

# MoLAN: A Unified Modality-Aware Noise Dynamic Editing Framework for Multimodal Sentiment Analysis

Xingle Xu<sup>1</sup>, Yongkang Liu<sup>2</sup>, Dexian Cai<sup>1</sup>, Shi Feng<sup>1\*</sup>,  
Xiaocui Yang<sup>1</sup>, Daling Wang<sup>1</sup>, Yifei Zhang<sup>1</sup>

<sup>1</sup>School of Computer Science and Engineering, Northeastern University,  
Shenyang 110819, China

<sup>2</sup>School of Computer and Communication Engineering, Northeastern University,  
Qinhuangdao 066004, China

xuxingle@stumail.neu.edu.cn, misonsky@163.com, 2301840@stu.neu.edu.cn  
{fengshi, yangxiaocui, wangdaling, zhangyifei}@cse.neu.edu.cn

## Abstract

Multimodal Sentiment Analysis aims to integrate information from various modalities to make complementary predictions. However, it often struggles with irrelevant or misleading visual and auditory information. Most existing approaches treat entire modality as an independent unit for feature enhancement or denoising, which often suppresses redundant noise at the cost of weakening critical information. To address this challenge, we propose **MoLAN**, a unified **ModaLity**-aware noise dynAmic editiNg framework. Specifically, MoLAN performs modality-aware block partitioning by dividing the features of each modality into multiple blocks. Each block is then dynamically assigned a distinct denoising strength based on its noise level and semantic relevance, enabling fine-grained noise suppression while preserving essential multimodal information. Notably, MoLAN is a unified and flexible framework that can be seamlessly integrated into a wide range of multimodal models. Building upon this framework, we further introduce **MoLAN<sup>+</sup>**, a new multimodal sentiment analysis approach. Experiments across five models and four datasets demonstrate the broad effectiveness of the MoLAN framework. Extensive evaluations show that MoLAN<sup>+</sup> achieves the state-of-the-art performance. The code is publicly available at <https://github.com/betterfly123/MoLAN-Framework>.

## 1 Introduction

Multimodal Sentiment Analysis (MSA) aims to integrate information from various modalities to achieve a more comprehensive and accurate understanding of the emotions (Zadeh et al., 2018b; Tsai et al., 2019a). MSA holds significant academic value in advancing multimodal learning, and offers broad industrial applications in areas such

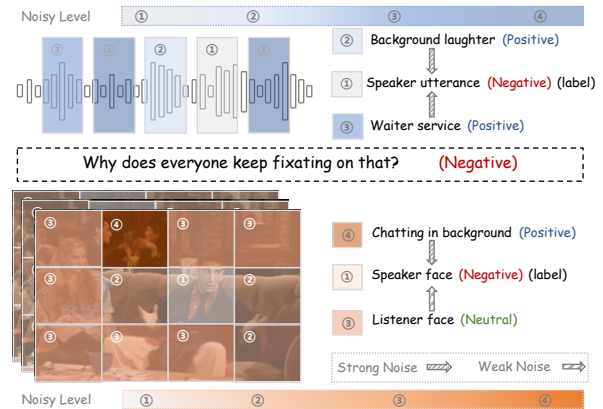


Figure 1: Distribution of noise. Lighter colors in the region mean more noise and less useful information.

as human-computer interaction and mental health monitoring (Zhu et al., 2025; Singh et al., 2024).

Existing MSA methods leverage multimodal synergy to achieve impressive improvements (Wu et al., 2024; Li and Liu, 2025; Zhang et al., 2026), but in real-world, noise interferes with representation learning and leads to performance degradation (Li and Li, 2025; Liu, 2024). To deal with multiple noise patterns, early solutions train individual models from scratch for each noise type (Yuan et al., 2021) or design a unified model to perceive noise adaptively (Zeng et al., 2022c). However, noise sources differ substantially across modalities, making cross-modality noise transfer prone to failure. Accordingly, researchers design pattern specific denoising methods (Yuan et al., 2024), yet they operate at the whole modality. Moreover, noise intensity varies across regions within the modality, so holistic denoising often suppresses noise at the cost of losing essential information.

As shown in Figure 1, the intensity of this noise varies across different regions. Specifically, background smiles that contradict Ross’s emotional state constitute strong noise in the visual modality, while the remaining background mostly represents

\*Corresponding author.

weak noise. In the audio modality, segments such as laughter that conflict with the sentiment label also form strong noise. The inconsistency of the noise distribution highlights the importance of fine-grained denoising. Therefore, the challenge of this paper is how to perform fine-grained noise dynamic editing on different modalities, so as to remove noise information while retaining information that is beneficial to MSA.

To address these issues, we propose **MoLAN**, a unified **ModaLity**-aware **noise dynAmic editiNg** framework. To achieve dynamic fine-grained denoising, MoLAN employs a block partitioning strategy that divides each modality into different sub-blocks. In each block, the denoising strength is dynamically computed based on the noise level of the block. This approach allows the model to apply varying degrees of denoising across different blocks, thereby enhancing its ability of selective noise editing while preserving essential information in each modality. In addition, considering the heterogeneity between different modalities (Fan et al., 2024; Wei et al., 2023), we design differentiated block partitioning strategies for each modality. Combining our experimental results and following the conclusions of previous studies (Li and Li, 2025; Zhang et al., 2023; Lin and Hu, 2022), we choose text modality as the main basis to calculate the denoising strength. Furthermore, MoLAN is a unified framework that can be flexibly integrated into different MSA models and Multimodal Large Language Models (MLLMs), thereby raising the upper limit of model performance. Based on the MoLAN, we propose **MoLAN<sup>+</sup>**, which uses denoised information to update the cross-attention between modalities and guide the model to focus on important information for MSA. MoLAN<sup>+</sup> further introduces denoising-driven contrastive learning to encourage the model to generate higher quality features, improving the performance of MSA task. The key contributions are as follows:

- To address the noise, we propose **MoLAN**, a unified **ModaLity**-aware **noise dynAmic editiNg** framework. It performs modality-aware block partitioning by dividing modality into multiple blocks. Each block is dynamically assigned a distinct denoising strength, enabling fine-grained noise editing. Additionally, MoLAN can be flexibly integrated into various models.
- Based on the denoising framework, we further introduce the noise suppression cross-attention mechanism and denoising-driven contrastive

learning, and design an MSA method **MoLAN<sup>+</sup>**. MoLAN<sup>+</sup> suppresses noise and guides the model to generate higher quality features.

- We conduct experiments on seven models and four datasets to demonstrate the broad effectiveness of the MoLAN framework. Additionally, extensive evaluations on four benchmark multimodal datasets show that MoLAN<sup>+</sup> achieves the state-of-the-art performance.

## 2 Related Work

### 2.1 Multimodal Sentiment Analysis (MSA)

As a fundamental task in multimodal affective computing (Zhang et al., 2024; Cheng et al., 2024), MSA leverages multimodal signals to understand emotions. Early studies mainly adopt fusion methods such as TFN (Zadeh et al., 2017) and LMF (Liu and Shen, 2018) to obtain joint representations. Subsequently, Transformer encoder architectures (Vaswani et al., 2017) and cross-modal attention become mainstream. For example, MulT (Tsai et al., 2019a) uses cross-modal attention to align and fuse modalities, and related work (Zhou et al., 2025; Wu et al., 2024; Guo et al., 2024) further explores more effective alignment strategies. More recently, knowledge is also incorporated. KuDA (Feng et al., 2024) leverages affective knowledge to dynamically select the dominant modality and adjust modality contributions, while KEBR (Zhu et al., 2024) injects non-verbal information from videos into textual semantics to enhance representations. Despite continuous progress in alignment and fusion, the impact of modality noise is often overlooked, which limits model performance. This work focuses on noise across modalities and performs noise editing to improve MSA.

### 2.2 Multimodal Sentiment Analysis Denoising

Recently, noise in MSA attracts increasing attention. t-HNE (Li and Li, 2025) removes visual and audio noise via text guidance and attention mechanisms. Meta-NA (Zeng et al., 2022c) simulates noise tasks through meta-learning to improve robustness. JOSFD (Jiang et al., 2024) introduces fuzzy logic into multimodal fusion and decision-making to model emotion uncertainty. Missing modality is also regarded as a type of noise. EMMR (Zeng et al., 2022b) reconstructs semantic features of key missing modalities, TATE (Zeng et al., 2022a) uses tags to guide the model to focus on missing information, IASE (Shi et al., 2024) ag-

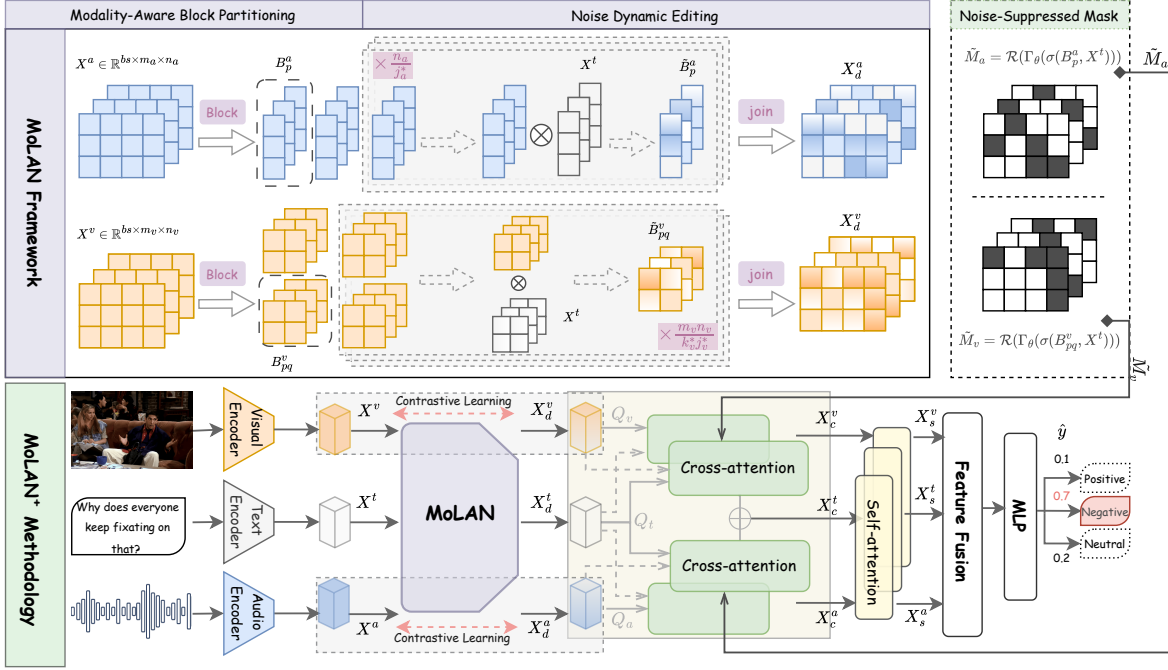


Figure 2: An illustration of MoLAN framework and MoLAN+ method. The purple box above represents the MoLAN framework, and the below represents the entire process of the MoLAN+ method. The MoLAN framework shown above provides a detailed description of the MoLAN block presented in MoLAN+ method.

gregates data with a bipartite-graph formulation to reduce the impact of missing modalities, UMDf (Li et al., 2024b) learns strong representations via a unified self-distillation mechanism, and Prompts are also used for missing modalities (Guo et al., 2024). Unfortunately, existing works adopt too coarse processing granularity in the process of noise removal. This defect may cause excessive denoising and loss of essential information, or may lead to incomplete noise removal. In contrast, our work performs fine-grained noise dynamic editing on modality features, ensuring that essential information is preserved while denoising.

### 3 MoLAN Framework

As shown in the purple box in the upper part of Figure 2, the MoLAN framework consists of two components: modality-aware block partitioning and noise dynamic editing. The pilot study is provided in Appendix A.

#### 3.1 Modality-Aware Block Partitioning

Since the distribution of noise is uneven, we introduce a block-level mechanism to enable fine-grained control over the denoising range. Through block partitioning, the minimum unit of the denoising operation shifts from the entire modality feature to a feature block, thereby achieving

more precise denoising. Considering the differences between modalities, visual information usually presents in regional forms, which is suitable for two-dimensional block partitioning. In contrast, audio information appears as continuous segments, so one-dimensional block partitioning is more effective. Feature representation of modalities as  $X^f \in \mathbb{R}^{bs \times m_f \times n_f}$ ,  $f \in t, a, v$ . Here,  $m_f$  and  $n_f$  denote the sequence length and hidden size of modality  $f$ , respectively.

$$P_{block} = \begin{cases} (k_v^*, j_v^*), & \text{if } f = v \\ (j_a^*), & \text{if } f = a \end{cases} \quad (1)$$

where  $P_{block}$  is the optimal block partitioning parameter, it represents the size of the block.  $bs$  is batch size. We use two-dimensional block partitioning as an example to illustrate the block partitioning process.  $(k_v^*, j_v^*)$  is calculated as:

$$(k_v^*, j_v^*) = \arg \min_{(k_v, j_v) \in \mathcal{D}_{m_v} \times \mathcal{D}_{n_v}} \left\| \frac{k_v}{\sqrt{m_v}} - 1 \right\|_2^2 + \left\| \frac{j_v}{\sqrt{n_v}} - 1 \right\|_2^2 \quad (2)$$

where  $\mathcal{D}_m$  and  $\mathcal{D}_n$  represent the sets of factors  $k_v$  and  $j_v$  of  $m_v$  and  $n_v$ , respectively. The  $\|\cdot\|_2^2$  represents the square of the L2 norm. Using the factor closest to the square root as the basis for block partitioning can achieve balanced segmentation: it can

avoid information loss or noise residue caused by too large blocks, and semantic loss caused by too small blocks. The modality feature  $X^v$  is reshaped according to the block partitioning parameters, and each sub-block is defined as:

$$B_{pq}^v = X^v [(p-1)k_v^*, (q-1)j_v^*], \quad (3)$$

$$\forall p \in \left[1, \frac{m_v}{k_v^*}\right], q \in \left[1, \frac{n_v}{j_v^*}\right]$$

where  $B_{pq}^v \in R^{bs \times k_v^* \times j_v^*}$ .  $p$  and  $q$  index the row and column positions of the sub-blocks. We conduct ablation studies to explain the choice of the block partitioning factor.

### 3.2 Noise Dynamic Editing

We first dynamically compute the adaptive denoising strength of each sub-block  $B_{pq}^v$  and then edit its noise accordingly. The denoising strength is calculated as follows:

$$\sigma(B_{pq}^v, X^t) = \frac{\langle \Phi(B_{pq}^v), \Psi(X^t) \rangle}{\|\Phi(B_{pq}^v)\|_2 \cdot \|\Psi(X^t)\|_2} \quad (4)$$

where  $X^t$  denotes the text vector.  $\Phi$  and  $\Psi$  are tensor reshaping mapping functions, which can be written as  $\Phi(o) = \Psi(o) = \text{reshape}(o)$ , where  $o$  denotes the feature tensor. These mapping functions alleviate the dimensional mismatch across modalities. In addition, we constrain the denoising strength  $\sigma$  to the non negative range  $[0, 1]$ , so negative values are not allowed during denoising. This avoids sign flipping that may reverse the editing direction and thus ensures the stability of the editing process. Based on this strength, we perform dynamic editing. The denoising operation for each sub-block  $B_{pq}^v$  can be represented as:

$$\tilde{B}_{pq}^v = \sigma(B_{pq}^v, X^t) \cdot B_{pq}^v \quad (5)$$

The denoised feature is obtained by recombining all blocks:

$$X_d^v = R(\{\tilde{B}_{pq}^v\}_{p=1, q=1}^{P, Q}) \quad (6)$$

where  $X_d^v \in R^{bs \times m_v \times n_v}$ .  $\mathcal{R}$  is the block reassembly operator.  $P, Q$  is total block numbers.  $P = m_v/k_v^*, Q = n_v/j_v^*$ .

## 4 MoLAN<sup>+</sup> Methodology

As shown in Figure 2, we propose the MoLAN<sup>+</sup> method built upon MoLAN framework. The MoLAN<sup>+</sup> method consists of three components:

the MoLAN framework, noise-suppressed cross attention, and denoising-driven contrastive learning. The detailed introduction of each module can be found in following subsections.

### 4.1 Problem Definition

In MSA task, the input signal consists of  $text(t)$ ,  $visual(v)$  and  $audio(a)$  modalities. The feature representation of these modalities can be denoted as  $X^f \in R^{bs \times m_f \times n_f}$ ,  $f \in t, a, v$ . The prediction is the sentiment score  $\hat{y}$ , which is a value.

To ensure a fair comparison, we use the same feature encoder as previous work (Tsai et al., 2019a; Wu et al., 2024; Sun and Tian, 2025a). After encoding, we feed the modality features into the MoLAN to obtain the denoised features:

$$X_d^v, X_d^a = MoLAN(X^t, X^v, X^a) \quad (7)$$

### 4.2 Noise-Suppressed Cross Attention

In the green module below of Figure 2, to further enhance noise suppression, we update the attention mechanism based on the denoising strength calculation information. Take visual modality as an example:

$$M_{pq}^v = \Gamma_\theta(\sigma(B_{pq}^v, X^t)) \in \{0, 1\} \quad (8)$$

$$\Gamma_\theta(\sigma(B_{pq}^v, X^t)) = I(\sigma(B_{pq}^v, X^t) \geq \theta) \odot J_{pq} \quad (9)$$

where  $I(\cdot)$  is the indicator function.  $J_{pq}$  is an all-ones matrix.  $\theta \in [0, 1]$  is the fixed similarity threshold. We aggregate the sub-blocks together to construct the mask matrix.

$$\tilde{M}_v = \mathcal{R}(\{M_{pq}^v\}_{p=1, q=1}^{P, Q}) \quad (10)$$

where  $\mathcal{R}$  is the block reassembly operator that joins the sub-blocks according to their original positions. The mask matrix derived from the MoLAN is combined to calculate the cross-modality attention score to reduce the impact of noise.

$$X_c^{f_q} = \Omega_{\text{Inter-M}}(X_d^{f_q}, X_d^{f_{kv}}, \tilde{M}_{f_q}) \quad (11)$$

$$\Omega_{\text{Inter-M}}(X_d^{f_q}, X_d^{f_{kv}}, \tilde{M}_{f_q}) = \frac{\exp_i[X_d^{f_q} X_d^{f_{kv}T} \cdot (\sqrt{d})^{-1} + \tilde{M}_{mq}] X_d^{f_{kv}}}{\sum \exp_j[X_d^{f_q} X_d^{f_{kv}T} \cdot (\sqrt{d})^{-1} + \tilde{M}_{mq}]} \quad (12)$$

where  $f_q, f_{kv} \in \{t, a, v\}$ .  $f_q$  represents the query source modality of the current modality, and  $f_{kv}$  represents the modality that provides key/value

pairs.  $\tilde{M}_{mq}$  denotes the attention mask corresponding to the  $mq$  modality.  $i$  and  $j$  denote the row and column position indices of the attention matrix, respectively. The text attention mask matrix is generated by the encoder.  $X_c^{fq}$  represents the feature after the attention mechanism.

### 4.3 Denoising-Driven Contrastive Learning

As shown in the light gray below of Figure 2, to enhance the encoder’s capability in distinguishing noise, we introduce a noise-driven contrastive learning loss. We perform contrastive learning between the denoised modal features and their corresponding original features. By minimizing the distance between positive pairs and maximizing the differences with other samples, the model is encouraged to learn more discriminative denoised representations. The contrastive learning loss function can be formulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{contrast}} = & \\ & -E \log \left[ \frac{\exp(\phi(X_d^v, X^v)/\tau)}{\sum_{j=1}^N \exp(\phi(X_d^v, X_j^v)/\tau)} \right] \\ & -E \log \left[ \frac{\exp(\psi(X_d^a, X^a)/\tau)}{\sum_{j=1}^N \exp(\psi(X_d^a, X_j^a)/\tau)} \right] \end{aligned} \quad (13)$$

where  $\phi$  and  $\psi$  denote similarity measurement functions, with cosine similarity adopted in this work.  $\tau$  is a temperature parameter that controls the smoothness of the distribution. During training, this module guides the encoder to focus on distinguishing modal noise, thereby improving the overall denoising quality and the effectiveness of fusion.

### 4.4 Sentiment Prediction

By introducing a cross-attention mechanism, information from different modalities is effectively interacted and integrated, enabling the initial fusion of multimodal features. Subsequently, the interacted representations are fed into a self-attention mechanism to further model the deep intra-modal dependencies. Finally, the information from different modalities is consolidated into a unified representation, which is used for the sentiment analysis.

$$X_s^f = \text{SelfAttn}(X_c^f, \tilde{M}_f) \quad (14)$$

$$\hat{y} = \text{MLP}_{\theta_{\text{FC}}}(\mathcal{F}_{\text{fuse}}(\text{Cat}(X_s^t, X_s^a, X_s^v))) \quad (15)$$

where  $\theta_{\text{FC}}$  denotes the parameters of the fully connected network.  $\mathcal{F}_{\text{fuse}}$  denotes linear layer, and  $\hat{y}$  is

the predicted sentiment value. The overall training of the MoLAN<sup>+</sup> is performed by minimizing the following loss:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \mathcal{L}_{\text{contrast}}^v + \mathcal{L}_{\text{contrast}}^a \quad (16)$$

where  $\mathcal{L}_{\text{task}}$  involves regression and classification tasks. For regression tasks, we adopt the L1 loss, following prior works (Mai et al., 2023, 2020), which measures the absolute difference. For classification tasks, the standard cross-entropy loss is employed to optimize the model. The loss of the predicted value  $\hat{y}$  and the ground truth  $y$  is:

$$\mathcal{L}_{\text{task}} = \begin{cases} \mathcal{L}_{\text{reg}} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| & (\text{Reg.}) \\ \mathcal{L}_{\text{cla}} = \frac{1}{N} \sum_{i=1}^N -y_i \log(\hat{y}_i) & (\text{Cls.}) \end{cases} \quad (17)$$

where  $N$  is the number of the samples. The model is trained based on the overall loss function.

## 5 Experiments

### 5.1 Datasets and Baselines

We use CMU-MOSI (Zadeh et al., 2016), CMU-MOSEI (Zadeh et al., 2018a), CH-SIMS (Yu et al., 2020), and IEMOCAP (Busso et al., 2008).

**Framework:** We compare our method with **MuT** (Tsai et al., 2019a), **SPECTRA** (Yu et al., 2023), **KuDA** (Feng et al., 2024), **SFTTR** (Sun and Tian, 2025b), **MMML** (Wu et al., 2024), as well as MLLM-based baselines **LLaVA-NeXT** (Li et al., 2024a) and **Qwen2.5-VL** (Bai et al., 2025).

**Models:** We further benchmark against representative MSA methods, including **TFN** (Zadeh et al., 2017), **LMF** (Liu et al., 2018), **MFN** (Tsai et al., 2019b), **Self-MM** (Yu et al., 2021), **UniMSE** (Hu et al., 2022b), **CHFV** (Guo et al., 2022), **ALMT** (Zhang et al., 2023), **EMT** (Sun et al., 2023), **GLoMo** (Zhuang et al., 2024), **JOSFD** (Jiang et al., 2024), **t-HNE** (Li and Li, 2025), and **MMML** (Wu et al., 2024). More details are in Appendix B.

### 5.2 Framework Overall Analysis

This section provides an overall analysis of the proposed framework to assess its performance across different evaluation dimensions. We first examine the framework’s effectiveness in enhancing MSA, followed by an analysis of its universality ability.

**Framework effectiveness.** We validate the proposed MoLAN framework on four standard MSA datasets, using five representative MSA models and

Model	CMU-MOSI					CMU-MOSEI				
	Acc <sub>2</sub>	F1	Acc <sub>7</sub>	MAE ↓	Corr	Acc <sub>2</sub>	F1	Acc <sub>7</sub>	MAE ↓	Corr
MuT	75.93/77.94	75.68/77.69	35.78	0.9494	65.91	77.69/79.19	77.99/79.60	48.65	0.6299	65.33
+MoLAN	<b>78.85/81.04</b>	<b>77.97/80.57</b>	<b>36.71</b>	<b>0.9140</b>	<b>68.56</b>	<b>78.50/81.19</b>	<b>78.32/81.16</b>	<b>49.52</b>	<b>0.6195</b>	<b>66.44</b>
SPECTRA	84.12/86.33	83.96/86.04	47.12	0.7941	75.69	83.11/85.34	82.16/85.71	51.37	0.5849	74.19
+MoLAN	<b>85.13/86.79</b>	<b>84.53/86.69</b>	<b>48.27</b>	<b>0.7848</b>	<b>76.41</b>	<b>84.36/85.79</b>	<b>83.56/86.60</b>	<b>52.69</b>	<b>0.5641</b>	<b>75.14</b>
KuDA	84.00/85.92	83.78/86.06	46.54	0.7110	78.35	82.95/86.14	82.61/86.24	51.54	0.5335	76.71
+MoLAN	<b>85.92/87.31</b>	<b>86.05/88.12</b>	<b>48.23</b>	<b>0.7011</b>	<b>80.15</b>	<b>84.55/87.35</b>	<b>84.2/88.00</b>	<b>53.33</b>	<b>0.5127</b>	<b>78.84</b>
SFTTR	81.73/83.15	82.54/84.05	45.95	0.7137	78.99	82.17/85.27	82.79/85.64	53.12	0.5395	76.33
+MoLAN	<b>83.65/85.62</b>	<b>84.04/86.25</b>	<b>47.17</b>	<b>0.6894</b>	<b>80.15</b>	<b>83.11/86.24</b>	<b>83.54/86.63</b>	<b>54.29</b>	<b>0.5197</b>	<b>78.36</b>
MMML	86.91/88.92	86.92/88.97	49.71	0.5820	87.05	86.43/87.96	86.45/87.76	53.39	0.5224	81.39
+MoLAN	<b>87.72/89.30</b>	<b>87.69/89.33</b>	<b>50.61</b>	<b>0.5827</b>	<b>87.46</b>	<b>86.85/88.10</b>	<b>86.53/87.88</b>	<b>55.11</b>	<b>0.5145</b>	<b>81.66</b>
LLaVA-NeXT§	79.49/81.24	79.02/80.54	43.17	0.8236	70.68	78.93/79.68	78.44/79.35	47.65	0.6077	70.00
+MoLAN	<b>80.48/83.14</b>	<b>80.97/82.05</b>	<b>44.60</b>	<b>0.8019</b>	<b>71.99</b>	<b>80.11/81.24</b>	<b>79.86/81.02</b>	<b>50.29</b>	<b>0.5869</b>	<b>72.53</b>
Qwen2.5-VL§	87.03/89.01	87.01/88.69	50.04	0.6306	86.39	86.73/87.45	86.70/86.98	54.06	0.5610	79.98
+MoLAN	<b>87.80/89.44</b>	<b>87.69/88.95</b>	<b>50.61</b>	<b>0.6074</b>	<b>87.00</b>	<b>87.09/87.87</b>	<b>87.00/87.49</b>	<b>56.02</b>	<b>0.5195</b>	<b>81.78</b>

Table 1: The performance of the MoLAN framework on the MOSI and MOSEI. The baseline results in the experiment are obtained through replication. Two evaluation metrics, ACC and F1, are adopted, specifically  $ACC_{2Has0} / ACC_{2Non0}$  and  $F1_{Has0} / F1_{Non0}$ . § denotes fine-tuning with LoRA. We perform significance testing on seven experimental groups, with p-value of  $5.33 \times 10^{-5}$ ,  $1.43 \times 10^{-6}$ ,  $4.87 \times 10^{-10}$ ,  $2.68 \times 10^{-6}$ ,  $2.16 \times 10^{-3}$ ,  $2.45 \times 10^{-5}$ , and  $2.66 \times 10^{-6}$ , all of which  $< 0.05$  indicate significant differences.

Model	CH-SIMS		IEMOCAP	
	ACC <sub>2</sub>	F1	Weighted-F1	Macro-F1
MuT	68.49	55.68	65.07	65.38
+MoLAN	<b>69.24</b>	<b>58.37</b>	<b>66.35</b>	<b>66.37</b>
SPECTRA	77.12	77.06	63.21	63.54
+MoLAN	<b>78.05</b>	<b>78.00</b>	<b>64.33</b>	<b>64.36</b>
KuDA	78.54	78.41	64.00	63.86
+MoLAN	<b>80.23</b>	<b>80.12</b>	<b>65.27</b>	<b>64.92</b>
SFTTR	79.22	79.20	64.05	61.54
+MoLAN	<b>80.96</b>	<b>80.39</b>	<b>65.61</b>	<b>62.94</b>
MMML	79.39	79.50	64.29	62.86
+MoLAN	<b>81.36</b>	<b>81.14</b>	<b>67.35</b>	<b>66.53</b>
LLaVA-NeXT§	76.64	76.02	64.61	62.19
+MoLAN	<b>77.93</b>	<b>77.11</b>	<b>65.87</b>	<b>65.08</b>
Qwen2.5-VL§	80.22	79.87	66.28	65.72
+MoLAN	<b>81.56</b>	<b>80.99</b>	<b>67.59</b>	<b>67.26</b>

Table 2: The performance of the MoLAN framework on the SIMS and IEMOCAP. § denotes fine-tuning with LoRA. Significance testing in Appendix C.

two MLLMs. Detailed experimental results are shown in Table 1 and Table 2. Five MSA models all show dramatically performance improvements after integrating the MoLAN framework. This trend demonstrates that MoLAN can be universally applied across diverse model architectures, effectively improving the quality of multimodal feature fusion. This further demonstrates that MoLAN, through fine-grained noise editing of modality information, can suppress irrelevant noise while retaining key information, thereby strengthening semantic alignment and collaborative representation across modalities. It is worth noting that even the current SOTA, MMML, achieves further improvement after integrating MoLAN. This observation indicates that existing methods still have limitations in modality denoising, and that MoLAN helps overcome these performance bottlenecks. Further-




Figure 3: Pixel-level heatmap. Color intensity indicates the magnitude of the pixel value, with brighter areas representing stronger image information.

more, experiments on MLLMs further validate the framework’s effectiveness. Both LLaVA-NeXT and Qwen2.5-VL exhibit consistent performance improvements after integrating MoLAN. Result demonstrates that MoLAN is not only applicable to traditional MSA models but also seamlessly integrates with mainstream MLLM architectures.

**Framework universality.** Experimental results show that whether MoLAN is integrated into dedicated MSA models or deployed on MLLMs, it consistently exhibits significant performance improvements. This consistent trend verifies the strong universality capability of MoLAN, indicating that it effectively adapts to different architectures, task settings, and data distributions. Overall, the results suggest that MoLAN serves as a general enhancement module for MSA tasks, providing a reliable solution for multimodal representation learning.

### 5.3 Effectiveness of Dynamic Strategy

The superior performance of the framework demonstrates the effectiveness of our proposed noise dynamic editing. Next, we illustrate the key role

Model	CMU-MOSI					CMU-MOSEI				
	ACC <sub>2</sub>	F1	ACC <sub>7</sub>	MAE ↓	Corr	ACC <sub>2</sub>	F1	ACC <sub>7</sub>	MAE ↓	Corr
TFN	-/80.80	-/80.70	34.90	0.9010	69.80	-/82.50	-/82.10	50.20	0.5930	67.70
LMF	-/82.50	-/82.40	33.20	0.9170	69.50	-/82.00	-/82.10	48.00	0.6230	70.00
MFM	-/81.70	-/81.60	35.40	0.8770	70.60	-/84.40	-/84.30	51.30	0.5680	70.30
MuT	75.93/77.94	75.68/77.69	35.78	0.9494	65.91	77.99/79.60	77.99/79.60	48.65	0.6299	65.33
Self-MM	84.00/84.42	85.98/85.95	-	0.7130	79.80	82.81/82.53	85.17/85.30	-	0.5300	76.50
UniMSE	85.85/86.90	85.83/86.42	48.68	0.6910	80.90	85.86/87.50	85.79/87.46	54.39	0.5230	77.30
CHFN	84.30/86.40	84.20/86.20	48.60	0.6890	80.90	83.70/86.20	83.90/86.10	54.30	0.5250	77.80
SPECTRA	84.12/86.33	83.96/86.04	47.12	0.7941	75.69	83.11/85.34	82.16/85.71	51.37	0.5849	74.19
ALMT	84.55/86.43	84.57/86.47	49.42	0.6830	80.50	84.78/86.79	85.19/86.86	54.28	0.5260	77.90
EMT	83.30/85.00	83.20/85.00	47.40	0.7050	79.80	83.40/86.00	83.70/86.00	54.50	0.5270	77.40
GLoMo	84.10/86.70	83.90/86.60	48.30	0.7180	78.20	83.70/86.50	84.00/86.40	55.00	0.5390	77.10
MMML◇	86.91/88.92	86.92/88.97	49.71	0.5820	87.05	86.43/87.96	86.45/87.76	53.39	0.5224	81.39
JOSFD†	-/89.80	-/89.70	52.10	0.5790	87.40	-/87.90	-/87.90	51.80	0.5150	79.70
t-HNe†	85.02/87.03	84.98/87.01	47.04	0.6800	81.00	85.20/87.14	85.32/87.59	54.05	0.5200	78.90
LLaVA-NeXT§	79.49/81.24	79.02/80.54	43.17	0.8236	70.68	78.93/79.68	78.44/79.35	47.65	0.6077	70.00
Qwen2.5-VL§	87.03/89.01	87.01/88.69	50.04	0.6306	86.39	86.73/87.45	86.70/86.98	54.06	0.5610	79.98
MoLAN+	<b>88.02/89.94</b>	<b>87.96/89.90</b>	<b>52.30</b>	<b>0.5700</b>	<b>87.72</b>	<b>87.79/88.29</b>	<b>87.46/88.21</b>	<b>56.86</b>	<b>0.4909</b>	<b>82.02</b>

Table 3: The performance of the MoLAN+ on the MOSI and MOSEI. ◇ indicates our reproduced results. † indicates MSA denoising method. § denotes fine-tuning with LoRA. We perform significance testing on MMML and MoLAN+, with p-value of  $7.52 \times 10^{-4} < 0.05$  indicates significant differences. The best results are in bold.

Model	CH-SIMS		IEMOCAP	
	ACC <sub>2</sub>	F1	Weighted-F1	Macro-F1
MuT◇	68.49	55.68	65.07	65.38
SPECTRA◇	77.12	77.06	63.21	63.54
KuDA◇	78.54	78.41	64.00	63.86
SFTTR◇	79.22	79.20	64.05	61.54
ALMT	81.19	81.57	-	-
MMML◇	79.39	79.50	64.29	62.86
LLaVA-NeXT§	76.64	76.02	64.61	62.19
Qwen2.5-VL§	80.22	79.87	66.28	65.72
MoLAN+	<b>82.24</b>	<b>81.63</b>	<b>68.79</b>	<b>67.69</b>

Table 4: Performance of MoLAN+ on the SIMS and IEMOCAP. Significance testing in Appendix C.

of noise dynamic editing from the perspectives of visualization. Figure 3 shows the heatmap comparison between the original image and the image processed by the MoLAN framework. It can be clearly observed that the energy distribution of the heatmap of the original image is uniform, and the emotional information (the target person’s face and hands) is not prominent enough due to the noise in the background area. After the noise dynamic editing of the MoLAN framework, the energy distribution of the heatmap on the right changes significantly. For the background noise area, the overall energy level is reduced, effectively suppressing the noise. In contrast, energy levels in areas such as the face remain high, forming clear highlights. This is because these brighter regions correspond to visual information that is more relevant to the current sentiment analysis, and are therefore assigned higher attention weights by the model. Meanwhile, the denoising strength is computed at the block level,

which enables the model to distinguish local key information from noise more precisely. As a result, adjacent blocks may receive different scores, leading to noticeable value changes at block boundaries. Such local block-wise discontinuity is a natural consequence of block-level modeling, and further demonstrates that our method can suppress noise while preserving essential information in a fine grained manner.

## 5.4 MoLAN+ Experiments

We conduct a comprehensive evaluation of MoLAN+ on four datasets. As presented in Table 3 and Table 4, MoLAN+ consistently achieves SOTA performance across all datasets, surpassing both existing baseline models and recent MLLM-based approaches. This superior performance demonstrates that the noise-suppression cross attention and noise-driven contrastive learning modules effectively enhance the model’s discriminative capability and semantic alignment across modalities. By emphasizing emotion-relevant multimodal representations during feature extraction, the model is able to mitigate the influence of noisy information and maintain stable performance across diverse conditions. Moreover, when compared with the framework experiments, we observe that although other models benefit from the integration of the MoLAN framework, their results still fall short of the overall performance achieved by MoLAN+. These findings suggest a strong synergy between noise-suppression cross attention, noise-driven contrastive learning, and the MoLAN frame-

Model	CMU-MOSI			CMU-MOSEI		
	ACC <sub>2</sub>	F1	Corr	ACC <sub>2</sub>	F1	Corr
MuT	77.94	77.69	65.91	79.19	79.60	65.33
MuT+MoLAN	81.04	80.57	68.56	81.19	81.16	66.44
w/o DE	78.25	78.20	66.49	79.29	79.93	65.47
w/o MoLAN(v)	78.91	78.56	67.19	80.32	80.25	65.98
w/o MoLAN(a)	79.15	78.95	67.95	80.65	80.76	66.10
w/o MB	78.65	78.43	66.97	79.82	79.73	65.61
v(2D), a(2D)	79.48	79.68	68.19	80.20	80.59	66.04
v(1D), a(2D)	78.17	78.02	66.57	79.41	79.69	65.63
MoLAN <sup>+</sup>	89.94	89.90	87.72	88.29	88.21	82.02
w/o NC	89.69	89.68	87.59	88.16	87.96	81.85
w/o DC	89.74	89.71	87.64	88.20	88.01	81.88

Table 5: Ablation studies.

work, jointly contributing to the superior adaptability and effectiveness of MoLAN<sup>+</sup> in MSA tasks.

We conduct a comparison between MoLAN<sup>+</sup> and two representative denoising-based MSA models, JOSFD and t-HNE. The significance tests yield p-values of 0.01 and  $4.85 \times 10^{-5}$ , both of which are far below the 0.05 threshold, indicating that the performance improvement achieved by MoLAN<sup>+</sup> is statistically significant. This result suggests that traditional denoising strategies often over-filter multimodality signals, inadvertently removing critical semantic information and thereby degrading model performance. In contrast, MoLAN<sup>+</sup> employs a noise dynamic editing mechanism that selectively suppresses irrelevant noise while preserving emotion-relevant features, leading to higher accuracy across multiple datasets.

## 5.5 Ablation Study

As shown in the upper part of Table 5, we first conduct ablation experiments on two parts of MoLAN. **Effectiveness of noise dynamic editing.** We apply a uniform denoising strength to the modality information. Experimental results (w/o DE) show that the model performance degrades after removing dynamic editing. This indicates that uniform denoising strength cannot handle the differences in noise levels between different regions. Such denoising may result in insufficient denoising in high-noise areas, while low-noise areas may lose valuable information due to over-denoising. The experimental results further verify the necessity of dynamic editing, which adaptively adjusts denoising strength based on local noise characteristics to achieve a better balance between noise suppression and information preservation.

We independently apply the denoising framework to the visual modality (MoLAN(v)) and the audio modality (MoLAN(a)). The results show that although the single-modality denoising does not reach the performance level of joint multimodal denoising, it still achieves a significant improve-

ment compared with the model without the denoising framework. This finding indicates that the denoising mechanism also plays a positive role in single-modality denoising, while the collaborative removal during joint multimodal denoising further enhances the overall performance.

**Effectiveness of modality-aware block partitioning.** We conduct an ablation study on the modality-aware block partitioning strategy to examine its effectiveness. Specifically, we first apply a uniform one-dimensional block partitioning to all modalities (w/o MB), and then a uniform two-dimensional block partitioning to both the visual and audio modalities (v(2D), a(2D)). The results show a noticeable decline in performance, indicating that the information distribution differs significantly across modalities. These findings suggest that only a targeted block partitioning strategy can effectively localize and suppress modality-specific noise. Furthermore, when we assign one-dimensional block partitioning to the visual modality and two-dimensional block partitioning to the audio modality (v(1D), a(2D)), the model performance also decreases. This further verifies that the adopted configuration, namely 2D block partitioning for visual and 1D block partitioning for audio, better aligns with the properties of each modality and thus enables more efficient denoising.

As shown in the lower part of Table 5, we perform ablation experiments for MoLAN<sup>+</sup>.

**Effectiveness of noise-suppressed cross attention.** We remove the noise-suppression cross-attention mechanism (w/o NC) to examine its effect. The experimental results show a performance drop, indicating that denoising information plays a crucial role in cross-attention computation. Specifically, when the mask matrix is removed, the model’s ability to perceive and suppress noise weakens, leading to an overall decline in performance. These results suggest that explicitly incorporating denoising information into the mask matrix effectively guides the cross-attention mechanism to focus on informative regions, reduce noise interference, and enhance the model’s noise resistance.

**Effectiveness of denoising-driven contrastive learning.** We further remove the denoising-driven contrastive learning module (w/o DC) to validate its contribution. The results reveal a drop in performance, suggesting that this module is essential for learning stable and discriminative modality representations. In particular, the contrastive objective establishes a semantic constraint between denoised

and original features, enabling the model to maintain both distinctiveness and consistency of feature distributions under noisy conditions. Once this mechanism is removed, such relational constraints are weakened, making it difficult for the model to separate noise from informative signals in the latent space. Therefore, the denoising-driven contrastive learning not only improves cross-modal alignment but also enhances the model’s universality ability in complex noisy environments. Additional analyses and case studies in Appendix D and Appendix E.

## 6 Conclusion

We propose MoLAN, a unified modality-aware noise dynamic editing framework that partitions modalities into blocks and dynamically assigns denoising strength based on each block’s noise level and semantic relevance. MoLAN is plug-and-play and can be integrated into various MSA models and MLLMs to improve performance. Built on MoLAN, MoLAN<sup>+</sup> further introduces noise suppression cross-attention mechanism and denoising-driven contrastive learning to emphasize essential information. Extensive experiments demonstrate the effectiveness of MoLAN and MoLAN<sup>+</sup>. Overall, our work provides a practical pathway for enhancing robustness in multimodal systems.

## Limitations

Our current evaluation mainly covers several representative MSA benchmarks to demonstrate the generality and integrability of MoLAN and MoLAN<sup>+</sup>. However, we don’t yet conduct systematic evaluations under more challenging settings, such as cross-domain generative tasks or more complex real-world scenarios. In addition, performance may depend on the quality of the guiding modalities. When these modalities are biased, noise editing may fail to provide stable and accurate regulation, reducing performance. This challenge becomes more pronounced in open environments with more complex noise patterns, requiring stronger robustness and adaptability. Therefore, comprehensive experiments and analyses on these broader scenarios remain to be further complemented.

## Acknowledgments

The work is supported by the National Natural Science Foundation of China (No. 62272092, No. 62172086), National Science Foundation for Young Scientists of China (No. 62502081), and the

Fundamental Research Funds for the Central Universities under Grants (N2523011, N25XQD004).

## References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Ebrahim (Abe) Kazemzadeh, Emily Mower Provost, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. 2008. *Iemocap: interactive emotional dyadic motion capture database*. *Language Resources and Evaluation*, 42:335–359.
- Zebang Cheng, Zhi-Qi Cheng, Jun-Yan He, Jingdong Sun, Kai Wang, Yuxiang Lin, Zheng Lian, Xiaojiang Peng, and Alexander G Hauptmann. 2024. Emotion-llama: Multimodal emotion recognition and reasoning with instruction tuning. *Advances in Neural Information Processing Systems*, 37:110805–110853.
- Cunhang Fan, Kang Zhu, Jianhua Tao, Guofeng Yi, Jun Xue, and Zhao Lv. 2024. Multi-level contrastive learning: Hierarchical alleviation of heterogeneity in multimodal sentiment analysis. *IEEE Transactions on Affective Computing*.
- Xinyu Feng, Yuming Lin, Lihua He, You Li, Liang Chang, and Ya Zhou. 2024. Knowledge-guided dynamic modality attention fusion framework for multimodal sentiment analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14755–14766.
- Jiwei Guo, Jiajia Tang, Weichen Dai, Yu Ding, and Wanzeng Kong. 2022. *Dynamically adjust word representations using unaligned multimodal information*. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM ’22, page 3394–3402, New York, NY, USA. Association for Computing Machinery.
- Zirun Guo, Tao Jin, and Zhou Zhao. 2024. Multimodal prompt learning with missing modalities for sentiment analysis and emotion recognition. In *ACL*, pages 1726–1736.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022a. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Guimin Hu, Ting-En Lin, Yi Zhao, Guangming Lu, Yuchuan Wu, and Yongbin Li. 2022b. *UniMSE: Towards unified multimodal sentiment analysis and emotion recognition*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7837–7851, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

- Xun Jiang, Xing Xu, Huimin Lu, Lianghua He, and Heng Tao Shen. 2024. Joint objective and subjective fuzziness denoising for multimodal sentiment analysis. *IEEE Transactions on Fuzzy Systems*.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024a. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*.
- Mingcheng Li, Dingkan Yang, Yuxuan Lei, Shunli Wang, Shuaibing Wang, Liuzhen Su, Kun Yang, Yuzheng Wang, Mingyang Sun, and Lihua Zhang. 2024b. A unified self-distillation framework for multimodal sentiment analysis with uncertain missing modalities. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 10074–10082.
- Tianyi Li and Daming Liu. 2025. Mpid: A modality-preserving and interaction-driven fusion network for multimodal sentiment analysis. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 4313–4322.
- Zuocheng Li and Lishuang Li. 2025. t-hne: A text-guided hierarchical noise eliminator for multimodal sentiment analysis. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 2834–2844.
- Ronghao Lin and Haifeng Hu. 2022. Multimodal contrastive learning via uni-modal coding and cross-modal prediction for multimodal sentiment analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 511–523.
- Yuanyuan et al. Liu. 2024. Noise-resistant multimodal transformer for emotion recognition. *International Journal of Computer Vision*, pages 1–21.
- Zhun Liu and Ying Shen. 2018. Efficient low-rank multimodal fusion with modality-specific factors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Long Papers)*.
- Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, AmirAli Bagher Zadeh, and Louis-Philippe Morency. 2018. [Efficient low-rank multimodal fusion with modality-specific factors](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2247–2256, Melbourne, Australia. Association for Computational Linguistics.
- Sijie Mai, Haifeng Hu, and Songlong Xing. 2020. Modality to modality translation: An adversarial representation learning and graph fusion network for multimodal fusion. In *AAAI*, volume 34, pages 164–172.
- Sijie Mai, Ying Zeng, and Haifeng Hu. 2023. Learning from the global view: Supervised contrastive learning of multimodal representation. *Information Fusion*, 100:101920.
- Liang Shi, Fuyong Xu, Ru Wang, Yongqing Wei, Guangjin Wang, Bao Wang, and Peiyu Liu. 2024. Information aggregate and sentiment enhance network to handle missing modalities for multimodal sentiment analysis. In *2024 ICME*, pages 1–6. IEEE.
- Gopendra Vikram Singh, Sai Vardhan Vemulapalli, Mauajama Firdaus, and Asif Ekbal. 2024. Deciphering cognitive distortions in patient-doctor mental health conversations: A multimodal llm-based detection and reasoning framework. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 22546–22570.
- Kaiwei Sun and Mi Tian. 2025a. Sequential fusion of text-close and text-far representations for multimodal sentiment analysis. In *Proceedings of the 31st international conference on computational linguistics*, pages 40–49.
- Kaiwei Sun and Mi Tian. 2025b. [Sequential fusion of text-close and text-far representations for multimodal sentiment analysis](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 40–49, Abu Dhabi, UAE. Association for Computational Linguistics.
- Licai Sun, Zheng Lian, Bin Liu, and Jianhua Tao. 2023. Efficient multimodal transformer with dual-level feature restoration for robust multimodal sentiment analysis. *IEEE Transactions on Affective Computing*, 15(1):309–325.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019a. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for computational linguistics. Meeting*, volume 2019, page 6558.
- Yao-Hung Hubert Tsai, Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019b. [Learning factorized multimodal representations](#). In *ICLR*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pages 9929–9939. PMLR.
- Yiwei Wei, Shaozu Yuan, Ruosong Yang, Lei Shen, Zhangmeizhi Li, Longbiao Wang, and Meng Chen. 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)*, pages 5240–5252.

- Zehui Wu, Ziwei Gong, Jaywon Koo, and Julia Hirschberg. 2024. Multimodal multi-loss fusion network for sentiment analysis. In *Proceedings of the 2024 NAACL*, pages 3588–3602.
- Tianshu Yu, Haoyu Gao, Ting-En Lin, Min Yang, Yuchuan Wu, Wentao Ma, Chao Wang, Fei Huang, and Yongbin Li. 2023. [Speech-text pre-training for spoken dialog understanding with explicit cross-modal alignment](#). In *Proceedings of the 61st ACL*, pages 7900–7913, Toronto, Canada. Association for Computational Linguistics.
- Wenmeng Yu, Hua Xu, Fanyang Meng, Yilin Zhu, Yixiao Ma, Jiele Wu, Jiyun Zou, and Kaicheng Yang. 2020. [Ch-sims: A chinese multimodal sentiment analysis dataset with fine-grained annotation of modality](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Wenmeng Yu, Hua Xu, Ziqi Yuan, and Jiele Wu. 2021. [Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):10790–10797.
- Ziqi Yuan, Wei Li, Hua Xu, and Wenmeng Yu. 2021. Transformer-based feature reconstruction network for robust multimodal sentiment analysis. In *Proceedings of the 29th ACM MM*, pages 4400–4407.
- Ziqi Yuan, Baozheng Zhang, Hua Xu, and Kai Gao. 2024. Meta noise adaption framework for multimodal sentiment analysis with feature noise. *IEEE Transactions on Multimedia*, 26:7265–7277.
- Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2017. [Tensor fusion network for multimodal sentiment analysis](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1103–1114, Copenhagen, Denmark. Association for Computational Linguistics.
- Amir Zadeh, Paul Pu Liang, Soujanya Poria, E. Cambria, and Louis-Philippe Morency. 2018a. [Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, and Louis-Philippe Morency. 2018b. Multi-attention recurrent network for human communication comprehension. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. 2016. [Mosi: Multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos](#). *ArXiv*, abs/1606.06259.
- Jiandian Zeng, Tianyi Liu, and Jiantao Zhou. 2022a. Tag-assisted multimodal sentiment analysis under uncertain missing modalities. In *Proceedings of the 45th ACM SIGIR*, pages 1545–1554.
- Jiandian Zeng, Jiantao Zhou, and Tianyi Liu. 2022b. Mitigating inconsistencies in multimodal sentiment analysis under uncertain missing modalities. In *Proceedings of the 2022 EMNLP*.
- Jiandian Zeng, Jiantao Zhou, and Tianyi Liu. 2022c. Robust multimodal sentiment analysis via tag encoding of uncertain missing modalities. *IEEE Transactions on Multimedia*.
- Haoyu Zhang, Yu Wang, Guanghao Yin, Kejun Liu, Yuanyuan Liu, and Tianshu Yu. 2023. Learning language-guided adaptive hyper-modality representation for multimodal sentiment analysis. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 756–767. Association for Computational Linguistics.
- Yiqun Zhang, Fanheng Kong, Peidong Wang, Shuang Sun, SWangLing SWangLing, Shi Feng, Daling Wang, Yifei Zhang, and Kaisong Song. 2024. Sticker-conv: Generating multimodal empathetic responses from scratch. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7707–7733.
- Yiqun Zhang, Xiaocui Yang, Xingle Xu, Zeran Gao, Yijie Huang, Shiyi Mu, Shi Feng, Daling Wang, Yifei Zhang, Kaisong Song, and 1 others. 2026. Affective computing in the era of large language models: A survey from the nlp perspective. *Knowledge-Based Systems*, page 115411.
- Miao Zhou, Lina Yang, Thomas Wu, Dongnan Yang, and Xinru Zhang. 2025. Dual-path dynamic fusion with learnable query for multimodal sentiment analysis. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 11366–11376.
- Aoqiang Zhu, Min Hu, Xiaohua Wang, Jiaoyun Yang, Yiming Tang, and Fuji Ren. 2024. Kebr: Knowledge enhanced self-supervised balanced representation for multimodal sentiment analysis. In *Proceedings of the 32nd ACM MM*, pages 5732–5741.
- Zhouan Zhu, Shangfei Wang, Yuxin Wang, and Jiaqiang Wu. 2025. Integrating visual modalities with large language models for mental health support. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 8939–8954.
- Yan Zhuang, Yanru Zhang, Zheng Hu, Xiaoyue Zhang, Jiawen Deng, and Fuji Ren. 2024. [Glomo: Global-local modal fusion for multimodal sentiment analysis](#). In *Proceedings of the 32nd ACM MM*, page 1800–1809.

## A Pilot Study

The initial modality feature embedding, serving as the encoded representations of each modality, form the foundation of MSA. However, existing studies generally overlook the denoising of these encoded features and fail to investigate the potential impact of noise within the initial feature representations. Noise within the initial features may interfere with the accuracy of MSA and hinder the model’s ability to capture critical information. By enhancing the quality of the initial features through effective denoising, the overall performance of the model can be further improved. Therefore, we design and conduct a series of validation experiments to systematically analyze the presence and influence of noise in the initial modality features.

Ratio	MulT		MMML	
	CMU-MOSI	CMU-MOSEI	CMU-MOSI	CMU-MOSEI
10%	77.89	79.36	88.56	88.03
20%	79.42 ↑	79.27 ↓	88.63 ↑	86.96 ↓
30%	77.49 ↓	79.44 ↑	89.02 ↑	87.89 ↑
40%	79.57 ↑	80.10 ↑	88.63 ↓	87.84 ↓

Table 6: Multimodal Noise. ↓ indicates a performance drop compared to the previous row, while ↑ indicates an improvement. The metric used is  $ACC_2$ .

### A.1 Random Masking Strategy

To investigate the impact of noise in the initial multimodal feature representations, we design a random masking strategy. Specifically, we randomly mask a portion of the elements in the feature embeddings according to a predefined masking ratio. In this way, we observe how the model performs under conditions of partial information loss.

$$M_{ij} = \begin{cases} 0, & \text{if } R_{ij} < p \\ 1, & \text{otherwise} \end{cases} \quad R_{ij} \sim \mathcal{U}(0, 1) \quad (18)$$

$$\tilde{\mathbf{F}} = \mathbf{F} \odot \mathbf{M} \quad (19)$$

where  $\mathbf{F}$  denote the feature embedding matrix, and  $\mathbf{M}$  the corresponding mask matrix. Each element  $R_{ij}$  is a random number sampled from a uniform distribution over the interval  $[0, 1]$ . The masking ratio is defined by  $p \in [0, 1]$ .

If the initial feature information contains little noise, an intuitive inference is that the model performance will degrade as the masking ratio increases. Therefore, we gradually increase the masking ratio and observe the trend in model performance, thereby indirectly reflecting the level of potential noise in the initial multimodal features.

Ratio	CMU-MOSI			CMU-MOSEI		
	Text	Visual	Audio	Text	Visual	Audio
10%	75.16	79.07	79.57	79.95	79.39	79.03
20%	71.91 ↓	80.16 ↑	77.77 ↓	75.06 ↓	79.42 ↑	79.13 ↑
30%	70.09 ↓	80.06 ↓	78.92 ↑	73.62 ↓	79.11 ↓	79.17 ↑
40%	66.83 ↓	79.41 ↓	78.43 ↓	73.46 ↓	78.78 ↓	77.98 ↓

Table 7: Unimodal Noise. Text, visual, and audio represent the modality being masked, while the other modality features remain unchanged.

### A.2 Multimodal Noise

We first apply the random masking strategy to the overall modality-level features and conduct experiments on both the visual and audio modalities with the same masking ratio. As shown in Table 6, the results show that the model performance improves as the masking ratio increases. This phenomenon indicates that the removed portions of the features may contain more noise than useful information. Therefore, eliminating these noisy components allows the model to focus on more critical representations, leading to better performance. These findings provide preliminary evidence that a considerable amount of noise exists in the initial modality features and that not all encoded information contributes positively to the task. Consequently, applying an effective denoising strategy improves the quality of modality representations and ultimately enhances the overall performance of MSA task.

### A.3 Unimodal Noise

To further investigate the impact of noise within different modalities, we apply masking to each individual modality separately. Using the MulT model as the base framework, we impose varying levels of masking on the text, audio, and visual modalities. We observe how model performance changes as the degree of masking increases for each single modality. This allows us to analyze the relative level of noise in the initial features of each modality and its effect on sentiment analysis.

As shown in Table 7, the audio and visual modalities exhibit performance trends that differ from those of the text modality. As the masking ratio increases, the performance of the audio and visual modalities shows fluctuating improvements. However, performance consistently declines in the text modality. This observation suggests that the initial text modality features are higher quality, containing less noise. Therefore, the text modality can serve as a reference standard. Moreover, the performance improvements in the audio and visual modalities are triggered at different masking ratios, further

validating the different distribution of noise across modalities. This observation suggests that the differences in noise characteristics between different modalities may lead to the loss of critical information in some modalities when a unified denoising strategy is applied to all modalities. Therefore, differentiated denoising strategies should be adopted for different modalities.

## B More Details

### B.1 Dataset and Metrics

These datasets encompass both Chinese and English corpora, incorporating text, audio, and visual modalities across diverse contexts such as monologues, dialogues, and film clips, with annotations covering both discrete and continuous emotional dimensions. For the CMU-MOSI and CMU-MOSEI datasets, we follow prior works (Wu et al., 2024; Jiang et al., 2024; Zhuang et al., 2024) to evaluate both regression and classification tasks. For regression, we report the **Mean Absolute Error (MAE)** and **Correlation coefficient (Corr)**. For classification, we calculate the **Acc<sub>2</sub>** and **F1** scores for both the including zero sentiment scores as positive ( $ACC_{2Has0} / F1_{Has0}$ ) and the ignoring zero sentiment scores ( $ACC_{2Non0} / F1_{Non0}$ ). Additionally, we report **ACC<sub>7</sub>**. For the CH-SIMS dataset, we adopt **ACC<sub>2</sub>** and **F1** scores as evaluation metrics for the classification task. For the IEMOCAP dataset, we use **Weighted-F1** and **Macro-F1** scores to assess the performance of the classification task.

### B.2 Baselines

**Framework:** We select influential and reproducible multimodal sentiment analysis models as comparative baselines: **MuT** (Tsai et al., 2019a), **SPECTRA** (Yu et al., 2023), **KuDA** (Feng et al., 2024), **SFTTR** (Sun and Tian, 2025b), **MMML** (Wu et al., 2024). These models cover different research paradigms such as multimodal feature alignment, cross-modality attention modeling, and knowledge enhancement. With the rapid rise of the Multimodal Large Language Model (MLLM), unified understanding and reasoning across modalities has become a key direction in multimodal research. To verify the applicability and stability of our framework within this new paradigm, we further introduce **LLaVA-NeXT** (Li et al., 2024a) and **Qwen2.5-VL** (Bai et al., 2025) as representative models for comparison, further comprehensively examining the framework’s universality per-

formance.

**Models:** In addition to the baseline models used in the framework experiments, we also compare a series of representative approaches for MSA: **TFN** (Zadeh et al., 2017): Models intra- and inter-modality dynamics in an end-to-end manner. **LMF** (Liu et al., 2018): Employs low-rank tensors for efficient multimodal fusion, reducing computational complexity. **MFMM** (Tsai et al., 2019b): Jointly optimizes generative and discriminative objectives on multimodal data and labels. **Self-MM** (Yu et al., 2021): Designs a self-supervised label generation module to automatically obtain unimodal supervision signals. **UniMSE** (Hu et al., 2022b): Unifies multimodal sentiment analysis and emotion recognition tasks from multiple perspectives. **CHFNN** (Guo et al., 2022): Based on a Transformer architecture, it efficiently fuses unaligned multimodal sequences. **ALMT** (Zhang et al., 2023): Introduces an adaptive hyper-modality learning module to suppress irrelevant or conflicting information under the guidance of language. **EMT** (Sun et al., 2023): Enhances model robustness in scenarios with incomplete modalities while maintaining efficiency and performance. **GLoMo** (Zhuang et al., 2024): Integrates local representations from each modality and combines them with global representations to enhance expressive power. **JOSFD** (Jiang et al., 2024): Incorporates fuzzy logic to model both subjective and objective fuzziness in sentiment information. **t-HNE** (Li and Li, 2025): Text-guided hierarchical denoiser improves sentiment analysis performance via two-stage denoising and contrastive learning mechanism. **t-HNE** is the **latest** MSA denoising method. **MMML** is the current **SOTA** MSA model.

### B.3 Implementation Detail

All experiments are conducted using the PyTorch framework on a hardware setup with 8 RTX A6000 GPUs. For the framework experiments, we strictly follow the hyperparameter settings reported in the original papers to ensure fair and consistent comparisons. In our methodological experiments, to ensure fair comparison, we maintain the same parameter size settings for our modality encoders as in previous studies (Wu et al., 2024). Model is trained using the AdamW optimizer to achieve stable and efficient convergence performance. Specifically, the learning rate is set to  $5e-6$  for the MOSI and MOSEI datasets,  $1e-5$  for the CH-SIMS dataset, and  $2e-8$  for the IEMOCAP dataset. In the frame-

Model	CMU-MOSI			CMU-MOSEI		
	ACC <sub>2</sub>	F1	Corr	ACC <sub>2</sub>	F1	Corr
MuT	77.94	77.69	65.91	79.19	79.60	65.33
Text guide	81.04	80.57	68.56	81.19	81.16	66.44
Visual guide	76.16	76.08	64.81	77.02	77.00	64.10
Audio guide	77.91	77.46	65.88	79.11	79.52	65.13

Table 8: Ablation study of denoising strength computation guidance.

work experiments, we first fine-tune the MLLM using LoRA (Hu et al., 2022a) on MSA datasets to adapt it to the task characteristics and stabilize model performance. We then integrate the fine-tuned model into the proposed framework to verify its effectiveness. In the method experiments, the MLLM results report the performance of the model after LoRA fine-tuning. To ensure the stability of the experimental results, we fix the random seed to 1 and run each experiment five times independently, reporting the average result.

## C Significance Testing

For Table 2, we perform significance testing on seven experimental groups, with p-value of  $4.26 \times 10^{-2}$ ,  $6.01 \times 10^{-4}$ ,  $2.95 \times 10^{-3}$ ,  $1.08 \times 10^{-3}$ ,  $1.19 \times 10^{-2}$ ,  $4.28 \times 10^{-3}$ , and  $4.45 \times 10^{-3}$ , all of which  $< 0.05$  indicate significant differences.

For Table 4, the significance testing on MMML and MoLAN<sup>+</sup>, with p-value ( $0.01 < 0.05$ ) indicating significant differences.

## D More Experiments

### D.1 Block Partitioning Strategy Selection

We adopt a block partitioning strategy based on the factor closest to the square root of the feature dimension, since not all feature dimensions are perfect squares. This design ensures a more balanced division of feature regions and avoids the two extremes of excessively large or excessively small blocks. Specifically, when the block is too large, it tends to mix effective information with noise within the same region, making it difficult for the denoising strength to accurately distinguish between the two. This may result in incomplete noise removal or excessive suppression of critical information. In contrast, when the block is too small, it over-segments semantic structures, weakens inter-block correlations. By comparison, the close to square root block partitioning strategy achieves a better trade-off between signal-noise separation and semantic integrity, allowing the denoising process to retain critical information while removing noise.

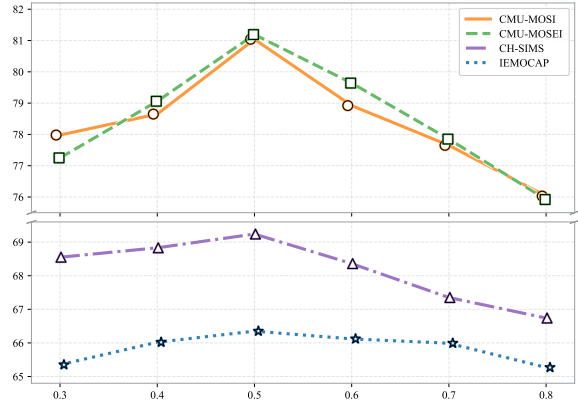


Figure 4: Performance comparison under different similarity thresholds  $\theta$ . The four curves represent experimental results on CMU-MOSI, CMU-MOSEI, CH-SIMS, and IEMOCAP datasets, respectively. The x-axis denotes the similarity threshold  $\theta$ , and the y-axis indicates the model performance.

As shown in Table 9, we design comparative experiments with different block sizes for both visual and audio modalities. The experimental results show that, in both visual and audio feature processing, excessively large or small block partitioning lead to a decline in model performance. In contrast, using the close to square root block partitioning strategy achieves the best balance between critical information preservation and noise suppression, allowing the model to capture critical features more effectively and maintain overall performance stability. Overall, the proposed close to square root block partitioning strategy aligns well with the characteristics of each modality and lays a solid foundation for subsequent denoising.

### D.2 Similarity Threshold Analysis

To determine the optimal similarity threshold  $\theta$  in the noise-suppressed cross attention mechanism, we conduct a series of comparative experiments across four benchmark datasets, as shown in Figure 4. The threshold  $\theta \in [0, 1]$  controls the activation of the denoising mask by determining which sub-blocks are preserved or suppressed based on similarity scores.

The experimental results demonstrate a consistent trend across all datasets: as  $\theta$  increases from 0.3 to 0.8, model performance initially improves and then declines. A smaller  $\theta$  (e.g., 0.3–0.4) allows excessive noisy features to pass through, resulting in incomplete denoising and suboptimal performance. Conversely, an overly large  $\theta$  (e.g., 0.7–0.8) overly suppresses critical information,

Model	CMU-MOSI			CMU-MOSEI		
	ACC <sub>2</sub>	F1	Corr	ACC <sub>2</sub>	F1	Corr
Visual(50,2)	78.10	77.81	66.05	79.65	79.36	65.20
Visual(25,4)	80.15	79.68	67.12	80.61	80.45	66.28
Visual(20,4)	81.04	80.57	68.56	81.19	81.16	66.44
Visual(10,5)	78.84	78.09	66.40	79.84	79.72	65.87
Visual(5,10)	77.13	77.06	65.69	79.12	79.50	65.27
Audio(3)	79.34	79.19	67.12	80.16	80.09	65.69
Audio(5)	80.56	80.44	68.29	80.88	81.00	66.12
Audio(15)	81.04	80.57	68.56	81.19	81.16	66.44
Audio(75)	80.55	80.48	68.17	80.11	80.05	65.57
Audio(125)	78.06	77.95	66.12	79.39	79.37	65.34

Table 9: Block Sizes Experiments. Take MulT+MoLAN on the CMU-MOSI and CMU-MOSEI dataset as an example. The audio feature dimension is  $[128, 375, 20]$ , and the visual feature dimension is  $[128, 500, 20]$ . The number after audio indicates the one-dimensional block size ( $j$ ), and the number after video indicates the two-dimensional block size ( $k, j$ ).

causing semantic loss and degraded performance. The performance peaks at ( $\theta = 0.5$ ), where the model achieves the best trade-off between noise suppression and information preservation. Therefore, we set  $\theta = 0.5$  as the default similarity threshold in all experiments.

### D.3 Denoising-Driven Contrastive Learning Analysis

To further verify the effectiveness of denoising-driven contrastive learning (DC) in modality denoising, we use two metrics in the field of contrastive learning: alignment and uniformity (Wang and Isola, 2020). Alignment measures the average distance between positive sample pairs in the embedding space, reflecting the model’s ability to cluster semantically similar samples. A smaller value indicates tighter intra-class compactness. Uniformity measures how evenly all samples are distributed on the unit hypersphere. A smaller value implies larger inter-class separation and better discrimination against noise.

As shown in Table 10, after incorporating DC, the model achieves lower alignment on both CMU-MOSI and CMU-MOSEI datasets, suggesting that positive samples cluster more tightly in the embedding space. Meanwhile, the uniformity score decreases significantly, indicating that negative samples become more separable and noise becomes easier to distinguish. These results indicate that DC effectively improves the model’s ability to distinguish critical information features from noisy redundant features, enabling the MSA model to learn purer modality representations.

Model	CMU-MOSI		CMU-MOSEI	
	Alignment ↓	Uniformity ↓	Alignment ↓	Uniformity ↓
w/o DC	0.9962	-0.5043	0.9993	-0.2674
with DC	0.9359	-2.2770	0.0198	-1.0030

Table 10: Ablation studies of Denoising-Driven Contrastive Learning. Alignment and Uniformity metrics on CMU-MOSI and CMU-MOSEI datasets.

### D.4 Denoising strength computation guidance

In our design, the denoising strength of each block is determined based on the text modality. To validate this choice, we conduct an ablation study comparing different guiding modalities, as shown in Table 8. The results show that using text as the guidance yields the best performance on two datasets. In contrast, using visual or audio features as the guidance leads to a performance drop, indicating that these modalities introduce more noise. These findings are consistent with the results observed in pilot study (Table 7), further confirming that text-guided denoising offers more reliable representations for MSA.

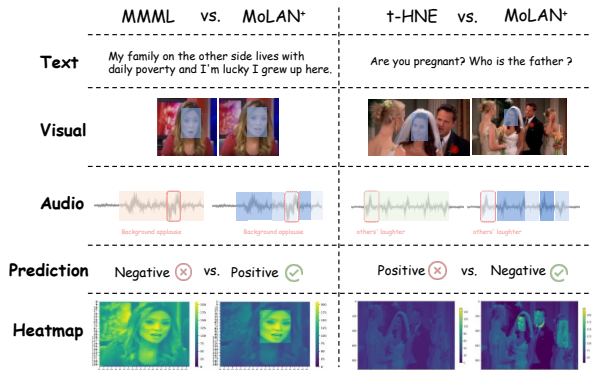


Figure 5: Case Study. The blue area in the visual modality represents the target person. The red, green, and blue colors in the audio modality represent the attention distribution of MMML, t-HNE, and MoLAN+ on different audio segments, respectively. The red boxes in the audio mark the noisy segments. The heatmap shows the model’s attention strength in different visual regions.

### E Case Study

As shown in Figure 5, we present two representative case studies comparing MoLAN+ with the original SOTA model MMML and the latest denoising model t-HNE. In the left, we compare MoLAN+ with MMML. It can be observed that MMML pays almost equal attention to different visual and audio segments without effectively suppressing noise interference. Consequently, the model fails to distinguish the emotionally relevant regions and mis-

interprets the sentiment. In contrast, MoLAN<sup>+</sup> focuses more accurately on the target speaker and the emotionally relevant regions, leading to a correct sentiment analysis. In the right, we compare MoLAN<sup>+</sup> with t-HNE. Although t-HNE performs denoising, it applies a uniform denoising intensity across all modalities, which results in the loss of critical information necessary for MSA. This over-smoothing effect causes the model to generate incorrect sentiment. Conversely, MoLAN<sup>+</sup> adaptively balances noise suppression and information retention through its dynamic noise-editing mechanism, thereby capturing emotion-relevant cues and achieving a more reliable analysis.