

# Bridging Visual Affective Gap: Borrowing Textual Knowledge by Learning from Noisy Image-Text Pairs

Anonymous Authors

## ABSTRACT

Visual emotion recognition (VER) is a longstanding field that has garnered increasing attention with the advancement of deep neural networks. Although recent studies have achieved notable improvements by leveraging the knowledge embedded within pre-trained visual models, the lack of direct association between factual-level features and emotional categories, called the “affective gap”, limits the applicability of pre-training knowledge for VER tasks. On the contrary, the explicit emotional expression and high information density in textual modality eliminate the “affective gap”. Therefore, we propose borrowing the knowledge from the pre-trained textual model to enhance the emotional perception of pre-trained visual models. We focus on the factual and emotional connections between images and texts in noisy social media data, and propose **Partitioned Adaptive Contrastive Learning (PACL)** to leverage these connections. Specifically, we manage to separate different types of samples and devise distinct contrastive learning strategies for each type. By dynamically constructing negative and positive pairs, we fully exploit the potential of noisy samples. Through comprehensive experiments, we demonstrate that bridging the “affective gap” significantly improves the performance of various pre-trained visual models in downstream emotion-related tasks.

## CCS CONCEPTS

• Information systems → Sentiment analysis; • Computing methodologies → Computer vision tasks.

## KEYWORDS

visual emotion recognition, affective gap, factual connection, emotional connection, contrastive learning

## ACM Reference Format:

Anonymous Authors. 2024. Bridging Visual Affective Gap: Borrowing Textual Knowledge by Learning from Noisy Image-Text Pairs. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM'24)*, October 28–November 1, 2024, Melbourne, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Visual emotion recognition (VER) aims at identifying human emotions towards different visual stimuli [41]. As a vital facet of human engagement with the world, perceiving emotions through visual

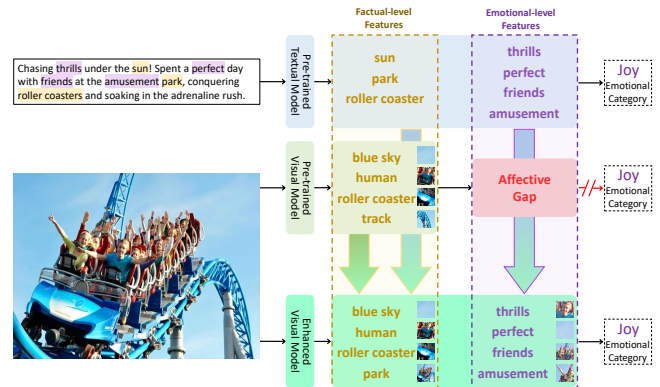


Figure 1: Illustration of the “affective gap” in the pre-trained visual model and how our proposed method bridges it using knowledge from the pre-trained textual model.

cues is progressively emerging as a pivotal challenge on the path toward the next generation of artificial general intelligence [22, 50]. Therefore, it has drawn increasing attention in recent years [9, 24] and exhibited broad applications in opinion mining [37], business intelligence [14], and autonomous driving [17].

Prior studies in VER have greatly enhanced their performance by leveraging the powerful generalizable knowledge of pre-trained models [61]. However, a misalignment exists between the objectives of their pre-training and downstream tasks, leading to the “affective gap” phenomenon, as depicted in Fig. 1. Specifically, the commonly adopted pre-training tasks, such as classification on ImageNet [4] and CLIP [28]-type vision-language pretraining, encourage models to acquire non-trivial factual-level features under semantic guidance. Yet, these features lack a direct association with the emotional categories, leading to the incomplete applicability of pre-training knowledge for downstream VER tasks. To bridge the “affective gap”, the pre-trained model needs the capability of encoding emotional-level features, which serves as an intermediate link between factual-level features and emotion categories.

An intuitive way to acquire such capability is learning from large-scale emotional annotated datasets, while it would encounter two primary challenges in practice. Firstly, the perception subjectivity of emotion [25, 63] results in high annotation costs, restricting the scale of high-quality datasets. Secondly, the sparse information density of visual data [3, 39] causes a dispersed distribution of emotional information, requiring models to extract features across multiple regions and scales, thereby increasing the demand for computational resources. However, language has natural advantages in these two aspects compared with vision, as it explicitly conveys emotions through specific words and possesses higher and more consistent information density [34]. These advantages allow pre-trained textual models to simultaneously learn generalizable factual-level and emotional-level features from large corpora in

Unpublished working draft. Not for distribution.

Permission to make digital or hard copies of all or part of this work for personal or professional use, is granted by ACM, provided that the copyright holder(s) is credited on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ACM MM, 2024, Melbourne, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58

59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112  
113  
114  
115  
116

unsupervised manners, eliminating the “affective gap” in textual modality.

Inspired by this discovery, we propose to bridge the visual “affective gap” by enhancing pre-trained visual models with the unified factual-level and emotional-level features from pre-trained textual models, as illustrated in Fig. 1. To transfer this knowledge from texts to images, we focus on image-text datasets where samples generally possess relatively strong factual and emotional connections between their images and texts. Specifically, we adopt TumEmo [45], an image-text dataset sourced from social media. Due to the inherent factual inconsistencies and emotional discrepancies of user-generated posts [46], the dataset inevitably incorporates certain noisy samples that are factual-mismatched, emotional-mismatched samples, or both. Given such characteristics, we propose **Partitioned Adaptive Contrastive Learning (PACL)** to leverage these connections within all samples. It comprises three stages. Firstly, we quantitatively assess the factual and emotional connections between the image and text of each sample, dividing the dataset into separate partitions. It allows us to formulate distinct learning strategies for different types of data. Secondly, we cluster on two unimodal subsets of the dataset, revealing the inter-sample relationships from both factual and emotional aspects. Thirdly, we design an adaptive contrastive strategy by utilizing the outcomes from the preceding stages. We dynamically filter out false negative pairs and reconstruct false positive pairs for different dataset partitions, thereby ensuring the consistent superiority of positive pairs over negative pairs, both factually and emotionally.

We follow the network architecture of CLIP [28], with its textual encoder initialized by SKEP [36] and its visual encoder initialized by the to-be-enhanced visual model. In this manner, the enhanced visual model can encode emotional-level features from an image while retaining the original perception of factual-level features. It effectively bridges the “affective gap” in visual models and significantly improves their performance in downstream VER tasks.

The overall contributions of this paper are three-fold:

- We harness the advantages of language over vision. As far as we know, this is the first attempt to bridge the “affective gap” in pre-trained visual models with pre-trained textual knowledge.
- We propose a method called PACL. It leverages the factual and emotional connections within noisy image-text pairs by dynamically constructing contrastive pairs for different types of samples.
- We conduct extensive experiments on **six** VER benchmarks and **four** kinds of downstream tasks with **four** pre-trained visual models. The results demonstrate the consistent improvements achieved by our method, underscoring the effectiveness and necessity of bridging the “affective gap”.

## 2 RELATED WORKS

### 2.1 Visual Emotion Recognition

Vision plays a crucial role in human emotional perception, and research on VER has been conducted for over two decades [64]. Early studies mainly focus on designing hand-crafted features based on psychology and art theories, including low-level elements [48], mid-level principals [54, 60], and high-level adjective-noun pairs [1]. Later on, the rapid advancements in deep learning reveal the superiority of learning-based features [29, 53], catalyzing the emergence

of numerous methods built upon it [43, 52]. As their performance gradually approaches the bottleneck, researchers shift focus towards perceiving emotions at multiple regions and scales [44, 51]. More recently, various correlative subfields witness increased attention, such as personalized prediction [47, 63], domain adaptation [59, 65] and zero-shot learning [49, 55].

The rapid development of VER in recent years is inseparable from the emergence of pre-trained models. The non-trivial knowledge embedded in these models has significantly aided them in perceiving emotions [58]. However, the misalignment between pre-training objectives and downstream VER tasks leads to the “affective gap”- the absence of intermediate emotional-level features between factual-level features and emotional categories. To tackle this problem, Feng *et al.* [10] imitate human’s visual sentiment perception mechanism and design six pre-training tasks to guide the model in learning factual-level and emotional-level features from scratch. Although achieving impressive performance gains for several backbones, they expand the original network size by threefold and demand extensive annotated data from diverse datasets, imposing significant computational costs on the training process. In contrast to them, we enhance the emotional perception of pre-trained visual models on top of the original factual-level knowledge, significantly reducing the demand for data.

### 2.2 Learning from Text Supervision

Language and vision have their respective pros and cons in different aspects. In recent years, many studies have explored harnessing the advantages of language to enhance visual models. Out of them, CLIP [28] demonstrates the efficacy of text supervision in guiding the acquisition of robust and transferable visual representations. Following works are proposed by considering such as fine-grained alignment [3, 30, 66], additional supervision [7, 66], external knowledge [11, 33] and non-contrastive objective [67].

Inspired by their success, EmotionCLIP [58] transfers this paradigm to VER, by guiding the visual model in learning emotion-level features with a subject-aware context encoding and a sentiment-guided contrastive learning. It incorporates the knowledge of a pre-trained textual model during training, similar to ours, but only employs it for computing relationships between samples. Additionally, their approach requires substantial video and text data strictly correlated in both factual and emotional aspects. However, they haven’t released their training data yet, and to our knowledge, no publicly available dataset meets this requirement. Consequently, subsequent works following them entail not only significant training costs, akin to [10], but also necessitate additional expenses for data collection. In comparison, our comprehensive utilization of pre-training textual knowledge and the noisy image-text pairs mitigates the demands on both the quantity and quality of training data and enables stronger performance on downstream tasks.

## 3 METHOD

The pipeline of our proposed PACL is demonstrated in Fig. 2. It comprises three stages: Dataset Partition (Section 3.1), Unimodal Sample Cluster (Section 3.2) and Adaptive Contrastive Learning (Section 3.3). We provide a detailed explanation of each stage below.

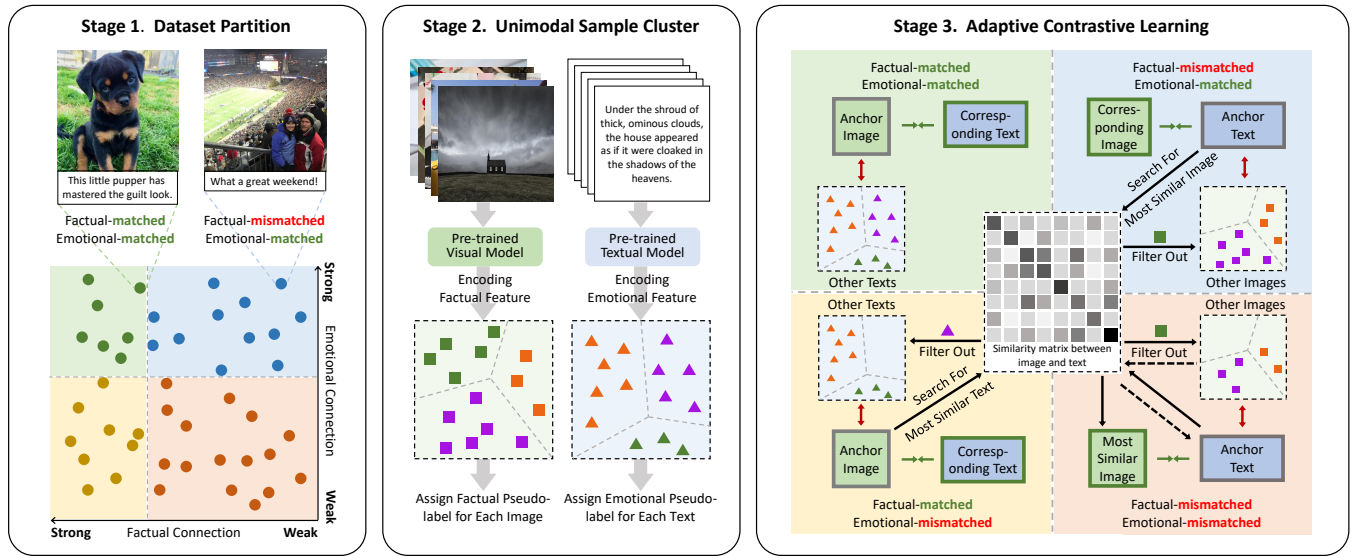


Figure 2: Overview of PACL’s three-stage pipeline. In the first stage, we divide the dataset into four partitions according to whether samples are factual-matched, emotional-matched, or not. In the second stage, we perform k-means clustering on image and text separately to obtain the factual and emotional pseudo-label for each of them. In the third stage, we guide the knowledge transfer from text to image with an adaptive contrastive learning. It dynamically constructs positive and negative pairs for different partitions by leveraging the pseudo-labels of samples.

### 3.1 Dataset Partition

We represent a total of  $N$  image-text pairs in the training dataset as  $\{(v_i, t_i) | i = 1, 2, \dots, N\}$ . Due to being collected from social media, images and texts from the same sample may suffer from inherent factual inconsistencies and emotional discrepancies. To prevent the model from learning misleading factual or emotional connections, we first divide the dataset into four partitions by quantitatively assessing the factual and emotional connections of each sample.

We employ CLIP [28] as the grounding evaluator of the factual connection, with its visual and textual encoder denoted by  $C_v(\cdot)$  and  $C_t(\cdot)$ . Given a sample  $(v_i, t_i)$ , we consider it is factual-matched if its CLIP-similarity score is larger than a threshold  $\sigma$ :

$$\frac{C_v(v_i) \odot C_t(t_i)}{\|C_v(v_i)\| \cdot \|C_t(t_i)\|} > \sigma. \quad (1)$$

Otherwise, we consider it factual-mismatched.

We employ DeepSentiBank [2], denoted by  $D(\cdot)$  and pre-trained BERTweet [26], denoted by  $B(\cdot)$ , as the grounding evaluator of the emotional connection. For sample  $(v_i, t_i)$ , we first transform its image  $v_i$  into the top-3 adjective-noun pairs  $D(v_i)$ . Then, we determine whether it is emotional-matched based on a calculation similar to the above:

$$\frac{B(D(v_i)) \odot B(t_i)}{\|B(D(v_i))\| \cdot \|B(t_i)\|} > \sigma. \quad (2)$$

Following this, we divide the dataset into four partitions: **1. strong-coupled samples** that are factual-matched and emotional-matched; **2. partial-coupled samples** that are emotional-matched but factual-mismatched; **3. partial-coupled samples** that are factual-matched but emotional-mismatched; **4. weak-coupled samples** that are factual-mismatched and emotional-mismatched. We



Figure 3: Samples from four partitions of TumEmo [45].

display a sample from each partition in Fig. 3 for intuitive understanding. It should be noted that we employ CLIP, DeepSentiBank, and BERTweet to simulate the partition conducted by humans, which can be replaced with other off-the-shelf tools.

### 3.2 Unimodal Sample Cluster

To optimize the construction of sample pairs in the subsequent contrastive learning, we utilize the knowledge embedded in the pre-trained visual and textual models to explore the unimodal inter-sample relationships from both factual and emotional perspectives. We adopt the network architecture of CLIP, as shown in Fig. 4. In this stage, none of the gradients are involved in the calculation.

For the visual encoder, we initialize it with parameters from a pre-trained visual model, denoted by  $V_\theta(\cdot)$ , which serves as the enhanced target of PACL. Leveraging its factual-level knowledge, we encode factual features for all images,  $\{v_i | i = 1, 2, \dots, N\}$ , and conduct k-means clustering inside its feature space. It yields  $K$

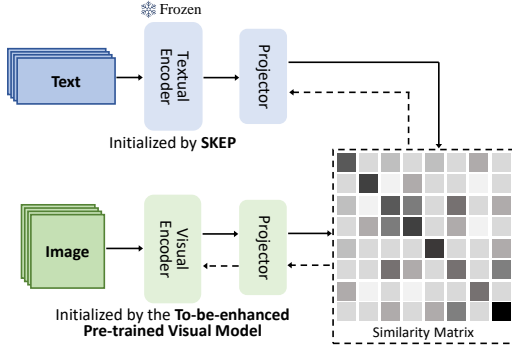


Figure 4: Network architecture of PACL.

clusters, where  $K$  is a hyperparameter describing the cluster granularity, and each image is assigned a cluster label. We denote the label of image  $v_i$  as its factual pseudo-label  $f_i$ . Thus, we obtain the factual relationship between images  $v_i, v_j$ : they are factually similar if  $f_i = f_j$ , and dissimilar if  $f_i \neq f_j$ .

For the textual encoder, we initialize it with parameters from SKEP[36], denoted by  $T_{\xi}(\cdot)$ . SKEP is an unsupervised sentiment knowledge-enhanced pre-trained model focusing more on emotional features. Leveraging this characteristic, we encode emotional features for all texts,  $\{t_i | i = 1, 2, \dots, N\}$ , and perform k-means clustering that also yields  $K$  clusters. Similarly, we denote the cluster label of text  $t_i$  as its emotional pseudo-label  $e_i$ .

### 3.3 Adaptive Contrastive Learning

With the outcomes of the two preparation stages, we devise an adaptive contrastive learning that constructs positive and negative pairs under different strategies for each dataset partition. In this stage, we treat SKEP as the knowledge source and keep the parameters of the textual encoder frozen. We update the parameters of the visual encoder through contrastive learning to guide it in learning the generalizable factual-level and emotional-level knowledge from textual modality.

**3.3.1 Preliminary.** We employ contrastive learning between images and texts multiple times. Denoting the features of  $(v_i, t_i)$  as  $z_i^v$  and  $z_i^t$ , here we provide a general form of contrastive loss:

$$\mathcal{L} = -\frac{1}{2} \log \frac{\sum_{j \in p^+(v_i)} \exp((z_i^v \odot z_j^t)/\tau)}{\sum_{j \in \{p^+(v_i) \cup p^-(v_i)\}} \exp((z_i^v \odot z_j^t)/\tau)} - \frac{1}{2} \log \frac{\sum_{j \in p^+(t_i)} \exp((z_j^v \odot z_i^t)/\tau)}{\sum_{j \in \{p^+(t_i) \cup p^-(t_i)\}} \exp((z_j^v \odot z_i^t)/\tau)}. \quad (3)$$

$p^+(\cdot), p^-(\cdot)$  represent the positive and negative samples of the anchor.  $\tau$  is the temperature hyperparameter. This loss comprises two components, with either image  $v_i$  or text  $t_i$  serving as the anchor. Taking the example of image  $v_i$ , the objective of Eq. (3) is to bring it closer to  $p^+(v_i)$  while pushing it away from  $p^-(v_i)$ .

**3.3.2 Strong-coupled Samples.** We first conduct contrastive learning solely on this partition to facilitate the alignment between the feature spaces of visual encoder  $V_{\theta}(\cdot)$  and textual encoder  $T_{\xi}(\cdot)$ . Since samples are both factual-matched and emotional-matched, we treat the images and texts from the same samples as positive

pairs, and those from different samples as negative pairs. It enables the model to assign higher cosine similarity scores to image-text pairs with strong factual and emotional connections. Denoting the projector of the visual branch as  $q_v(\cdot)$  and it of the textual branch as  $q_t(\cdot)$ , the features of image  $v_i$  and text  $t_i$  are computed as:

$$z_i^v = q_v(V_{\theta}(v_i)), \quad z_i^t = q_t(T_{\xi}(t_i)). \quad (4)$$

By substituting Eq. (4) into Eq. (3), we obtain the contrastive loss of this dataset partition:

$$\mathcal{L}_1 = \mathcal{L}(p^+(v_i) = p^+(t_i) = \{i\}, p^-(v_i) = p^-(t_i) = \{j | j \neq i\}). \quad (5)$$

**3.3.3 Partial-coupled Samples.** These samples are from two dataset partitions. We discuss them together due to their common properties, that their images and texts have strong connections only in one of the factual and emotional aspects. To leverage these connections, we construct positive pairs in the same way as strong-coupled samples. However, since partial-coupled samples no longer guarantee the advantages of positive pairs over negative pairs in both factual and emotional aspects, the former construction of negative pairs is sub-optimal.

For the partition where samples are emotional-matched but factual-mismatched, we take  $(v_i, t_i)$  as an example. Compared to  $v_i$ , the image  $v_j$  from another sample may demonstrate a stronger factual connection with  $t_i$ , and in this case, treating  $(v_i, t_i)$  as a positive pair while  $(v_j, t_i)$  as a negative pair compromises the factual-level learning of the visual model. Therefore, we devise a mapping  $m_f(\cdot)$  that searches for images that might possess strong factual connections with text  $t_i$  to avoid inappropriate negative pairs:

$$m_f(t_i) = \{v_j | l = \underset{k}{\operatorname{argmax}} z_k^v \odot z_i^t, f_j = f_l\}. \quad (6)$$

It first finds the most matching image, denoted by  $v_l$ , based on the cosine similarity scores and then selects images with factual pseudo-label same as  $f_l$ . We also define an inversion of  $m_f(\cdot)$  that searches for texts based on image  $v_i$ :

$$m_f^{-1}(v_i) = \{t_j | v_i \in m_f(t_j)\}. \quad (7)$$

Subsequently, the negative samples of  $v_i$  are all other texts with the exclusion of  $m_f^{-1}(v_i)$ , and the negative samples of  $t_i$  are all other images with the exclusion of  $m_f(t_i)$ . Thereby, the contrastive loss of this partition is:

$$\mathcal{L}_2 = \mathcal{L}(p^+(v_i) = p^+(t_i) = \{i\}, p^-(v_i) = \{j | j \neq i, t_j \notin m_f^{-1}(v_i)\}, p^-(t_i) = \{j | j \neq i, v_j \notin m_f(t_i)\}). \quad (8)$$

For the partition where samples are factual-matched but emotional-mismatched, we devise a symmetric  $m_e(\cdot)$  that searches for texts that possess strong emotional connections with image  $v_i$ :

$$m_e(v_i) = \{t_j | l = \underset{k}{\operatorname{argmax}} z_i^v \odot z_k^t, e_j = e_l\}, \quad (9)$$

and also an inversion of  $m_e(\cdot)$  that searches for images based on text  $t_i$ :

$$m_e^{-1}(t_i) = \{v_j | t_i \in m_e(v_j)\} \quad (10)$$

Similarly, the negative samples of  $v_i$  ( $t_i$ ) are all other texts (images) with the exclusion of  $m_e(v_i)$  ( $m_e^{-1}(t_i)$ ). The contrastive loss of this partition is:

$$\begin{aligned} \mathcal{L}_3 &= \mathcal{L}(p^+(v_i) = p^-(t_i) = \{i\}, \\ p^-(v_i) &= \{j | j \neq i, t_j \notin m_e(v_i)\}, \\ p^-(t_i) &= \{j | j \neq i, v_j \notin m_e^{-1}(t_i)\}. \end{aligned} \quad (11)$$

**3.3.4 Weak-coupled Samples.** Since samples of this partition are factual-mismatched and emotional-mismatched, constructing positive pairs as before would disrupt the semantic relationships learned by the model [19] and lead to degenerations in downstream performance. Therefore, we reconstruct positive pairs by leveraging the learned correlation between images and texts in previous partitions. Taking sample  $(v_i, t_i)$  as an example, we select positive samples according to the cosine similarity scores:

$$p^+(v_i) = \{\operatorname{argmax}_k z_i^v \odot z_k^t\}, \quad (12)$$

$$p^+(t_i) = \{\operatorname{argmax}_k z_k^v \odot z_i^t\}. \quad (13)$$

For negative pairs, we follow the strategies adopted for partial-coupled samples. We simultaneously apply  $m_f(\cdot)$  and  $m_e(\cdot)$  to filter out potential false negative pairs with either strong factual or emotional connections. As a result, the contrastive loss is:

$$\begin{aligned} \mathcal{L}_4 &= \mathcal{L}(p^+(v_i) = \{\operatorname{argmax}_k z_i^v \odot z_k^t\}, \\ p^+(t_i) &= \{\operatorname{argmax}_k z_k^v \odot z_i^t\}, \\ p^-(v_i) &= \{j | j \neq i, t_j \notin m_f^{-1}(v_i) \cup m_e(v_i)\}, \\ p^-(t_i) &= \{j | j \neq i, v_j \notin m_f(t_i) \cup m_e^{-1}(t_i)\}. \end{aligned} \quad (14)$$

**3.3.5 Total Loss.** We devise a progressive learning schedule to guide the model to learn from strong-coupled samples to weak-coupled samples step by step. In the first one-third of epochs, we only incorporate the strong-coupled samples into training:

$$\mathcal{L}_{total} = \mathcal{L}_1. \quad (15)$$

In the second one-third of epochs, we incorporate both the strong-coupled and partial-coupled samples into training:

$$\mathcal{L}_{total} = \mathcal{L}_1 + \mathcal{L}_2 + \mathcal{L}_3. \quad (16)$$

In the last one-third of epochs, we incorporate all samples into training:

$$\mathcal{L}_{total} = \mathcal{L}_1 + \mathcal{L}_2 + \mathcal{L}_3 + \mathcal{L}_4. \quad (17)$$

## 4 EXPERIMENT

### 4.1 Datasets

We conduct training of PACL on **TumEmo** [45]. It contains 200k image-text pairs crawled from social media. Its samples are annotated based on their emotions. Since we emphasize leveraging the factual and emotional connections between images and texts, we avoid using knowledge from the annotations.

In the downstream testing, we adopt six publicly available VER datasets that are annotated according to diverse emotion theories. **FI** [53] is labeled according to Mikels' psychological study [23]. **Emotion-6** and **UnbiasedEmo** [24] contain 6 common emotion

categories. **WebEmo** [24] is labeled according to Parrot's hierarchical emotion model [27], which contains 2 basic categories at the first level, 7 categories at the second level and 25 fine-grained categories at the third level. **Emotic** [16] comprises 26 fine-grained emotion categories, with each image assigned multiple emotion labels, and 10-scale VAD (Valence-Arousal-Dominance) [32] ratings. **IAPS** [18] is labeled by 9-scale VAD ratings.

### 4.2 Implementation Details

All of our implementations are based on Pytorch. We employ three visual models pre-trained on ImageNet [4], including ResNet-50 [13], ViT-base [8], Swin-base [21], and one visual model pre-trained by CLIP [28] (ViT-clip) as the target models for enhancement. Most of our experiments adopt SKEP [36] as the pre-trained textual model. In Section 5, we also evaluate the enhancement effect of other textual models [6, 20, 31]. In the first stage of PACL, we set the partition threshold  $\sigma$  to 0.7. In the second stage, we set cluster granularity  $K$  to 2. In the third stage, we set the contrastive temperature  $\tau$  to 0.07, the initial learning rate of ResNet-50 to 1e-3, ViT-base and ViT-clip to 1e-5, Swin-base to 2e-5, projection layers  $p_v(\cdot)$ ,  $p_t(\cdot)$  to 1e-3. We train the model for 30 epochs with a minibatch size of 64. We adopt AdamW optimizer with a cosine schedule to decay the learning rate. In the downstream testing, we adopt the original partition for training, validation, and testing if provided. Otherwise, we randomly divide the dataset in a ratio of 8:1:1. All the experiments are conducted on four NVIDIA 4090 GPUs.

### 4.3 Compared Baselines

Two other works also focus on enhancing pre-trained visual models in perceiving emotions. **Feng et al.** [10] factorize human perception into three steps and devise a pre-training method to imitate it. Their pre-training tasks include colorization, super resolution, scene recognition, jigsaw puzzles, ANP prediction, and image captioning, requiring 10.5M images with annotations. We reproduce it under all of our settings for comparison. **EmotionCLIP** [58] adopts vision-language pre-training to extract emotional representations from verbal and nonverbal communication. It leverages the pre-trained knowledge from CLIP [28] and DistilBert [31]. It is trained on 1M video-text pairs with the annotation of sentiment score and human bounding boxes. Since the training data is not publicly available yet, we compare PACL with it under pre-trained ViT-clip, which is the only model released. **PACL** is our proposed model, it leverages the pre-trained knowledge from SKEP [36] and is trained on 200K image-text pairs without any additional annotations. PACL has **the loosest requirements for training data** among them.

### 4.4 Enhancement Comparison

We adopt four kinds of VER downstream tasks. **1. In single-label learning**, models are required to classify each image into a single emotional category. **2. In hierarchical-label learning**, models are required to predict categories for three emotion levels simultaneously. We report accuracy (Acc) and weighted-F1 (F1) for the first two tasks. **3. In multi-label learning**, models are required to predict multiple emotion categories for each image. We report the mean of average precision (mAP) and area under roc curve (AUC) of categories. **4. In VAD learning**, models are required to predict

**Table 1: Emotional enhancement of visual models on single-label learning datasets including FI [53], Emotion-6 [24], UnbiasedEmo [24], hierarchical-label learning dataset WebEmo [24], multi-label learning dataset Emotic [16], and VAD learning datasets including Emotic [16] and IAPS [18]. Accuracy (Acc $\uparrow$ ), weighted-F1 (F1 $\uparrow$ ), mean of average precision (mAP $\uparrow$ ), area under roc curve (AUC $\uparrow$ ), mean of mean square error (MSE $\downarrow$ ), mean of R square ( $R^2$   $\uparrow$ ) are reported for these datasets. Acc and F1 are followed by the number of total categories, and MSE is multiplied by  $10^2$ . The best results are marked in bold.**

Model	Single-label						Hierarchical-label						Multi-label		VAD			
	FI		Emotion-6		UnbiasedEmo		WebEmo						Emotic		Emotic		IAPS	
	Acc-8	F1-8	Acc-6	F1-6	Acc-6	F1-6	Acc-2	F1-2	Acc-6	F1-6	Acc-25	F1-25	mAP	AUC	MSE	$R^2$	MSE	$R^2$
<i>Linear Evaluation:</i>																		
ResNet-50	57.72	56.72	46.47	45.18	60.20	60.08	66.25	66.21	41.08	39.03	24.94	22.80	25.85	66.36	2.825	0.1267	2.130	0.2790
ResNet-50 + Feng <i>et al.</i>	58.39	57.88	47.54	46.48	62.17	61.27	65.82	65.71	40.63	38.34	24.57	22.14	<b>27.32</b>	67.49	2.867	0.1279	2.253	0.2685
ResNet-50 + PACL	<b>59.50</b>	<b>59.03</b>	<b>51.14</b>	<b>50.59</b>	<b>66.78</b>	<b>66.39</b>	<b>67.53</b>	<b>67.42</b>	<b>43.60</b>	<b>40.95</b>	<b>25.85</b>	<b>24.32</b>	27.02	<b>67.60</b>	<b>2.646</b>	<b>0.1404</b>	<b>1.982</b>	<b>0.3104</b>
ViT-base	62.08	62.06	54.85	54.24	74.67	74.24	70.72	70.67	45.90	43.67	27.69	26.10	30.16	68.26	2.763	0.1265	2.010	0.2936
ViT-base + Feng <i>et al.</i>	64.12	64.06	57.02	56.73	77.71	77.54	70.55	70.51	45.34	45.09	27.37	27.14	33.01	69.30	2.733	0.1293	1.991	0.3088
ViT-base + PACL	<b>64.77</b>	<b>64.64</b>	<b>58.25</b>	<b>58.03</b>	<b>78.95</b>	<b>78.91</b>	<b>71.98</b>	<b>71.88</b>	<b>48.32</b>	<b>46.01</b>	<b>28.64</b>	<b>27.89</b>	<b>33.40</b>	<b>70.68</b>	<b>2.435</b>	<b>0.1466</b>	<b>1.903</b>	<b>0.3319</b>
Swin-base	63.63	63.50	53.29	51.27	74.30	73.98	70.39	70.23	46.39	44.25	28.82	27.41	31.27	68.09	2.715	0.1308	1.966	0.2872
Swin-base + Feng <i>et al.</i>	65.03	64.74	55.47	52.84	76.30	75.67	70.12	70.03	46.37	44.36	28.56	27.17	32.00	69.08	2.665	0.1335	1.909	0.2948
Swin-base + PACL	<b>65.70</b>	<b>65.58</b>	<b>57.13</b>	<b>56.79</b>	<b>79.93</b>	<b>79.67</b>	<b>72.14</b>	<b>72.13</b>	<b>47.96</b>	<b>45.81</b>	<b>29.89</b>	<b>27.95</b>	<b>32.91</b>	<b>70.19</b>	<b>2.451</b>	<b>0.1449</b>	<b>1.794</b>	<b>0.3285</b>
ViT-clip	71.74	71.61	64.43	64.15	85.53	85.24	78.41	77.85	54.86	53.68	37.36	36.12	38.50	72.67	2.313	0.1681	1.560	0.3473
ViT-clip + Feng <i>et al.</i>	64.12	64.06	57.02	56.73	77.71	77.54	70.55	70.51	45.34	45.09	27.37	27.14	33.01	69.30	2.733	0.1293	1.991	0.3088
ViT-clip + EmotionCLIP	72.63	72.32	64.72	64.57	84.36	84.06	78.77	78.22	55.28	54.35	37.12	35.82	34.63	72.33	2.303	0.1711	1.538	0.3478
ViT-clip + PACL	<b>73.17</b>	<b>72.53</b>	<b>66.08</b>	<b>66.01</b>	<b>86.72</b>	<b>85.97</b>	<b>79.51</b>	<b>79.33</b>	<b>56.13</b>	<b>55.59</b>	<b>37.40</b>	<b>37.03</b>	<b>39.05</b>	<b>72.97</b>	<b>2.249</b>	<b>0.1744</b>	<b>1.506</b>	<b>0.3535</b>
<i>Fine-tuning:</i>																		
ResNet-50	65.14	64.84	53.46	53.25	73.68	73.15	73.81	73.80	48.88	48.34	30.85	29.89	28.36	68.43	2.757	0.1337	2.114	0.2842
ResNet-50 + Feng <i>et al.</i>	66.69	66.55	58.21	58.05	<b>77.47</b>	<b>76.90</b>	74.43	74.38	49.66	49.39	32.63	31.01	28.44	68.20	2.646	0.1408	2.082	0.2954
ResNet-50 + PACL	<b>67.11</b>	<b>66.79</b>	<b>58.60</b>	<b>58.16</b>	<b>77.30</b>	<b>76.65</b>	<b>74.76</b>	<b>74.75</b>	<b>50.18</b>	<b>49.21</b>	<b>32.72</b>	<b>31.42</b>	<b>30.39</b>	<b>70.12</b>	<b>2.531</b>	<b>0.1460</b>	<b>1.927</b>	<b>0.3130</b>
ViT-base	69.57	69.56	57.45	57.09	77.63	77.51	76.25	76.08	52.49	51.01	33.88	32.52	32.52	71.73	2.633	0.1374	1.975	0.3016
ViT-base + Feng <i>et al.</i>	69.73	69.60	60.15	59.38	81.03	81.00	76.34	76.11	53.39	52.41	33.98	33.62	33.12	71.87	2.603	0.1451	1.964	0.3129
ViT-base + PACL	<b>70.91</b>	<b>70.68</b>	<b>60.92</b>	<b>60.65</b>	<b>82.57</b>	<b>82.26</b>	<b>77.61</b>	<b>77.74</b>	<b>54.87</b>	<b>54.50</b>	<b>34.75</b>	<b>34.61</b>	<b>35.09</b>	<b>72.24</b>	<b>2.378</b>	<b>0.1511</b>	<b>1.860</b>	<b>0.3325</b>
Swin-base	70.24	70.33	55.57	55.27	78.29	78.02	75.83	75.75	51.57	51.05	34.19	33.10	33.11	71.61	2.607	0.1369	1.922	0.2961
Swin-base + Feng <i>et al.</i>	72.04	71.68	58.62	58.41	82.29	82.17	76.71	76.53	53.90	53.62	34.50	34.08	32.97	71.53	2.580	0.1461	1.861	0.3115
Swin-base + PACL	<b>73.21</b>	<b>73.17</b>	<b>60.04</b>	<b>59.72</b>	<b>83.55</b>	<b>83.34</b>	<b>77.19</b>	<b>77.06</b>	<b>54.69</b>	<b>54.32</b>	<b>35.33</b>	<b>35.27</b>	<b>36.11</b>	<b>72.83</b>	<b>2.407</b>	<b>0.1530</b>	<b>1.779</b>	<b>0.3344</b>
ViT-clip	77.33	77.25	64.89	64.60	88.16	88.04	79.32	79.25	55.47	54.99	39.71	38.82	38.78	72.90	2.246	0.1729	1.523	0.3562
ViT-clip + Feng <i>et al.</i>	69.73	69.60	60.15	59.38	81.03	81.00	76.34	76.11	53.39	52.41	33.98	33.62	33.12	71.87	2.603	0.1451	1.964	0.3129
ViT-clip + EmotionCLIP	77.62	77.40	65.09	64.77	89.44	88.81	80.20	80.16	56.14	55.83	39.67	38.70	36.54	72.52	2.264	0.1719	1.476	0.3575
ViT-clip + PACL	<b>80.55</b>	<b>80.28</b>	<b>66.91</b>	<b>66.65</b>	<b>90.14</b>	<b>90.13</b>	<b>81.84</b>	<b>81.57</b>	<b>56.35</b>	<b>56.02</b>	<b>39.99</b>	<b>39.21</b>	<b>39.36</b>	<b>73.03</b>	<b>2.213</b>	<b>0.1784</b>	<b>1.425</b>	<b>0.3664</b>

the continuous values in the ‘Valence’, ‘Arousal’, and ‘Dominance’ dimensions. We report the mean of mean squared error (MSE) and R squared ( $R^2$ ) of the three dimensions. Besides, we conduct experiments under two evaluation protocols: linear evaluation and fine-tuning. In both protocols, we set the learning rate of the linear classification head to  $1e-2$  and train for 40 epochs. For parameters of the backbone network, we keep them frozen under linear evaluation and set the initial learning rate of ResNet-50 to  $1e-4$ , ViT-base, Swin-base, and ViT-clip to  $1e-5$  under fine-tuning. The other settings of downstream testing are aligned with Section 4.2.

The experimental results are presented in Table 1. We can observe that PACL consistently improves all the pre-trained visual models on all tasks by a large margin. In contrast, the effectiveness of Feng *et al.* is limited under linear evaluation, especially in hierarchical-label learning. It highlights the advantages of harnessing textual knowledge in perceiving emotions at multiple levels. Additionally, since they start the pre-training from scratch, they can not leverage the generalizable factual knowledge in the pre-trained visual models. For instance, PACL achieves significant performance gains by starting from ViT-clip instead of ViT-base, underscoring the effectiveness and scalability of bridging the “affective gap”. EmotionCLIP, on the other hand, also undergoes poor performance on some tasks, specifically, multi-label learning. We attribute this to

the misalignment between its pre-training and testing. Although it explicitly models the dependencies between context and humans, it may be misled by the massive communications between humans in training data and learns emotional representations heavily dependent on humans. However, the multi-label learning dataset Emotic emphasizes the emotion perception from both context and humans, resulting in the undesirable performance of EmotionCLIP. On the contrary, PACL leverages diverse image-text pairs and the unified factual and emotional features embedded in texts, naturally encouraging visual models to perceive emotions in different regions and scales, thereby promoting universal enhancements on all tasks.

## 4.5 Comparison with SOTA

To further demonstrate the necessity and effectiveness of bridging the “affective gap”, we present a brief view of the current SOTA methods in two downstream tasks.

**4.5.1 Classification on FI, WebEmo.** Following previous methods [40], we group eight categories of FI into two coarse-grained parent categories. We report the classification accuracy of SOTA methods on FI and WebEmo at multiple emotion levels in Table 2. It should be noted that the models are tested on each level separately, different from the settings in hierarchical-label learning in 4.4. As shown in Table 2, ViT-clip enhanced by PACL outperforms other

**Table 2: Accuracy of SOTA methods in classification at different emotion levels on FI [53] and WebEmo [24]. The best results are marked in bold, and the second-best results are underlined.**

Model	FI		WebEmo		
	Acc-2	Acc-8	Acc-2	Acc-6	Acc-25
Sentibank [1]	56.47	44.49	-	-	-
DeepSentibank [2]	64.39	53.16	-	-	-
PCNN [52]	75.34	56.16	-	-	-
Zhu <i>et al.</i> [68]	84.26	73.03	-	-	-
Rao <i>et al.</i> [43]	87.51	75.46	-	-	-
WSCNet [42]	86.74	70.07	79.43	52.61	32.75
PDANet [62]	87.25	72.13	80.96	53.46	32.82
Zhang and Xu [56]	90.97	75.91	82.47	53.88	33.01
MDAN [40]	91.08	76.41	82.72	55.65	34.28
SimEmotion [5]	<u>95.42</u>	80.33	-	-	-
ViT-clip + PACL	95.27	<b>80.55</b>	<u>84.06</u>	<u>57.42</u>	<u>40.31</u>
ViT-clip + PACL + PDANet	<b>96.06</b>	<b>80.83</b>	<b>84.74</b>	<b>58.11</b>	<b>40.50</b>

methods in most cases. Additionally, since our enhanced model does not incorporate additional parameters besides the backbone network, it can be combined with other methods for further improvements. By adding a polarity-consistent regression loss [62], our enhanced model achieves the best performances across all scenarios. These results validate the efficacy and scalability of PACL.

**Table 3: Accuracy of SOTA methods in zero-shot classification on FI [53] and Emotion-6 [24]. [C] denotes the names of emotion categories.**

Model	Text Prompt	FI	Emotion-6
		Acc-8	Acc-6
Zhang <i>et al.</i> [57]	—	21.73	22.74
Kodirov <i>et al.</i> [15]	—	19.35	26.53
Sung <i>et al.</i> [35]	—	18.49	28.10
Zhan <i>et al.</i> [55]	—	24.51	30.49
Ye <i>et al.</i> [49]	—	<b>36.03</b>	<b>33.57</b>
CLIP [28]	A photo with [C] emotion.	20.71	10.16
ResNet-50 + PACL	A photo with [C] emotion.	25.58	29.46
CLIP [28]	The photo makes me feel so [C]!	17.45	13.78
ResNet-50 + PACL	The photo makes me feel so [C]!	<u>31.72</u>	<u>30.23</u>

**4.5.2 Zero-shot on FI, Emotion-6.** Since we adopt the network architecture of PACL the same as CLIP [28], our enhanced models are naturally capable of conducting zero-shot emotion classification by leveraging the learned emotional connections between images and texts. We report the zero-shot classification accuracy of the enhanced ResNet-50 and the SOTA methods in Table 3. Similar to CLIP, our enhanced model is also sensitive to text prompts. Therefore, we test its performances under both regular and manually designed prompts. From the results, we observe that the enhanced ResNet-50 significantly outperforms CLIP by learning from much fewer image-text pairs. We attribute this to the difference between the types of training data. PACL learns from social media posts, which contain generally strong emotional connections between images and texts, as well as more frequent occurrences of textual emotional expressions. Compared with other methods, the enhanced

**Table 4: Ablation experiments on dataset partition and sample usage. For factual (emotional) partition, we conduct contrastive learning by treating factual-matched (emotional-matched) samples as strong-coupled and factual-mismatched (emotional-mismatched) samples as partial-coupled. “Both” denotes factual + emotional.**

Partition	Sample Usage	UnbiasedEmo	
		Acc-6	F1-6
Initial weight		60.20	60.08
None	All	60.53	60.01
Factual	Factual-matched	63.49	62.91
Factual	Factual-matched + Factual-mismatched	65.13	64.71
Emotional	Emotional-matched	64.47	63.77
Emotional	Emotional-matched + Emotional-mismatched	65.46	64.83
Both	Strong-coupled	64.14	64.10
Both	Strong-coupled + Partial-coupled	65.79	65.38
Both	Strong-coupled + Partial-coupled + Weak-coupled	<b>66.78</b>	<b>66.39</b>

ResNet-50 achieves the second-best performance on both datasets. Although zero-shot emotion classification is not our primary goal, the satisfactory performance achieved by the enhanced ResNet-50 provides us with the intuitive advantages of learning from social media posts.

## 5 ANALYSIS

In this section, we validate the effectiveness of each component of PACL. All the experiments are conducted by adopting pre-trained ResNet-50 as the backbone under linear evaluation.

### 5.1 Ablation Study

We conduct ablation experiments in Table 4 to explore the influences of different dataset partitions and sample usages. Firstly, we do not apply partition strategies and conduct contrastive learning by treating all the samples as strong-coupled. It results in minor enhancements for the pre-trained model since TumEmo contains certain noisy samples with either weak factual or emotional connections. In these cases, the vanilla contrastive learning is sub-optimal. Next, we apply partition based only on factual connections, treating factual-matched samples as strong-coupled and factual-mismatched samples as partial-coupled. PACL enhances the pre-trained model by a large margin with solely factual-matched samples and further improves its performance by utilizing factual-mismatched samples. We also apply partition based only on emotional connections with the same settings, which brings a similar trend. By dynamically constructing contrastive pairs, we succeed in harnessing these noisy samples. Notably, although the downstream task is emotion-related, applying partition based on either emotional or factual connections can both lead to significant enhancements, demonstrating the cruciality of proper handling of both factual-mismatched and emotional-mismatched samples. Finally, by simultaneously applying factual and emotional partition as PACL, the pre-trained model progressively learns from factual and emotional connections of different kinds of data, achieving the best result. The above results indicate the effectiveness of PACL in fully exploiting the dataset.

**Table 5: Experiments of hyperparameter selection. Each experiment for a specific hyperparameter is carried out independently, with other non-tested hyperparameters set to their adopted values.  $\sigma, K, \tau$  are from Eq. (1), Section 3.2, Eq. (3), respectively.**


Partition Threshold $\sigma$	FI		Emotion-6		WebEmo	
	Acc-8	F1-8	Acc-6	F1-6	Acc-2	F1-2
0	57.08	56.74	46.59	45.11	65.51	65.21
0.3	58.50	58.03	49.10	47.74	66.36	66.33
0.5	58.58	57.96	49.46	48.00	67.02	66.94
0.7 (adopted)	<b>59.59</b>	<b>59.03</b>	<b>51.14</b>	<b>50.59</b>	<b>67.53</b>	<b>67.42</b>
0.9	56.64	56.06	45.75	44.31	65.08	64.94
1	39.56	37.77	36.17	33.72	58.54	58.11
Cluster Granularity $K$	FI		Emotion-6		WebEmo	
	Acc-8	F1-8	Acc-6	F1-6	Acc-2	F1-2
2 (adopted)	<b>59.59</b>	<b>59.03</b>	51.14	50.59	<b>67.53</b>	<b>67.42</b>
3	59.23	58.51	51.02	50.34	67.30	67.12
6	58.46	57.66	<b>51.98</b>	<b>51.78</b>	67.44	67.35
10	57.98	57.82	49.34	49.08	66.82	66.50
15	57.33	56.67	47.90	47.13	66.60	66.38
25	56.08	55.82	45.63	43.56	65.21	64.79
Contrastive Temperature $\tau$	FI		Emotion-6		WebEmo	
	Acc-8	F1-8	Acc-6	F1-6	Acc-2	F1-2
0.01	59.43	59.12	49.34	48.89	66.53	66.44
0.03	<b>59.74</b>	<b>59.38</b>	50.19	49.83	67.22	67.10
0.05	59.45	58.77	50.96	<b>50.62</b>	67.03	66.79
0.07 (adopted)	59.59	59.03	<b>51.14</b>	50.59	<b>67.53</b>	<b>67.42</b>
0.5	58.81	58.53	50.27	49.80	66.91	66.73
1	58.45	58.13	49.61	49.29	66.80	66.55

## 5.2 Hyperparameter Selections

To probe the influence of hyperparameters, we conduct the experiments in Table 5.

**5.2.1 Partition Threshold  $\sigma$ .** It controls the precision and size of each partition. By adopting a small threshold, the strong-coupled samples can not guarantee the generally strong factual and emotional connections between images and texts. It disrupts the alignment between the visual and textual feature spaces [19]. On the other hand, a large threshold limits the number of strong-coupled samples, which is also not conducive to the alignment. Consequently, we adopt a medium threshold that provides sufficient high-quality strong-coupled samples. Specifically, 22% samples are strong-coupled, 51% samples are partial-coupled, and 27% samples are weak-coupled under  $\sigma = 0.7$ .

**5.2.2 Cluster Granularity  $K$ .** It balances the distances between intra-cluster and inter-cluster samples, thereby influencing the reconstruction of negative pairs for partial-coupled and weak-coupled samples. Fine-grained clustering guarantees the correctness of filtered negative pairs, while coarse-grained clustering emphasizes the correctness of the remaining. According to the experimental results, we adopt a small  $K$ . This indicates that the advantages of filtering out additional true negative pairs outweigh the potential drawbacks of ignoring false negative pairs. Interestingly, the enhanced model achieves the best results at  $K = 6$  on Emotion-6. We attribute this to the alignment in the granularities of clustering and downstream classification. However, we still adopt  $K = 2$  due to its general advantages across all datasets.



**Figure 5: Classification accuracy of the enhanced visual model on Emotion-6 [24] and UnbiasedEmo [24] under text supervisions from different pre-trained textual models. Specifically, we employ SKEP [36], DistilBert [31], Bert-base [6], RoBerta-base [20], Bert-large [6] and ReBerta-large [20].**

**5.2.3 Contrastive Temperature  $\tau$ .** It controls the strength of penalties on negative samples [38]. As shown in the results, PACL is not sensitive to its selection. The enhanced model achieves slightly superior results under smaller temperatures compared to larger selections. Therefore, we adopt  $\tau = 0.07$  according to results and related works [12, 28].

## 5.3 Different Text Supervisions

We evaluate the generalizability of PACL by employing the knowledge from various pre-trained textual models to enhance the visual model. As shown in Fig. 5, different pre-trained textual models bring consistent performance improvements on both Emotion-6 and UnbiasedEmo. It validates the universal advantages of textual models in encoding emotional-level features compared to visual models. Additionally, we can observe a positive correlation between the representative capability of the textual model and its enhancement effectiveness through PACL on the visual model. It demonstrates the scalability of PACL relative to the textual models. With the evolution of large language models in recent years, the potential of PACL can be further exploited.

## 6 CONCLUSION

In this paper, we focus on bridging the visual “affective gap” of the pre-trained visual models. It arises from the misalignment between the objectives of pre-training and downstream tasks. Inspired by the advantages of language, we propose Partitioned Adaptive Contrastive Learning (PACL) that enhances visual models by leveraging the unified factual-level and emotional-level features from the pre-trained textual model. PACL learns from the factual and emotional connections from noisy image-text pairs collected from social media by dataset partition, unimodal sample cluster, and adaptive contrastive learning. Through extensive experiments, we demonstrate the effectiveness and scalability of PACL in enhancing the emotional perception of visual models, validating the importance of bridging the “affective gap”.



## REFERENCES

- [1] Damian Borth, Rongrong Ji, Tao Chen, Thomas M. Breuel, and Shih-Fu Chang. 2013. Large-Scale Visual Sentiment Ontology and Detectors using Adjective Noun Pairs. In *MM 2013, Barcelona, Spain, October 21-25, 2013*. ACM, 223–232. <https://doi.org/10.1145/2502081.2502282>
- [2] Tao Chen, Damian Borth, Trevor Darrell, and Shih-Fu Chang. 2014. DeepSentimentBank: Visual Sentiment Concept Classification with Deep Convolutional Neural Networks. *CoRR abs/1410.8586* (2014). arXiv:1410.8586 <http://arxiv.org/abs/1410.8586>
- [3] Yuxiao Chen, Jianbo Yuan, Yu Tian, Shijie Geng, Xinyu Li, Ding Zhou, Dimitris N. Metaxas, and Hongxia Yang. 2023. Revisiting Multimodal Representation in Contrastive Learning: From Patch and Token Embeddings to Finite Discrete Tokens. In *CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*. Computer Vision Foundation / IEEE, 15095–15104. <https://doi.org/10.1109/CVPR52729.2023.01449>
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR 2009, 20-25 June 2009, Miami, Florida, USA*. Computer Vision Foundation / IEEE, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [5] Sinuo Deng, Lifang Wu, Ge Shi, Lehao Xing, Wenjin Hu, Heng Zhang, and Ye Xiang. 2023. Simple But Powerful, a Language-Supervised Method for Image Emotion Classification. *IEEE Trans. Affect. Comput.* 14, 4 (2023), 3317–3331. <https://doi.org/10.1109/TAFFC.2022.3225049>
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [7] Xiaoyi Dong, Jianmin Bao, Yinglin Zheng, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. 2023. MaskCLIP: Masked Self-Distillation Advances Contrastive Language-Image Pretraining. In *CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*. Computer Vision Foundation / IEEE, 10995–11005. <https://doi.org/10.1109/CVPR52729.2023.01058>
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/forum?id=YicbFdNTTy>
- [9] Shaojing Fan, Zhiqi Shen, Ming Jiang, Bryan L. Koenig, Juan Xu, Mohan S. Kankanhalli, and Qi Zhao. 2018. Emotional Attention: A Study of Image Sentiment and Visual Attention. In *CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. Computer Vision Foundation / IEEE, 7521–7531. <https://doi.org/10.1109/CVPR.2018.00785>
- [10] Tinglei Feng, Jiakuan Liu, and Jufeng Yang. 2023. Probing Sentiment-Oriented PreTraining Inspired by Human Sentiment Perception Mechanism. In *CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*. Computer Vision Foundation / IEEE, 2850–2860. <https://doi.org/10.1109/CVPR52729.2023.00279>
- [11] Yuting Gao, Jinfeng Liu, Zihan Xu, Jun Zhang, Ke Li, Rongrong Ji, and Chunhua Shen. 2022. PyramidCLIP: Hierarchical Feature Alignment for Vision-Language Model Pretraining. In *NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*. [http://papers.nips.cc/paper\\_files/paper/2022/hash/e9882f77c44a10acc01132302bac9d8-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/e9882f77c44a10acc01132302bac9d8-Abstract-Conference.html)
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 9726–9735. <https://doi.org/10.1109/CVPR42600.2020.00975>
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. Computer Vision Foundation / IEEE, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [14] Morris B Holbrook and John O’Shaughnessy. 1984. The Role of Emotion in Advertising. *Psychology & Marketing* 1, 2 (1984), 45–64.
- [15] Elyor Kodirov, Tao Xiang, and Shaogang Gong. 2017. Semantic Autoencoder for Zero-Shot Learning. In *CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*. Computer Vision Foundation / IEEE, 4447–4456. <https://doi.org/10.1109/CVPR.2017.473>
- [16] Ronak Kosti, José M. Álvarez, Adrià Recasens, and Àgata Lapedriza. 2017. EMOTIC: Emotions in Context Dataset. In *CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*. Computer Vision Foundation / IEEE, 2309–2317. <https://doi.org/10.1109/CVPRW.2017.285>
- [17] Sven Kraus, Matthias Althoff, Bernd Heifing, and Martin Buss. 2009. Cognition and Emotion in Autonomous Cars. In *Intelligent Vehicles Symposium 2009*. IEEE, 635–640.
- [18] Peter J Lang, Margaret M Bradley, Bruce N Cuthbert, et al. 1997. International Affective Picture System (IAPS): Technical Manual and Affective Ratings. *NIMH Center for the Study of Emotion and Attention* 1, 39-58 (1997), 3.
- [19] Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y. Zou. 2022. Mind the Gap: Understanding the Modality Gap in Multi-modal Contrastive Representation Learning. In *NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*. [http://papers.nips.cc/paper\\_files/paper/2022/hash/702f4db7543a7432431df588d57bc7c9-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/702f4db7543a7432431df588d57bc7c9-Abstract-Conference.html)
- [20] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR abs/1907.11692* (2019). arXiv:1907.11692 <http://arxiv.org/abs/1907.11692>
- [21] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*. IEEE, 9992–10002. <https://doi.org/10.1109/ICCV48922.2021.00986>
- [22] Yoshihiro Maruyama. 2020. The Conditions of Artificial General Intelligence: Logic, Autonomy, Resilience, Integrity, Morality, Emotion, Embodiment, and Embeddedness. In *13th AGI 2020, St. Petersburg, Russia, September 16-19, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12177)*. Springer, 242–251. [https://doi.org/10.1007/978-3-030-52152-3\\_25](https://doi.org/10.1007/978-3-030-52152-3_25)
- [23] Joseph A Mikels, Barbara L Fredrickson, Gregory R Larkin, Casey M Lindberg, Sam J Maglio, and Patricia A Reuter-Lorenz. 2005. Emotional Category Data on Images from the International Affective Picture System. *Behavior Research Methods* 37 (2005), 626–630.
- [24] Rameswar Panda, Jianming Zhang, Haoxiang Li, Joon-Young Lee, Xin Lu, and Amit K. Roy-Chowdhury. 2018. Contemplating Visual Emotions: Understanding and Overcoming Dataset Bias. In *ECCV 2018, Munich, Germany, September 8-14, 2018, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 11206)*. Springer, 594–612. [https://doi.org/10.1007/978-3-030-01216-8\\_36](https://doi.org/10.1007/978-3-030-01216-8_36)
- [25] Kuan-Chuan Peng, Tsuhan Chen, Amir Sadovnik, and Andrew C. Gallagher. 2015. A Mixed Bag of Emotions: Model, Predict, and Transfer Emotion Distributions. In *CVPR 2015, Boston, MA, USA, June 7-12, 2015*. Computer Vision Foundation / IEEE, 860–868. <https://doi.org/10.1109/CVPR.2015.7298687>
- [26] Juan Manuel Pérez, Juan Carlos Giudici, and Franco M. Luque. 2021. Pysentimiento: A Python Toolkit for Sentiment Analysis and SocialNLP tasks. *CoRR abs/2106.09462* (2021). arXiv:2106.09462 <https://arxiv.org/abs/2106.09462>
- [27] Robert Plutchik and Henry Kellerman. 1980. *Emotion, Theory, Research, and Experience: Theory, Research and Experience*. Academic press.
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 8748–8763. <http://proceedings.mlr.press/v139/radford21a.html>
- [29] Tianrong Rao, Xiaoxu Li, and Min Xu. 2020. Learning Multi-level Deep Representations for Image Emotion Classification. *Neural Process. Lett.* 51, 3 (2020), 2043–2061. <https://doi.org/10.1007/S11063-019-10033-9>
- [30] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. 2022. DenseCLIP: Language-Guided Dense Prediction with Context-Aware Prompting. In *CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*. Computer Vision Foundation / IEEE, 18061–18070. <https://doi.org/10.1109/CVPR52688.2022.01755>
- [31] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *CoRR abs/1910.01108* (2019). arXiv:1910.01108 <http://arxiv.org/abs/1910.01108>
- [32] Harold Schlosberg. 1954. Three Dimensions of Emotion. *Psychological review* 61, 2 (1954), 81.
- [33] Sheng Shen, Chunyuan Li, Xiaowei Hu, Yujia Xie, Jianwei Yang, Pengchuan Zhang, Zhe Gan, Lijuan Wang, Lu Yuan, Ce Liu, Kurt Keutzer, Trevor Darrell, Anna Rohrbach, and Jianfeng Gao. 2022. K-LITE: Learning Transferable Visual Models with External Knowledge. In *NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*. [http://papers.nips.cc/paper\\_files/paper/2022/hash/63fef082863f4775c3563e18cbb17-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/63fef082863f4775c3563e18cbb17-Abstract-Conference.html)
- [34] Yang Sun, Nan Yu, and Guohong Fu. 2021. A Discourse-Aware Graph Neural Network for Emotion Recognition in Multi-Party Conversation. In *EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*. Association for Computational Linguistics, 2949–2958. <https://doi.org/10.18653/v1/2021.FINDINGS-EMNLP.252>
- [35] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. 2018. Learning to Compare: Relation Network for Few-Shot Learning. In *CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. Computer Vision Foundation / IEEE, 1199–1208. <https://doi.org/10.1109/CVPR.2018.00131>
- [36] Hao Tian, Can Gao, Xinyan Xiao, Hao Liu, Bolei He, Hua Wu, Haifeng Wang, and Feng Wu. 2020. SKEP: Sentiment Knowledge Enhanced Pre-training for Sentiment Analysis. In *ACL 2020, Online, July 5-10, 2020*. Association for Computational Linguistics, 4067–4076. <https://doi.org/10.18653/v1/2020.ACL-MAIN.374>
- [37] Quoc-Tuan Truong and Hady W. Lauw. 2017. Visual Sentiment Analysis for Review Images with Item-Oriented and User-Oriented CNN. In *MM 2017, Mountain View, CA, USA, October 23-27, 2017*. ACM, 1274–1282. <https://doi.org/10.1145/3123266.3123374>

929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971  
972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044

- 1045 [38] Feng Wang and Huaping Liu. 2021. Understanding the Behaviour of Contrastive  
1046 Loss. In *CVPR 2021, virtual, June 19-25, 2021*. Computer Vision Foundation / IEEE,  
2495–2504. <https://doi.org/10.1109/CVPR46437.2021.00252>
- 1047 [39] Yiwei Wei, Shaozu Yuan, Ruosong Yang, Lei Shen, Zhangmeizhi Li, Longbiao  
1048 Wang, and Meng Chen. 2023. Tackling Modality Heterogeneity with Multi-View  
1049 Calibration Network for Multimodal Sentiment Detection. In *ACL 2023, Toronto,  
1050 Canada, July 9-14, 2023*. Association for Computational Linguistics, 5240–5252.  
<https://doi.org/10.18653/V1/2023.ACL-LONG.287>
- 1051 [40] Liwen Xu, Zhengtao Wang, Bin Wu, and Simon Lui. 2022. MDAN: Multi-level  
1052 Dependent Attention Network for Visual Emotion Analysis. In *CVPR 2022, New  
1053 Orleans, LA, USA, June 18-24, 2022*. IEEE, 9469–9478. <https://doi.org/10.1109/CVPR52688.2022.00926>
- 1054 [41] Jingyuan Yang, Jie Li, Leida Li, Xiumei Wang, and Xinbo Gao. 2021. A Circular-  
1055 Structured Representation for Visual Emotion Distribution Learning. In *CVPR  
1056 2021, virtual, June 19-25, 2021*. Computer Vision Foundation / IEEE, 4237–4246.  
<https://doi.org/10.1109/CVPR46437.2021.00422>
- 1057 [42] Jufeng Yang, Dongyu She, Yu-Kun Lai, Paul L. Rosin, and Ming-Hsuan Yang. 2018.  
1058 Weakly Supervised Coupled Networks for Visual Sentiment Analysis. In *CVPR  
1059 2018, Salt Lake City, UT, USA, June 18-22, 2018*. Computer Vision Foundation /  
1060 IEEE, 7584–7592. <https://doi.org/10.1109/CVPR.2018.00791>
- 1061 [43] Jufeng Yang, Dongyu She, and Ming Sun. 2017. Joint Image Emotion Classification  
1062 and Distribution Learning via Deep Convolutional Neural Network. In *IJCAI  
1063 2017, Melbourne, Australia, August 19-25, 2017*. ijcai.org, 3266–3272. <https://doi.org/10.24963/IJCAI.2017/456>
- 1064 [44] Jufeng Yang, Dongyu She, Ming Sun, Ming-Ming Cheng, Paul L. Rosin, and Liang  
1065 Wang. 2018. Visual Sentiment Prediction Based on Automatic Discovery of  
1066 Affective Regions. *IEEE Trans. Multim.* 20, 9 (2018), 2513–2525. <https://doi.org/10.1109/TMM.2018.2803520>
- 1067 [45] Xiaocui Yang, Shi Feng, Daling Wang, and Yifei Zhang. 2021. Image-Text  
1068 Multimodal Emotion Classification via Multi-View Attentional Network. *IEEE Trans.  
1069 Multim.* 23 (2021), 4014–4026. <https://doi.org/10.1109/TMM.2020.3035277>
- 1070 [46] Xiaocui Yang, Shi Feng, Yifei Zhang, and Daling Wang. 2021. Multimodal Sentiment  
1071 Detection Based on Multi-channel Graph Neural Networks. In *ACL/IJCNLP  
1072 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*. Association for Com-  
1073 putational Linguistics, 2021. <https://doi.org/10.18653/V1/2021.ACL-LONG.28>
- 1074 [47] Yun Yang, Peng Cui, Wenwu Zhu, and Shiqiang Yang. 2013. User Interest and  
1075 Social Influence Based Emotion Prediction for Individuals. In *MM 2013, Barcelona,  
1076 Spain, October 21-25, 2013*. ACM, 785–788. <https://doi.org/10.1145/2502081.2502204>
- 1077 [48] Victoria Yanulevskaya, Jan C. van Gemert, Katharina Roth, Ann-Katrin Herbold,  
1078 Nicu Sebe, and Jan-Mark Geusebroek. 2008. Emotional Valence Categorization using  
1079 Holistic Image Features. In *ICIP 2008, October 12-15, 2008, San Diego, California,  
1080 USA*. IEEE, 101–104. <https://doi.org/10.1109/ICIP.2008.4711701>
- 1081 [49] Yingrui Ye, Yuya Moroto, Keisuke Maeda, Takahiro Ogawa, and Miki Haseyama.  
1082 2022. Affective Embedding Framework with Semantic Representations from  
1083 Tweets for Zero-Shot Visual Sentiment Prediction. In *MMAsia 2022, Tokyo, Japan,  
1084 December 13-16, 2022*. ACM, 6:1–6:7. <https://doi.org/10.1145/3551626.3564950>
- 1085 [50] Richard Yonck. 2020. *Heart of the Machine: Our Future in a World of Artificial  
1086 Emotional Intelligence*. Arcade.
- 1087 [51] Quanzeng You, Hailin Jin, and Jiebo Luo. 2017. Visual Sentiment Analysis by  
1088 Attending on Local Image Regions. In *AAAI 2017, February 4-9, 2017, San Francisco,  
1089 California, USA*. AAAI Press, 231–237. <https://doi.org/10.1609/AAAI.V31I1.10501>
- 1090 [52] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. 2015. Robust Image  
1091 Sentiment Analysis Using Progressively Trained and Domain Transferred Deep  
1092 Networks. In *AAAI 2015, January 25-30, 2015, Austin, Texas, USA*. AAAI Press,  
1093 381–388. <https://doi.org/10.1609/AAAI.V29I1.9179>
- 1094 [53] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. 2016. Building a Large  
1095 Scale Dataset for Image Emotion Recognition: The Fine Print and The Benchmark.  
1096 In *AAAI 2016, February 12-17, 2016, Phoenix, Arizona, USA*. AAAI Press, 308–314.  
<https://doi.org/10.1609/AAAI.V30I1.9987>
- 1097 [54] Jianbo Yuan, Sean McDonough, Quanzeng You, and Jiebo Luo. 2013. Stribute:  
1098 Image Sentiment Analysis from a Mid-Level Perspective. In *WISDOM 2013,  
1099 Chicago, IL, USA, August 11, 2013*. ACM, 10:1–10:8. <https://doi.org/10.1145/2502069.2502079>
- 1100 [55] Chi Zhan, Dongyu She, Sicheng Zhao, Ming-Ming Cheng, and Jufeng Yang.  
1101 2019. Zero-Shot Emotion Recognition via Affective Structural Embedding. In  
1102 *ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. IEEE, 1151–1160.  
<https://doi.org/10.1109/ICCV.2019.00124>
- 1103 [56] Haimin Zhang and Min Xu. 2021. Weakly Supervised Emotion Intensity Pre-  
1104 diction for Recognition of Emotions in Images. *IEEE Trans. Multim.* 23 (2021),  
1105 2033–2044. <https://doi.org/10.1109/TMM.2020.3007352>
- 1106 [57] Li Zhang, Tao Xiang, and Shaogang Gong. 2017. Learning a Deep Embedding  
1107 Model for Zero-Shot Learning. In *CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*.  
1108 Computer Vision Foundation / IEEE, 3010–3019. <https://doi.org/10.1109/CVPR.2017.321>
- 1109 [58] Sitao Zhang, Yimu Pan, and James Z. Wang. 2023. Learning Emotion Representations  
1110 from Verbal and Nonverbal Communication. In *CVPR 2023, Vancouver, BC,  
1111 Canada, June 17-24, 2023*. Computer Vision Foundation / IEEE, 18993–19004.  
<https://doi.org/10.1109/CVPR52729.2023.01821>
- 1112 [59] Sicheng Zhao, Xuanbai Chen, Xiangyu Yue, Chuang Lin, Pengfei Xu, Ravi Kr-  
1113 ishna, Jufeng Yang, Guiguang Ding, Alberto L. Sangiovanni-Vincentelli, and Kurt  
1114 Keutzer. 2022. Emotional Semantics-Preserved and Feature-Aligned CycleGAN  
1115 for Visual Emotion Adaptation. *IEEE Trans. Cybern.* 52, 10 (2022), 10000–10013.  
<https://doi.org/10.1109/TCYB.2021.3062750>
- 1116 [60] Sicheng Zhao, Yue Gao, Xiaolei Jiang, Hongxun Yao, Tat-Seng Chua, and Xi-  
1117 aoshuai Sun. 2014. Exploring Principles-of-Art Features For Image Emotion  
1118 Recognition. In *MM 2014, Orlando, FL, USA, November 03 - 07, 2014*. ACM, 47–56.  
<https://doi.org/10.1145/2647868.2654930>
- 1119 [61] Sicheng Zhao, Guoli Jia, Jufeng Yang, Guiguang Ding, and Kurt Keutzer. 2021.  
1120 Emotion Recognition From Multiple Modalities: Fundamentals and Methodologies.  
1121 *IEEE Signal Process. Mag.* 38, 6 (2021), 59–73. <https://doi.org/10.1109/MSP.2021.3106895>
- 1122 [62] Sicheng Zhao, Zizhou Jia, Hui Chen, Leida Li, Guiguang Ding, and Kurt Keutzer.  
1123 2019. PDANet: Polarity-consistent Deep Attention Network for Fine-grained  
1124 Visual Emotion Regression. In *MM 2019, Nice, France, October 21-25, 2019*. ACM,  
1125 192–201. <https://doi.org/10.1145/3343031.3351062>
- 1126 [63] Sicheng Zhao, Hongxun Yao, Yue Gao, Rongrong Ji, Wenlong Xie, Xiaolei Jiang,  
1127 and Tat-Seng Chua. 2016. Predicting Personalized Emotion Perceptions of Social  
1128 Images. In *MM 2016, Amsterdam, The Netherlands, October 15-19, 2016*. ACM,  
1129 1385–1394. <https://doi.org/10.1145/2964284.2964289>
- 1130 [64] Sicheng Zhao, Xingxu Yao, Jufeng Yang, Guoli Jia, Guiguang Ding, Tat-Seng Chua,  
1131 Björn W. Schuller, and Kurt Keutzer. 2022. Affective Image Content Analysis:  
1132 Two Decades Review and New Perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 10 (2022), 6729–6751. <https://doi.org/10.1109/TPAMI.2021.3094362>
- 1133 [65] Sicheng Zhao, Xin Zhao, Guiguang Ding, and Kurt Keutzer. 2018. EmotionGAN:  
1134 Unsupervised Domain Adaptation for Learning Discrete Probability Distributions  
1135 of Image Emotions. In *MM 2018, Seoul, Republic of Korea, October 22-26, 2018*.  
1136 ACM, 1319–1327. <https://doi.org/10.1145/3240508.3240591>
- 1137 [66] Chong Zhou, Chen Change Loy, and Bo Dai. 2022. Extract Free Dense Labels  
1138 from CLIP. In *ECCV 2022, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part  
1139 XXVIII (Lecture Notes in Computer Science, Vol. 13688)*. Springer, 696–712. [https://doi.org/10.1007/978-3-031-19815-1\\_40](https://doi.org/10.1007/978-3-031-19815-1_40)
- 1140 [67] Jinghao Zhou, Li Dong, Zhe Gan, Lijuan Wang, and Furu Wei. 2023. Non-  
1141 Contrastive Learning Meets Language-Image Pre-Training. In *CVPR 2023, Van-  
1142 couver, BC, Canada, June 17-24, 2023*. Computer Vision Foundation / IEEE, 11028–  
1143 11038. <https://doi.org/10.1109/CVPR52729.2023.01061>
- 1144 [68] Xinge Zhu, Liang Li, Weigang Zhang, Tianrong Rao, Min Xu, Qingming Huang,  
1145 and Dong Xu. 2017. Dependency Exploitation: A Unified CNN-RNN Approach  
1146 for Visual Emotion Recognition. In *IJCAI 2017, Melbourne, Australia, August 19-25,  
1147 2017*. ijcai.org, 3595–3601. <https://doi.org/10.24963/IJCAI.2017/503>