathbf{x}} = D_\omega(\mathbf{h}_s, \mathbf{h}_n)$.

3.2.2 Semantic-Visual Relationship Learning

A relation network (RN) [30] is adopted to learn the semantic-consistent representations \mathbf{h}_s by maximizing the compatibility score (CS) between a latent representation \mathbf{h}_s and the corresponding semantic embeddings \mathbf{a}^s . The relation module R_κ learns the pair-wise relationship between the latent representations and semantic vectors. The input of R_κ is the pairs of a latent representation \mathbf{h}_s and N^c unique semantic embeddings $\mathcal{A}^{batch} = \{\mathbf{a}_{(c)}\}_{c=1}^{N^c}$ from a training batch with B training instances. The ground-truth CS of the matched pairs are set to be 1; the mismatched pairs are 0, which can be formulated as:

$$CS(\mathbf{h}_{(t)}, \mathbf{a}_{(c)}) = \begin{cases} 0, & y_{(t)} \neq y_{(c)} \\ 1, & y_{(t)} = y_{(c)} \end{cases}, \quad (3)$$

where t and c refer to the t -th semantic-consistent representation and c -th unique semantic embedding in a training batch, $y_{(t)}$ and $y_{(c)}$ denote the class label of $\mathbf{h}_{(t)}$ and $\mathbf{a}_{(c)}$.

With the CS defined in Equation 3, we now formulate the loss function for the semantics embeddings \mathcal{A}^{batch} and semantic-consistent information. The relation module R_κ with a Sigmoid activation function outputs the learned compatibility score between 0 and 1 for each pair. The loss function for optimizing \mathbf{h}_s can then be formulated as:

$$\mathcal{L}_{\mathbf{h}_s} = \sum_{t=1}^B \sum_{c=1}^{N^c} \|R_\kappa(\mathbf{h}_{s(t)}, \mathbf{a}_{(c)}) - CS(\mathbf{h}_{s(t)}, \mathbf{a}_{(c)})\|^2, \quad (4)$$

where we calculate the mean squared error between the output compatibility score of each pair of $\mathbf{h}_{s(t)}$ and $\mathbf{a}_{(c)}$ and the ground-truth in every single training batch. Normally, we have $N^c \leq B$ as one class that may contain many sampled visual feature vectors in one batch.

3.2.3 Disentanglement by Total Correlation Penalty

To encourage the disentanglement between semantic-consistent \mathbf{h}_s and semantic-unrelated \mathbf{h}_n representations, we introduce a *total correlation* penalty in our proposed method. Within the encoding procedure, the latent representation $\mathbf{h} \sim \gamma(\mathbf{h} | \mathbf{x})$ is expected to contain both the semantic-consistent and the semantic-unrelated information. Therefore, the disentanglement between these two

streams of information is crucial for meaningful representation learning. Along with the relation network that facilitates the learning of semantic-consistent representation \mathbf{h}_s , we aim to make \mathbf{h}_n semantic-unrelated by encouraging the disentanglement between \mathbf{h}_n and \mathbf{h}_s . From a probabilistic point of view, we can think both of them come from different conditional distributions:

$$\mathbf{h}_s \sim \gamma_1(\mathbf{h}_s | \mathbf{x}), \quad \mathbf{h}_n \sim \gamma_2(\mathbf{h}_n | \mathbf{x}). \quad (5)$$

Hence, for the two latent representations, the total correlation can be formulated as:

$$TC = \text{KL}(\gamma || \gamma_1 \cdot \gamma_2), \quad (6)$$

where $\gamma := \gamma(\mathbf{h}_s, \mathbf{h}_n | \mathbf{x})$ is the joint conditional probability of \mathbf{h}_s and \mathbf{h}_n . In order to efficiently approximate the total correlation, we apply the density ratio estimation to distinguish the samples from the two distributions in an adversarial manner. A discriminator Dis_φ is constructed to output an estimate of the probability $Dis_\varphi(\mathbf{h})$ whose input is independent. Thus,

$$TC = \mathbb{E}_\gamma \left(\log \frac{\gamma}{\gamma_1 \cdot \gamma_2} \right) \approx \mathbb{E}_\gamma \left(\log \frac{Dis(\mathbf{h})}{1 - Dis(\mathbf{h})} \right), \quad (7)$$

where the approximation is derived in Appendix A. Simultaneously, we train the discriminator Dis_φ to maximize the probability of assigning the correct label to \mathbf{h} and $\tilde{\mathbf{h}}$:

$$\mathcal{L}_{dis} = \log Dis_\varphi(\mathbf{h}) + \log(1 - Dis_\varphi(\tilde{\mathbf{h}})), \quad (8)$$

where $\tilde{\mathbf{h}}$ is the result by randomly permuting each \mathbf{h}_s and \mathbf{h}_n across the batch. The permutation process is described as following: (1) Given a batch of latent representations $\{\mathbf{h}_{(t)}\}_{t=1}^B$; (2) they are split into $\{\mathbf{h}_{s(t)}\}_{t=1}^B$ and $\{\mathbf{h}_{n(t)}\}_{t=1}^B$; (3) for both \mathbf{h}_s and \mathbf{h}_n , we permute the batch indices twice on $\mathcal{B} = \{1, \dots, B\}$, yielding \mathcal{B}' and \mathcal{B}'' ; (4) with the permuted indices, reorder the latent representations $\{\mathbf{h}_{s\mathcal{B}'(t)}\}_{t=1}^B$, $\{\mathbf{h}_{n\mathcal{B}''(t)}\}_{t=1}^B$ and concatenate them into $\{\tilde{\mathbf{h}}_{(t)}\}_{t=1}^B$.

3.3. Visual Feature Generation

To model the distribution of visual features conditioned on the semantic information, we leverage the conditional variational auto-encoder (cVAE) [29] as the generative model. In GZSL, we aim to transfer the knowledge from seen classes to some other unseen classes. Hence, we represent the category information in cVAE as the class embeddings \mathbf{a} to enable parameter sharing between classes. The objective function of the cVAE in our framework can then be written as:

$$\begin{aligned} \mathcal{L}_{cVAE} = & -KL[q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{a})||p_\theta(\mathbf{z}|\mathbf{a})] \\ & + \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{a})}[\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{a})], \end{aligned} \quad (9)$$

Algorithm 1 SDGZSL training

Input: $\{\mathcal{X}^s, \mathcal{Y}^s\}$, \mathcal{A}^s , learning rate λ

Initialize: φ and $W = \{\phi, \theta, \psi, \omega, \kappa\}$

```

1: while not converged do
2:   Randomly select a batch  $\{\mathbf{x}_{(t)}^s, \mathbf{y}_{(t)}^s\}_{t=1}^B, \{\mathbf{a}_{(c)}^s\}_{c=1}^{N^c}$ 
3:   for step = 0, ...,  $n_{dis}$  do
4:     Compute  $\tilde{\mathbf{h}}$  and  $\mathbf{h}$  with  $\phi, \theta, \psi$  by Eq. 1
5:     Vector permutation for  $\tilde{\mathbf{h}}$  and  $\mathbf{h}$ 
6:     Compute  $\lambda_3 \mathcal{L}_{dis}$  by Eq. 8
7:     Update  $\varphi \leftarrow \varphi + \lambda \nabla_\varphi \lambda_3 \mathcal{L}_{dis}$ 
8:     Compute  $\mathcal{L}_{overall1} = \mathcal{L}_{cVAE} + \mathcal{L}_{rec} + \lambda_1 \mathcal{L}_{h_s}$ 
       with Eq. 9, 2, 4
9:     Update  $W \leftarrow W + \lambda \nabla_W \mathcal{L}_{overall1}$ 
10:   end for
11:   Randomly select a batch  $\{\mathbf{x}_{(t)}^s, \mathbf{y}_{(t)}^s\}_{t=1}^B, \{\mathbf{a}_{(c)}^s\}_{c=1}^{N^c}$ 
12:   Compute  $\mathcal{L}_{overall2} = \mathcal{L}_{cVAE} + \mathcal{L}_{rec} + \lambda_1 \mathcal{L}_{h_s} +$ 
        $\lambda_2 TC$  with Eq. 9, 2, 4, 7
13:   Update  $W \leftarrow W + \lambda \nabla_W \mathcal{L}_{overall2}$ 
14: end while

```

Output: trained generative network P_θ and encoder E_ψ

where the first term is the KL divergence between two distributions $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{a})$ and $p_\theta(\mathbf{z}|\mathbf{a})$ and the second term is the reconstruction loss. As illustrated in Figure 2, we use Q_ϕ and P_θ to denote the inference and generator networks to model $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{a})$ and $p_\theta(\mathbf{z}|\mathbf{a})$, respectively. Specifically, given the visual features \mathbf{x} and the semantic embeddings \mathbf{a} , the inference network Q_ϕ produces the latent variables \mathbf{z} . The generator network P_θ leverages the inferred latent variable \mathbf{z} and the class embeddings \mathbf{a} to reconstruct the visual features $\hat{\mathbf{x}}$. The reconstructed and original visual features $\hat{\mathbf{x}}, \mathbf{x}$ are fed into the disentangling module.

3.4. Training and Inference

Algorithm 1 shows the pseudocode of the model training. We iteratively train the Dis_φ with the overall framework for n_{dis} steps and then fix the weights in Dis_φ to train other components. In Algorithm 1, the weights for \mathcal{L}_{h_s}, TC , and \mathcal{L}_{dis} are denoted as λ_1, λ_2 and λ_3 . Once the training of SDGZSL is converged, the semantic-consistent representations $\tilde{\mathbf{h}}_s^u$ of unseen classes can be generated by the generative network P_θ and disentangling encoder E_ψ from the Gaussian noise \mathbf{z} and the unseen semantic embeddings \mathbf{a}^u , i.e., $P_\theta: \mathbb{R}^z \times \mathbb{R}^k \rightarrow \mathbb{R}^d$, $E_\psi: \mathbb{R}^d \rightarrow \mathbb{R}^{l+m}$. z represents the dimensions of the latent variables \mathbf{z} . Then, we disentangle the training seen features $\{\mathbf{x}_{(i)}^s\}_{i=1}^{N^s}$ into $\{\mathbf{h}_{s(i)}^s\}_{i=1}^{N^s}$, together with the generated unseen semantic-consistent representations $\tilde{\mathbf{h}}_s^u$ we can simply train a supervised classifier. Further prediction for either seen or unseen objects can be conducted by supervised classification. In our paper, a Softmax classifier is adopted for evaluation.

Table 1: Performance comparison in accuracy (%) on four datasets. We report the accuracies of unseen, seen classes and their harmonic mean for GZSL, which are denoted as U, S and H. For ZSL, performance results are reported with the average top-1 classification accuracy (T1). The top two results of the T1 and H are highlighted in bold. † and ‡ represent embedding-based and generative methods, respectively. * means a fine-tuned backbone is used.

Methods	aPaY				AWA				CUB				FLO				
	T1	U	S	H	T1	U	S	H	T1	U	S	H	T1	U	S	H	
†	LFGAA [22]	-	-	-	-	68.1	27.0	93.4	41.9	67.6	36.2	80.9	50.0	-	-	-	-
	DCN [21]	43.6	14.2	75.0	23.9	65.2	25.5	84.2	39.1	56.2	28.4	60.7	38.7	-	-	-	-
	TCN [13]	38.9	24.1	64.0	35.1	71.2	61.2	65.8	63.4	59.5	52.6	52.0	52.3	-	-	-	-
	DVBE [23]	-	32.6	58.3	41.8	-	63.6	70.8	67.0	-	53.2	60.2	56.5	-	-	-	-
‡	f-CLSWGAN [35]	40.5	32.9	61.7	42.9	65.3	56.1	65.5	60.4	57.3	43.7	57.7	49.7	69.6	59.0	73.8	65.6
	CANZSL [6]	-	-	-	-	68.9	49.7	70.2	58.2	60.6	47.9	58.1	52.5	69.7	58.2	77.6	66.5
	LisGAN [20]	43.1	34.3	68.2	45.7	70.6	52.6	76.3	62.3	58.8	46.5	57.9	51.6	69.6	57.7	83.8	68.3
	CADA-VAE [28]	-	31.7	55.1	40.3	64.0	55.8	75.0	63.9	60.4	51.6	53.5	52.4	65.2	51.6	75.6	61.3
	f-VAEGAN-D2 [36]	-	-	-	-	71.1	57.6	70.6	63.5	61.0	48.4	60.1	53.6	67.7	56.8	74.9	64.6
	DLFZRL [31]	46.7	-	-	38.5	70.3	-	-	60.9	61.8	-	-	51.9	-	-	-	-
	TF-VAEGAN [24]	-	-	-	-	72.2	59.8	75.1	66.6	64.9	52.8	64.7	58.1	70.8	62.5	84.1	71.7
	TF-VAEGAN* [24]	-	-	-	-	73.4	55.5	83.6	66.7	74.3	63.8	79.3	70.7	74.7	69.5	92.5	79.4
	OCD-CVAE [15]	-	-	-	-	71.3	59.5	73.4	65.7	60.3	44.8	59.9	51.3	-	-	-	-
	E-PGN [37]	-	-	-	-	73.4	52.6	83.5	64.6	72.4	52.0	61.1	56.2	85.7	71.5	82.2	76.5
	AGZSL [37]	41.0	35.1	65.5	45.7	73.8	65.1	78.9	71.3	57.2	41.4	49.7	45.2	82.7	63.5	94.0	75.7
	AGZSL* [37]	43.7	36.2	58.6	44.8	76.4	69.0	86.5	76.8	77.2	69.2	76.4	72.6	85.2	73.7	91.9	81.7
cVAE	39.2	30.2	55.3	39.0	65.4	54.4	72.6	62.2	61.4	47.0	59.9	52.7	68.7	60.1	89.6	71.9	
SDGZSL w/o RN&TC	20.3	14.6	37.0	21.0	57.5	41.3	70.0	51.9	42.3	25.1	54.7	34.5	67.5	56.6	87.6	68.7	
SDGZSL w/o TC	39.8	33.8	49.1	40.0	54.0	45.7	79.4	58.0	27.4	23.6	46.2	31.2	57.5	44.8	65.7	53.3	
SDGZSL w/o RN	30.9	27.3	41.9	33.1	55.8	48.5	64.6	55.5	30.4	20.8	31.4	25.0	53.2	40.6	67.6	50.7	
SDGZSL	45.4	38.0	57.4	45.7	72.1	64.6	73.6	68.8	75.5	59.9	66.4	63.0	85.4	83.3	90.2	86.6	
SDGZSL*	47.0	39.1	60.7	47.5	74.3	69.6	78.2	73.7	78.5	73.0	77.5	75.1	86.9	86.1	89.1	87.8	

4. Experiments

4.1. Experimental Setting

Datasets The proposed framework is evaluated on four widely used benchmark datasets of image classification, including two coarse-grained datasets (Attribute Pascal and Yahoo (aPaY) [8] and Animals with Attributes 2 (AWA) [19]) and two fine-grained datasets (Caltech-UCSD Birds-200-2011 (CUB) [32] and Oxford Flowers (FLO) [25]). aPaY contains 18,627 images from 42 classes and is annotated with 64 attributes. It combines datasets a-Pascal and a-Yahoo, which has 30 and 12 classes respectively. AWA is a relatively larger coarse-grained dataset with 30,475 images from 50 animal species, in which 40 are selected as seen classes and the rest is unseen. Each species in the dataset is annotated with 85 attributes. CUB consists of 11,788 images from fine-grained bird species with 150 seen and 50 unseen classes. FLO contains 102 flower categories with 82 seen and 20 unseen classes. The semantic embeddings of FLO and CUB are 1,024-dimensional character-based CNN-RNN features [26] extracted from the fine-grained visual descriptions (10 sentences per image).

Evaluation Protocol The metric used to evaluate the GZSL task is harmonic mean, which calculates the joint accuracy of the seen and unseen classes. The formula used to calculate the harmonic mean H can be written as: $H = (2 \times U \times S) / (U + S)$, where U and S denote the av-




Figure 3: GZSL accuracy (in %) comparison between h_n , h , and h_s .

erage per-class top-1 accuracy of unseen and seen classes, respectively. A high harmonic mean indicates the good performance of both seen and unseen classes.

4.2. Comparison with State-of-the-art Methods

Table 1 shows the GZSL performance of compared methods and ours with and without fine-tuning the backbone on the datasets. We choose the recent state-of-the-art embedding-based and generative methods. They are marked with † and ‡ respectively. Generally, our pro-

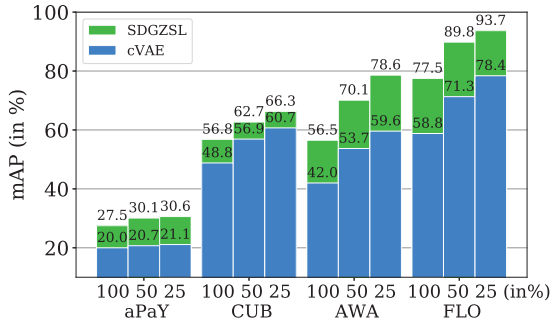


Figure 4: Zero-shot image retrieval result comparison between cVAE and SDGZSL.

posed method consistently performs better than all the compared methods, except on AWA dataset. These methods directly use the visual features extracted from a pre-trained or fine-tuned ResNet101 model, which limits semantic-visual alignment learning. The intuition of this work is in agreement with DLFZRL [31], we both aim to disentangle more effective representations from visual features. DLFZRL proposes to disentangle the discriminative features from visual features, both semantic and non-semantic. However, we argue that the non-semantic discriminative features cannot be transferred from seen classes to unseen classes. The conditional generative models learn to project the semantic vectors to the visual space, but intuitively there is no way to generalize the non-semantic discriminative features from the semantic perspective. As can be seen from the performance comparison, our method surpasses DLFZRL on all the reported datasets. Our method is also applicable to the generative methods mentioned above. It is worth mentioning that, unlike DLFZRL that is designed in a two-stage process, we carefully design our framework by incorporating the disentangling modules into the generative model so that the overall framework can be trained in an end-to-end manner.

4.3. Conventional Zero-shot Learning Results

We aim to solve the generalized zero-shot learning problem in this work, but it is necessary to demonstrate that our approach can also achieve state-of-the-art performance on conventional zero-shot learning that only aims to classify unseen class samples over unseen classes. The results shown in Table 1 demonstrate the performance comparison between our proposed SDGZSL method and existing state-of-the-art models. It can be seen that our method performs better than all other methods on CUB and FLO datasets. Even if DLFZRL is slightly higher than our method on aPaY dataset, we can surpass it on all other datasets. This is the same case for E-PGN on dataset AWA.



Figure 5: Qualitative results of our approach on AWA, where six random unseen class labels are shown on top. The top-3 retrieved images are highlighted in green and the top-3 retrieved false examples are highlighted in red.

4.4. Zero-shot Retrieval Results

We conduct the image retrieval task to illustrate the qualitative results of the proposed framework. Given the semantic embeddings of a specific class, we synthesize a certain number of semantic-consistent representations for a specific class and compute the centroid point as a retrieval query, which is further used to retrieve the nearest samples. To evaluate the performance on the retrieved samples, the mean average precision (mAP) score is adopted. In Figure 4, we compare our proposed SDGZSL method with the base generative model cVAE when retrieving 100%, 50%, and 25% of the images from all the unseen classes on aPaY, AWA, CUB, and FLO. It can be seen that the disentangling module can significantly boost the retrieval performance among all settings, which can also demonstrate the effectiveness of the disentangling modules from the retrieval perspective. Figure 5 illustrates the retrieved examples on the AWA dataset. The class name is given on top, following by the top-3 true positive retrieved images with green boxes and the top-3 false-positive retrieved examples with red boxes. It can be seen that all the false-positive images look very similar to the groundtruth examples. For example, the top-3 failed retrieved samples for the rat are all bats, as the two species have many visual patterns in common. Thus, the nearest neighbor-based retrieval may fail to distinguish the false positive cases. The results demonstrate that the synthesized semantic-consistent representations are close to the samples of the same class in the semantic-consistent feature space.

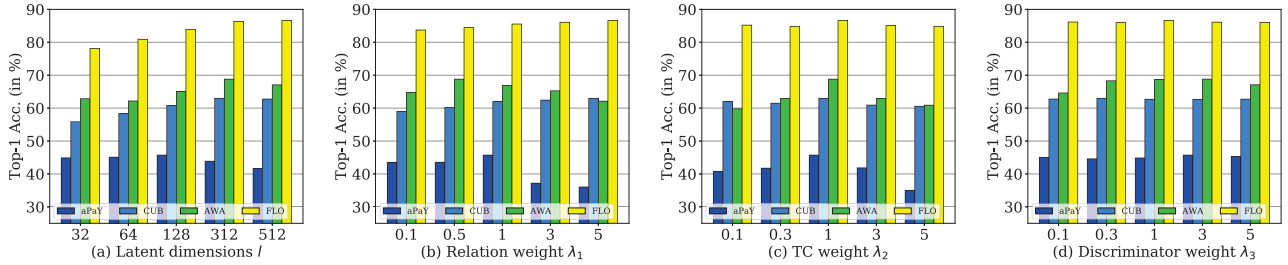


Figure 6: Hyper-parameter study w.r.t. latent dimensions l , relation weight λ_1 , TC weight λ_2 , and discriminator weight λ_3 .




Figure 7: t-SNE visualization of different representations of 50 unseen classes on CUB: (a) semantic-consistent representations h_s ; (b) semantic-unrelated representations h_n .

4.5. Model Analysis

Ablation Study. In this ablation study, we evaluate various stripped-down versions of our full proposed model to validate the contributions of the key disentangling components of the model. In Table 1, we report the GZSL performance of each version on the four benchmark datasets. The best performance is achieved when RN and TC are both applied. In Figure 3, we show the performance comparison between h_s , h_n , and h . From the observations on all datasets, semantic-unrelated features h_n performs much worse than the original latent features h , which proves the ineffectiveness of h_n in GZSL. In contrast, using semantic-consistent features h_s can further improve the performance, which validates the significant impact of h_s when transferring the connection between semantic to semantic-unrelated representations from seen to unseen classes in GZSL. It is worth noting that when taking both h_s and h_n (*i.e.*, h) to classify test samples, the performance is similar to cVAE.

Hyper-Parameter Sensitivity. There are mainly four hyper-parameters that control the objective function, including the number of the disentangled feature dimension l , the weight of *total correlation* term λ_1 , the weight of loss in the relation module λ_2 , and the weight of the discriminator loss λ_3 . To better understand the effect of the disentangling components, we report the sensitivity of the four hyper-parameters in Figure 6.

t-SNE Visualization. To further validate the properties of the disentanglement, we visualize semantic-consistent representations h_s from x in Figure 7 (a) and semantic-unrelated representations h_n from x in (b). We choose all of 50 unseen classes from CUB dataset that have enough classes to show the class-wise comparison. Clearly, as we expected, the semantic-consistent representations h_s is much more discriminative than the semantic-unrelated representations h_n . However, we can still see discriminative patterns from h_n as there are remaining discriminative features even if they are semantic-unrelated, *e.g.*, characteristics that are not annotated in the attributes. We argue that these discriminative features may help classify between these classes, but as the non-semantic discriminative features are not annotated in the attributes, it is intuitively impossible to transfer the semantic-visual relationship from seen classes to unseen classes.

5. Conclusion

In this paper, we propose a novel semantics disentangling approach for generalized zero-shot learning. Specifically, the visual features of an image extracted from pre-trained ResNet101 are further factorized into two independent representations that are semantic-consistent and semantic-unrelated. In our approach, an encoder-decoder architecture is coupled with a relation module to learn the visual-semantic interaction. Further, we leverage the total correlation term to encourage the disentanglement between the two representations. The disentangling encoder-decoder model is incorporated into a conditional variational autoencoder and trained in an end-to-end manner. The generation ability trained on seen classes is transferred to the unseen classes and synthesizes the missing visual samples. We evaluate our proposed method on four image classification datasets. Extensive experiments show that our approach consistently performs better than other state-of-the-arts.

Acknowledgments: This work was partially supported by Australian Research Council DP190101985, CE200100025, DE200101610 and Sichuan Science and Technology Program under Grant 2020YFG0080.

References

- [1] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid. Label-embedding for image classification. *TPAMI*, pages 1425–1438, 2015. [1](#), [10](#), [11](#)
- [2] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, pages 2927–2936, 2015. [11](#)
- [3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, pages 214–223, 2017. [2](#)
- [4] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *TPAMI*, pages 1798–1828, 2013. [2](#)
- [5] R. T. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *NeurIPS*, 2018. [2](#)
- [6] Z. Chen, J. Li, Y. Luo, Z. Huang, and Y. Yang. Canzsl: Cycle-consistent adversarial networks for zero-shot learning from natural language. In *WACV*, pages 874–883, 2020. [6](#)
- [7] C. Eastwood and C. K. Williams. A framework for the quantitative evaluation of disentangled representations. In *ICLR*, 2018. [2](#)
- [8] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *CVPR*, pages 1778–1785, 2009. [6](#)
- [9] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, and T. Mikolov. Devise: A deep visual-semantic embedding model. In *NeurIPS*, pages 2121–2129, 2013. [1](#), [11](#)
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NeurIPS*, volume 27, 2014. [2](#)
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [2](#)
- [12] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016. [2](#)
- [13] H. Jiang, R. Wang, S. Shan, and X. Chen. Transferable contrastive network for generalized zero-shot learning. In *ICCV*, pages 9765–9774, 2019. [1](#), [6](#)
- [14] Diederik P. K. and Max W. Auto-encoding variational bayes. In *ICLR*, 2014. [2](#)
- [15] R. Keshari, R. Singh, and M. Vatsa. Generalized zero-shot learning via over-complete distribution. In *CVPR*, pages 13300–13308, 2020. [1](#), [6](#)
- [16] H. Kim and A. Mnih. Disentangling by factorising. In *ICML*, pages 2649–2658, 2018. [2](#)
- [17] E. Kodirov, T. Xiang, and S. Gong. Semantic autoencoder for zero-shot learning. In *CVPR*, pages 3174–3183, 2017. [1](#), [11](#)
- [18] V. Kumar Verma, G. Arora, A. Mishra, and P. Rai. Generalized zero-shot learning via synthesized examples. In *CVPR*, pages 4281–4289, 2018. [2](#)
- [19] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *TPAMI*, pages 453–465, 2013. [1](#), [6](#), [11](#)
- [20] J. Li, M. Jin, K. Lu, Z. Ding, L. Zhu, and Z. Huang. Leveraging the invariant side of generative zero-shot learning. In *CVPR*, pages 7402–7411, 2019. [6](#)
- [21] S. Liu, M. Long, J. Wang, and M. I. Jordan. Generalized zero-shot learning with deep calibration network. In *NeurIPS*, pages 2005–2015, 2018. [1](#), [6](#)
- [22] Y. Liu, J. Guo, D. Cai, and X. He. Attribute attention for semantic disambiguation in zero-shot learning. In *CVPR*, pages 6698–6707, 2019. [1](#), [6](#)
- [23] S. Min, H. Yao, H. Xie, C. Wang, Z. J. Zha, and Y. Zhang. Domain-aware visual bias eliminating for generalized zero-shot learning. In *CVPR*, pages 12664–12673, 2020. [6](#)
- [24] S. Narayan, A. Gupta, F. S. Khan, C. G. Snoek, and L. Shao. Latent embedding feedback and discriminative features for zero-shot classification. In *ECCV*, pages 479–495, 2020. [1](#), [2](#), [6](#)
- [25] M. E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, 2008. [6](#)
- [26] S. Reed, Z. Akata, H. Lee, and B. Schiele. Learning deep representations of fine-grained visual descriptions. In *CVPR*, pages 49–58, 2016. [6](#)
- [27] B. Romera P. and P. Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, pages 2152–2161, 2015. [11](#)
- [28] E. Schonfeld, S. Ebrahimi, S. Sinha, T. Darrell, and Z. Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *CVPR*, pages 8247–8255, 2019. [1](#), [2](#), [6](#)
- [29] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In *NeurIPS*, pages 3483–3491, 2015. [5](#)
- [30] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018. [4](#)
- [31] B. Tong, C. Wang, M. Klinkigt, Y. Kobayashi, and Y. Nonaka. Hierarchical disentanglement of discriminative latent features for zero-shot learning. In *CVPR*, pages 11467–11476, 2019. [2](#), [3](#), [6](#), [7](#)
- [32] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [6](#)
- [33] S. Watanabe. Information theoretical analysis of multivariate correlation. *IBM J. Res. Dev.*, pages 66–82, 1960. [2](#)
- [34] Y. Xian, Z. Akata, G. Sharma, Q. Nguyen, M. Hein, and B. Schiele. Latent embeddings for zero-shot classification. In *CVPR*, pages 69–77, 2016. [1](#), [11](#)
- [35] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata. Feature generating networks for zero-shot learning. In *CVPR*, pages 5542–5551, 2018. [1](#), [2](#), [6](#)
- [36] Y. Xian, S. Sharma, B. Schiele, and Z. Akata. f-vaegan-d2: A feature generating framework for any-shot learning. In *CVPR*, pages 10275–10284, 2019. [6](#)
- [37] Y. Yu, Z. Ji, J. Han, and Z. Zhang. Episode-based prototype generating network for zero-shot learning. In *CVPR*, pages 14035–14044, 2020. [2](#), [6](#)
- [38] Y. Zhu, M. Elhoseiny, B. Liu, X. Peng, and A. Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *CVPR*, pages 1004–1013, 2018. [1](#)