

Zero-Shot Emotion Recognition via Affective Structural Embedding

Chi Zhan¹, Dongyu She¹, Sicheng Zhao²✉, Ming-Ming Cheng¹, Jufeng Yang¹✉

¹Nankai University ²University of California, Berkeley

{chizhan_nt, sherry6656}@163.com, schzhao@gmail.com, {cmm, yangjufeng}@nankai.edu.cn

Abstract

Image emotion recognition attracts much attention in recent years due to its wide applications. It aims to understand the emotional response of humans, where candidate emotion categories are generally defined by specific psychological theories. However, with the development of psychological theories, emotion categories become increasingly diverse, fine-grained, and difficult to collect samples. In this paper, we investigate zero-shot learning (ZSL) problem in the emotion recognition task, which aims to recognize the new unseen emotions. Specifically, we propose an affective structural embedding framework, utilizing mid-level semantic representation, i.e., adjective-noun pairs (ANP) features, to construct an intermediate embedding space. By doing this, the learned intermediate space can bridge the affective gap between low-level visual features and high-level semantics. In addition, we introduce an adversarial constraint to combine the visual and affective embeddings so as to retain the discriminative capacity of visual features and the affective structural information of semantic features during training process. Our method is evaluated on five widely-used affective datasets and the experimental results show that the proposed algorithm outperforms the state-of-the-art approaches.

1. Introduction

With the rapid development of social media, more and more people prefer to record their lives and express opinions via visual contents, e.g., images and videos [47]. In particular, computational understanding emotions of online images has attracted increasing attention from academia and industry due to its various applications, e.g., opinion mining [30], online advertisement [15] and social networks [18].

In the past few years, many methods [38, 39, 43, 49] have made huge progress for image emotion recognition, which aims to classify emotions evoked by image content. Most existing methods follow the general view in psychology that a specific emotion can be recognized as a fixed

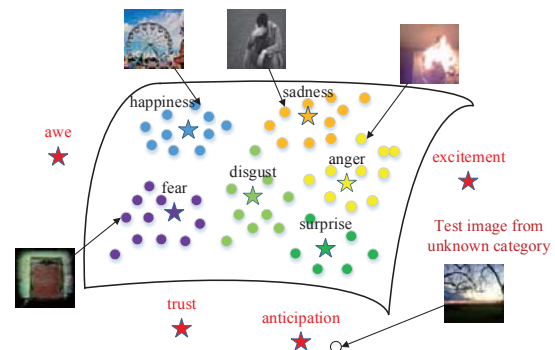


Figure 1. Overview of zero-shot emotion recognition problem. Each class has a prototype (denoted by a star). Training classes are on the manifold with different colors. Test unseen classes are in red and side information is provided to determine where an unseen classes locates. To classify an unseen image, we assign it a label that corresponds to the nearest unseen prototypes.

number of basic emotions. Peng *et al.* [28] train regression models to predict the probabilities of Ekman’s six basic emotions [11, 10], including happiness, sadness, disgust, anger, fear, and surprise. There are also methods employing different psychological theory for emotion modeling. For example, Yang *et al.* [40] jointly optimize emotion classification and distribution learning task according to Mikels’ wheel [24], which replace happiness and surprise in Ekman’s basic emotions with amusement, content, awe, and excitement. However, with the development of psychological theories, basic emotion categories become increasingly fine-grained. Traditional supervised learning methods can only recognize the seen classes, e.g., happiness, anger and other four emotions on the manifold as shown in Fig. 1. Such recognition model trained on the pre-defined categories cannot recognize emotions dynamically, when new categories are explored according to different psychological theories. In addition, it is labor-intensive and time-consuming to collect samples for rare emotion categories.

Zero-shot learning (ZSL) [37] aims to recognize new categories that are not exists in the training set, which has been widely used in various vision tasks [9, 34, 51]. The conventional zero-shot learning methods usually build a common

space based on the correspondence between the seen images and their class semantic representations. The space are also shared by both seen and unseen classes, which rely on the side information (*e.g.* attributes and Word2vec) about how unseen classes are semantically related to the seen classes. Then zero-shot learning can be simplified into a nearest neighbor search task, and test images will be assigned to the nearest unseen class in the common embedding space. Such ZSL paradigm relies on the cross-modality similarity between visual features and class semantic representations. There exists an affective gap between low-level image features and high-level emotional semantics [23, 46, 50] and directly computing similarities is hard to describe the similarity relationship correctly between them. Thus, zero-shot emotion recognition becomes more challenging.

In this paper, we propose an affective structural embedding framework using the mid-level semantic representations, *i.e.* adjective-noun pairs (ANP) [6] features, to construct an intermediate embedding space. Both visual and class semantic features are embedded into the learned embedding space and aligned with the affective structure of the ANP features. Therefore, our method can effectively bridge the affective gap, which addresses the zero-shot as well as generalized zero-shot learning problems in emotion recognition. Note that in the zero-shot setting, training and test classes are disjoint and in the more realistic generalized zero-shot setting, training classes are present at test time. In the training process, both visual embedding and affective embedding are dynamically changing, and it is difficult to combine them directly and effectively. We further introduce an affective adversarial constraint to force the visual embedding to choose an embedding space that preserves the affective structural information.

Our contributions are summarized as follows: 1. We propose an end-to-end affective structural embedding framework to learn an intermediate space and preserve emotion-related information, in which both visual and class semantic features are learned. To the best of our knowledge, this is the first zero-shot learning work on image emotion recognition. 2. We apply an affective adversarial constraint to retain the discriminative capacity of visual features and the affective structural information of semantic features during training process. The experimental results show the superiority of the proposed method over the state-of-the-art methods on five public datasets.

2. Related Work

2.1. Image Emotion Recognition

Previous approaches for image emotion recognition mainly focus on the classification problem utilizing hand-crafted features or deep learning features. In the early years, many methods design hand-crafted features with different

levels to recognize image emotion. For low-level features, Machajdik *et al.* [23] define a combination of hand-crafted features according to aesthetics and psychology theory, including color, texture, and composition. Zhao *et al.* [48] further investigate more robust visual features related to art principles as the mid-level representation. In another research [7], adjective-noun pairs are regarded as mid-level semantic features, and a bank of visual sentiment classifiers (SentiBank) is proposed for image affective analysis.

Recently, Convolutional Neural Network (CNN) [17, 14] has been applied to image emotion recognition tasks and achieved satisfactory results. Inspired by the research [7], DeepSentiBank [8] adopts a deep CNN model to construct a detector for visual sentiment concept based on the adjective-noun pairs. You *et al.* [42] propose a novel progressive CNN architecture PCNN, to make use of large noisy web data for binary sentiment classification. Yang *et al.* [38] explore the relation between emotions via deep metric learning and employ a multi-task framework to optimize retrieval and classification simultaneously. Later several methods [41, 39] consider both global and local information for image emotion recognition.

All the above methods employ a supervised manner to learn the relationship between image visual content and emotions, which depends on the pre-defined psychological theories. In addition, many recent methods [5, 22] suggest the types of emotions are much more various than previously assumed. Because of the diversity of emotional descriptions, it is difficult to assign an emotional image to an existing stereotypical label practically. The focus of our research is to classify a novel emotion class which does not appear in the training set.

2.2. Zero-Shot Learning

Zero-shot learning aims to classify unseen classes without any training data. To cope with the challenge, most methods [12, 45, 1, 35] utilize semantic attributes describing cross-class properties to transfer the semantic knowledge from the seen classes to the novel unseen classes. However, semantic attributes need manual definition and annotation, which limits the scalability of the above approaches. Several methods [13, 2, 31] explore zero-shot learning using word vector representations [25], which is constructed by the large-scale text corpora in an unsupervised way. For the affective datasets without attribute annotations, we choose word vector representations as the class semantic features.

Many zero-shot learning methods use the embedding-based method seek to measure the similarities between the visual features and class semantic features in different embedding space. For example, DeViSE [13] directly learns a linear mapping from the image space to the semantic space using a ranking loss function. SJE [2] optimizes the struc-

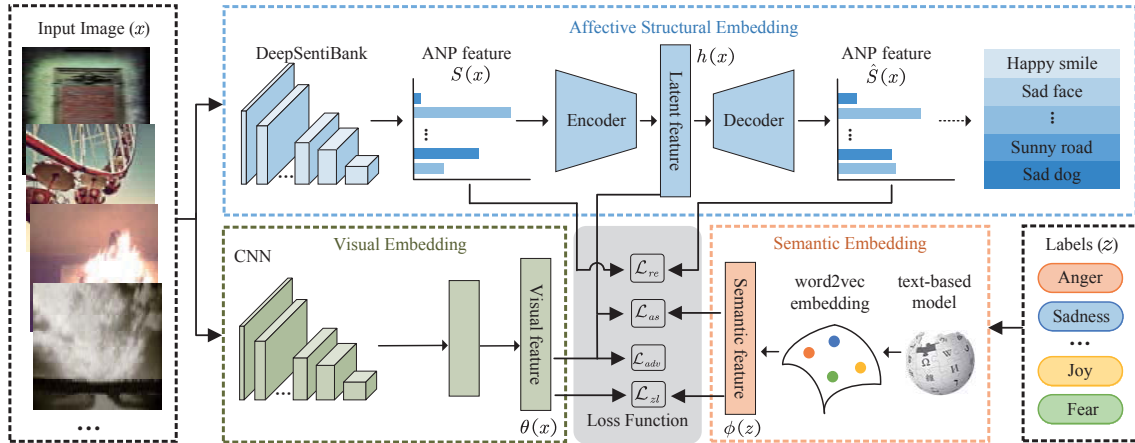


Figure 2. Pipeline of the proposed approach for zero-shot emotion recognition. Given the training image, we first extract the ANP features using the pre-trained DeepSentiBank detector and feed them into an auto-encoder to conduct the latent ANP space. Meanwhile, visual features are also embedded in the latent ANP space to align with the embedded semantic features and measure the similarities for zero-shot emotion recognition. The whole framework is trained by optimizing the multi-loss function in an end-to-end manner.

tural SVM loss function to learn the bilinear compatibility between visual and semantic space. SAE [21] proposes a semantic autoencoder to regularize the model. It firstly projects the image features to the semantic space and further reconstructs them back to the visual space. DEM [44] chooses to embed the semantic features to the visual space. PSR [3] further considers the inter-class semantic relationships during the mapping process. Moreover, many zero-shot learning approaches [19, 20, 32] learn to embed both visual and semantic features into a latent intermediate space.

However, all the above ZSL methods fail to capture the specific emotion information for emotion recognition problem. The visual and class semantic features are located in different structural spaces, both of which are independent of emotions. Our model utilizes the mid-level semantic representations to construct an intermediate space. It can reserve emotion-related information and effectively bridge the affective gap.

There are also a few recent methods that tackle the zero-shot learning problem utilizing adversarial learning methods and generative adversarial networks (GAN). GAZSL [52] leverages GANs to imagine the visual features given the noisy textual descriptions from Wikipedia. CVAE-ZSL [27] proposes to use conditional variational autoencoder to generate samples for unseen classes. f-CLSWGAN [36] applies GAN to generate image features conditional on class attributes. The idea of GAN and adversarial learning methods are to train a generator that can fool a discriminator to confuse the distributions of the generated and true samples. The max-min training procedure can lead the generator to model the data distribution. Our method is similar to the GAN applied in the feature level. In this paper, we employ adversarial learning to bridge the

gap between visual and affective features.

3. Methodology

In this section, we first formalize the zero-shot emotion task and then introduce the proposed affective structural embedding model. As shown in Fig. 2, we propose an independent affective structural embedding with traditional visual-semantic embedding. Specifically, the extracted ANP features are fed into an auto-encoder to learn the latent ANP space, and then both visual and class semantic features are embedded into the learned ANP space so as to effectively bridge the affective gap. In addition, we introduce an affective adversarial constraint to effectively combine visual and ANP features so as to retain the discriminative capacity and the affective structural information.

3.1. Problem Definition

Following conventional zero-shot learning problem, we split the affective dataset with s seen classes and u unseen classes. The training set is then defined as $D_S = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, where $x_i^s \in \mathcal{X}_S$ denotes the i -th image of the seen class and $y_i^s \in \mathcal{Y}_S$ is the corresponding class label. We define the test set as $D_U = \{(x_j^u, y_j^u)\}_{j=1}^{n_u}$, where $x_j^u \in \mathcal{X}_U$ denotes the j -th unseen image and $y_j^u \in \mathcal{Y}_U$ is the class label. The seen and unseen classes are disjoint, *i.e.* $\mathcal{Y}_S \cap \mathcal{Y}_U = \emptyset$. Additionally, we choose the word vector z_i^s and z_j^u obtained by the NLP model [29] as the class semantic features. Note that during the training stage, only seen class images are used to learn the classifier model with the assistance of semantic information z^s . Given a test image x^u and the semantic feature z^u , we aim to predict the corresponding class y^u .

3.2. Affective Structural Alignment

ZSL problems are usually addressed by measuring similarities between the visual and class semantic features. Since affective datasets, *e.g.*, the FI dataset, lack the attribute annotations for class semantic features, we utilize a text-based model, *i.e.* Word2vec [26], which is learned from a large-scale text corpus in an unsupervised manner requiring little or no human labor to annotate. However, they only capture the weakly semantic relationship between different classes and are not discriminative enough to classify. What’s more, visual features directly extracted from pre-trained CNN model are also limited by the affective gap. In order to align visual and class semantic features with more emotional structure in the latent intermediate space, we introduce an independent affective structural embedding.

First, we utilize the mid-level semantic representations ANP features to construct an intermediate latent space. Given a training image x , we choose the pre-trained ANP detector DeepSentiBank [8] $S(\cdot)$ to extract the ANP feature $S(x) \in \mathbb{R}^d$. To learn an effective latent space for compact affective representation of original affective features $S(x)$, we adopt an auto-encoder model. Suppose the input of the auto-encoder as the d -dimension ANP feature $S(x) \in \mathbb{R}^{d \times n}$ which contains n samples. The encoding part of the auto-encoder embeds the input into the l -dimension latent space $h(x)$ using FC layers, which can be defined as: $h(x) = f(W_1 S(x))$. Similarly, the decoder aims to reconstruct the input as $\hat{S}(x) \in \mathbb{R}^{d \times n}$: $\hat{S}(x) = f(W_2 h(x))$. W_1 and W_2 are the weight matrices of the FC layers and $f(\cdot)$ is the activation function. To learn the auto-encoder parameters, the input and output of the auto-encoder should be close enough by optimizing the following loss:

$$\mathcal{L}_{re} = \|S(\hat{x}) - S(x)\|_2^2. \quad (1)$$

Meanwhile, the class semantic feature z corresponding to the training image x is also projected into the learned latent ANP space by the non-linear embedding $\phi(\cdot)$. In other words, we hope to minimize the distance between embedded class semantic features $\phi(z)$ and learned latent ANP features $h(x)$. So the first part of loss function is defined as $\|h(x) - \phi(z)\|_2^2$. On the other hand, class semantic space and learned latent space have different inter-class structures. As the auto-encoder is used to reconstruct the ANP features and make the latent features preserving emotion-relevant information, we seek to match the structures of class semantic features and the ANP features in the learned latent space. Inspired by [16, 33], we project the class semantic features to the mean of the ANP features of the corresponding classes. Thus the second part of loss function is defined as $\|C^h - \phi(z)\|_2^2$, where C^h denotes the mean vector of the latent ANP features $h(x)$ for each class. By optimizing all the above constraints, embedded class semantic

features could learn emotion-relevant class representations, which are better associated with the latent emotional concepts. The affective structural alignment loss is formulated as:

$$\mathcal{L}_{as} = \|h(x) - \phi(z)\|_2^2 + \|C^h - \phi(z)\|_2^2. \quad (2)$$

The total affective structural embedding is optimized by the combination of reconstruction loss and the affective structural alignment loss:

$$\mathcal{L}_{ae} = \mathcal{L}_{re} + \mathcal{L}_{as}. \quad (3)$$

3.3. Affective Adversarial Constraint

To address the zero-shot learning problem, we also embed visual and class semantic features to construct visual-semantic embedding model and measure similarities. Suppose $\theta(\cdot)$ denotes the visual features embedding process, the loss to align the visual and semantic features is defined as:

$$\mathcal{L}_{zl} = \|\theta(x) - \phi(z)\|_2^2. \quad (4)$$

Currently, both the traditional visual-semantic embedding and proposed affective structural embedding contribute to recognizing unknown emotions. Visual features have a better discriminative capacity while ANP features contain some affective structural information, which bridges the affective gap and useful for the zero-shot emotion learning. However, it is difficult to combine both embeddings effectively during training process as both visual and affective structural embedding are dynamically changing. Our goal is to retain the discriminative capacity of visual features $\theta(x)$ and combine the rich affective structural information preserved in $h(x)$. To this end, we apply an adversarial constraint, which try to fool a discriminator network D so that the output visual features are as similar as embedded ANP features:

$$\mathcal{L}_{adv} = \mathbb{E}_x (\log D(h(x))) + \mathbb{E}_x (\log [1 - D(\theta(x))]), \quad (5)$$

where $\theta(\cdot)$ tries to minimize \mathcal{L}_{adv} against D that tries to maximize it. Considering this kind of adversarial learning is tricky to optimize, in order to obtain better train stability, we adopt the strategy of WGAN [4]. Please refer to [4] for more details.

Combining all above mentioned constraints, the whole model is trained by the following loss function:

$$\mathcal{L} = \mathcal{L}_{zl} + \mathcal{L}_{ae} + \mathcal{L}_{adv}. \quad (6)$$

3.4. ZSL Prediction

Given the test image x and the set of class semantic features \mathbf{Z} of candidate emotion classes, we can classify the unseen emotion class via the simple nearest-neighbor research. More specifically, the test image and candidate class

semantic features are fed into the visual and semantic embedding branch separately to get $\theta(x)$ and $\phi(\mathbf{Z})$. Then the test image is recognized by calculating its distance to the class semantic embedding features in the latent space:

$$\hat{y}^t = \min_{y \in \mathcal{Y}_U} \|\theta(x) - \phi(\mathbf{Z}^y)\|_2^2, \quad (7)$$

where \mathbf{Z} denotes the semantic features associated with the emotion label y . For the generalized ZSL setting, we only need to modify the candidate space of labels as $y \in \mathcal{Y}_U \cup \mathcal{Y}_S$.

4. Experiments

In this section, we first introduce the detailed experimental setup, including the datasets, implementation details, and evaluation metrics. And then we compare with the state-of-the-art approaches and analyze the results.

4.1. Datasets

We perform our experiments on five datasets, including Flickr and Instagram (FI) [43], IAPSA [24], ArtPhoto [23], Emotion6 [28] and Abstract Paintings [23]. The FI dataset is collected from 3 million weakly labeled web images Flickr and Instagram by labeling with Mikels' eight emotion categories. A group of 225 Amazon Mechanical Turk workers was employed to label the images. In total, 23,308 images receiving at least three agreements between workers are included in the FI dataset. The International Affective Picture System (IAPS) is widely used in visual sentiment analysis research, which contains 395 pictures from IAPS and is also labeled with Mikel's eight emotion categories. ArtPhoto includes 806 art photographs from photo sharing sites and the owner of each image provides the ground truth labels. Abstract Paintings includes 228 abstract paintings consisting of texture and color. Emotion6 is collected from Flickr for the sentiment prediction, which contains 1980 images and is annotated by seven emotional categories.

4.2. Implementation Details

We employ the ResNet-50 model as the backbone CNN network and initialize our framework with the weights from the pre-trained model on ImageNet. In addition, we apply the pre-trained DeepSentiBank [8] detector to extract 2089-dimension ANP features. For the auto-encoder, the dimension of the latent ANP feature is fixed to 1024. We utilize a fully connected layer before the ReLU layer in order to embed both the visual and semantic features to the latent ANP space. The discriminator D is composed of two fully connected layers and a ReLU layer, and takes 1024-d features as input. The learning rate of stochastic gradient descent (SGD) is $1e-4$ and the weight decay is $1e-3$. The momentum is set as 0.9. We implement our model with Pytorch and run all experiments on an NVIDIA GTX 1080Ti GPU.

For the class semantic feature, we choose to use Word2vec [26], where each instance is represented by a 300-dimensional vector. The features are constructed automatically from large unlabeled text corpora without additional manual annotation.

4.3. Evaluation Metrics

Following previous ZSL methods [37], we employ the average per-class accuracy as the evaluation metric. For the generalized ZSL setting [37], we compute the average per-class accuracy on unseen classes ($A_{U \rightarrow T}$) and average per-class accuracy on seen classes ($A_{S \rightarrow T}$) when the prediction label set is the union of seen and unseen classes. We also compute the harmonic mean (H) on seen classes and unseen classes, *i.e.*, $H = 2 * (A_{U \rightarrow T} * A_{S \rightarrow T}) / (A_{U \rightarrow T} + A_{S \rightarrow T})$.

4.4. Results and Analysis

To evaluate the effectiveness of our model for zero-shot emotion recognition, we compare with a variety of ZSL methods, including common ZSL methods (*i.e.* LATEM [35], SSE [45], SAE [21] and DEM [44]) and the recent ZSL methods (*i.e.* LAD [20], CDL [19] and RN [32]). Since the ANP features extracted by the DeepSentiBank detector can be used as visual features in image emotion recognition, we also report the zero-shot recognition results using extracted ANP features.

We evaluate the performance of the proposed zero-shot emotion recognition method on five affective datasets. As shown in Table 1, we conduct the experiment on each dataset with two kinds of split strategies in order to prove the effectiveness and robustness of the proposed method. In detail, for the datasets with 8 emotions, the splits of training classes and testing classes are 6:2 and 4:4; for the Emotion6 dataset with 6 emotions, the splits of training classes and testing classes are 4:2 and 3:3. For the FI dataset with 6 : 2 split setting, our method attains 68.87%, which is slightly higher than the state-of-the-art reported by CDL (67.07%). Among the comparison methods, CDL gets the best performance followed by DEM (65.49%) while SSE gets the worst performance. Our method and CDL both consider the structural information of different space during the embedding process, so the performance is much better than other methods. Besides, our method further considers the specific affective structural information which is relevant to the emotion recognition. What's more, we could utilize the affective adversarial constraint to automatically find the optimal solution of combining the two features when both features are changing during the training process. Thus, our method could obtain the best performance. For the 4 : 4 split setting, our method has achieved gains up to 3.77% than DEM (50.96%). For the other small datasets, our method can still achieve 1.29% ~ 3.97% improvements.

We also observe that the ANP features (S) have a similar

Table 1. Zero shot emotion recognition accuracy (%) of all methods on the FI, ArtPhoto, Abstract, IAPSA and Emotion6 datasets. We evaluate the proposed model with several baseline zero-shot learning methods. S denotes the DeepSentiBank features, while D denotes the CNN-based features, and M denotes the concatenation of the DeepSentiBank and CNN-based features.

Setting	FI		ArtPhoto		Abstract		IAPSA		Emotion6		
	6 : 2	4 : 4	6 : 2	4 : 4	6 : 2	4 : 4	6 : 2	4 : 4	4 : 2	3 : 3	
[35]	D	53.32±0.42	32.73±0.32	38.57±0.25	27.66±0.76	56.71±0.54	24.64±0.48	49.81±0.31	25.57±0.22	52.88±0.52	32.42±0.44
	S	57.79±0.44	38.59±0.53	42.98±0.32	25.26±0.86	49.04±0.22	22.39±0.30	43.74±0.74	25.26±0.63	55.91±0.15	29.89±0.32
	M	58.25±0.36	37.57±0.42	44.35±0.88	24.17±0.19	42.37±0.79	25.68±0.11	41.24±0.52	29.35±0.43	55.12±0.21	34.82±0.31
[45]	D	42.67±0.04	33.61±0.04	45.57±0.02	26.55±0.01	47.36±0.05	21.41±0.04	41.34±0.01	22.97±0.07	42.12±0.11	28.99±0.07
	S	43.37±0.03	21.04±0.04	42.85±0.04	26.83±0.06	44.73±0.05	17.64±0.02	50.73±0.07	27.47±0.02	53.18±0.04	27.98±0.07
	M	42.02±0.01	33.55±0.05	40.63±0.08	20.51±0.05	49.64±0.04	22.33±0.02	44.51±0.09	30.62±0.01	52.85±0.05	31.21±0.06
[21]	D	61.12±0.00	37.34±0.00	49.66±0.00	23.45±0.00	60.53±0.00	29.41±0.00	49.04±0.00	29.73±0.00	44.24±0.00	38.89±0.00
	S	54.66±0.00	31.60±0.00	51.70±0.00	22.88±0.00	42.11±0.00	25.49±0.00	53.85±0.00	32.43±0.00	45.45±0.00	37.88±0.00
	M	57.82±0.00	26.57±0.00	45.58±0.00	23.45±0.00	52.63±0.00	28.43±0.00	48.08±0.00	29.28±0.00	54.55±0.00	35.86±0.00
[20]	D	51.44±0.13	34.94±0.31	41.27±0.24	20.53±0.22	43.21±0.18	13.75±0.42	38.42±0.25	27.66±0.17	45.36±0.33	33.65±0.29
	S	43.97±0.17	29.36±0.34	43.51±0.20	22.34±0.16	45.44±0.10	15.24±0.32	35.71±0.14	22.68±0.41	40.04±0.26	28.34±0.21
	M	44.18±0.11	26.53±0.25	42.18±0.18	22.03±0.28	42.11±0.27	23.53±0.33	50.96±0.14	18.83±0.25	47.27±0.25	32.93±0.18
[44]	D	65.49±0.52	49.37±0.73	48.30±0.39	29.66±0.48	63.42±0.88	32.35±0.63	47.12±0.71	34.69±0.83	50.61±0.57	33.74±0.49
	S	62.06±0.48	34.70±0.63	48.97±0.53	31.35±0.66	60.53±0.74	32.35±0.60	46.15±0.53	30.18±0.77	51.66±0.61	36.36±0.51
	M	64.73±0.43	50.96±0.78	50.33±0.49	27.97±0.64	60.71±0.88	30.72±0.53	44.52±0.76	31.38±0.78	50.43±0.62	31.28±0.46
[32]	D	64.97±0.83	47.83±0.92	40.45±0.87	28.86±0.68	57.85±0.69	26.69±0.74	40.57±0.85	32.43±0.67	56.51±0.82	39.49±0.93
	S	50.34±0.79	31.57±1.02	44.23±0.88	23.79±0.71	42.51±0.72	16.55±0.65	43.31±0.61	30.76±0.77	52.45±0.98	29.21±0.91
	M	63.99±0.80	49.31±0.89	47.14±0.91	25.32±0.53	56.16±0.82	29.52±0.64	43.98±0.76	33.62±0.88	49.85±1.07	33.64±0.84
[19]	D	67.03±0.54	41.28±0.45	50.52±0.72	30.46±0.61	52.11±0.53	24.11±0.51	52.88±0.76	30.56±0.69	51.67±0.70	36.36±0.48
	S	61.44±0.51	36.88±0.42	48.57±0.79	25.81±0.67	53.64±0.62	22.62±0.46	51.00±0.59	35.74±0.60	56.97±0.78	40.10±0.67
	M	67.07±0.48	41.37±0.52	48.35±0.70	29.22±0.72	55.19±0.64	27.42±0.53	48.73±0.73	33.92±0.78	53.86±0.64	37.91±0.60
Ours	68.87 ±0.79	54.73 ±0.64	53.22 ±0.97	35.58 ±1.03	64.71 ±1.17	34.45 ±0.93	57.82 ±0.82	38.30 ±0.75	59.94 ±0.86	42.83 ±1.14	

Table 2. Generalized ZSL recognition accuracy on the FI dataset following 6 : 2 split setting and cross dataset recognition accuracy of all methods between the FI dataset and Emotion6 dataset. $A_{U \rightarrow T}$ denotes the recognition accuracy on unseen classes, while $A_{S \rightarrow T}$ denotes the average recognition accuracy on seen classes. H denotes the harmonic mean. 'FI→Emo6' denotes using the images of FI dataset belong to the common categories as training set and the images of Emotion6 dataset belong to the other categories as the test set. 'Emo6→FI' denotes the similar setting but change the dataset.

Method	$A_{U \rightarrow T}$	$A_{S \rightarrow T}$	H	FI→Emo6	Emo6→FI
LATEM [35]	1.82	55.31	3.54	51.21	26.43
RN [32]	3.23	62.56	6.14	59.85	29.22
SSE [45]	7.51	53.57	13.17	51.32	29.40
DEM [44]	13.43	56.25	21.68	56.52	22.36
LAD [20]	20.83	59.46	30.85	52.27	21.03
SAE [21]	24.25	65.59	35.42	54.70	30.03
CDL [19]	26.48	54.87	35.72	55.15	32.34
Ours	28.12	66.57	39.54	61.94	34.48

performance with the deep features (D) extracted by deep learning models pre-trained on ImageNet. For example, on the Emotion6 dataset, the results of using ANP features are generally better than using deep features, *e.g.* LATEM, SSE, SAE, DEM, and CDL. For other datasets such as the FI dataset, the performance of using ANP features is slightly lower in most cases. This demonstrates that the ANP features contain some affective structural information, which may be not discriminative for classification tasks, but use-

ful for the zero-shot learning problem. Our method utilizes deep features with the ANP features as a supplement and shows the effectiveness of considering both discriminative visual features and affective structural information. We have also validated all the comparison methods using both ANP and deep features and reported the results in the rows marked with M in Table 1. That also shows that applying both two features directly induces no performance gain, while our method effectively improves zero-shot emotion recognition performance.

We further report the generalized ZSL recognition accuracy of all methods on the FI dataset following 6 : 2 split setting in Table 2. For the accuracy on unseen classes $A_{U \rightarrow T}$, CDL obtains the best performance 26.48% among the comparison methods, while SAE achieves 65.59% when it comes to the accuracy on seen classes $A_{S \rightarrow T}$, which is much better than CDL. For $A_{U \rightarrow T}$ and $A_{S \rightarrow T}$, our method obtains 28.12% and 66.57%. Compared with the most competitive CDL, our method improves the harmonic mean by 3.82% on the FI dataset. Our method outperforms all compared methods in all three cases.

4.5. Ablation Study

We conduct ablation experiments to illustrate the effectiveness of the affective structural alignment and the affective adversarial constraint in Table 3. In particular, "Base" denotes the basic visual-semantic embedding to conduct

Table 3. Ablation experiment on the FI dataset. ‘Base’ denotes using the basic visual-semantic embedding to conduct zero-shot emotion recognition. \mathcal{L}_{as} and \mathcal{L}_{re} denote two parts of affective structural embedding. \mathcal{L}_{as}^* denotes the structural alignment loss without the second part. \mathcal{L}_{adv} denotes using the affective adversarial constraint.

Base	\mathcal{L}_{adv}	\mathcal{L}_{as}^*	\mathcal{L}_{as}	\mathcal{L}_{re}	FI
✓					65.12
✓		✓			65.46
✓			✓		66.41
✓	✓		✓		67.53
✓			✓	✓	67.28
✓	✓		✓	✓	68.87

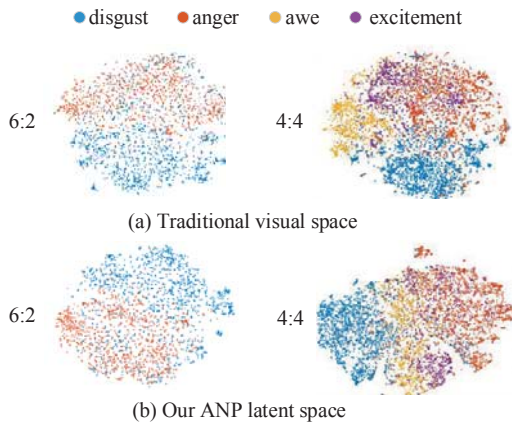


Figure 3. t-SNE plots of the distribution of unseen class visual samples from the FI dataset in traditional visual space and our ANP latent space. ‘6:2’ and ‘4:4’ denote the two splits for zero-shot emotion recognition.

zero-shot emotion recognition, where visual and class semantic features are directly embedded into the common space and measured similarities. From the results, we can clearly see that the affective adversarial constraint plays a significant role in improving the zero-shot recognition accuracy by 1.12% on the FI dataset. The results validate that adversarial learning combining two embedded visual and ANP features is superior to those direct combination or in a manual combination rule.

Affective structural embedding optimized by \mathcal{L}_{as} and both \mathcal{L}_{as} and \mathcal{L}_{re} further boost the ZSL accuracy by 1.29% and 2.16%. Among the affective structural embedding, aligning the class semantic features with the center of the latent ANP features also results in 0.95% performance improvement. The results demonstrate that considering emotional structural information to align the visual and semantic features contributes a lot to zero-shot emotion recognition. Fig. 3 illustrates the distribution of unseen class visual samples in traditional visual space and our ANP latent space. It clearly shows that the embedded visual features are more separated from other classes in the ANP space. Finally, we obtain the best performance by training all the components,

Table 4. Zero-shot emotion recognition accuracy (%) on the FI dataset with different choices of testing classes. Note that * means that we choose different testing classes under the same train/test ratio corresponding to Table 1. Here, we take ‘excitement’ and ‘sadness’ as test classes for the 6:2 split setting and all the negative emotions as test classes for the 4:4 split setting.

Setting	LATEM	SSE	SAE	LAD	CDL	DEM	RN	Ours
6 : 2*	70.47	63.73	67.78	61.84	68.70	71.98	72.41	74.03
4 : 4*	27.09	37.94	42.24	38.30	41.95	37.84	39.87	42.92

which shows the complementarity of these contributions.

We also explore the influence of different testing classes in Table 4 under the same train/test ratio in Table 1. For the 6 : 2* split setting, we take the ‘excitement’ and ‘sadness’ as test classes and others as training classes while in Table 1 we choose two negative emotions as test classes. For the 4 : 4* split setting, we take all the negative emotions as test classes (*i.e.* ‘fear’, ‘sad’, ‘disgust’ and ‘anger’) and all the positive emotions as training classes while in Table 1 training and test classes are both two positive emotions and two negative emotions. We can see that, the performance of almost all the compared methods is improved when test emotions are in two polarities and is decreased when test emotions are all negative and close to each other. Our method achieves the state-of-the-art results consistently with different emotion prediction configurations.

4.6. Cross Datasets Recognition

To better evaluate the performance of zero-shot learning on image emotion recognition tasks, we further conduct the cross datasets experiment for emotion recognition. Conventional zero-shot learning methods assume the testing images are sampled from the same distribution with the training images, which is inconsistent with the real situation. Thus we try to recognize unseen emotion categories in different datasets. The FI and Emotion6 datasets share the four emotion categories including sadness, disgust, fear and anger. Specifically, we take the four emotion categories from one dataset as the training set and other categories from another dataset as the test set. Ideally, we focus on the image emotion recognition, which means the more emotion-related features and embedding, the higher the performance.

Table 2 shows the results of the cross datasets recognition. The proposed method obtains the best results compared with the other zero-shot methods in both two experiments. In particular, the proposed method achieves 2.08% improvements under FI→Emo6 setting and improves the recognition accuracy by 2.14% improvements under Emo6→FI setting. Affective structural information provided by ANP features are shared by different datasets. Considering both the affective structural information and affective enhanced visual features ensure the ability of our method to cross datasets recognition.



Figure 4. Qualitative results of CDL and our method on FI (a) and Emotion6 (b). Predictions of unseen class labels for both two datasets are listed from top to bottom in probability and the most probable prediction is at the top. The ground truth labels are in red. “ex, aw, di, an, jo, sa and su” denote excitement, awe, disgust, anger, joy, sadness and surprise respectively.



Figure 5. Qualitative results for cross dataset recognition on the FI and Emotion6 datasets. Misclassified images are marked with red bounding boxes.

4.7. Qualitative Results

We present some qualitative analysis for our proposed algorithm and CDL on FI and Emotion6 dataset in Fig. 4. As shown in Fig. 4 (a), the image of the class “excitement” and “anger” can be correctly predicted among the four unseen classes. While the example image of the class “disgust” is predicted to the class “anger”. Both emotion classes are similar in the emotion theories and the same image may evoke two emotions in different situations. Actually, the ground truth “disgust” comes in second just following “disgust” in the predictions. For the last examples in the Emotion6 dataset Fig. 4 (b), our model predictions confuse the class “sadness” and “surprise”. These results further prove the ambiguity of emotions that even people will feel difficult to distinguish such similar emotions. Compared with the most competitive method, we can also see that our method significantly outperforms CDL in these examples. Although CDL still gives some correct results in cases, the predic-

tions do not show the relationship between emotions. On the other hand, our method can not only output the correct prediction, but also output higher prediction probability for the emotions that are close to the ground truth, while the probability of the opposite emotion is lower.

Fig. 5 shows some results of cross datasets recognition. When unseen classes belong to the Emotion6 dataset, our model distinguishes the “joy” and “surprise” classes successfully. For the last examples, when unseen classes belong to the FI dataset, our model predicts the “awe” class to the “excitement”. This is probably because such military images could evoke different emotions among different people. Interestingly, we train the model with images of four negative emotions and our model could recognize different positive emotions in the FI and Emotion6 dataset.

5. Conclusion

In this paper, we propose a novel affective structural embedding framework for the zero-shot emotion recognition problem. By utilizing ANP features to construct an affective embedding space, the affective gap between visual and semantic features can be effectively bridged. In addition, we introduce an affective adversarial constraint to force the visual embedding to choose an embedding space that preserves the affective structural information and retains the discriminative capacity simultaneously. Experiments on five widely-used affective datasets show that our method significantly outperforms the state-of-the-art approaches for zero-shot emotion recognition.

Acknowledgment

This work was supported by the NSFC (NO.61876094, 61701273, U1933114), Natural Science Foundation of Tianjin, China (NO.18JCYBJC15400, 18ZXZNGX00110), the Open Project Program of the National Laboratory of Pattern Recognition (NLPR), the Fundamental Research Funds for the Central Universities, and Berkeley DeepDrive.

References

- [1] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for image classification. *IEEE transactions on pattern analysis and machine intelligence*, 38(7):1425–1438, 2016.
- [2] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, 2015.
- [3] Yashas Annadani and Soma Biswas. Preserving semantic relations for zero-shot learning. In *CVPR*, 2018.
- [4] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [5] Lisa Feldman Barrett. Are emotions natural kinds? *Perspectives on psychological science*, 1(1):28–58, 2006.
- [6] Damian Borth, Tao Chen, Rongrong Ji, and Shih-Fu Chang. Sentibank: large-scale ontology and classifiers for detecting sentiment and emotions in visual content. In *ACM MM*, 2013.
- [7] Damian Borth, Rongrong Ji, Tao Chen, Thomas Breuel, and Shih-Fu Chang. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *ACM MM*, 2013.
- [8] Tao Chen, Damian Borth, Trevor Darrell, and Shih-Fu Chang. DeepSentibank: Visual sentiment concept classification with deep convolutional neural networks. *arXiv preprint arXiv:1410.8586*, 2014.
- [9] Ming Lin Chuang Gan, Yi Yang, Yueting Zhuang, and Alexander G Hauptmann. Exploring semantic interclass relationships (sir) for zero-shot action recognition. In *AAAI*, 2015.
- [10] Paul Ekman. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200, 1992.
- [11] Paul Ekman, Wallace V Friesen, Maureen O’sullivan, Anthony Chan, Irene Diacoyanni-Tarlatzis, Karl Heider, Rainer Krause, William Ayhan LeCompte, Tom Pitcairn, Pio E Ricci-Bitti, et al. Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of personality and social psychology*, 53(4):712, 1987.
- [12] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *CVPR*, 2009.
- [13] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NIPS*, 2013.
- [14] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE TPAMI*, 2019.
- [15] Himaanshu Gauba, Pradeep Kumar, Partha Pratim Roy, Priyanka Singh, Debi Prosad Dogra, and Balasubramanian Raman. Prediction of advertisement preference by fusing eeg response and sentiment analysis. *Neural Networks*, 92:77–88, 2017.
- [16] Omkar Gune, Biplob Banerjee, and Subhasis Chaudhuri. Structure aligning discriminative latent embedding for zero-shot learning. In *BMVC*, 2018.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [18] Jia Jia, Sen Wu, Xiaohui Wang, Peiyun Hu, Lianhong Cai, and Jie Tang. Can we understand van gogh’s mood? learning to infer affects from images in social networks. In *ACM MM*, 2012.
- [19] Huajie Jiang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Learning class prototypes via structure alignment for zero-shot recognition. *ECCV*, 2018.
- [20] Huajie Jiang, Ruiping Wang, Shiguang Shan, Yi Yang, and Xilin Chen. Learning discriminative latent attributes for zero-shot classification. In *ICCV*, 2017.
- [21] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *CVPR*, 2017.
- [22] Kristen A Lindquist, Tor D Wager, Hedy Kober, Eliza Bliss-Moreau, and Lisa Feldman Barrett. The brain basis of emotion: a meta-analytic review. *Behavioral and brain sciences*, 35(3):121–143, 2012.
- [23] Jana Machajdik and Allan Hanbury. Affective image classification using features inspired by psychology and art theory. In *ACM MM*, 2010.
- [24] Joseph A Mikels, Barbara L Fredrickson, Gregory R Larkin, Casey M Lindberg, Sam J Maglio, and Patricia A Reuter-Lorenz. Emotional category data on images from the international affective picture system. *Behavior research methods*, 37(4):626–630, 2005.
- [25] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [26] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [27] Ashish Mishra, Shiva Krishna Reddy, Anurag Mittal, and Hema A. Murthy. A generative model for zero shot learning using conditional variational autoencoders. In *CVPR*, 2018.
- [28] Kuan-Chuan Peng, Tsuhan Chen, Amir Sadovnik, and Andrew C Gallagher. A mixed bag of emotions: Model, predict, and transfer emotion distributions. In *CVPR*, 2015.
- [29] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014.
- [30] Shengsheng Qian, Tianzhu Zhang, and Changsheng Xu. Multi-modal multi-view topic-opinion mining for social event analysis. In *ACM MM*, 2016.
- [31] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, 2013.
- [32] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: relation network for few-shot learning. In *CVPR*, 2018.
- [33] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016.
- [34] Shuang Wu, Sravanthi Bondugula, Florian Luisier, Xiaodan Zhuang, and Pradeep Natarajan. Zero-shot event detection using multi-modal fusion of weakly supervised concepts. In *CVPR*, 2014.
- [35] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *CVPR*, 2016.

- [36] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *CVPR*, 2018.
- [37] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning - the good, the bad and the ugly. In *CVPR*, 2017.
- [38] Jufeng Yang, Dongyu She, Yukun Lai, and Ming-Hsuan Yang. Retrieving and classifying affective images via deep metric learning. In *AAAI*, 2018.
- [39] Jufeng Yang, Dongyu She, Yu-Kun Lai, Paul L Rosin, and Ming-Hsuan Yang. Weakly supervised coupled networks for visual sentiment analysis. In *CVPR*, 2018.
- [40] Jufeng Yang, Dongyu She, and Ming Sun. Joint image emotion classification and distribution learning via deep convolutional neural network. In *IJCAI*, 2017.
- [41] Quanzen You, Hailin Jin, and Jiebo Luo. Visual sentiment analysis by attending on local image regions. In *AAAI*, 2017.
- [42] Quanzen You, Jiebo Luo, Hailin Jin, and Jianchao Yang. Robust image sentiment analysis using progressively trained and domain transferred deep networks. In *AAAI*, 2015.
- [43] Quanzen You, Jiebo Luo, Hailin Jin, and Jianchao Yang. Building a large scale dataset for image emotion recognition: The fine print and the benchmark. In *AAAI*, 2016.
- [44] Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *CVPR*, 2017.
- [45] Ziming Zhang and Venkatesh Saligrama. Zero-shot learning via semantic similarity embedding. In *ICCV*, 2015.
- [46] Sicheng Zhao, Guiguang Ding, Yue Gao, and Jungong Han. Approximating discrete probability distribution of image emotions by multi-modal features fusion. In *IJCAI*, 2017.
- [47] Sicheng Zhao, Yue Gao, Guiguang Ding, and Tat-Seng Chua. Real-time multimedia social event detection in microblog. *IEEE Transactions on Cybernetics*, 48(11):3218–3231, 2018.
- [48] Sicheng Zhao, Yue Gao, Xiaolei Jiang, Hongxun Yao, Tat-Seng Chua, and Xiaoshuai Sun. Exploring principles-of-art features for image emotion recognition. In *ACM MM*, 2014.
- [49] Sicheng Zhao, Hongxun Yao, Yue Gao, Guiguang Ding, and Tat-Seng Chua. Predicting personalized image emotion perceptions in social networks. *IEEE transactions on Affective Computing*, 9(4):526–540, 2018.
- [50] Sicheng Zhao, Hongxun Yao, Yue Gao, Rongrong Ji, and Guiguang Ding. Continuous probability distribution prediction of image emotions via multitask shared sparse regression. *IEEE Transactions on Multimedia*, 19(3):632–645, 2017.
- [51] Pengkai Zhu, Hanxiao Wang, Tolga Bolukbasi, and Venkatesh Saligrama. Zero-shot detection. *arXiv preprint arXiv:1803.07113*, 2018.
- [52] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *CVPR*, 2018.