

# SPP-SCL: Semi-Push-Pull Supervised Contrastive Learning for Image-Text Sentiment Analysis and Beyond

Jiesheng Wu<sup>1</sup>, Shengrong Li<sup>2</sup> \*

<sup>1</sup>School of Computer and Information, Anhui Normal University, Wuhu, 241003, China

<sup>2</sup>College of Artificial Intelligence, Nanjing University of Aeronautics and Astronautics, Nanjing, 211106, China  
jasonwu@mail.nankai.edu.cn, lisrong@nuaa.edu.cn

## Abstract

Existing Image-Text Sentiment Analysis (ITSA) methods may suffer from inconsistent intra-modal and inter-modal sentiment relationships. Therefore, we develop a method that balances before fusing to solve the issue of vision-language imbalance intra-modal and inter-modal sentiment relationships; that is, a Semi-Push-Pull Supervised Contrastive Learning (SPP-SCL) method is proposed. Specifically, the method is implemented using a novel two-step strategy, namely first using the proposed intra-modal supervised contrastive learning to pull the relationships between the intra-modal and then performing a well-designed conditional execution statement. If the statement result is false, our method will perform the second step, which is inter-modal supervised contrastive learning to push away the relationships between inter-modal. The two-step strategy will balance the intra-modal and inter-modal relationships to achieve the purpose of relationship consistency and finally perform cross-modal feature fusion for sentiment analysis and detection. Experimental studies on three public image-text sentiment and sarcasm detection datasets demonstrate that SPP-SCL significantly outperforms state-of-the-art methods by a large margin and is more discriminative in sentiment.

## Introduction

With the explosive growth of social media, users increasingly express sentiments through images accompanied by textual descriptions. This has spurred extensive interest in **image-text sentiment analysis (ITSA)** (Das and Singh 2023; Xue et al. 2022; Zadeh et al. 2017). The core challenge of ITSA lies in accurately modeling and integrating emotional cues from different modalities to make consistent and robust predictions.

Early methods typically concatenate high-level features from each modality for fusion (Xu and Mao 2017; Xu 2017), but such strategies often fall short in capturing nuanced or conflicting emotional signals across modalities. To address this, recent studies have introduced interaction-based fusion modules (Yang et al. 2021a; Yu and Jiang 2019; Zhu et al. 2023b) to better model fine-grained cross-modal sentiment relationships. Despite progress, a critical issue remains un-

derexplored: *the imbalance and inconsistency among intra-modal and inter-modal sentiment representations.*

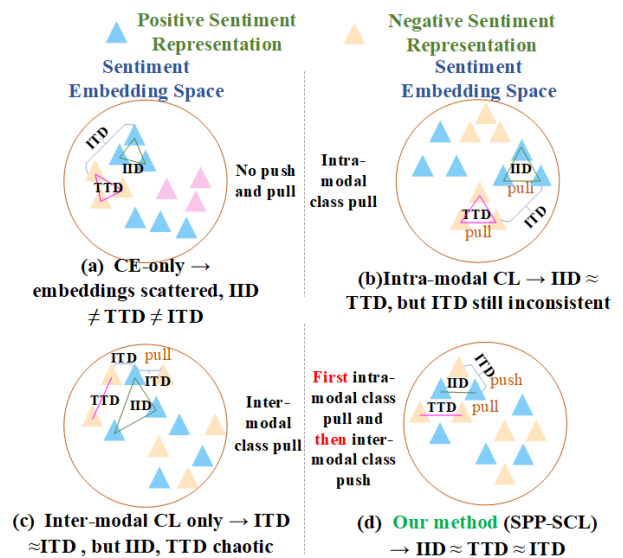


Figure 1: Comparison of intra-modal and inter-modal sentiment distances under different training strategies. (a) Cross-entropy loss only: image and text sentiment representations are scattered, leading to inconsistent intra- and inter-modal sentiment relationships. (b) Intra-modal contrastive learning aligns IID and TTD, but fails on ITD. (c) Inter-modal contrastive learning aligns ITD but ignores intra-modal structure. (d) Our proposed SPP-SCL balances all three sentiment distances, yielding consistent sentiment embeddings.

As illustrated in Fig. 1, existing methods primarily rely on cross-entropy loss, which encourages discriminative classification but fails to align sentiment representations across modalities. Specifically, we observe that the average pairwise distances within image sentiment features (IID), within text sentiment features (TTD), and between image-text pairs sentiment features (ITD) often differ significantly, even for samples sharing the same sentiment label. **Here, IID (Image-to-Image sentiment embedding Distance), TTD (Text-to-Text sentiment embedding Distance), and ITD (Image-to-Text sentiment embedding Distance) denote**

\*Corresponding author.

the average Euclidean distances between sentiment embeddings of samples with the same label, computed respectively within the image modality, within the text modality, and across modalities. Such inconsistencies in sentiment-level embedding space undermine generalization and hinder effective fusion. While some works employ contrastive learning (CL) to address this issue (Li et al. 2022b), most adopt a *fuse-then-contrast* strategy, applying CL on fused multi-modal features. This approach risks semantic noise propagation and lacks targeted alignment of unimodal emotional representations.

To tackle these challenges, we propose a novel **Balance-before-Fuse** framework that aligns sentiment representations within and across modalities before fusion. Inspired by alignment-then-fusion paradigms in vision-language tasks (Li et al. 2021, 2022a; Yang, Bisk, and Gao 2021), we introduce a two-step **Semi-Push-Pull Supervised Contrastive Learning (SPP-SCL)** approach. In the first step, we apply supervised contrastive losses to image and text branches independently, pulling same-sentiment samples closer within each modality and achieving intra-modal sentiment alignment (*i.e.*,  $\text{IID} \approx \text{IID}$  and  $\text{TTD} \approx \text{TTD}$ ). In the second step, we conditionally apply inter-modal contrastive learning to align sentiment representations across modalities, based on a diagonal similarity threshold. **We aim to simultaneously minimize the discrepancy among intra-image, intra-text, and image-text sentiment distances under a unified supervised contrastive learning framework**, thereby constructing a modality-balanced and sentiment-consistent embedding space that facilitates robust multi-modal fusion. To further improve fusion quality, we introduce two lightweight yet effective components: a **Hierarchical Attention (HA)** module for context-aware sentiment extraction from text, and a **Cross-Modal Fusion (CMF)** module for dynamic integration of aligned features. *Unlike general-purpose vision-language alignment frameworks, our method is specifically tailored for sentiment analysis by aligning emotion-aware representations across modalities, rather than generic semantic features.*

**Our main contributions are summarized as follows:**

- We propose a novel two-step **Semi-Push-Pull Supervised Contrastive Learning (SPP-SCL)** framework that aligns intra- and inter-modal *sentiment* relationships prior to fusion, yielding a well-structured space.
- We design three supervised contrastive objectives—two intra-modal (for image and text) and one inter-modal—that enable fine-grained sentiment-level alignment without requiring data augmentation or large-scale pretraining.
- We develop a Hierarchical Attention module for efficient sentiment representation from text, and a Cross-Modal Fusion module for robust visual-linguistic interaction.
- Extensive experiments on three public datasets (MVSA-S, MVSA-M, and HFM) demonstrate that our method significantly outperforms state-of-the-art baselines, especially in fine-grained or ambiguous sentiment scenarios.

## Related Work

### Image-Text Sentiment Analysis

Image-Text Sentiment Analysis (ITSA) aims to classify the sentiment of paired visual and textual content, particularly in social media contexts (Xie et al. 2024; He et al. 2023; Tang et al. 2023). Early approaches such as HSAN and MultiSentiNet (Xu 2017; Xu and Mao 2017) extracted features independently from each modality and fused them via simple concatenation, which limited cross-modal interaction. To address this, subsequent works introduced interaction-aware architectures. For example, Co-MN-Hop6 (Xu, Mao, and Chen 2018) and MVAN (Yang et al. 2021a) leveraged memory-based networks to enable iterative alignment between modalities, while MGNNs (Yang et al. 2021b) used sentiment-aware graph neural networks. More recent methods such as ITIN and MULSER (Zhu et al. 2023a,b) focused on deep fusion but often overlooked token-level alignment or sentiment consistency. To improve alignment, CLMLF (Li et al. 2022b) adopted contrastive learning using both label- and data-based signals. However, it applied contrastive objectives directly on fused representations, without explicitly balancing intra- and inter-modal features. Inspired by this, our method introduces a *balance-before-fuse* strategy to align intra- and inter-modal sentiment embeddings prior to fusion.

### Contrastive Learning

Contrastive Learning (CL) has shown strong representation power in vision and language domains through instance discrimination techniques, as demonstrated in MoCo (He et al. 2020), SimCLR (Chen et al. 2020), and SimCSE (Gao, Yao, and Chen 2021). In the multi-modal setting, CLIP (Radford et al. 2021) paved the way for joint vision-language representation learning. Supervised contrastive learning (SCL) (Khosla et al. 2020) extends this idea by incorporating label information, enabling multi-positive alignment. Recent applications of SCL include text classification (Gunel et al. 2020) and trimodal sentiment analysis (Mai et al. 2022; Yang et al. 2023). However, most methods treat contrastive learning as an auxiliary module or apply it after feature fusion, which may propagate semantic noise across modalities. In contrast, we propose a two-step supervised contrastive framework that explicitly aligns intra-modal and inter-modal sentiment embeddings, ensuring a consistent and discriminative feature space for downstream fusion.

## Methodology

### Overview

The overall architecture of our proposed **Semi-Push-Pull Supervised Contrastive Learning (SPP-SCL)** framework is illustrated in Fig. 2. The model operates in two sequential steps designed to align sentiment representations both within and across modalities, prior to final fusion. We first adopt two frozen feature encoders—a ResNet-50 pretrained on ImageNet for image inputs and a BERT-base model pretrained on general text corpora for textual inputs—to extract high-level semantic features. These encoders are **kept**

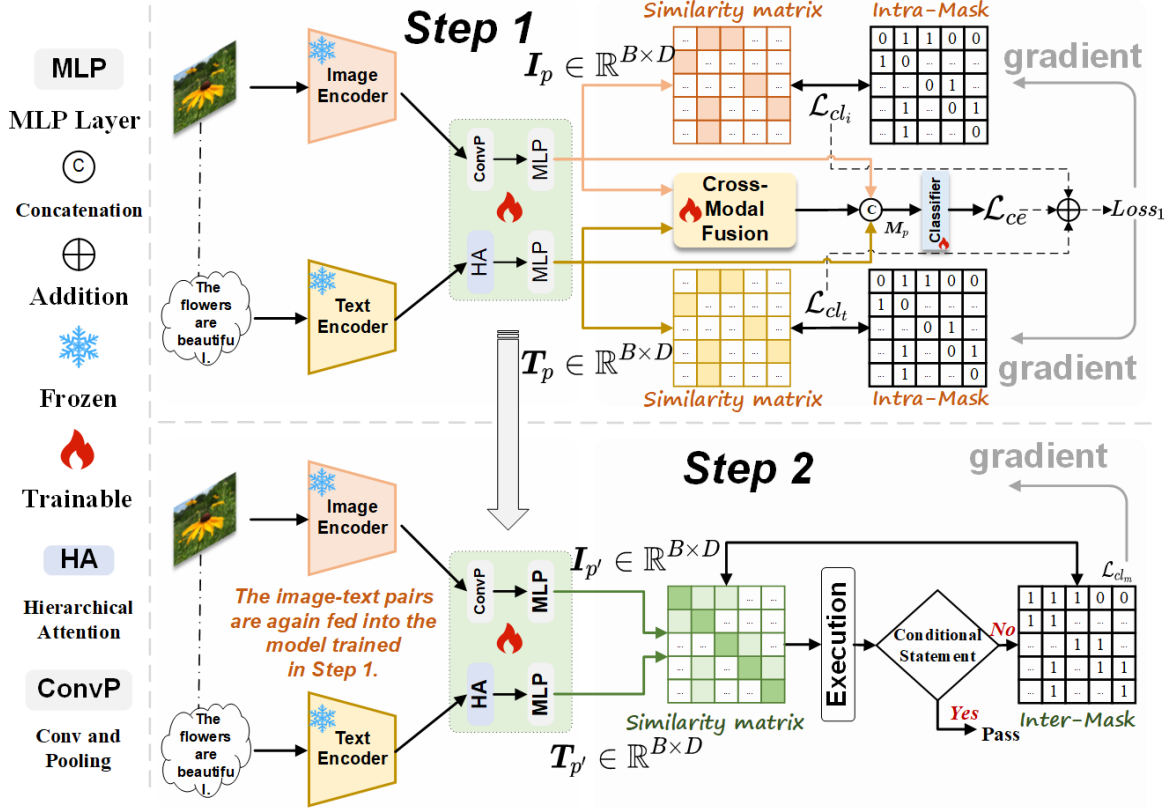


Figure 2: Overall architecture of SPP-SCL. The framework includes two main steps: intra-modal sentiment alignment via supervised contrastive learning ( $\mathcal{L}_{cl_i}$  and  $\mathcal{L}_{cl_t}$ ), and conditional inter-modal sentiment alignment ( $\mathcal{L}_{cl_m}$ ).

**fixed during training** to prevent overfitting on small-scale sentiment datasets and to ensure that performance improvements stem primarily from our proposed contrastive alignment strategy rather than feature extractor fine-tuning. In the **first step**, intra-modal sentiment alignment is performed using two supervised contrastive losses:  $\mathcal{L}_{cl_i}$  for the image branch and  $\mathcal{L}_{cl_t}$  for the text branch. These losses encourage sentiment representations of samples from the same class to be pulled closer within each modality. Both image and text embeddings are projected into a shared sentiment embedding space, where the contrastive objectives operate jointly with a standard cross-entropy loss for sentiment classification. This step ensures intra-modal consistency, minimizing discrepancies such as IID and TTD. In the **second step**, we conduct a conditional judgment on the alignment quality between modalities. Specifically, we compute the similarity matrix between paired image and text sentiment embeddings. If the diagonal consistency score falls below a predefined threshold, we activate an inter-modal supervised contrastive loss  $\mathcal{L}_{cl_m}$ , which further pulls image-text pairs of the same sentiment class closer together, thus reducing ITD and achieving global sentiment-level alignment.

By decoupling intra- and inter-modal sentiment alignment into a two-step training process, SPP-SCL ensures a well-structured sentiment embedding space where **intra-image,**

**intra-text, and image-text sentiment distances are harmonized**, resulting in more robust and discriminative fused representations.

### Multi-modal Feature Extraction

Given an image-text pair sample from an ITSA dataset, denoted as  $S = \{I, T\}$ , where  $I$  and  $T$  represent the image and text modalities, respectively, let the corresponding sentiment label be  $Y = y_i$ , indicating that the sample belongs to the  $i$ -th sentiment class. Our goal is to propose a framework for balancing sentiment consistency, independent of encoder scale. SPP-SCL is backbone-agnostic and can replace other encoders directly. We focus on widely adopted backbones (BERT-base, ResNet-50) to ensure fair comparison.

**Image Encoder** We adopt ResNet-50 (He et al. 2016) as the image encoder. Given an input image  $I$ , we extract feature maps from the layer preceding the final global average pooling. These features are then passed through a  $1 \times 1$  convolutional layer with batch normalization and ReLU activation, followed by an adaptive average pooling layer and a fully connected projection layer with non-linear activation. The output is the final image representation  $I_p \in \mathbb{R}^{B \times D}$ , where  $B$  denotes the batch size and  $D$  the embedding dimension. The overall encoding process is defined as:

$$I_p = \mathcal{M}(\mathcal{P}(c(\text{ResNet}(I)))) , \quad (1)$$

where  $c(\cdot)$  denotes the  $1 \times 1$  convolutional layer,  $\mathcal{P}(\cdot)$  the adaptive average pooling operation, and  $\mathcal{M}(\cdot)$  the fully connected transformation for dimensionality reduction.

**Text Encoder** We adopt BERT (Devlin et al. 2018) as the text encoder. Given a text input  $T$ , we follow prior findings (Jawahar, Sagot, and Seddah 2019; Sun et al. 2019) that fine-tuning only the last four transformer layers (layers 9–12) yields optimal performance for text classification tasks. Accordingly, we extract two types of features from these layers: (1) the [CLS] token embeddings, which capture contextual semantics, and (2) the token-level word embeddings from the corresponding hidden states, which preserve fine-grained lexical information.

To effectively combine these two feature types, we propose a **Hierarchical Attention (HA)** module that fuses contextual and word-level representations into a unified sentiment embedding. Specifically, let  $\mathbf{T}_c \in \mathbb{R}^{B \times 4 \times 1 \times C}$  denote the [CLS] embeddings and  $\mathbf{T}_w \in \mathbb{R}^{B \times 4 \times N \times C}$  denote the word embeddings, where  $B$  is the batch size,  $N$  the number of words, and  $C$  the hidden dimension.

The HA module first applies one-dimensional global average pooling to  $\mathbf{T}_c$ , reducing it to  $\mathbb{R}^{B \times 4 \times 1 \times 1}$ , and then passes it through a fully connected layer followed by a sigmoid activation  $\mathcal{S}(\cdot)$  to generate normalized attention weights across layers. These weights are then element-wise multiplied with  $\mathbf{T}_w$  to produce the weighted word representation  $\mathbf{T}'_w$ . Afterward, we apply summation and channel-wise pooling operations to aggregate token features across layers and time steps. Finally, the fused features are fed into an LSTM (Hochreiter and Schmidhuber 1997) followed by a projection layer  $\mathcal{M}(\cdot)$  to obtain the final text representation  $\mathbf{T}_p \in \mathbb{R}^{B \times D}$ . The entire process can be formalized as:

$$\begin{aligned} \mathbf{T}'_c &= \mathcal{S}(\mathcal{M}(\mathcal{P}_g(\mathbf{T}_c))), \mathbf{T}'_w = \mathbf{T}'_c \odot \mathbf{T}_w, \\ \mathbf{T}_p &= \mathcal{M}(\text{LSTM}(\mathcal{P}_c(\mathcal{U}(\mathbf{T}'_w)))) \end{aligned} \quad (2)$$

where  $\mathcal{S}(\cdot)$  denotes the sigmoid activation,  $\mathcal{P}_g(\cdot)$  is global average pooling,  $\mathcal{P}_c(\cdot)$  is channel-wise pooling,  $\mathcal{U}(\cdot)$  denotes summation over layers,  $\odot$  represents element-wise multiplication, and  $\text{LSTM}(\cdot)$  is the LSTM network.

### Cross-Modal Fusion Module

Effective fusion of visual and textual sentiment cues is essential for sentiment classification in multi-modal contexts (Tang et al. 2022). To this end, we propose a Cross-Modal Fusion (CMF) module that integrates the visual representation  $\mathbf{I}_p$  and the textual representation  $\mathbf{T}_p$  into a unified, sentiment-aware feature. Inspired by DFAF (Gao et al. 2019), CMF combines two key components: **intra-modal feature enhancement** via Self-Attention (SA) (Vaswani et al. 2017) and **inter-modal interaction** via dynamic attention modulation. The intuition is to dynamically adjust intra-modal attention weights based on complementary information from the other modality, enabling fine-grained sentiment-level interaction rather than simple semantic matching. Specifically, for both  $\mathbf{I}_p$  and  $\mathbf{T}_p$ , we construct respective Query, Key, and Value vectors:  $\mathbf{Q}_I, \mathbf{K}_I, \mathbf{V}_I$  for images and  $\mathbf{Q}_T, \mathbf{K}_T, \mathbf{V}_T$  for text. During self-attention, we inject inter-modal cues by modulating the query and key vec-

tors via element-wise multiplication with the other modality’s global representation (e.g.,  $\mathbf{T}_p$  modulates  $\mathbf{Q}_I, \mathbf{K}_I$ ). This allows the model to dynamically adjust its attention focus based on sentiment-relevant context from the other modality. The fusion process is defined as:

$$\begin{aligned} \mathbf{I}_m &= \text{Softmax} \left( \frac{(\mathbf{Q}_I \odot \mathbf{T}_p) \otimes (\mathbf{K}_I \odot \mathbf{T}'_p)}{\sqrt{D}} \right) \otimes \mathbf{V}_I, \\ \mathbf{T}_m &= \text{Softmax} \left( \frac{(\mathbf{Q}_T \odot \mathbf{I}_p) \otimes (\mathbf{K}_T \odot \mathbf{I}'_p)}{\sqrt{D}} \right) \otimes \mathbf{V}_T, \\ \mathbf{M}_p &= [\mathbf{I}_m, \mathbf{T}_m, \mathbf{I}_p, \mathbf{T}_p], \end{aligned} \quad (3)$$

where  $\mathbf{I}_m, \mathbf{T}_m \in \mathbb{R}^{B \times D}$  are the enhanced modality-specific features,  $\odot$  denotes element-wise multiplication,  $\otimes$  is matrix multiplication, and  $[\cdot]$  denotes concatenation.  $\text{Softmax}(\cdot)$  denotes the standard attention normalization. The resulting fused representation  $\mathbf{M}_p$  is sentiment-aware and jointly informed by both intra- and inter-modal contextual signals.

### Semi-Push-Pull Supervised Contrastive Learning

As previously discussed, robust image-text sentiment analysis (ITSA) requires a well-structured embedding space where image and text representations of the same sentiment class are closely aligned. This entails minimizing the distance and relationship discrepancies across three perspectives: Image-to-Image (IID), Text-to-Text (TTD), and Image-to-Text (ITD). To this end, we propose a **Semi-Push-Pull Supervised Contrastive Learning (SPP-SCL)** framework that incorporates both intra-modal and inter-modal alignment strategies.

**Intra-Modal Supervised Contrastive Learning** To enforce sentiment consistency within each modality, we introduce two supervised contrastive losses:  $\mathcal{L}_{cl_i}$  for the image branch and  $\mathcal{L}_{cl_t}$  for the text branch. Following (Khosla et al. 2020), we first construct an **Intra-Mask matrix**  $\mathbb{1}^{ij} = \mathbb{1}[y_i = y_j, i \neq j]$  to identify positive pairs based on sentiment labels.

Let  $\mathbf{I}_p$  and  $\mathbf{T}_p \in \mathbb{R}^{B \times D}$  denote the  $\ell_2$ -normalized image and text representations. The intra-modal similarity matrix is computed using cosine similarity. The supervised contrastive loss for the image branch is defined as:

$$\mathcal{L}_{cl_i} = -\frac{1}{B} \sum_{i=1}^B \frac{1}{B} \sum_{j=1}^B \mathbb{1}^{ij} \log \left( \frac{\exp(s(i, j)/\tau)}{\sum_{k=2}^B \exp(s(i, k)/\tau)} \right), \quad (4)$$

where  $s(i, j)$  is the cosine similarity between samples  $i$  and  $j$ , and  $\tau$  is the temperature scaling factor. The diagonal is excluded since self-similarity is trivial. The loss  $\mathcal{L}_{cl_t}$  for the text branch follows the same formulation:

$$\mathcal{L}_{cl_t} = -\frac{1}{B} \sum_{i=1}^B \frac{1}{B} \sum_{j=1}^B \mathbb{1}^{ij} \log \left( \frac{\exp(s(i, j)/\tau)}{\sum_{k=2}^B \exp(s(i, k)/\tau)} \right). \quad (5)$$

These losses effectively reduce intra-modal sentiment variance and ensure that IID  $\simeq$  TTD within the sentiment embedding space.

**Inter-modal Supervised Contrastive Learning** While intra-modal alignment achieves IID  $\simeq$  TTD, it does not guarantee cross-modal consistency (*i.e.*, ITD alignment). To address this, we propose an inter-modal supervised contrastive loss  $\mathcal{L}_{cl_m}$ .

We define an **Inter-Mask matrix**  $\mathbb{1}^{mn} = \mathbb{1}[y_m = y_n, m \neq n]$  across image-text pairs. Let  $\mathbf{I}_{p'}$  and  $\mathbf{T}_{p'}$  denote  $\ell_2$ -normalized features extracted after the first training step. The inter-modal similarity matrix is then computed between all image-text pairs. The inter-modal contrastive loss is defined as:

$$\mathcal{L}_{cl_m} = -\frac{1}{B} \sum_{m=1}^B \frac{1}{B} \sum_{n=1}^B \mathbb{1}^{mn} \log \left( \frac{\exp(s(m, n) / \tau)}{\sum_{t=1}^B \exp(s(m, t) / \tau)} \right). \quad (6)$$

This loss encourages image-text samples to align in sentiment space, thereby enforcing IID  $\simeq$  TTD  $\simeq$  ITD.

**Training Strategy: Two-Step Optimization** To balance efficiency and alignment quality, we adopt a two-step training strategy. In the first step (SPP-SCL<sub>1</sub>), we jointly optimize  $\mathcal{L}_{cl_i}$ ,  $\mathcal{L}_{cl_t}$ , and the cross-entropy loss  $\mathcal{L}_{ce}$ . After training, we evaluate whether sentiment distances among all three types (IID, TTD, ITD) are sufficiently consistent based on a threshold  $\alpha$ . If not, we proceed to the second step (SPP-SCL<sub>2</sub>), where we fine-tune the model using the inter-modal contrastive loss  $\mathcal{L}_{cl_m}$  to close any remaining cross-modal gaps. The full procedure is outlined in Algorithm 1.

## Loss Function

The overall training loss integrates four components:

$$\mathcal{L} = \lambda \mathcal{L}_{ce}(\mathbf{M}_p, Y) + \mathcal{L}_{cl_i}(\mathbf{I}_p, Y) + \mathcal{L}_{cl_t}(\mathbf{T}_p, Y) + \mathcal{L}_{cl_m}(\mathbf{T}_{p'}, \mathbf{I}_{p'}, Y), \quad (7)$$

where  $\mathcal{L}_{ce}$  represents the cross-entropy loss function, and  $\lambda$  is a hyperparameter used to balance the different losses.

## Experiments

### Datasets

We evaluate on three widely used ITSA benchmarks: MVSA-S, MVSA-M (Niu et al. 2016), and the sarcasm-oriented HFM dataset (Cai, Cai, and Wan 2019; Liang et al. 2022). Details are listed in Table 1.

Datasets	Numbers	Train sets	Validation sets	Test sets
MVSA-S	4,511	3,611	450	450
MVSA-M	17,024	13,624	1,700	1,700
HFM	24,635	19,816	2,410	2,409

Table 1: Statistics of three datasets.

### Experimental Details

Our SPP-SCL model is implemented using PyTorch and HuggingFace Transformers, and trained on an NVIDIA V100 GPU (16GB)<sup>1</sup>. We use ResNet-50 (He et al. 2016) and BERT-base-uncased (Devlin et al. 2018) as frozen image

<sup>1</sup>Codes: <https://github.com/TomorrowJW/SPP-SCL>

### Algorithm 1: # Two-Step Training Strategy of SPP-SCL (PyTorch Style)

```

Require: Input image-text pair  $(I, T)$ , sentiment label  $Y$ ,
model  $f(\cdot)$ , batch size  $B$ , threshold  $\alpha$ 
1: for each batch  $(I, T, Y)$  from dataloader do
2:   # # Step 1: Intra-modal alignment
   (SPP-SCL1)
3:    $M_p, I_p, T_p \leftarrow f(I, T)$ 
4:    $\mathcal{L} = \lambda \cdot \mathcal{L}_{ce}(M_p, Y) + \mathcal{L}_{cl_i}(I_p, Y) + \mathcal{L}_{cl_t}(T_p, Y)$ 
5:   Backpropagate and update: optimizer.step()
6:   # # Step 2: Inter-modal check
   (SPP-SCL2)
7:    $I'_p, T'_p \leftarrow f(I, T)$  # # extract embeddings
   again
8:   Compute similarity matrix:  $\mathbf{S} = \text{sim}(I'_p, T'_p)$ 
9:   Compute average diagonal:  $d = \text{mean}(\text{diag}(\mathbf{S}))$ 
10:  Count:  $c = \text{sum}(\mathbf{S} < d)$ 
11:   $c_{\text{mask}} = \text{sum}(\text{InterMask})$ 
12:  if  $c < \alpha \cdot c_{\text{mask}}$  then
13:     $\mathcal{L}_{cl_m} = \mathcal{L}_{cl_m}(I'_p, T'_p, Y)$ 
14:    Backpropagate and update:
      optimizer.step()
15:  end if
16: end for

```

Hyperparameters	Dataset		
	MVSA-S	MVSA-M	HFM
$\lambda$	5	5	10
Initial learning rate	1e-4	1e-4	5e-5
Learning rate decay/multiple	10/0.5	15/0.5	5/0.5
Weight decay	1e-6	1e-6	1e-6
DropOut	0.2	0.2	0.2
Batch size	64	128	256
Epochs	150	150	150

Table 2: Statistics for all hyperparameters.

and text encoders. Training is performed using the AdamW optimizer (Kingma and Ba 2014) with StepLR scheduling. Key hyperparameters include temperature  $\tau = 0.07$ , threshold  $\alpha = 2/3$ , feature dimension  $D = 32$ , max text length  $N = 200$ , and image size  $224 \times 224$ . Additional settings are listed in Table 2. To make a fair comparison, we follow the existing evaluation. For the MVSA-S and MVSA-M datasets, the Accuracy (Acc) and Weighted F1 value are used for evaluation, and for the HFM dataset, the Accuracy (Acc) and Macro F1 are used for evaluation.

### Comparison Methods

- **Uni-modal baselines:** For text, we include CNN (Kim 2014), Bi-LSTM (Zhou et al. 2016), and BERT (Devlin et al. 2018); for images, ResNet (He et al. 2016) and ViT (Dosovitskiy et al. 2020). OSDA (Yang et al. 2021a), a multi-view visual classifier, was designed by them.
- **Multi-modal baselines:** On MVSA-S and MVSA-M, we compare with MultiSentiNet (Xu and Mao 2017), HSAN (Xu 2017), Co-MN-Hop6 (Xu, Mao, and Chen 2018), MGNNS (Yang et al. 2021b), CLMLF (Li et al. 2022b), ITIN (Zhu et al. 2023a), MULSER (Zhu et al. 2023b), CLMLF<sup>1</sup>, MVCN (Wei et al. 2023), CTMWA (Zhang

modal	Model	Publication	MVSA-S		MVSA-M		Model	Publication	HFM	
			ACC	F1	ACC	F1			ACC	F1
Text	CNN	EMNLP2014	0.6819	0.5590	0.6564	0.5766	CNN	EMNLP2014	0.8003	0.7532
	Bi-LSTM	ACL2016	0.7012	0.6506	0.6790	0.6790	Bi-LSTM	ACL2016	0.8190	0.7753
	BERT	NAACL2018	0.7111	0.6970	0.6759	0.6624	BERT	NAACL2018	0.8389	0.8326
Image	ResNet-50	CVPR2016	0.6467	0.6155	0.6188	0.6098	ResNet-50	CVPR2016	0.7277	0.7138
	OSDA	TMM2021	0.6675	0.6651	0.6662	0.6623	OSDA	TMM2021	-	-
	ViT	ICLR2021	0.6378	0.6226	0.6194	0.6119	ViT	ICLR2021	0.7309	0.7152
Multi-modal	MultiSentiNet	CIKM2017	0.6984	0.6984	0.6886	0.6811	Concat(2)	ACM MM2016	0.8103	0.7799
	HSAN	ISI2017	0.6988	0.6690	0.6796	0.6776	Concat(3)	ACM MM2016	0.8174	0.7874
	Co-MN-Hop6	SIGIR2018	0.7051	0.7001	0.6892	0.6883	MMSD	ACL2019	0.8344	0.8018
	MGNNs	ACL2021	0.7377	0.7270	0.7249	0.6934	D&R Net	ACL2020	0.8402	0.8060
	CLMLF	NAACL2022	0.7533	0.7346	0.7200	0.6983	CLMLF	NAACL2022	0.8543	0.8487
	ITIN	TMM2023	0.7519	0.7497	0.7352	0.7349	HKE	EMNLP2022	0.8702	-
	MULSER	TMM2023	0.7560	0.7534	0.7375	0.7371	DIP	CVPR2023	0.8820	0.8767
	CLMLF <sup>1</sup>	ACL2023	0.7378	0.7291	0.7112	0.6863	CLMLF <sup>1</sup>	ACL2023	0.8489	0.8446
	MVCN	ACL2023	0.7606	0.7455	0.7207	0.7001	MVCN	ACL2023	0.8568	0.8523
	CTMWA	TMM2024	0.7196	0.7143	0.7256	0.7157	MMGCL	TAC2024	0.8857	0.8870
	MSFN	TOMM2025	0.7898	0.7848	0.7475	0.7262	ESAM	TMM2025	0.9011	0.8819
	<b>SPP-SCL</b>	<b>2025</b>	<b>0.8133</b>	<b>0.8015</b>	<b>0.7871</b>	<b>0.7753</b>	<b>SPP-SCL</b>	<b>2025</b>	<b>0.9469</b>	<b>0.9450</b>

Table 3: Quantitative results of the proposed SPP-SCL and these state-of-the-art methods on MVSA-S, MVSA-M, and HFM datasets. The best results are highlighted in **Bold** in each column.

et al. 2024), and MSFN (Zhang et al. 2025). On HFM, we evaluate against Concat(2), Concat(3) (Schifanella et al. 2016), MMSD (Cai, Cai, and Wan 2019), D&R Net (Xu, Zeng, and Mao 2020), CLMLF (Li et al. 2022b), HKE (Liu, Wang, and Li 2022), DIP (Wen, Jia, and Yang 2023), CLMLF<sup>1</sup>, MVCN (Wei et al. 2023), MMGCL (Liang et al. 2024), and ESAM (Yuan et al. 2025).

## Quantitative Comparison

The quantitative results in Table 3 show that SPP-SCL consistently outperforms all baseline methods across the three datasets. Overall, multi-modal methods perform better than uni-modal ones, with text-based models outperforming image-based models—likely due to the sparse and noisy nature of visual sentiment cues, which limits the effectiveness of image-only classifiers. In a word, SPP-SCL first aligns intra-modal sentiment features independently and then balances inter-modal distributions, leading to more robust sentiment representations and better overall performance.

## Ablation Study

**Effectiveness of the two-step strategy** We evaluate a one-step variant of our model—where all contrastive losses are applied simultaneously—to assess the benefit of the proposed two-step training strategy. Table 4 shows that the one-

Method	MVSA-S		MVSA-M		HFM	
	ACC	F1	ACC	F1	ACC	F1
<b>SPP-SCL (Two-step)</b>	<b>0.8133</b>	<b>0.8015</b>	<b>0.7871</b>	<b>0.7753</b>	<b>0.9469</b>	<b>0.9450</b>
SPP-SCL (One-step)	0.7956	0.7684	0.7276	0.6929	0.8800	0.8765

Table 4: Ablation results of two-step strategy.

step strategy performs worse than the two-step strategy, especially on the HFM dataset. This demonstrates that our approach better balances intra- and inter-modal sentiment embedding distances, leading to improved performance.

**Effectiveness of semi-push-pull supervised contrastive learning** We remove  $\mathcal{L}_{cl_i}$ ,  $\mathcal{L}_{cl_t}$  and  $\mathcal{L}_{cl_m}$  respectively

from the entire SPP-SCL. Table 5 shows that removing ei-

Method	MVSA-S		MVSA-M		HFM	
	ACC	F1	ACC	F1	ACC	F1
<b>SPP-SCL</b>	<b>0.8133</b>	<b>0.8015</b>	<b>0.7871</b>	<b>0.7753</b>	<b>0.9469</b>	<b>0.9450</b>
w/o $\mathcal{L}_{cl_m}$	0.7289	0.7040	0.7688	0.7500	0.7812	0.7753
w/o $\mathcal{L}_{cl_i}, \mathcal{L}_{cl_t}$	0.7244	0.7170	0.7459	0.7411	0.9427	0.9409

Table 5:  $\mathcal{L}_{cl_i}$ ,  $\mathcal{L}_{cl_t}$  and  $\mathcal{L}_{cl_m}$  ablation results. “w/o  $\mathcal{L}_{cl_i}, \mathcal{L}_{cl_t}$ ” indicates the removal of  $\mathcal{L}_{cl_i}, \mathcal{L}_{cl_t}$ , while “w/o  $\mathcal{L}_{cl_m}$ ” indicates the removal of  $\mathcal{L}_{cl_m}$ .

ther component degrades SPP-SCL’s performance, confirming their importance. On MVSA-S and MVSA-M, intra-modal losses are more critical than inter-modal losses. However, on HFM, the model still performs well without intra-modal loss (F1=0.9409), indicating a larger gap between text and visual features and highlighting the need for  $\mathcal{L}_{cl_m}$  to balance them.

To further verify the effectiveness of the two components, t-SNE (Van der Maaten and Hinton 2008) visualizes fusion features before classification in Fig. 3. The full SPP-SCL shows well-separated clusters with minimal overlap, while removing  $\mathcal{L}_{cl_m}$  or  $\mathcal{L}_{cl_i}, \mathcal{L}_{cl_t}$  results in larger cluster overlaps. This confirms the components’ role in balancing intra- and inter-modal sentiment consistency. Fig. 4 visualizes the distance distribution, showing that average distances (IID, TTD, ITD) are nearly equal and same-category samples from both modalities cluster together. This indicates the sentiment alignment already has notable discriminative power before cross-modal fusion.

**Effectiveness of hierarchical attention module** To validate the HA module, we conduct ablation experiments extracting [CLS] embeddings from the last 1 to 4 BERT layers and word embeddings from hidden layers. Results are shown in Table 6. The results show performance varies when using features from BERT’s last four layers across datasets. Generally, fusing all four layers in the HA module gives the best results. MVSA-S, MVSA-M, and HFM are most sensitive to fusing the last two or three layers, but still perform

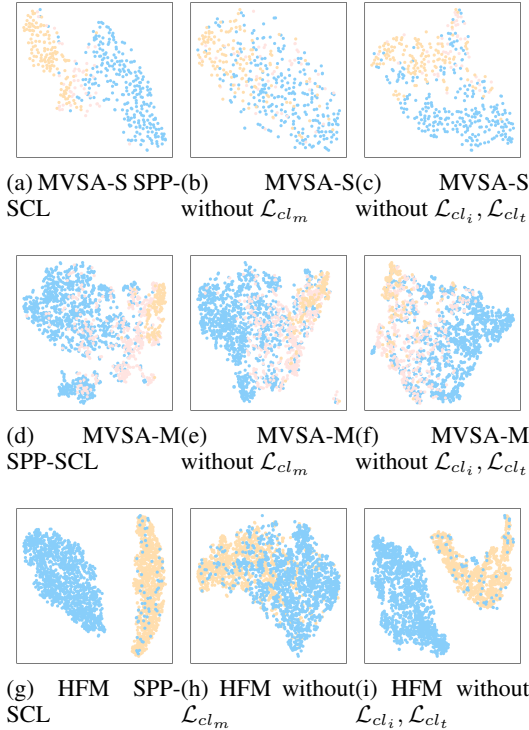


Figure 3: Visualization of the fusion feature distribution on the three datasets.

well. This suggests dataset differences, possibly related to size—MVSA-M and HFM are larger. Overall, using the last four layers yields the best performance.

Method	MVSA-S		MVSA-M		HFM	
	ACC	F1	ACC	F1	ACC	F1
w: last three layers	0.7667	0.7565	0.7612	0.7275	0.9402	0.9385
w: last two layers	<b>0.8222</b>	<b>0.8181</b>	0.5918	0.5847	0.9311	0.9293
w: last one layers	0.7133	0.7025	0.7282	0.6840	0.6617	0.6512
<b>SPP-SCL</b>	<b>0.8133</b>	<b>0.8015</b>	<b>0.7871</b>	<b>0.7753</b>	<b>0.9469</b>	<b>0.9450</b>

Table 6: The ablation experiment results of the HA module. “w: last three Layers, w: last two Layers and w: last one Layer” respectively indicate the ablation results from the last three layers, last two layers, and last one layer of BERT.

**Effectiveness of cross-modal fusion module** To validate the CMF module, we remove it from SPP-SCL and replace fusion with simple concatenation of  $I_p$  and  $T_p$ . Results are in Table 7. The results show a performance drop without the CMF module, proving its effectiveness. Simple concatenation fails to capture semantic interactions for effective fusion. Despite this, “w/o CMF” still outperforms image-only models, confirming the advantage of multi-modal methods.

**Sensitivity of hyperparameters** We observe the best performance at  $\alpha = 2/3$ , as shown in Fig. 5. A lower value (e.g.,  $1/3$ ) triggers inter-modal alignment too early, possibly disrupting already consistent intra-modal structures. In

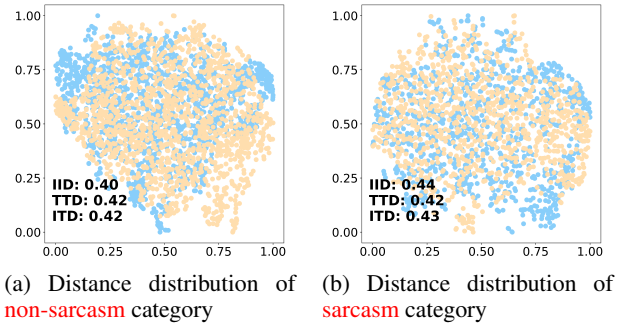


Figure 4: Visualization of the sentiment distance distribution.

Method	MVSA-S		MVSA-M		Method	HFM	
	ACC	F1	ACC	F1		ACC	F1
ResNet-50	0.6467	0.6155	0.6188	0.6098	ResNet-50	0.7277	0.7138
ViT	0.6378	0.6226	0.6194	0.6119	ViT	0.7309	0.7152
w/o CMF	0.7289	0.7482	0.7453	0.7389	w/o CMF	0.8117	0.8260
<b>SPP-SCL</b>	<b>0.8133</b>	<b>0.8015</b>	<b>0.7871</b>	<b>0.7753</b>	<b>SPP-SCL</b>	<b>0.9469</b>	<b>0.9450</b>

Table 7: The ablation experiment results of the CMF module. “w/o CMF” indicates the removal of the CMF module.

contrast, a higher value (e.g., 1) may skip necessary inter-modal alignment, leading to suboptimal sentiment fusion. Thus,  $\alpha = 2/3$  achieves the best trade-off between sensitivity and stability.

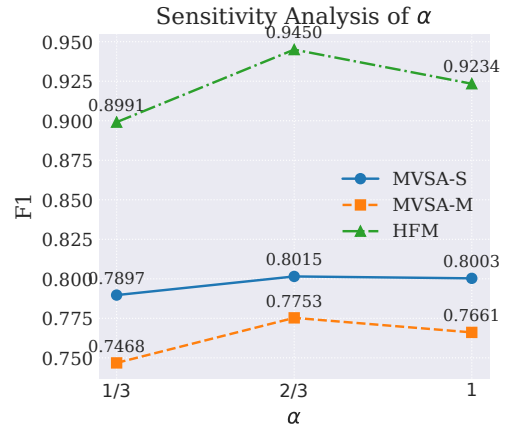


Figure 5: Sensitivity of hyperparameter  $\alpha$ .

## Conclusion

We propose SPP-SCL, a model that balances visual and textual modalities before fusion for image-text sentiment analysis using a novel two-step supervised contrastive learning strategy, combining intra- and inter-modal learning. Without data augmentation, this approach achieves a balanced sentiment embedding space. We also introduce a hierarchical attention module and a cross-modal fusion module to enhance feature extraction and fusion. Experiments on three datasets validate our method’s effectiveness. Future work will explore semi-supervised image-text sentiment analysis.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No.62541601, No.62306010), Natural Science Foundation of Anhui Province under Grant No. 2408085MF169.

## References

- Cai, Y.; Cai, H.; and Wan, X. 2019. Multi-modal sarcasm detection in twitter with hierarchical fusion model. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, 2506–2515.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Das, R.; and Singh, T. D. 2023. Multimodal sentiment analysis: a survey of methods, trends, and challenges. *ACM Computing Surveys*, 55(13s): 1–38.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Gao, P.; Jiang, Z.; You, H.; Lu, P.; Hoi, S. C. H.; Wang, X.; and Li, H. 2019. Dynamic Fusion With Intra- and Inter-Modality Attention Flow for Visual Question Answering. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6632–6641.
- Gao, T.; Yao, X.; and Chen, D. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Gunel, B.; Du, J.; Conneau, A.; and Stoyanov, V. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. *arXiv preprint arXiv:2011.01403*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, L.; Wang, Z.; Wang, L.; and Li, F. 2023. Multi-modal Mutual Attention-Based Sentiment Analysis Framework Adapted to Complicated Contexts. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(12): 7131–7143.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Jawahar, G.; Sagot, B.; and Seddah, D. 2019. What does BERT learn about the structure of language? In *ACL 2019-57th Annual Meeting of the Association for Computational Linguistics*.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33: 18661–18673.
- Kim, Y. 2014. Convolutional Neural Networks for Sentence Classification. In Moschitti, A.; Pang, B.; and Daelemans, W., eds., *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, 1746–1751. ACL.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022a. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34: 9694–9705.
- Li, Z.; Xu, B.; Zhu, C.; and Zhao, T. 2022b. CLMLF: A Contrastive Learning and Multi-Layer Fusion Method for Multimodal Sentiment Detection. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 2282–2294.
- Liang, B.; Gui, L.; He, Y.; Cambria, E.; and Xu, R. 2024. Fusion and Discrimination: A Multimodal Graph Contrastive Learning Framework for Multimodal Sarcasm Detection. *IEEE Transactions on Affective Computing*, 1–15.
- Liang, B.; Lou, C.; Li, X.; Yang, M.; Gui, L.; He, Y.; Pei, W.; and Xu, R. 2022. Multi-modal sarcasm detection via cross-modal graph convolutional network. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 1767–1777. Association for Computational Linguistics.
- Liu, H.; Wang, W.; and Li, H. 2022. Towards multi-modal sarcasm detection via hierarchical congruity modeling with knowledge enhancement. *arXiv preprint arXiv:2210.03501*.
- Mai, S.; Zeng, Y.; Zheng, S.; and Hu, H. 2022. Hybrid contrastive learning of tri-modal representation for multimodal sentiment analysis. *IEEE Transactions on Affective Computing*, 14(3): 2276–2289.
- Niu, T.; Zhu, S.; Pang, L.; and El Saddik, A. 2016. Sentiment analysis on multi-view social data. In *MultiMedia Modeling: 22nd International Conference, MMM 2016, Miami, FL, USA, January 4-6, 2016, Proceedings, Part II 22*, 15–27. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Schifanella, R.; De Juan, P.; Tetreault, J.; and Cao, L. 2016. Detecting sarcasm in multimodal social platforms. In *Proceedings of the 24th ACM international conference on Multimedia*, 1136–1145.

- Sun, C.; Qiu, X.; Xu, Y.; and Huang, X. 2019. How to fine-tune bert for text classification? In *Chinese Computational Linguistics: 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings 18*, 194–206. Springer.
- Tang, J.; Liu, D.; Jin, X.; Peng, Y.; Zhao, Q.; Ding, Y.; and Kong, W. 2022. BAFN: Bi-direction attention based fusion network for multimodal sentiment analysis. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(4): 1966–1978.
- Tang, J.; Liu, D.; Jin, X.; Peng, Y.; Zhao, Q.; Ding, Y.; and Kong, W. 2023. BAFN: Bi-Direction Attention Based Fusion Network for Multimodal Sentiment Analysis. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(4): 1966–1978.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wei, Y.; Yuan, S.; Yang, R.; Shen, L.; Li, Z.; Wang, L.; and Chen, M. 2023. Tackling Modality Heterogeneity with Multi-View Calibration Network for Multimodal Sentiment Detection. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5240–5252.
- Wen, C.; Jia, G.; and Yang, J. 2023. Dip: Dual incongruity perceiving network for sarcasm detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2540–2550.
- Xie, Z.; Yang, Y.; Wang, J.; Liu, X.; and Li, X. 2024. Trustworthy Multimodal Fusion for Sentiment Analysis in Ordinal Sentiment Space. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(8): 7657–7670.
- Xu, N. 2017. Analyzing multimodal public sentiment based on hierarchical semantic attentional network. In *2017 IEEE international conference on intelligence and security informatics (ISI)*, 152–154. IEEE.
- Xu, N.; and Mao, W. 2017. Multisentinet: A deep semantic network for multimodal sentiment analysis. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2399–2402.
- Xu, N.; Mao, W.; and Chen, G. 2018. A co-memory network for multimodal sentiment analysis. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, 929–932.
- Xu, N.; Zeng, Z.; and Mao, W. 2020. Reasoning with multimodal sarcastic tweets via modeling cross-modality contrast and semantic association. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, 3777–3786.
- Xue, X.; Zhang, C.; Niu, Z.; and Wu, X. 2022. Multi-level attention map network for multimodal sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering*, 35(5): 5105–5118.
- Yang, J.; Bisk, Y.; and Gao, J. 2021. Taco: Token-aware cascade contrastive learning for video-text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11562–11572.
- Yang, X.; Feng, S.; Wang, D.; Hong, P.; and Poria, S. 2023. Multiple Contrastive Learning for Multimodal Sentiment Analysis. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Yang, X.; Feng, S.; Wang, D.; and Zhang, Y. 2021a. Image-Text Multimodal Emotion Classification via Multi-View Attentional Network. *IEEE Transactions on Multimedia*, 23: 4014–4026.
- Yang, X.; Feng, S.; Zhang, Y.; and Wang, D. 2021b. Multimodal sentiment detection based on multi-channel graph neural networks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 328–339.
- Yu, J.; and Jiang, J. 2019. Adapting bert for target-oriented multimodal sentiment classification. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 5408–5414.
- Yuan, S.; Wei, Y.; Zhou, H.; Xu, Q.; Chen, M.; and He, X. 2025. Enhancing Semantic Awareness by Sentimental Constraint with Automatic Outlier Masking for Multimodal Sarcasm Detection. *IEEE Transactions on Multimedia*.
- Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; and Morency, L.-P. 2017. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*.
- Zhang, B.; Yuan, Z.; Xu, H.; and Gao, K. 2024. Crossmodal Translation Based Meta Weight Adaption for Robust Image-Text Sentiment Analysis. *IEEE Transactions on Multimedia*, 26: 9949–9961.
- Zhang, S.; Liu, J.; Jiao, Y.; Zhang, Y.; Chen, L.; and Li, K. 2025. A Multimodal Semantic Fusion Network with Cross-Modal Alignment for Multimodal Sentiment Analysis. *ACM Transactions on Multimedia Computing, Communications and Applications*.
- Zhou, P.; Shi, W.; Tian, J.; Qi, Z.; Li, B.; Hao, H.; and Xu, B. 2016. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 2: Short papers)*, 207–212.
- Zhu, T.; Li, L.; Yang, J.; Zhao, S.; Liu, H.; and Qian, J. 2023a. Multimodal Sentiment Analysis With Image-Text Interaction Network. *IEEE Transactions on Multimedia*, 25: 3375–3385.
- Zhu, T.; Li, L.; Yang, J.; Zhao, S.; and Xiao, X. 2023b. Multimodal Emotion Classification With Multi-Level Semantic Reasoning Network. *IEEE Transactions on Multimedia*, 25: 6868–6880.