

# TiCAL: Typicality-Based Consistency-Aware Learning for Multimodal Emotion Recognition

Wen Yin<sup>1</sup>, Siyu Zhan<sup>1</sup>, Cencen Liu<sup>1</sup>, Xin Hu<sup>1</sup>, Guiduo Duan<sup>1,2</sup>, Xiurui Xie<sup>1</sup>,  
Yuan-Fang Li<sup>3</sup>, Tao He<sup>1,2,\*</sup>

<sup>1</sup>The Laboratory of Intelligent Collaborative Computing of University of Electronic Science and Technology of China

<sup>2</sup>Ubiquitous Intelligence and Trusted Services Key Laboratory of Sichuan Province

<sup>3</sup>Faculty of Information Technology, Monash University

{yinwen1999,zhansy,202411900402,202411900415,guiduo.duan,xiexiurui}@uestc.edu.cn

yuanfang.li@monash.edu, tao.he01@hotmail.com

## Abstract

Multimodal Emotion Recognition (MER) aims to accurately identify human emotional states by integrating heterogeneous modalities such as visual, auditory, and textual data. Existing approaches predominantly rely on unified emotion labels to supervise model training, often overlooking a critical challenge: *inter-modal emotion conflicts*, wherein different modalities within the same sample may express divergent emotional tendencies. In this work, we address this overlooked issue by proposing a novel framework, Typicality-based Consistent-aware Multimodal Emotion Recognition (TiCAL), inspired by the stage-wise nature of human emotion perception. TiCAL dynamically assesses the consistency of each training sample by leveraging pseudo unimodal emotion labels alongside a typicality estimation. To further enhance emotion representation, we embed features in a hyperbolic space, enabling the capture of fine-grained distinctions among emotional categories. By incorporating consistency estimates into the learning process, our method improves model performance, particularly on samples exhibiting high modality inconsistency. Extensive experiments on benchmark datasets, e.g. MOSEI and MER2023, validate the effectiveness of TiCAL in mitigating inter-modal emotional conflicts and enhancing overall recognition accuracy, e.g., with about 2.6% improvements over the state-of-the-art DMD.

## 1 Introduction

Multimodal Emotion Recognition (MER) aims to infer human emotional states by integrating complementary information from multiple modalities such as visual expressions, acoustic signals, and textual content (Tsai et al. 2019; Liang et al. 2021; He et al. 2022; Wen et al. 2025). Compared to unimodal approaches (Abbaschian, Sierra-Sosa, and Elmaghraby 2021; Matsumoto et al. 2008), which rely on a single source, MER benefits from cross-modal cues to achieve a strong understanding of affective behaviors.

However, most existing MER methods assume that all modalities within a sample convey consistent emotional signals, supervising them with a single unified label (Yang et al. 2024a; Li, Wang, and Cui 2023; Sun et al. 2024). This assumption is often invalid, as modalities can express con-

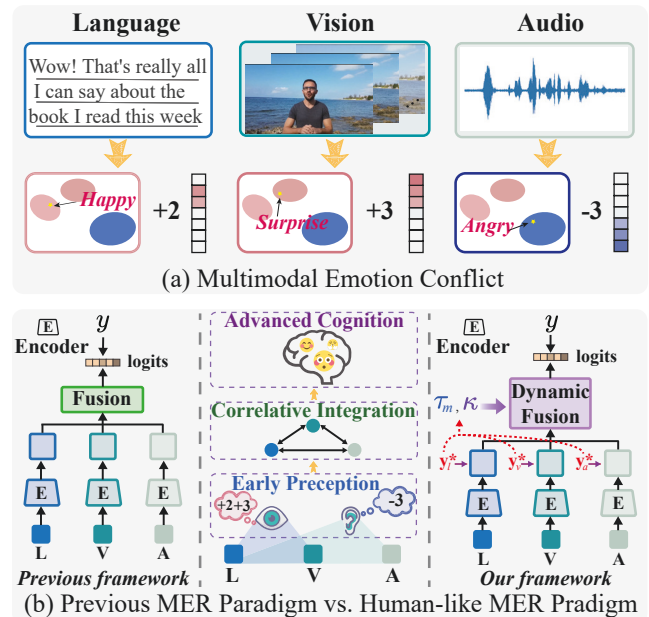


Figure 1: (a) Example of inter-modal emotional conflict where modalities show divergent affective cues. (b) Comparison of our paradigm with prior work: we model modality-specific emotional tendencies and enforce consistency-aware dynamic fusion, inspired by human perception.  $\tau_m$  and  $\kappa$  represent typicality and consistency calculated using our method, respectively.

flicting emotional tendencies—a challenge we refer to as *inter-modal emotional conflict*. For example, Fig. 1(a) shows a MOSEI sample (Zadeh et al. 2018) labeled as “happy”, where the linguistic modality exhibits strong positive sentiment (polarity +2), while the acoustic modality reflects anger-like tendency (polarity -3). Such divergence arises from overlapping acoustic traits, e.g., high vocal intensity—shared by various emotions like happiness and anger.

Treating such multimodal inputs as if they share a single emotion label (e.g., “happy”) introduces ambiguity, leading to suboptimal feature learning and poor generalization. Although recent MER studies have improved fusion mecha-

\*Corresponding author.

nisms through advanced encoders (Cheng et al. 2024; Tong et al. 2022; Zhang, Li, and Bing 2023) or contrastive objectives (Li, Wang, and Cui 2023; Yang et al. 2023; Sun et al. 2024), they rarely address the fundamental cause of performance degradation—the absence of explicit modeling for inter-modal consistency and conflict. While Xia et al. (2025) attempts to address this by optimizing single-modal predictions, their approach does not capture inter-modal consistency, which is critical for effective fusion.

In contrast, human emotion perception is inherently stage-wise and consistency-aware. Neuroscience studies suggest that we first process unimodal cues independently (Early Perception (Choo et al. 2025)), then integrate correlated signals (Correlative Integration (Choi, Lee, and Lee 2018)), and finally reconcile conflicting information through high-level reasoning (Advanced Cognition (Lange, Heerdink, and Van Kleef 2022)). Moreover, recent studies (Wei and Hu 2024; Zhang, Latham, and Saxe 2023; Li et al. 2023) show that the degree of inter-modal consistency directly impacts the optimization of fusion networks—samples with conflicting signals require deeper and more adaptive integration. This raises a research question: *Can we explicitly quantify and leverage inter-modal consistency to achieve human-like, conflict-aware multimodal fusion for emotion recognition?*

To answer this, we propose **TiCAL** (Typicality-based Consistency-aware Learning), a novel framework for multimodal emotion recognition (MER). TiCAL begins by generating intermediate pseudo-labels for each modality using high-confidence anchor samples list (HASL), which capture unimodal emotional tendencies. These pseudo-labels are then embedded into a hyperbolic space, providing a richer structure to model hierarchical and nuanced emotional cues. To ensure reliable cross-modal consistency, we introduce a typicality metric—inspired by the notion of typicality (Liu et al. 2024b)—to assess pseudo-label confidence and selectively emphasize trustworthy instances. Finally, we design a stage-wise perception framework that leverages both typicality and consistency, mimicking the human process of forming modality-specific impressions before integrating them into a robust multimodal understanding.

Experimental results on benchmark datasets, including CMU-MOSI, CMU-MOSEI, and MER2023, demonstrate that TiCAL provides an interpretable and robust solution for resolving inter-modal emotional conflicts in MER. Specifically, on CMU-MOSI, our method achieves state-of-the-art results with  $\sim 2.6\%$  improvement over strong baselines such as DMD (Li, Wang, and Cui 2023).

In summary, our contributions are as follows:

- **Human-like Stage-wise Framework:** We propose **TiCAL** (Typicality-based Consistency-aware Learning), a novel framework that performs dynamic multi-stage fusion by leveraging inter-modal consistency and unimodal typicality, mimicking human-like emotion perception.
- **Reliable Consistency Estimation:** TiCAL captures unimodal emotional tendencies through pseudo-labels and introduces a typicality metric to assess their reliability, enabling robust and accurate consistency estimation across modalities.

- **State-of-the-art Performance:** We conduct extensive experiments on publicly available MER and MSA datasets, achieving superior performance across various SOTA models. Our results demonstrate the feasibility of emulating the human emotion perception mechanism and validate the effectiveness of multi-stage perception guided by consistent dynamics.

## 2 Related Work

**Multimodal Emotion Recognition.** Multimodal Emotion Recognition (MER) seeks to infer human emotional states by integrating signals from multiple modalities and has been widely used in image retrieval (He et al. 2021b) and visual reasoning (Hu et al. 2025; Yang et al. 2024c). Recent work has focused on developing advanced fusion strategies and interaction mechanisms to improve representation learning and model performance (Yang et al. 2023; Sun et al. 2024; Ou, de Bruijn, and Schulz 2025; Li, Wang, and Cui 2023; Yang et al. 2022; Lv et al. 2021). For instance, Yang et al. (2023) introduced a unified contrastive learning framework that combines intra-sample modality decomposition with inter-sample supervised contrastive objectives. Other methods, such as Sun et al. (2024) and Li, Wang, and Cui (2023), aim to disentangle modality-specific and modality-invariant features to facilitate more robust multimodal representations. Although a recent study (Xia et al. 2025) optimized multimodal heterogeneous information by minimizing unimodal losses, it still failed to effectively quantify and leverage inter-modal consistency, leading to suboptimal modality fusion and impaired branch optimization.

**Imbalanced Multimodal Learning.** Imbalanced Multimodal Learning addresses the challenge of integrating diverse modalities that offer complementary yet asymmetrically signals. A common issue in this setting is the model’s tendency to overfit or over-rely on dominant modalities. To mitigate this, prior work has proposed various strategies to balance the learning process across modalities (Zhang, Latham, and Saxe 2023; Dai et al. 2025; Zhou et al. 2024; Fan et al. 2023; Li et al. 2023, 2024; He et al. 2021a; Li et al. 2025). For example, Hu et al. (2024) and Li et al. (2023) identified gradient conflicts in multitask-like training and proposed adaptive gradient modulation techniques to resolve them. Zhang et al. (2024) approached the problem by introducing alternating single-modal optimization, enabling independent updates for each modality. While effective, these methods primarily operate at the gradient level, focusing on adjusting optimization dynamics without explicitly modeling unimodal learning objectives. Even recent approaches such as Fan et al. (2023) and Hua et al. (2024), which refine objectives via prototype-based or regularization-based techniques, largely overlook the significance of the inter-modal consistency in guiding the different modal fusion stages.

## 3 Methodology

This section presents our proposed framework, **TiCAL**, as illustrated in Fig. 2. In Sec. 3.1, we generate pseudo-labels for each unimodal input and regularize the corresponding emotion features in a hyperbolic space. This encourages the

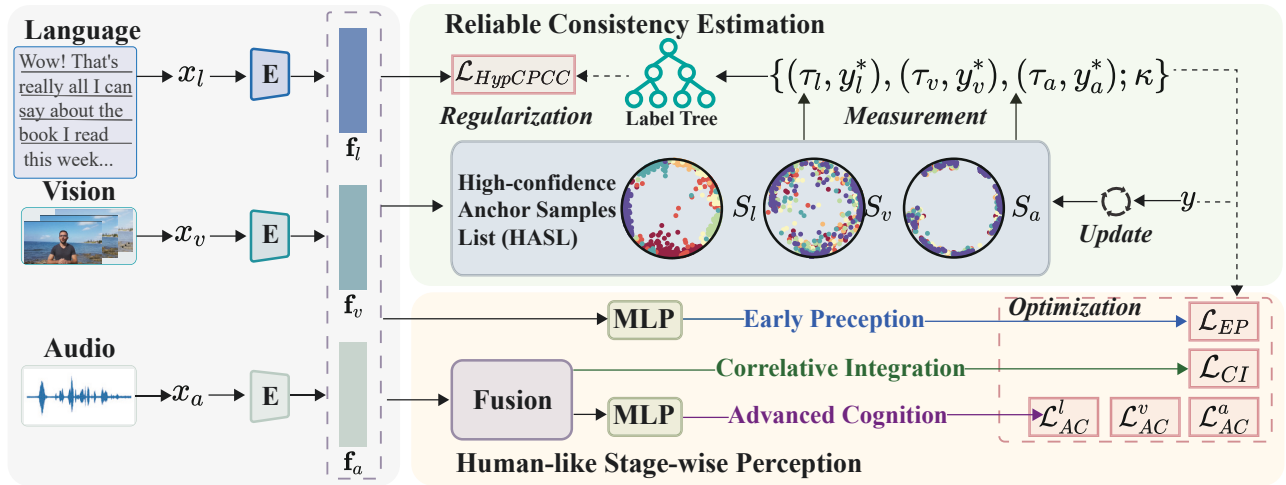


Figure 2: The overview of our proposed TiCAL. TiCAL comprises Reliable Consistency Estimation and Human-like Stage-wise Perception. Specifically, we generate unimodal pseudo labels  $y_m^*$  using a high-confidence anchor samples list (HASL) and regularize them in the hyperbolic space (In § 3.1). Then we estimate the reliable inter-modal consistency  $\kappa$  based on typicality  $\tau_m$  and unimodal pseudo labels  $y_m^*$  (In § 3.2). Inspired by human emotion perception, we design a three-stage prediction structure—*EP*, *CI*, and *AC*—and conduct dynamic typicality-based consistency-aware optimization (In § 3.3).

learning of modality-specific emotional cues while preserving the inherent hierarchical structure of emotional representations. In Sec. 3.2, we introduce a consistency metric to quantify inter-modal consistency for each sample, leveraging the pseudo-labels and a notion of typicality. Finally, in Sec. 3.3, we draw inspiration from human emotional perception to adaptively fuse multimodal features.

**Problem Formulation.** In the MER task, we consider a triplet of input modalities  $\mathbf{x} = [x_m]$ , where  $m \in l, v, a$  denotes the language, visual, and acoustic modalities, respectively. Each input  $x_m$  is processed by a modality-specific encoder to extract a corresponding feature representation  $\mathbf{f}_m$ . The resulting unimodal features are then integrated through a multimodal fusion module, which outputs the final emotion prediction  $y$ .

### 3.1 Unimodal Labeling and Feature Structuring

We consider generating the pseudo unimodal label for modality-specific emotion by anchor samples that the model predicts correctly with high confidence. Subsequently, we incorporate a hierarchically structured label into hyperbolic space to for unimodal features regularization.

**Generating Pseudo Unimodal Labels from HASL** A naïve approach to generating pseudo unimodal labels is to mask two modalities and predict from the remaining one. However, this often yields noisy labels due to the lack of multimodal context. Inspired by Liu et al. (2024b), we instead maintain a high-confidence anchor samples list (HASL) to estimate pseudo labels via feature similarity in hyperbolic space. We define an HASL for each modality as

$$S_m = \{(\mathbf{f}_{m_i}, y)\}, \quad (1)$$

where  $\mathbf{f}_{m_i}$  denotes the feature of the  $i$ -th sample in modality  $m$ . The length  $n$  of HASL is fixed and can be adjusted

as a hyperparameter. During early training (epoch  $\leq \lambda$ ), we populate  $S_m$  with correctly predicted samples ( $\hat{y}_i = y$ ) having confidence  $u_{\hat{y}_i} > \theta$ . After initialization,  $S_m$  is updated using a First-In-First-Out (FIFO) strategy. That is, the high-confidence samples saved earlier during training will be the first to be removed.

After  $\lambda$  epochs, we assign pseudo unimodal labels by measuring feature similarity. Specifically, for a given sample with multimodal feature  $\mathbf{f}_m$ , we compute its distance to all stored features  $\{\mathbf{f}_{m_j}\}_{j=1}^{S_m}$  in the HASL. The label of the closest anchor is treated as the pseudo label:

$$d_m = \min_{\mathbf{f}_{m_j} \in S_m} d(\mathbf{f}_m, \mathbf{f}_{m_j}); y_m^* = y_{min}, \quad (2)$$

where  $y_{min}$  represents the label with the minimum distance from the sample feature in HASL,  $d(\cdot, \cdot)$  is any distance metric in the feature space. We interpret  $y_m^*$  as the unimodal emotional tendency of the sample inferred via similarity in representation space.

**Hyperbolic Regularization for Hierarchical Emotion Features** As discussed, emotional states often share overlapping features—e.g., both happiness and anger may involve heightened vocal intensity—leading to misjudgments of pseudo-labels. Thus, to generate high-quality pseudo labels  $y_m^*$ , we enhance their emotional discriminability through fine-grained unimodal representation learning (Shi and Huang 2023; Hu et al. 2023). Specifically, we embed features in a hyperbolic space using the Poincaré Ball model (Ungar 2001), which effectively captures the hierarchical and fine-grained structure of emotional states (Sarkar 2011).

**Theorem 1 (Poincaré Ball Model)** *The Poincaré Ball Model (Ungar 2001) is a classical formulation of hyperbolic space in non-Euclidean geometry. It represents a  $d$ -dimensional hyperbolic space as an open unit ball:*

$\mathbb{B}^d = \{\mathbf{x} \in \mathbb{R}^d \mid |\mathbf{x}| < 1\}$ , where  $|\cdot|$  denotes the Euclidean norm. Given two points (e.g., feature vectors)  $\mathbf{f}_1, \mathbf{f}_2 \in \mathbb{B}^d$ , their hyperbolic distance is defined as:

$$d_{\mathbb{B}}(\mathbf{f}_1, \mathbf{f}_2) = \operatorname{arcosh} \left( 1 + 2 \frac{|\mathbf{f}_1 - \mathbf{f}_2|^2}{(1 - |\mathbf{f}_1|^2)(1 - |\mathbf{f}_2|^2)} \right). \quad (3)$$

For pseudo label generation from HASL, we adopt the hyperbolic distance (Eq. 3) to compare features and further optimize unimodal representations in hyperbolic space. To incorporate emotion hierarchies, we construct a weighted tree structure over emotion classes, enhancing one-hot labels into hierarchical ones. This structure captures class relationships more effectively. For details on the tree construction, see Appendix 1.1. During training, we apply a hyperbolic regularization that aligns feature embeddings with their corresponding hierarchical labels by minimizing the distance between feature-label node pairs in the tree.

**Definition 1 (Tree Node Pair Distance)** Given a weighted tree  $\mathbb{T} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ , where  $\mathcal{V}$  is the set of nodes,  $\mathcal{E}$  the set of edges, and  $\mathcal{W}$  the corresponding edge weights, the tree distance  $d_{\mathbb{T}}(v, v')$  is defined as the total weight of the shortest path connecting two nodes  $v, v' \in \mathcal{V}$  in  $\mathbb{T}$ .

To ensure that the unimodal features is empowered with the hierarchical structure, we compute tree-based distances  $d_{\mathbb{T}}$  between label pairs in a batch. Simultaneously, we measure the hyperbolic distances  $d_{\mathbb{B}}$  between corresponding feature representations in the Poincaré Ball model. To assess the alignment between the label hierarchy and the features, we employ the Hyperbolic Cophenetic Correlation Coefficient (HypCPCC) (Sokal and Rohlf 1962), which quantifies the correlation between tree-based and hyperbolic distances as:

$$\text{HypCPCC} = \frac{\sum_{i < j} (d_{\mathbb{T}} - \bar{d}_{\mathbb{T}}) (d_{\mathbb{B}} - \bar{d}_{\mathbb{B}})}{\sqrt{\sum_{i < j} (d_{\mathbb{T}} - \bar{d}_{\mathbb{T}})^2 \cdot \sum_{i < j} (d_{\mathbb{B}} - \bar{d}_{\mathbb{B}})^2}}, \quad (4)$$

where  $\bar{d}_{\mathbb{T}}$  and  $\bar{d}_{\mathbb{B}}$  denote the mean distances across all node pairs in the batch for the tree and hyperbolic spaces, respectively. By maximizing the HypCPCC value during training, we encourage the unimodal features to respect the hierarchical relationships encoded by the pseudo labels, thus enhancing their fine-grained emotional discrimination capability.

### 3.2 Reliable Inter-modal Consistency Estimation

In this section, we quantify the reliable inter-modal consistency of each sample based on pseudo unimodal labels and their associated typicality. To assess the reliability of each pseudo label, we compute a typicality score that reflects its alignment with the anchor distribution. We hypothesize that inter-modal consistency captures the degree of discrepancy across modalities, which directly influences prediction quality. Moreover, we conduct an analytical study to validate this hypothesis in Appendix 3.3 & 3.4.

**Assessing Typicality for Pseudo Unimodal Labels** Although pseudo unimodal labels can be derived via Eq. 2, their reliability may vary. To assess this, we adopt typicality (Yuksekgonul et al. 2023; Liu et al. 2024b), which suggests

that typical samples are semantically representative, easy to learn, and generalize well. These samples tend to resemble the majority in structure and behavior. We define typicality based on the relative hyperbolic distance to anchor samples within a batch. Intuitively, samples closer to anchors are more typical and thus more trustworthy.

**Definition 2 (Unimodal typicality)** We define the typicality of unimodal feature  $\mathbf{f}_m$  as <sup>1</sup>

$$\tau_m = \frac{\max D_m - d_m}{\max D_m - \min D_m}, \quad (5)$$

where  $d_m$  is the sample’s minimum distance to HASL, and  $D_m$  is the set of all such distances in the current batch. Here,  $\max$  and  $\min$  refer to the maximum and minimum values in  $D_m$ , respectively.

A high typicality score  $\tau_m$  indicates that the pseudo unimodal label is reliable, while a low  $\tau_m$  suggests the sample is atypical or ambiguous. At this stage, we obtain a set of pseudo-label and typicality pairs for each modality:  $\{y_m^*, \tau_m\}$ .

**Typicality-based Consistency Quantification** In this work, we define inter-modal consistency by jointly considering pseudo labels and typicality, as relying on pseudo labels alone may overlook critical discrepancies. For instance, even if pseudo labels across modalities are aligned, large variations in typicality—where some modalities are highly typical while others are ambiguous—suggest underlying inconsistency in emotional expression. Specifically, we define the calculation of inter-modal consistency as:

**Definition 3 (Inter-modal Consistency)** With the pseudo unimodal labels and its corresponding typicality  $(y_m^*, \tau_m)$ , we define the inter-modal consistency as

$$\kappa = \sqrt{(\tau_l \tau_v \tau_a)^t \cdot e^{-k \cdot d_{label}}}, \quad (6)$$

where  $t$  and  $k$  are temperature and scaling hyperparameters. And  $d_{label}$  is to measure the discrepancy between the estimated labels, defined as

$$d_{label} = \left( \frac{|y_l^* - \mu_{label}| + |y_v^* - \mu_{label}| + |y_a^* - \mu_{label}|}{3} \right)^\rho, \quad (7)$$

where  $\mu_{label}$  is the mean value of estimated labels.  $\rho$  is a density hyperparameter.

### 3.3 Human-Like Stage-wise Perception for MER

Inspired by human-like emotion processing, we decompose our model into three stages—Early Perception (EP), Correlative Integration (CI), and Advanced Cognition

<sup>1</sup>Here typicality differs from “typical distribution samples” in probability theory (Yuksekgonul et al. 2023), which refers to “more valuable samples”, but we focus on the reliability of the estimation labels of different samples in batches.

<sup>2</sup>Different from Liu et al. (2024b), we directly calculate hyperbolic distances based on features  $\mathbf{f}_m$  instead of the mean  $\mu$  and variance  $\sigma^2$ . More analysis can be found in the Appendix 3.1.3.

(AC)—guided by neuroscience insights (see Fig. 2). EP performs unimodal emotion prediction from raw features; CI fuses multimodal information; and AC conducts more in-depth predictions. Each stage is supervised with the class-weighted cross-entropy losses  $\mathcal{L}$  to alleviate the long-tailed problem in datasets. We design lightweight network structures, such as MLP, to ensure high portability. Implementation details are provided in the Appendix 2. Specifically, we integrate the typicality (Eq. 5) and consistency (Eq. 6) metrics into the training objective, dynamically modulating optimization across stages.

**Unbiased Unimodal Weights with  $\tau$ .** Based on the typicality in Eq.(5), we can determine which modality can be dominant. To mitigate the risk of over-relying on any single dominant modality under conflicts or unreliable channels, we introduce unbiased unimodal weights that assign lower relative weights to more typical (and hence dominant) modalities as:

$$\mathcal{L}_{AC} = \varphi(\tau_i)\mathcal{L}_{AC}^l(\hat{y}, y) + \varphi(\tau_v)\mathcal{L}_{AC}^v(\hat{y}, y) + \varphi(\tau_a)\mathcal{L}_{AC}^a(\hat{y}, y), \quad (8)$$

where  $\varphi(\tau) = e^{(1-\tau)}$  is a weight adjustment function that inversely scales with the typicality score.

**Dynamic Optimization with  $\kappa$ .** Prior work (Wei and Hu 2024; Zhang, Latham, and Saxe 2023; Li et al. 2023) has shown that samples with high consistency should be prioritized during early training, while more inconsistent samples benefit from later-stage refinement. Following this principle, we dynamically weight the losses for the EP and AC stages based on the inter-modal consistency score  $\kappa$ . The CI stage is assigned a fixed weight in order to ensure the basic integration of the modal. The overall task loss is defined as:

$$\mathcal{L}_{task} = \kappa\mathcal{L}_{EP}(\hat{y}, y) + \mathcal{L}_{CI}(\hat{y}, y) + (1 - \kappa)\mathcal{L}_{AC}(\hat{y}, y), \quad (9)$$

where  $\hat{y}$  is the predicted label and  $y$  is the ground truth. The EP stage is trained directly on original features, while the CI stage uses a cross-attention for basic integration and prediction (see Fig. 2). High-consistency samples ( $\kappa \rightarrow 1$ ) emphasize learning in the EP stage, while low-consistency samples ( $\kappa \rightarrow 0$ ) shift focus toward the AC stage.

**The Overall Optimization.** In the early stage of training ( $epoch \leq \lambda$ ), we only train the model with the task loss  $\mathcal{L}_{task}$  without dynamic optimization by  $\kappa$  and  $\tau$ , but in the later stage of training ( $epoch > \lambda$ ), we introduce hyperbolic regularization to optimize the original unimodal features as:

$$\mathcal{L}_{all} = \mathcal{L}_{task} - \mathcal{L}_{HypCPCC}(\mathbf{f}_m, \mathbf{y}_m^*). \quad (10)$$

**The Model Inference.** During the model inference stage, we can obtain three stages of emotion predictions ( $\hat{y}_{EP}$ ,  $\hat{y}_{CI}$ , and  $\hat{y}_{AC}$ ) and multimodal consistency  $\kappa$ . The final prediction label of the model is

$$\hat{y}_{final} = \kappa\hat{y}_{EP} + \hat{y}_{CI} + (1 - \kappa)\hat{y}_{AC}. \quad (11)$$

## 4 Experiments

### 4.1 Experimental Settings

**Datasets.** We evaluate TiCAL on four widely used multimodal datasets: CMU-MOSI (Zadeh et al. 2016) and CMU-MOSEI (Zadeh et al. 2018) for the Multimodal Sentiment

Analysis (MSA) task, and DFEW (Jiang et al. 2020) and MER2023 (Lian et al. 2023) for the Multimodal Emotion Recognition (MER) task. CMU-MOSI comprises 2,199 short monologue video clips, while CMU-MOSEI includes 22,856 movie review clips sourced from YouTube. Both DFEW and MER2023 consist of video clips from movies and TV series, each annotated with discrete emotion categories. In MSA, each sample is labelled with a sentiment score ranging from -3 to +3, representing seven sentiment levels. For MER, samples are labelled with six emotion classes.

**Implementation Details.** Our experiments were conducted using the PyTorch (Paszke et al. 2019) on a single NVIDIA RTX 4090 GPU with 24GB memory. We set the initial learning rate, batch size, and decay interval to 1e-4, 16, 0.005, respectively, and trained the model for 30 epochs using the Adam optimizer. For the initialization epoch and the hyperparameters related to consistency, we set:  $\lambda = 5$ ,  $\rho = 4$ ,  $k = 0.5$ , and  $t = 1/5$ . We use Bert-base-uncased (Devlin et al. 2019) to obtain the original text features. We use Facet (Baltrušaitis, Robinson, and Morency 2016) and COVAREP (Degottex et al. 2014) as visual encoder and acoustic encoder respectively. For the MSA task, we use accuracy of 2 class (Acc2), accuracy of 7 class (Acc7) and F1 scores as evaluation metrics, where Acc2 only focuses on the positive and negative for sentiment, while for the MER task, we use Unweighted Average Recall (UAR) and Weighted Average Recall (WAR) scores.

### 4.2 Comparison with State-of-the-Art Methods

We compare our model to three types of models: MSA or MER models, multimodal learning models, and language models (LLMs). The SOTA MSA models include CAGC (Sun et al. 2024), DMD (Li, Wang, and Cui 2023), PMR (Lv et al. 2021), FDMER (Yang et al. 2022), and CONFED (Yang et al. 2023), while the SOTA MER model includes Emotion-LLaMA (Cheng et al. 2024), VideoMAE (Tong et al. 2022), MAE-DFER (Sun et al. 2023), Former-DFER (Zhao and Liu 2021), FBP (Cheng et al. 2023), MER-Baseline (Lian et al. 2023). Specifically, for the MSA task, we compare with the most advanced multimodal learning methods, namely CGGM (Guo et al. 2024), ReconBoost (Hua et al. 2024), AGM (Li et al. 2023), and LFM (Yang et al. 2024b). In addition, we have also selected advanced and well-known LLMs for comparison: GPT-4V (Lian et al. 2024), Video-Llama (Zhang, Li, and Bing 2023), MiniGPT-v2 (Chen et al. 2023), LLaVA-NEXT (Liu et al. 2024a), and Qwen-Audio (Chu et al. 2023).

**Multimodal Sentiment Analysis Results.** Tab. 1 presents the experimental results for the MSA task, where the baseline method contains two types: multimodal learning (ML) and task-specific multimodal sentiment analysis (MSA). The results indicate that TiCAL outperforms all previous models. For the MOSI dataset, our method achieves an average improvement of 2.1%, 2.09%, 1.19% over the SOTA MSA model DMD in terms of Acc-2, F1, and Acc-7, respectively. Our method on MOSEI also outperforms the baseline model, apart from F1. We attribute this improvement

Model	Type	MOSI			MOSEI		
		Acc-2	F1	Acc-7	Acc-2	F1	Acc-7
AGM* (ICCV'23)	ML	78.23	77.56	31.48	72.14	71.46	41.04
ReconBoost* (ICML'24)		77.96	76.89	32.73	68.61	68.54	39.57
LFM* (NIPS'24)		79.41	79.60	34.02	74.88	74.27	42.59
CGGM* (NIPS'24)		82.84	82.74	33.73	79.67	78.97	45.76
PMR (CVPR'21)	MSA	83.60	83.40	40.60	83.30	82.60	52.20
FDMER (MM'22)		84.60	84.70	44.10	86.10	85.80	54.10
ConFEDE (ACL'23)		85.52	85.52	42.27	85.82	85.83	54.86
DMD (CVPR'23)		86.00	86.00	45.60	86.60	86.60	54.50
CAGC (CVPR'24)		85.70	85.60	44.80	-	-	-
TiCAL (Ours)	MSA	<b>88.10</b>	<b>88.09</b>	<b>46.79</b>	<b>87.03</b>	<b>87.05</b>	<b>55.23</b>

Table 1: Experimental results of state-of-the-art multimodal sentiment analysis (MSA) and multimodal learning (ML) models on two benchmark datasets: **MOSI** and **MOSEI**. Models marked with \* are reproduced using the publicly available code.

Method	Type	UAR	WAR
Qwen-Audio	LLM	25.23	31.74
LLaVA-NEXT		25.12	33.75
MiniGPT-v2		35.36	34.47
Video-Llama		26.17	35.24
GPT-4V		26.77	35.75
Former-DFER	MER	53.69	65.70
MAE-DFER		60.99	74.43
VideoMAE		63.41	74.60
Emotion-LLaMA		64.21	77.06
TiCAL (Ours)	MER	<b>65.87</b>	<b>78.29</b>

Table 2: Experimental results of MER on the LLM and MER models on DFEW dataset.

to the dynamic integration of multi-stage fusion in TiCAL, which leverages both the temporal and contextual dependencies more effectively. These results highlight the superiority of our method in multimodal sentiment analysis tasks.

**Multimodal Emotion Recognition Results.** Tab. 2 and 3 present the experimental results on the MER task, following the evaluation protocol of Emotion-LLaMA (Cheng et al. 2024). Our proposed method achieves new SOTA performance, surpassing both large language model (LLM)-based approaches and task-specific baselines across both the DFEW and MER2023 datasets. These empirical results demonstrate that TiCAL effectively captures and integrates multimodal emotional cues, whether for polarity-based sentiment classification or discrete emotion category recognition. This highlights the model’s strong adaptability and its capacity to perceive nuanced emotional signals across modalities.

### 4.3 Ablation Study

**Effectiveness of Key Components.** We perform ablation studies on the MOSI dataset to evaluate the contributions of individual components in our TiCAL framework. The results are presented in Tab. 4. First, removing the typicality-based dynamic optimization (see Sec. 3.3) results in a no-

Method	Modality	F1
MER2023-Baseline	A, V	86.75
MER2023-Baseline	A, V, T	86.40
FBP	A, V, T	88.55
VAT	A, V, T	89.11
Emotion-LLaMA	A, V	89.05
Emotion-LLaMA	A, V, T	90.36
TiCAL (Ours)	A, V, T	<b>91.56</b>

Table 3: The experimental results of MER on the MER2023 dataset in terms of F1 score. Notably, the “Modality” column indicates the input modality configuration, where L, A, and V represent the language, audio, and visual modalities.

Components	Acc-2	F1	Acc-7
TiCAL	<b>88.10</b>	<b>88.09</b>	<b>46.79</b>
Eq. 8 w/o $\tau$	87.18(-0.92)	87.20(-0.89)	44.77(-1.02)
Eq. 9 w/o $\kappa$	86.59(-1.51)	86.55(-1.54)	44.68(-2.11)
Eq. 5 w $d_{\mathbb{R}}$	87.24 (-0.86)	87.18(-0.91)	45.17(-1.62)
Eq. 10 w/o $\mathcal{L}_{Hyp}$	86.03(-2.07)	86.17(-1.92)	44.24(-2.55)

Table 4: Ablation study results of the key components. “w/o  $\tau$ ” and “w/o  $\kappa$ ” represent removing typicality and consistency, respectively.  $\mathcal{L}_{Hyp}$  is *HypCPCC* loss in Sec. 3.1.  $d_{\mathbb{R}}$  is Euclidean distance.

ticeable performance drop, indicating its importance in guiding modality integration. Excluding the consistency module leads to a further decrease in accuracy, demonstrating its critical role in capturing inter-modal alignment. Moreover, replacing the hyperbolic distance used in HASL with Euclidean distance degrades model performance, highlighting the effectiveness of hierarchical feature regularization in hyperbolic space. Finally, delete the unimodal optimization strategy (see Sec. 3.3) also negatively impacts performance, confirming the necessity of Estimation optimization for pseudo labels. More ablation experiments can be found in Appendix 3.1.

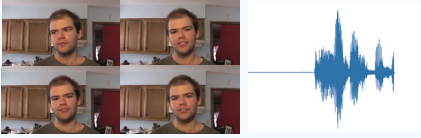
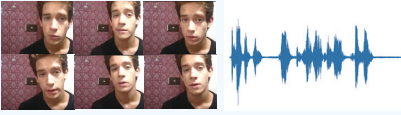
<b>Input</b>	<p>I really liked david finch.</p> 	<p>The soundtrack is good it's really good some really great songs.</p> 
<b>DMD (Li, Wang, and Cui 2023) CAGC (Sun et al. 2024)</b>	<p><b>Positive +3</b> ✗ <b>Positive +1</b> ✓</p>	<p><b>Negative -1</b> ✗ <b>Positive +1</b> ✗</p>
<b>TiCAL</b>	<p><math>\tau_l = 0.91, \tau_v = 0.73, \tau_a = 0.61</math>  <math>y_i^* = +2, y_v^* = +1, y_a^* = +1</math>  <math>\kappa = 0.8173</math>  <b>Positive +1</b> ✓</p>	<p><math>\tau_l = 0.68, \tau_v = 0.21, \tau_a = 0.11</math>  <math>y_i^* = +2, y_v^* = -1, y_a^* = -1</math>  <math>\kappa = 0.4712</math>  <b>Positive +2</b> ✓</p>

Table 5: Two real-case examples from the MOSI test set comparing predictions across different state-of-the-art MSA models and our proposed TiCAL method. Each example includes visualized input modalities to aid interpretability parameters. These examples highlight TiCAL’s ability to make accurate predictions, even when baseline models fail, particularly in the presence of modality conflicts or ambiguous cues.




Figure 3: The performance of our TiCAL method with different confidence threshold  $\theta$  in HASL initialization. The left side is on MOSI and the right side is on MOSEI.

**Impact of Confidence Threshold.** To examine the effect of the threshold  $\theta$  on model performance, we evaluate TiCAL under five different  $\theta$  values and report the results on the MOSI and MOSEI dataset in Fig. 3. The results show that the model achieves the highest accuracy when  $\theta = 0.8$ . Intuitively, a larger  $\theta$  enforces higher typicality among stored samples. However, when  $\theta = 0.9$ , the threshold may become overly restrictive, preventing proper initialization of the HASL module, which in turn leads to performance degradation. More hyperparameter analyses are provided in Appendix 3.2.

#### 4.4 Case Analysis

As illustrated in Table 5, we present a case study by visualizing representative examples of both successful and failed predictions made by different multimodal sentiment analysis (MSA) methods on the MOSI dataset. For each sample, we include the pseudo unimodal labels, their corresponding typicality scores, and the consistency measure computed using our proposed framework. In the first example, where the multimodal signals are relatively consistent, DMD fails to

make a correct prediction, while DMD succeeds. However, in the second example—characterized by significant cross-modal inconsistency—both DMD and CAGC fail to classify the emotion correctly. In contrast, our TiCAL framework successfully predicts both samples, regardless of their consistency level.

These examples highlight TiCAL’s ability to adaptively interpret unimodal cues and leverage typicality-weighted consistency to guide learning. This case study underscores the model’s robustness in handling emotionally conflicting modalities, which is critical for real-world applications of multimodal emotion recognition.

## 5 Conclusion

In this work, we proposed TiCAL (Typicality-based Consistency-Aware Learning), a novel framework inspired by human emotion perception mechanisms, to address the long-overlooked issue of inter-modal emotional conflicts in multimodal emotion recognition (MER). By generating pseudo unimodal emotion labels and embedding them in a hyperbolic space, TiCAL captures modality-specific emotional tendencies with structured semantic representation. Moreover, our typicality-based consistency metric enables dynamic, stage-wise optimization that reflects human-like emotional integration. Extensive experiments across diverse MER and MSA benchmarks confirm the effectiveness of TiCAL, achieving state-of-the-art performance while offering greater interpretability and robustness in handling conflicting modality cues.

Despite these promising results, our work has several **limitations**. For example, TiCAL assumes full modality availability during both training and inference, which may not be practical in real-world applications where certain modalities can be noisy or missing. Future work will explore the extension of TiCAL to handle incomplete modalities, enabling developing more general and robust multimodal emotion recognition models.

## Acknowledgments

This research was partially supported by the National Natural Science Foundation of China (NSFC) (62306064 and U19A2059) and the Sichuan Science and Technology Program (2022ZHCG0008, 2023ZYD0165 and 2024ZDZX0011). We appreciate all the authors for their fruitful discussions. In addition, thanks are extended to anonymous reviewers for their insightful comments.

## References

- Abbaschian, B. J.; Sierra-Sosa, D.; and Elmaghraby, A. 2021. Deep learning techniques for speech emotion recognition, from databases to models. *Sensors*, 21(4): 1249.
- Baltrušaitis, T.; Robinson, P.; and Morency, L.-P. 2016. Openface: an open source facial behavior analysis toolkit. In *WACV*, 1–10. IEEE.
- Chen, J.; Zhu, D.; Shen, X.; Li, X.; Liu, Z.; Zhang, P.; Krishnamoorthi, R.; Chandra, V.; Xiong, Y.; and Elhoseiny, M. 2023. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*.
- Cheng, Z.; Cheng, Z.-Q.; He, J.-Y.; Wang, K.; Lin, Y.; Lian, Z.; Peng, X.; and Hauptmann, A. 2024. Emotion-llama: Multimodal emotion recognition and reasoning with instruction tuning. *NerulIPS*, 37: 110805–110853.
- Cheng, Z.; Lin, Y.; Chen, Z.; Li, X.; Mao, S.; Zhang, F.; Ding, D.; Zhang, B.; and Peng, X. 2023. Semi-supervised multimodal emotion recognition with expression mae. In *ACMMM*, 9436–9440.
- Choi, I.; Lee, J.-Y.; and Lee, S.-H. 2018. Bottom-up and top-down modulation of multisensory integration. *Current Opinion in Neurobiology*, 52: 115–122.
- Choo, C. M.; Bai, S.; Privitera, A. J.; and Chen, S.-H. A. 2025. Brain Imaging Studies of Multisensory Integration in Emotion Perception: A Scoping Review. *Neuroscience & Biobehavioral Reviews*, 106118.
- Chu, Y.; Xu, J.; Zhou, X.; Yang, Q.; Zhang, S.; Yan, Z.; Zhou, C.; and Zhou, J. 2023. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*.
- Dai, R.; Li, C.; Yan, Y.; Mo, L.; Qin, K.; and He, T. 2025. Unbiased Missing-modality Multimodal Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 24507–24517.
- Degottex, G.; Kane, J.; Drugman, T.; Raitio, T.; and Scherer, S. 2014. COVAREP—A collaborative voice analysis repository for speech technologies. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 960–964. IEEE.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NACL*, 4171–4186.
- Fan, Y.; Xu, W.; Wang, H.; Wang, J.; and Guo, S. 2023. Pmr: Prototypical modal rebalance for multimodal learning. In *CVPR*, 20029–20038.
- Guo, Z.; Jin, T.; Chen, J.; and Zhao, Z. 2024. Classifier-guided gradient modulation for enhanced multimodal learning. *NerulIPS*, 37: 133328–133344.
- He, T.; Gao, L.; Song, J.; Cai, J.; and Li, Y.-F. 2021a. Semantic compositional learning for low-shot scene graph generation. *arXiv preprint arXiv:2108.08600*.
- He, T.; Gao, L.; Song, J.; and Li, Y.-F. 2021b. Semisupervised network embedding with differentiable deep quantization. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8): 4791–4802.
- He, T.; Gao, L.; Song, J.; and Li, Y.-F. 2022. Towards open-vocabulary scene graph generation with prompt-based fine-tuning. In *European conference on computer vision*, 56–73. Springer.
- Hu, D.; Bao, Y.; Wei, L.; Zhou, W.; and Hu, S. 2023. Supervised adversarial contrastive learning for emotion recognition in conversations. *arXiv preprint arXiv:2306.01505*.
- Hu, X.; Qin, K.; Duan, G.; Li, M.; Li, Y.-F.; and He, T. 2025. SPADE: Spatial-Aware Denoising Network for Open-vocabulary Panoptic Scene Graph Generation with Long- and Local-range Context Reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15562–15572.
- Hu, Z.; Zhou, S.; Zhao, S.; and Yuan, Z. 2024. MVCTrack: Boosting 3D Point Cloud Tracking via Multimodal-Guided Virtual Cues.
- Hua, C.; Xu, Q.; Bao, S.; Yang, Z.; and Huang, Q. 2024. Reconboost: Boosting can achieve modality reconciliation. *arXiv preprint arXiv:2405.09321*.
- Jiang, X.; Zong, Y.; Zheng, W.; Tang, C.; Xia, W.; Lu, C.; and Liu, J. 2020. Dfaw: A large-scale database for recognizing dynamic facial expressions in the wild. In *ACMMM*, 2881–2889.
- Lange, J.; Heerdink, M. W.; and Van Kleef, G. A. 2022. Reading emotions, reading people: Emotion perception and inferences drawn from perceived emotions. *Current opinion in psychology*, 43: 85–90.
- Li, H.; Li, X.; Hu, P.; Lei, Y.; Li, C.; and Zhou, Y. 2023. Boosting multi-modal model performance with adaptive gradient modulation. In *ICCV*, 22214–22224.
- Li, J.; Qiu, X.; Xu, L.; Guo, L.; Qu, D.; Long, T.; Fan, C.; and Li, M. 2025. UniF2ace: Fine-grained Face Understanding and Generation with Unified Multimodal Models. *arXiv preprint arXiv:2503.08120*.
- Li, M.; Zhou, P.; Liu, J.-W.; Keppo, J.; Lin, M.; Yan, S.; and Xu, X. 2024. Instant3d: instant text-to-3d generation. *IJCV*.
- Li, Y.; Wang, Y.; and Cui, Z. 2023. Decoupled multimodal distilling for emotion recognition. In *CVPR*, 6631–6640.
- Lian, Z.; Sun, H.; Sun, L.; Chen, K.; Xu, M.; Wang, K.; Xu, K.; He, Y.; Li, Y.; Zhao, J.; et al. 2023. Mer 2023: Multi-label learning, modality robustness, and semi-supervised learning. In *ACMMM*, 9610–9614.
- Lian, Z.; Sun, L.; Sun, H.; Chen, K.; Wen, Z.; Gu, H.; Liu, B.; and Tao, J. 2024. Gpt-4v with emotion: A zero-shot benchmark for generalized emotion recognition. *Information Fusion*, 108: 102367.

- Liang, T.; Lin, G.; Feng, L.; Zhang, Y.; and Lv, F. 2021. Attention is not enough: Mitigating the distribution discrepancy in asynchronous multimodal sequence fusion. In *ICCV*, 8148–8156.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024a. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Liu, Y.; Cui, J.; Tian, Z.; Yang, S.; He, Q.; Wang, X.; and Su, J. 2024b. Typicalness-Aware Learning for Failure Detection. *arXiv preprint arXiv:2411.01981*.
- Lv, F.; Chen, X.; Huang, Y.; Duan, L.; and Lin, G. 2021. Progressive modality reinforcement for human multimodal emotion recognition from unaligned multimodal sequences. In *CVPR*, 2554–2562.
- Matsumoto, D.; Keltner, D.; Shiota, M. N.; O’Sullivan, M.; and Frank, M. 2008. Facial expressions of emotion. *Handbook of emotions*, 3: 211–234.
- Ou, Y.; de Bruijn, G.-J.; and Schulz, P. J. 2025. Social Media as an Emotional Barometer: Bidirectional Encoder Representations From Transformers–Long Short-Term Memory Sentiment Analysis on the Evolution of Public Sentiments During Influenza A on Sina Weibo. *Journal of Medical Internet Research*, 27: e68205.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32.
- Sarkar, R. 2011. Low distortion delaunay embedding of trees in hyperbolic plane. In *International symposium on graph drawing*, 355–366. Springer.
- Shi, T.; and Huang, S.-L. 2023. MultiEMO: An attention-based correlation-aware multimodal fusion framework for emotion recognition in conversations. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 14752–14766.
- Sokal, R. R.; and Rohlf, F. J. 1962. The comparison of dendrograms by objective methods. *Taxon*, 33–40.
- Sun, K.; Xie, Z.; Ye, M.; and Zhang, H. 2024. Contextual augmented global contrast for multimodal intent recognition. In *CVPR*, 26963–26973.
- Sun, L.; Lian, Z.; Liu, B.; and Tao, J. 2023. Mae-dfer: Efficient masked autoencoder for self-supervised dynamic facial expression recognition. In *ACMMM*, 6110–6121.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Video-mae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NeurIPS*, 35: 10078–10093.
- Tsai, Y.-H. H.; Bai, S.; Liang, P. P.; Kolter, J. Z.; Morency, L.-P.; and Salakhutdinov, R. 2019. Multimodal transformer for unaligned multimodal language sequences. In *ACL*, volume 2019, 6558.
- Ungar, A. A. 2001. Hyperbolic trigonometry and its application in the poincaré ball model of hyperbolic geometry. *Computers & Mathematics with Applications*, 41(1-2): 135–147.
- Wei, Y.; and Hu, D. 2024. Mmpareto: boosting multimodal learning with innocent unimodal assistance. *ICML*.
- Wen, Y.; Yong, W.; Guiduo, D.; Dongyang, Z.; XIN, H.; Yuan-Fang, L.; and Tao, H. 2025. Knowledge-Aligned Counterfactual-Enhancement Diffusion Perception for Un-supervised Cross-Domain Visual Emotion Recognition. In *CVPR*.
- Xia, W.; Jia, G.; Zhao, S.; and Yang, J. 2025. Seek Common Ground While Reserving Differences: Semi-Supervised Image-Text Sentiment Recognition. In *CVPR*, 29601–29611.
- Yang, D.; Huang, S.; Kuang, H.; Du, Y.; and Zhang, L. 2022. Disentangled representation learning for multimodal emotion recognition. In *Proceedings of the 30th ACM international conference on multimedia*, 1642–1651.
- Yang, D.; Li, M.; Xiao, D.; Liu, Y.; Yang, K.; Chen, Z.; Wang, Y.; Zhai, P.; Li, K.; and Zhang, L. 2024a. Towards multimodal sentiment analysis debiasing via bias purification. In *ECCV*, 464–481. Springer.
- Yang, J.; Yu, Y.; Niu, D.; Guo, W.; and Xu, Y. 2023. Confede: Contrastive feature decomposition for multimodal sentiment analysis. In *ACL*, 7617–7630.
- Yang, Y.; Wan, F.; Jiang, Q.-Y.; and Xu, Y. 2024b. Facilitating multimodal classification via dynamically learning modality gap. *NeurIPS*, 37: 62108–62122.
- Yang, Z.; Liu, X.; Ouyang, D.; Duan, G.; Zhang, D.; He, T.; and Li, Y.-F. 2024c. Towards Open-vocabulary HOI Detection with Calibrated Vision-language Models and Locality-aware Queries. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1495–1504.
- Yuksekgonul, M.; Zhang, L.; Zou, J. Y.; and Guestrin, C. 2023. Beyond confidence: Reliable models should also consider atypicality. *NeurIPS*, 36: 38420–38453.
- Zadeh, A.; Zellers, R.; Pincus, E.; and Morency, L.-P. 2016. Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6): 82–88.
- Zadeh, A. B.; Liang, P. P.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *ACL*, 2236–2246.
- Zhang, H.; Li, X.; and Bing, L. 2023. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Zhang, X.; Yoon, J.; Bansal, M.; and Yao, H. 2024. Multimodal representation learning by alternating unimodal adaptation. In *CVPR*, 27456–27466.
- Zhang, Y.; Latham, P. E.; and Saxe, A. 2023. Understanding unimodal bias in multimodal deep linear networks. *arXiv preprint arXiv:2312.00935*.
- Zhao, Z.; and Liu, Q. 2021. Former-dfer: Dynamic facial expression recognition transformer. In *ACMMM*, 1553–1561.
- Zhou, S.; Li, L.; Zhang, X.; Zhang, B.; Bai, S.; Sun, M.; Zhao, Z.; Lu, X.; and Chu, X. 2024. LiDAR-PTQ: Post-Training Quantization for Point Cloud 3D Object Detection. In *ICLR*.