

Enhancing Emotion Recognition in Incomplete Data: A Novel Cross-Modal Alignment, Reconstruction, and Refinement Framework

Anonymous ACL submission

Abstract

Multimodal emotion recognition systems rely heavily on the full availability of modalities, suffering significant performance declines when modal data is incomplete. To tackle this issue, we present the Cross-Modal Alignment, Reconstruction, and Refinement (CM-ARR) framework, an innovative approach that sequentially engages in cross-modal alignment, reconstruction, and refinement phases to handle missing modalities and enhance emotion recognition. This framework utilizes unsupervised distribution-based contrastive learning to align heterogeneous modal distributions, reducing discrepancies and modeling semantic uncertainty effectively. The reconstruction phase applies normalizing flow models to transform these aligned distributions and recover missing modalities. The refinement phase employs supervised point-based contrastive learning to disrupt semantic correlations and accentuate emotional traits, thereby enriching the affective content of the reconstructed representations. Extensive experiments on the IEMOCAP and MSP-IMPROV datasets confirm the superior performance of CM-ARR under conditions of both missing and complete modalities. Notably, averaged across six scenarios of missing modalities, CM-ARR achieves absolute improvements of 2.11% in WAR and 2.12% in UAR on the IEMOCAP dataset, and 1.71% and 1.96% in WAR and UAR, respectively, on the MSP-IMPROV dataset. Our code is available at [this url](#).

1 Introduction

Multimodal emotion recognition (MMER) entails the analysis of emotional cues across various modalities, including speech, text, and body language, among others. These modalities serve complementary functions in the expression and interpretation of human emotions. However, in practical applications, the availability of these modalities is frequently compromised; specific modalities may



Modality	Content	Availability	Prediction
Speech		×	
Text	I did exactly what they told me to do.	√	Neutral (×)
Video		√	

Figure 1: An example of missing modalities: when the speech modality is missing, emotion recognition is guided by the text and video modalities, leading to incorrect predictions. The ground truth is “angry”.

be absent or inaccessible due to various factors. For example, text data may be unavailable due to errors in automatic speech recognition systems, speech may be obscured by excessive background noise, and visual data may be impaired by poor lighting or occlusions. These challenges underscore the need for MMER systems to be highly adaptable and robust, capable of effectively functioning even with incomplete modal information.

Conventional multimodal learning paradigms, as documented in the literature (Yoon et al., 2018; Liu et al., 2022; Sun et al., 2024), typically operate under the assumption of complete modality presence. These approaches are dedicated to constructing fusion models optimized for scenarios where all modalities are fully available, a presumption that can undermine their utility in situations where modalities are partially missing. Illustratively, as depicted in Fig. 1, an emotion classified as “angry” in a fully modal context—primarily due to pronounced tones in the speech modality—may be reinterpreted as “neutral” in the absence of the speech component, with the text and video modalities becoming the guide.

The research community has developed innovative methodologies aimed at enhancing the resilience of MMER systems faced with incomplete modalities, primarily focusing on predicting missing data across modalities. However, the signifi-

cant distribution gaps between different modalities present substantial challenges. For instance, Wang et al. (2023) utilize flow models to map heterogeneous modality representations into a latent space following a Gaussian distribution, aiming to ensure consistency. Nevertheless, this approach does not adequately address the distribution gaps between modalities. Alternatively, Zuo et al. (2023) suggest the use of modality-invariant features to aid in reconstructing modality-specific characteristics. While this method is promising, it has been critiqued for its limited effectiveness, particularly due to difficulties in accurately predicting modality-specific features. These strategies underscore the critical need for effective cross-modal alignment before attempting to predict missing data across modalities.

To address the aforementioned problems, this paper introduces a novel framework for cross-modal alignment, reconstruction, and refinement, designated as **CM-ARR**. The initial alignment phase aims to bridge the distributional divergences between modalities, which facilitates subsequent reconstruction efforts. Specifically, inspired by MAP (Ji et al., 2023), we employ an **unsupervised¹ distribution-based contrastive learning** approach that replaces point representations with their Gaussian distributional counterparts. This method can convey richer multimodal semantic information by effectively encoding uncertainty. Next in the reconstruction phase, we deploy a network based on normalizing flow models. This method transforms aligned modality representations into a Gaussian latent space (Wang et al., 2023). Gaussian distributions for missing modalities are estimated by transferring characteristics from available modalities. Subsequently, the representation of the missing modality is obtained through inverse normalization of the estimated Gaussian distribution. In the final refinement phase, we refine modality representations to better capture emotional characteristics. The initial phases prioritize semantic alignment, sidelining emotional attributes. Consequently, we implement **supervised point-based contrastive learning**. This method considers modalities from different instances of the same class as positive samples and thus disrupts the semantic correlation between modalities. Doing so enables the model to capture emotional attributes and related features be-

¹Modalities from the same instance are treated as positive samples without relying on explicit class labels.

yond mere semantics, enhancing the reconstruction of emotional information.

Overall, CM-ARR begins by aligning modalities to harmonize disparate modal distributions, a step that facilitates the effective estimation of missing data in the subsequent reconstruction phase, and concludes with refinement to accentuate emotional traits. The key contributions of this paper are summarized as follows:

- We introduce the CM-ARR framework, a pioneering approach for cross-modal alignment, reconstruction, and refinement, designed to enhance emotion recognition in scenarios characterized by incomplete data.
- We present two contrastive learning strategies: unsupervised distribution-based contrastive learning for effective uncertainty modeling and mitigation of distributional disparities, alongside supervised point-based contrastive learning that disrupts strong semantic intermodality correlations, facilitating a deeper understanding of emotional consistency.
- Our empirical investigations, conducted on the IEMOCAP and MSP-IMPROV datasets under both missing and full modality conditions, affirm the superior performance of our proposed CM-ARR framework.

2 Related Work

2.1 Incomplete Multimodal Learning

In MMER, there have been remarkable advances in research addressing the modality absence problem. These methods fall into two main categories: missing modality generation (Cai et al., 2018; Suo et al., 2019; Du et al., 2018) and multimodal joint representation learning methods (Pham et al., 2019; Han et al., 2019; Yuan et al., 2021).

Missing Modality Generation aims to utilize available modalities to predict or reconstruct missing ones. Tran et al. (2017) introduce a Cascaded Residual Autoencoder (CRA) to fill in data with missing modalities. This method effectively recovers incomplete data by integrating a series of autoencoders in a cascaded structure and leveraging a residual mechanism to address corrupted data. Similarly, Cai et al. (2018) develop a 3D encoder-decoder network that captures the intermodal relationships and compensates for missing modalities through adversarial and classification losses.

Multimodal Joint Representation Learning seeks to learn latent representations in a common feature space from available data that remain robust even when some modalities are missing. [Pham et al. \(2019\)](#) introduce a method for learning robust joint representations through cyclic translation between modalities, thereby enhancing the model’s capability to comprehend and represent multimodal data. [Zhao et al. \(2021\)](#) propose the Missing Modality Imagination Network (MMIN), a unified model designed to address the issue of uncertain missing modalities. [Zeng et al. \(2022b\)](#) employ a Tag-Assisted Transformer Encoder (TATE) network, which guides the network to focus on different missing cases by encoding specific tags for the missing modalities. Furthermore, they ([Zeng et al., 2022a](#)) propose an Ensemble-based Missing Modality Reconstruction (EMMR) framework to detect and recover semantic features of the key missing modality. However, these methods do not consider the effect of heterogeneous modal gaps on missing modality reconstruction and emotion recognition. IF-MMIN ([Zuo et al., 2023](#)) and DiCMoR ([Wang et al., 2023](#)) work on this problem. The former learns modality-invariant features, and the latter transfers distributions from available modalities to missing modalities to maintain distribution consistency in the recovered data. Nevertheless, these approaches only partially bridge modal gaps and overlook semantic uncertainties across modalities. To overcome these limitations, we introduce the CM-ARR framework, which leverages Gaussian distributions to both align modalities and model semantic uncertainty.

2.2 Contrastive Learning

Contrastive learning (CL) ([Khosla et al., 2020](#); [He et al., 2020](#)) aims to foster efficient data representations by drawing similar samples closer and distancing dissimilar ones. In recent years, CL has become a cornerstone in the field of representation learning ([Radford et al., 2021](#); [Wang et al., 2021](#); [Ghosh et al., 2022](#); [Sun et al., 2023](#)). Notably, [Ji et al. \(2023\)](#) address the heterogeneity between image and text modalities using unsupervised CL. [Pan et al. \(2023\)](#) employ supervised contrastive learning to enhance emotional representation learning by clustering similar text and speech modality samples. Building on these principles, our work introduces both unsupervised distribution-based contrastive learning and supervised point-based contrastive learning. These approaches are designed

to bridge the gaps between heterogeneous modalities and decipher common emotional patterns for improved prediction accuracy.

3 CM-ARR

In this section, we detail the proposed CM-ARR framework. Fig. 2 illustrates the architecture of CM-ARR, which comprises three main phases: alignment, reconstruction, and refinement. Without loss of generality, we consider a multimodal dataset consisting of three modalities: text, speech, and video.

3.1 Alignment Phase

3.1.1 Feature Extraction

Given a speech signal, video segment and its corresponding transcribed text, we extract high-level features for each modality as follows:

Text Representation: For each text sequence, we obtain high-level text features R_t using the pre-trained Bert-base model ([Devlin et al., 2018](#)), which has 12 encoder layers, each with 12 self-attention heads and 768 hidden units.

Speech Representation: For each speech signal, we obtain high-level speech features R_s using the pre-trained Wav2vec2-base model ([Baevski et al., 2020](#)), where the pre-trained Wav2vec2-base model has 12 encoder layers, each with 8 self-attention heads and 768 hidden units.

Video Representation: For each video segment, we utilize a pre-trained DenseNet model ([Huang et al., 2017](#)) to extract facial expression features R_v , trained on the Facial Expression Recognition Plus (FER+) dataset ([Barsoum et al., 2016](#)). These features, referred to as "Denseface," are frame-level sequential features derived from detected faces in video frames, with each feature vector comprising 342 dimensions.

3.1.2 Unsupervised Distribution-based Contrastive Learning

In the alignment phase, to mitigate the gaps between heterogeneous modalities while also modeling the uncertainty, we introduce an uncertainty modeling component (UMC) that employs Gaussian distributions to capture semantic uncertainty. This is coupled with unsupervised distribution-based contrastive learning to bring the modal distributions closer.

Fig. 3 illustrates the architecture of the UMC, tasked with learning a Gaussian distribution for each point-based modality representation, R_n ,

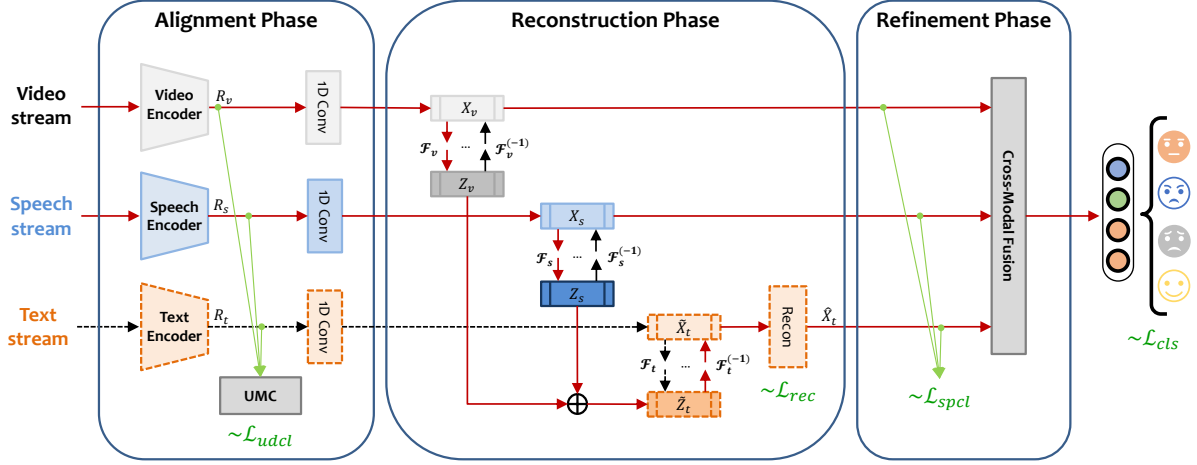


Figure 2: The framework of CM-ARR consists of three phases: the alignment phase employs unsupervised distribution-based contrastive learning to semantically align the video, speech, and text modalities (see UMC in Fig. 3); the reconstruction phase applies normalizing flow models to each modality; the refinement phase utilizes supervised point-based contrastive learning to accentuate emotional traits. The red arrows denote the inference process assuming the text modality is missing.




Figure 3: The overall structure of the proposed UMC, where $g(\cdot)$ denotes the gelu function, LN signifies the LayerNorm operation, and MLP indicates the feed forward layer.

where $n \in \{s, v, t\}$. We specifically employ the MLP and the multi-head attention mechanism to enhance feature-level and sequence-level interactions, respectively. The UMC learns a mean vector μ_n and a variance vector Σ_n for each R_n , transforming the point-based modality representations into Gaussian distributions.

We then implement an unsupervised distribution-based contrastive learning approach to align heterogeneous modal distributions effectively, utilizing the 2-Wasserstein distance (Kim et al., 2021) to measure the distance between the Gaussian distributions of three modalities:

$$D_{2W} = \|\mu_n - \mu_m\|_2^2 + \|\Sigma_n - \Sigma_m\|_2^2, \quad (1)$$

where $m, n \in \{s, v, t\}$ and $n \cap m = \emptyset$.

Suppose there are N speech-text, text-video, and speech-video pairs in each batch, where modalities from the same instance are treated as positive samples and those from different instances as negatives.

Taking the example of speech-text pairs, we utilize InfoNCE loss (He et al., 2020) to compute the loss \mathcal{L}_{udcl} :

$$\mathcal{L}_{udcl} = \mathcal{L}_{nce}^{s2t} + \mathcal{L}_{nce}^{t2s}, \quad (2)$$

$$\mathcal{L}_{nce}^{s2t} = -\log \frac{\exp(S(s_i, t_i) / \tau)}{\sum_{n=1}^N \exp(S(s_i, t_n) / \tau)}, \quad (3)$$

$$\mathcal{L}_{nce}^{t2s} = -\log \frac{\exp(S(t_i, s_i) / \tau)}{\sum_{n=1}^N \exp(S(t_i, s_n) / \tau)}, \quad (4)$$

$$S(t, s) = a \cdot D_{2W} + b, \quad (5)$$

where τ represents a learned temperature parameter. $S(\cdot, \cdot)$ denotes the similarity between a speech-text pair. a is a negative scale factor and b is a shift value.

3.2 Reconstruction Phase

Using representations from different feature spaces directly may lead to inefficacy due to scale, distribution, and semantic inconsistencies. Hence, we initially project each modality's features into a common dimensional space using a 1D convolutional layer, obtaining X_s , X_v , and X_t for further analysis and processing.

Subsequently, we define the normalizing flow model for each modality as \mathcal{F}_n , with $n \in \{s, v, t\}$, and $\mathcal{F}_n^{(-1)}$ representing its inverse transformation. The features of each modality are fed into their respective normalizing flow model, translating the input features from their original complex distribu-

Datasets	Languages	Year	Subjects	Type	Access	Sample size of emotions used				Total
						ang	hap	neu	sad	
IEMOCAP	English	2008	5 male, 5 female	Acted	Licensed	1103	1636	1708	1084	5531
MSP- IMPROV	English	2017	6 male, 6 female	Acted	Licensed	460	999	1733	627	3819

Table 1: Statistics of benchmark datasets. Note that the IEMOCAP dataset combines happiness and excited emotions into “hap”.

tions to a manageable Gaussian distribution Z_n .

$$Z_n = \mathcal{F}_n(X_n). \quad (6)$$

Conversely, the Gaussian distribution Z_n can be transformed back into the complex distribution of the input features via the inverse transformation $\mathcal{F}_n^{(-1)}$, giving \tilde{X}_n . Assuming text modality is missing and speech and video modalities are available, we input X_s and X_v into respective flow models, \mathcal{F}_s and \mathcal{F}_v , to obtain Z_s and Z_v . Z_t for the missing text modality can then be computed as:

$$\tilde{Z}_t \leftarrow (Z_s + Z_v) / 2 \sim \mathcal{N}(\mu_t, \Sigma_t), \quad (7)$$

where \mathcal{N} denotes the Gaussian distribution. μ_t and Σ_t represent the mean and covariance of the text Gaussian distribution, respectively. At this point, \tilde{Z}_t represents an estimated Gaussian consistent with missing text. \tilde{X}_t featuring the original text distribution is then generated through the inverse process of the text-specific flow:

$$\tilde{X}_t = \mathcal{F}_t^{(-1)}(\tilde{Z}_t). \quad (8)$$

Finally, the text-specific reconstruction module is used to recover the text features \hat{X}_t . The module consists of multiple residual channel attention blocks (Wei et al., 2022), where we replace the 2D convolutional layer with 1D. The reconstruction loss is computed as:

$$\mathcal{L}_{\text{rec}} = \left\| \hat{X}_t - X_t \right\|_F^2. \quad (9)$$

Similarly, when speech and video modalities is missing, we follow analogous steps, recovering speech and video features by transferring from available text.

3.3 Refinement Phase

To further promote the emotional information in the modality representations, we incorporate supervised point-based contrastive learning to refine the modality representations of CM-ARR.

3.3.1 Supervised Point-based Contrastive Learning

The learned representations still exhibit strong semantic correlations between various modality pairs. To address this, we utilize supervised point-based contrastive learning. Specifically, we treat modalities from different instances but with the same emotion labels (where different instances may have distinct semantics but share similar emotion characteristics) as positive samples, and those from different labels as negatives. This method transforms one-to-one modality relationships into many-to-many, emotion-centric relationships, thereby enabling the network to learn enhanced emotion representations beyond mere semantics.

3.3.2 Cross-Modal Fusion

Finally, after obtaining the recovered text features \hat{X}_t and the available speech and video features X_s and X_v , we fuse them into the multimodal representation H using three cross-modal attention blocks and one self-attention block (Vaswani et al., 2017) for the emotion recognition task. The process is as follows:

$$H_{tv} = \text{Cross-Attention}_{tv}(Q_t, K_v, V_v), \quad (10)$$

$$H_{ts} = \text{Cross-Attention}_{ts}(Q_t, K_s, V_s), \quad (11)$$

$$H' = \text{Concat}[H_{tv}, H_{ts}], \quad (12)$$

$$H = \text{Self-Attention}(H'), \quad (13)$$

where $K_s, V_s = X_s, K_v, V_v = X_v$ and $Q_t = \hat{X}_t$.

3.4 Optimization Objective

To train our proposed CM-ARR, four loss functions are required: unsupervised distribution-based contrastive learning loss \mathcal{L}_{udcl} , supervised point-based contrastive learning loss \mathcal{L}_{spcl} , modality reconstruction loss \mathcal{L}_{rec} , and emotion recognition loss \mathcal{L}_{cls} . In summary, our training loss is defined as:

$$\mathcal{L} = \alpha \mathcal{L}_{udcl} + \beta \mathcal{L}_{spcl} + \lambda \mathcal{L}_{rec} + \mathcal{L}_{cls}, \quad (14)$$

where α, β , and λ represent the trade-off factors.

Dataset	Avail.	CRA	MMIN	IF-MMIN	CIF-MMIN	DiCMoR ⁺	Ours
IEMOCAP	{ <i>t</i> }	31.15 / 27.96	67.49 / 68.46	66.58 / 67.51	67.97 / 68.93	68.83 / 70.15	69.74 / 70.92
	{ <i>s</i> }	54.58 / 56.79	54.84 / 56.87	56.06 / 58.38	56.26 / 58.46	71.02 / 72.30	74.15 / 75.38
	{ <i>v</i> }	53.31 / 51.17	53.17 / 50.28	53.11 / 51.28	51.40 / 51.39	51.60 / 49.03	54.06 / 52.42
	{ <i>v, t</i> }	31.66 / 28.42	72.98 / 73.60	72.61 / 73.14	73.00 / 74.12	72.14 / 73.36	73.04 / 74.09
	{ <i>s, v</i> }	63.12 / 63.93	64.03 / 64.71	64.80 / 66.49	66.03 / 67.17	72.28 / 73.49	75.51 / 76.38
	{ <i>s, t</i> }	32.88 / 30.13	74.07 / 75.49	73.77 / 75.48	74.50 / 75.72	76.89 / 77.83	78.95 / 79.70
	Avg.	44.45 / 43.07	64.43 / 64.90	64.48 / 65.38	64.86 / 65.96	68.79 / 69.36	70.90 / 71.48
	{ <i>s, v, t</i> }	-	77.17 / 77.89	77.97 / 78.58	79.26 / 80.34	78.36 / 79.80	79.86 / 81.06
Dataset	Avail.	CRA	MMIN	IF-MMIN	CIF-MMIN	DiCMoR ⁺	Ours
MSP-IMPROV	{ <i>t</i> }	46.78 / 28.37	62.08 / 57.30	61.97 / 58.23	62.42 / 58.71	62.13 / 59.85	63.57 / 61.49
	{ <i>s</i> }	37.90 / 38.96	51.60 / 43.35	50.46 / 40.45	50.66 / 40.37	52.78 / 47.17	55.55 / 50.87
	{ <i>v</i> }	59.46 / 42.42	60.09 / 45.75	61.68 / 45.29	61.10 / 46.15	59.84 / 49.33	62.04 / 51.01
	{ <i>v, t</i> }	54.96 / 38.84	69.37 / 63.94	67.49 / 63.40	69.90 / 65.36	67.30 / 64.21	68.47 / 65.88
	{ <i>s, v</i> }	57.85 / 47.70	63.74 / 55.91	62.42 / 53.14	63.75 / 55.21	62.70 / 52.65	63.84 / 54.25
	{ <i>s, t</i> }	48.57 / 37.97	64.00 / 60.98	63.25 / 59.91	63.78 / 60.80	67.92 / 65.37	69.49 / 66.85
	Avg.	52.59 / 39.04	61.81 / 54.53	61.21 / 53.40	61.93 / 54.43	62.11 / 56.43	63.82 / 58.39
	{ <i>s, v, t</i> }	-	69.70 / 64.89	69.03 / 63.84	72.02 / 67.12	71.69 / 68.08	72.37 / 69.69

Table 2: Performance comparison across testing conditions. The values reported in each cell denote WAR / UAR. “Avail.” indicates the available modalities. “Avg.” indicates average performance across all conditions. The best results are marked in boldface.

4 Experiments

4.1 Datasets

Our experiments utilize two widely adopted datasets: IEMOCAP (Busso et al., 2008) and MSP-IMPROV (Busso et al., 2016). IEMOCAP, collected by the University of Southern California, is a multi-modal emotion corpus comprising 10,039 utterances from 10 actors who express a range of specific emotions. MSP-IMPROV features 7,798 utterances from six sessions with 12 actors, focusing on the exploration of emotional behaviors during spontaneous dyadic improvisations. Additional details about these datasets are available in Table 1.

4.2 Experimental Setup

We evaluate all comparative methods on IEMOCAP and MSP-IMPROV using 5-fold cross-validation and 6-fold cross-validation, respectively. We employ evaluation metrics such as weighted average recall (WAR) and unweighted average recall (UAR) to assess the performance. In all experiments, parameters are configured as: $\alpha=1.0$, $\beta=0.1$, $\lambda=10$, learning rate= $1e-5$, batch size= 16 , epochs= 100 .

We benchmark CM-ARR against several state-of-the-art (SOTA) frameworks for incomplete multimodal emotion recognition, including CRA (Tran et al., 2017), MMIN (Zhao et al., 2021), IF-MMIN

(Zuo et al., 2023), CIF-MMIN (Liu et al., 2024), and DiCMoR⁺ (Wang et al., 2023). DiCMoR⁺ denotes our enhanced version of the DiCMoR framework, where we substitute the original speech encoder with a pre-trained Wav2vec2 model. This improves reproduction quality, creating a more robust framework. Using Wav2vec2 also enables a fairer comparison to our method, which similarly utilizes Wav2vec2 for speech encoding.

4.3 Comparison with SOTA Methods

Table 2 presents the performance of CM-ARR against SOTA models in terms of WAR and UAR under full and missing modality testing conditions. Across all evaluation metrics, CM-ARR consistently outperforms the competing models, indicating its superior performance.

Further analysis of CM-ARR’s recognition performance on the IEMOCAP dataset under various missing modality conditions reveals notable findings. Specifically, when compared to the best SOTA model, DiCMoR⁺, CM-ARR achieves a relative enhancement of 3.07% and 3.06% in WAR and UAR, respectively, on average. Notably, CM-ARR demonstrates exceptional performance when at least one modality is present, particularly with the availability of the speech modality, showing a significant relative improvements of 4.41% and 4.26% across WAR and UAR, respectively. Con-

Dataset	Avail.	Baseline	w/o \mathcal{L}_{udcl}	w/o \mathcal{L}_{spcl}	w/o attention	w/ Point	Ours
IEMOCAP	{t}	66.63 / 67.15	69.44 / 70.20	69.49 / 70.45	68.71 / 69.75	66.71 / 67.96	69.74 / 70.92
	{s}	55.57 / 57.75	71.90 / 72.67	71.89 / 72.99	71.68 / 72.72	70.48 / 71.74	74.15 / 75.38
	{v}	42.91 / 37.54	50.24 / 48.35	51.14 / 48.56	50.00 / 47.76	50.35 / 48.99	54.06 / 52.42
	{v, t}	68.71 / 68.94	71.59 / 72.90	72.55 / 73.33	70.72 / 71.53	71.06 / 72.19	73.04 / 74.09
	{s, v}	61.10 / 62.84	72.06 / 72.96	73.37 / 73.86	72.10 / 73.47	72.65 / 73.26	75.51 / 76.38
	{s, t}	75.34 / 76.61	76.51 / 78.04	77.80 / 78.83	76.77 / 77.91	75.53 / 76.67	78.95 / 79.70
	Avg.	61.71 / 61.80	68.62 / 69.18	69.37 / 69.67	68.33 / 68.85	67.79 / 68.46	70.90 / 71.48

Table 3: Emotion recognition results of the ablation study evaluating CM-ARR components on IEMOCAP. The values reported in each cell denote WAR / UAR.

versely, the performance gains are more modest when only the text modality is present. This discrepancy underscores the speech modality’s capacity to encapsulate substantial textual information, facilitating the effective reconstruction of textual modality representations from speech. In contrast, the text modality’s limited encapsulation of speech-related information results in less effective speech modality reconstruction. In addition, the performance of our method is also optimized compared to SOTA under full modality testing condition. In conclusion, CM-ARR’s ability to leverage available modalities for reconstructing missing modalities significantly mitigates the challenges posed by missing modalities, affirming its effectiveness in addressing the missing modality problem in multi-modal emotion recognition.

The right side of Table 2 shows the performance comparison between CM-ARR and SOTA methods on the MSP-IMPROV corpus. Given the corpus’s complexity as a challenging sentiment analysis dataset, the performance of these methods is generally modest. However, experimental results indicate that CM-ARR consistently surpasses SOTA methods across various scenarios, demonstrating its superior effectiveness and robust generalization capabilities.

4.4 Ablation Study

In Table 3, ablation experiments are conducted on each component of the CM-ARR framework. To illustrate the limitations of models trained exclusively on full modalities in addressing missing modality scenarios, we establish a full modality baseline model, denoted as ‘Baseline’, which includes feature extraction and cross-modal fusion components. Results from the IEMOCAP dataset indicate a significant performance decline in the Baseline model when faced with missing modalities, underscoring its vulnerability to conditions

of modality absence, given its training on the presumption of modality completeness.

Effects of Unsupervised Distribution-based Contrastive Learning: To verify the effectiveness of the alignment phase, we perform an ablation experiment (w/o \mathcal{L}_{udcl}) to evaluate the performance, as shown in Table 3. The results show that the models with unsupervised distribution-based contrastive learning achieve better performance. This suggests that distribution-based representations could learn richer semantic information from modal uncertainty and help bridge the distributional divergences between modalities, which facilitates subsequent reconstruction. Additionally, replacing distribution-based contrast learning with point-based representation (w/ Point) further demonstrates that leveraging modal uncertainty to gather diverse semantic information offers added advantages. Consequently, Gaussian distribution-based representations prove superior to instance-based representations.

Effects of Supervised Point-based Contrastive Learning: To validate the effectiveness of the refinement phase, we present the results of an ablation experiment (w/o \mathcal{L}_{spcl}). The results show that it is helpful to improve the performance by disrupting the semantic correlation between modalities through our point-based supervised contrast learning. This method allows the model to capture more generalized patterns within modal information, reducing the risk of overfitting to specific semantic content. Consequently, this approach emphasizes emotionally significant attributes, thereby enhancing the representation’s robustness.

4.5 Parameter Analysis

To thoroughly investigate the impact of various parameter settings on model performance, we conduct the comparative analysis focusing on the effects of \mathcal{L}_{udcl} , \mathcal{L}_{rec} , and \mathcal{L}_{spcl} weights.




Figure 4: The effect of weights α , β , and λ on performance.

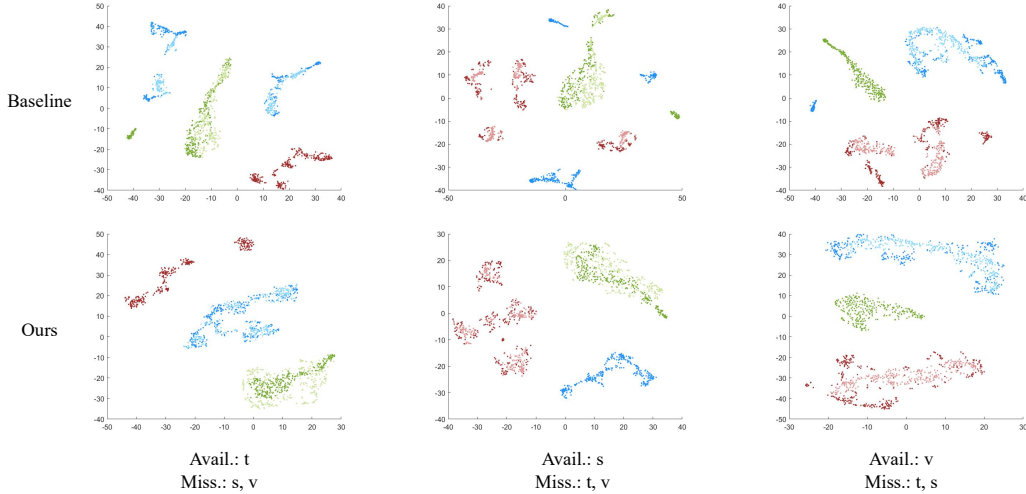


Figure 5: Visualization of the representations from different methods on the IEMOCAP corpus test set. Light blue represents speech reconstruction representations, while light red and light green depicts text and video reconstruction representations, respectively, with their corresponding darker shades indicating ground truth.

In Fig. 4 (a), the UAR (orange line) demonstrates relatively stable performance compared to WAR. This indicates that a moderate loss weight (1.0) yields the best performance for both WAR and UAR in unsupervised distribution-based contrastive learning. Similarly, the impacts of reconstruction loss and supervised point-based contrastive learning weights on performance are illustrated in Fig. 4 (b) and (c).

In summary, all three loss weights significantly influence model performance. The optimal weights are 1.0 for \mathcal{L}_{udcl} , 10.0 for \mathcal{L}_{rec} , and 0.5 for \mathcal{L}_{spl} .

4.6 Visualization Analysis

In Fig. 5, we use t-SNE (Van der Maaten and Hinton, 2008) to visualize the distribution of the modality representations for Baseline and our CM-ARR.

Fig. 5 illustrate the impact of CM-ARR on the reconstructed representation in scenarios of modality absence. In baseline, there is noticeably less overlap between the reconstructed modality and its ground-truth representations, with the distribution shape of the reconstructed representa-

tions markedly differing from that of the ground-truth representations. In contrast, ours (CM-ARR) demonstrates that the distributional similarity of reconstructed representations to the ground-truth representations is significantly enhanced, particularly evident in the overlap of clusters and the distribution shapes.

5 Conclusion

In this paper, we introduce the Cross-Modal Alignment, Reconstruction, and Refinement (CM-ARR) framework, designed to improve multimodal emotion recognition in incomplete data scenarios. CM-ARR effectively models uncertainty within the semantic space using unsupervised distribution-based contrastive learning, reducing the distributional gap. The reconstruction phase utilizes a normalizing flow model to transform aligned distributions, while the refinement phase augments the emotional content of the reconstructed representations. Extensive validation on the widely recognized IEMOCAP and MSP-IMPROV datasets confirms the superior effectiveness of our approach.

562 Limitations

563 The efficacy of the CM-ARR framework primarily
564 hinges on its ability to leverage available modalities
565 for reconstructing missing ones, thereby mitigat-
566 ing the adverse impacts of modality absence. The
567 experiments with missing modalities show that dif-
568 ferent modalities contribute to emotion recognition
569 to different degrees. For example, the video modal-
570 ity in the iemocap corpus is weak with a low degree
571 of its contribution to emotion recognition. There-
572 fore, how to deal with the transformation between
573 weak and strong modalities and measure the im-
574 portance of these modalities is a more interesting
575 issue, which will motivate us to further optimize
576 our proposed CM-ARR.

577 References

578 Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed,
579 and Michael Auli. 2020. wav2vec 2.0: A framework
580 for self-supervised learning of speech representations.
581 *Advances in Neural Information Processing Systems*,
582 33:12449–12460.

583 Emad Barsoum, Cha Zhang, Cristian Canton Ferrer,
584 and Zhengyou Zhang. 2016. Training deep networks
585 for facial expression recognition with crowd-sourced
586 label distribution. In *Proceedings of the 18th ACM
587 international conference on multimodal interaction*,
588 pages 279–283.

589 Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe
590 Kazemzadeh, Emily Mower, Samuel Kim, Jean-
591 nette N Chang, Sungbok Lee, and Shrikanth S
592 Narayanan. 2008. Iemocap: Interactive emotional
593 dyadic motion capture database. *Language Re-
594 sources and Evaluation*, 42:335–359.

595 Carlos Busso, Srinivas Parthasarathy, Alec Burmania,
596 Mohammed AbdelWahab, Najmeh Sadoughi, and
597 Emily Mower Provost. 2016. Msp-improv: An acted
598 corpus of dyadic interactions to study emotion per-
599 ception. *IEEE Transactions on Affective Computing*,
600 8(1):67–80.

601 Lei Cai, Zhengyang Wang, Hongyang Gao, Dinggang
602 Shen, and Shuiwang Ji. 2018. Deep adversarial learn-
603 ing for multi-modality missing data completion. In
604 *Proceedings of the 24th ACM SIGKDD international
605 conference on knowledge discovery & data mining*,
606 pages 1158–1166.

607 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and
608 Kristina Toutanova. 2018. Bert: Pre-training of deep
609 bidirectional transformers for language understand-
610 ing. *arXiv preprint arXiv:1810.04805*.

611 Changde Du, Changying Du, Hao Wang, Jinpeng Li,
612 Wei-Long Zheng, Bao-Liang Lu, and Huiguang He.
613 2018. Semi-supervised deep generative modelling of

incomplete multi-modality emotional data. In *Pro-
ceedings of the 26th ACM International Conference
on Multimedia*, pages 108–116. 614
615
616

Sreyan Ghosh, Utkarsh Tyagi, S Ramaneswaran,
Harshvardhan Srivastava, and Dinesh Manocha.
2022. Mmer: Multimodal multi-task learning
for speech emotion recognition. *arXiv preprint
arXiv:2203.16794*. 617
618
619
620
621

Jing Han, Zixing Zhang, Zhao Ren, and Björn Schuller.
2019. Implicit fusion by joint audiovisual training for
emotion recognition in mono modality. In *ICASSP
2019-2019 IEEE International Conference on Acous-
tics, Speech and Signal Processing (ICASSP)*, pages
5861–5865. IEEE. 622
623
624
625
626
627

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and
Ross Girshick. 2020. Momentum contrast for un-
supervised visual representation learning. In *Pro-
ceedings of the IEEE/CVF Conference on Computer
Vision and Pattern Recognition*, pages 9729–9738. 628
629
630
631
632

Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and
Kilian Q Weinberger. 2017. Densely connected con-
volutional networks. In *Proceedings of the IEEE
conference on computer vision and pattern recogni-
tion*, pages 4700–4708. 633
634
635
636
637

Yatai Ji, Junjie Wang, Yuan Gong, Lin Zhang, Yanru
Zhu, Hongfa Wang, Jiaying Zhang, Tetsuya Sakai,
and Yujiu Yang. 2023. Map: Multimodal uncertainty-
aware vision-language pre-training model. In *Pro-
ceedings of the IEEE/CVF Conference on Computer
Vision and Pattern Recognition*, pages 23262–23271. 638
639
640
641
642
643

Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron
Sarna, Yonglong Tian, Phillip Isola, Aaron
Maschinot, Ce Liu, and Dilip Krishnan. 2020. Su-
pervised contrastive learning. *Advances in Neural
Information Processing Systems*, 33:18661–18673. 644
645
646
647
648

Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. Vilt:
Vision-and-language transformer without convolu-
tion or region supervision. In *International Con-
ference on Machine Learning*, pages 5583–5594.
PMLR. 649
650
651
652
653

Rui Liu, Haolin Zuo, Zheng Lian, Bjorn W Schuller,
and Haizhou Li. 2024. Contrastive learning based
modality-invariant feature acquisition for robust mul-
timodal emotion recognition with missing modalities.
IEEE Transactions on Affective Computing. 654
655
656
657
658

Yang Liu, Haoqin Sun, Wenbo Guan, Yuqi Xia,
and Zhen Zhao. 2022. Multi-modal speech emo-
tion recognition using self-attention mechanism and
multi-scale fusion framework. *Speech Communica-
tion*, 139:1–9. 659
660
661
662
663

Yu Pan, Yanni Hu, Yuguang Yang, Jixun Yao, Wen Fei,
Lei Ma, and Heng Lu. 2023. Gemo-clap: Gender-
attribute-enhanced contrastive language-audio pre-
training for speech emotion recognition. *arXiv
preprint arXiv:2306.07848*. 664
665
666
667
668

669	Hai Pham, Paul Pu Liang, Thomas Manzini, Louis-Philippe Morency, and Barnabás Póczos. 2019. Found in translation: Learning robust joint representations by cyclic translations between modalities. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 33, pages 6892–6899.	
670		
671		
672		
673		
674		
675	Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In <i>International Conference on Machine Learning</i> , pages 8748–8763. PMLR.	
676		
677		
678		
679		
680		
681		
682	Dekai Sun, Yancheng He, and Jiqing Han. 2023. Using auxiliary tasks in multimodal fusion of wav2vec 2.0 and bert for multimodal emotion recognition. In <i>ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 1–5. IEEE.	
683		
684		
685		
686		
687		
688	Haoqin Sun, Shiwan Zhao, Xuechen Wang, Wenjia Zeng, Yong Chen, and Yong Qin. 2024. Fine-grained disentangled representation learning for multimodal emotion recognition. In <i>ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 11051–11055. IEEE.	
689		
690		
691		
692		
693		
694		
695	Qiuling Suo, Weida Zhong, Fenglong Ma, Ye Yuan, Jing Gao, and Aidong Zhang. 2019. Metric learning on healthcare data with incomplete modalities. In <i>IJCAI</i> , volume 3534, page 3540.	
696		
697		
698		
699	Luan Tran, Xiaoming Liu, Jiayu Zhou, and Rong Jin. 2017. Missing modalities imputation via cascaded residual autoencoder. In <i>Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition</i> , pages 1405–1414.	
700		
701		
702		
703		
704	Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. <i>Journal of Machine Learning Research</i> , 9(11).	
705		
706		
707	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. <i>Advances in Neural Information Processing Systems</i> , 30.	
708		
709		
710		
711		
712	Haobo Wang, Ruixuan Xiao, Yixuan Li, Lei Feng, Gang Niu, Gang Chen, and Junbo Zhao. 2021. Pico: Contrastive label disambiguation for partial label learning. In <i>International Conference on Learning Representations</i> .	
713		
714		
715		
716		
717	Yuanzhi Wang, Zhen Cui, and Yong Li. 2023. Distribution-consistent modal recovering for incomplete multimodal learning. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> , pages 22025–22034.	
718		
719		
720		
721		
722	Xiu-Shen Wei, Shu-Lin Xu, Hao Chen, Liang Xiao, and Yuxin Peng. 2022. Prototype-based classifier learning for long-tailed visual recognition. <i>Science China Information Sciences</i> , 65(6):160105.	
723		
724		
725		
	Seunghyun Yoon, Seokhyun Byun, and Kyomin Jung. 2018. Multimodal speech emotion recognition using audio and text. In <i>2018 IEEE Spoken Language Technology Workshop (SLT)</i> , pages 112–118. IEEE.	726 727 728 729
	Ziqi Yuan, Wei Li, Hua Xu, and Wenmeng Yu. 2021. Transformer-based feature reconstruction network for robust multimodal sentiment analysis. In <i>Proceedings of the 29th ACM International Conference on Multimedia</i> , pages 4400–4407.	730 731 732 733 734
	Jiandian Zeng, Jiantao Zhou, and Tianyi Liu. 2022a. Mitigating inconsistencies in multimodal sentiment analysis under uncertain missing modalities. In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 2924–2934.	735 736 737 738 739 740
	Jiandian Zeng, Jiantao Zhou, and Tianyi Liu. 2022b. Robust multimodal sentiment analysis via tag encoding of uncertain missing modalities. <i>IEEE Transactions on Multimedia</i> .	741 742 743 744
	Jinming Zhao, Ruichen Li, and Qin Jin. 2021. Missing modality imagination network for emotion recognition with uncertain missing modalities. In <i>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)</i> , pages 2608–2618.	745 746 747 748 749 750 751
	Haolin Zuo, Rui Liu, Jinming Zhao, Guanglai Gao, and Haizhou Li. 2023. Exploiting modality-invariant feature for robust multimodal emotion recognition with missing modalities. In <i>ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 1–5. IEEE.	752 753 754 755 756 757