

# Learning Noise-Robust Joint Representation for Multimodal Emotion Recognition under Incomplete Data Scenarios

Qi Fan\*  
Inner Mongolia University  
Hohhot, China  
fanqi1203@foxmail.com

Haolin Zuo\*  
Inner Mongolia University  
Hohhot, China  
zuohaolin\_0613@163.com

Rui Liu†  
Inner Mongolia University  
Hohhot, China  
liurui\_imu@163.com

Zheng Lian  
Institute of Automation, Chinese  
Academy of Sciences  
Beijing, China  
lianzheng2016@ia.ac.cn

Guanglai Gao  
Inner Mongolia University  
Hohhot, China  
csggl@imu.edu.cn

## Abstract

Multimodal emotion recognition (MER) in practical scenarios is significantly challenged by the presence of missing or incomplete data across different modalities. To overcome these challenges, researchers have aimed to simulate incomplete conditions during the training phase to enhance the system's overall robustness. Traditional methods have often involved discarding data or substituting data segments with zero vectors to approximate these incompletenesses. However, such approaches neither accurately represent real-world conditions nor adequately address the issue of noisy data availability. For instance, a blurry image cannot be simply replaced with zero vectors, while still retaining information. To tackle this issue and develop a more precise MER system, we introduce a novel noise-robust MER model that effectively learns robust multimodal joint representations from noisy data. This approach includes two pivotal components: firstly, a noise scheduler that adjusts the type and level of noise in the data to emulate various realistic incomplete situations. Secondly, a Variational AutoEncoder (VAE)-based module is employed to reconstruct these robust multimodal joint representations from the noisy inputs. Notably, the introduction of the noise scheduler enables the exploration of an entirely new type of incomplete data condition, which is impossible with existing methods. Extensive experimental evaluations on the benchmark datasets IEMOCAP and CMU-MOSEI demonstrate the effectiveness of the noise scheduler and the excellent performance of our proposed model. Our project is publicly available on [https://github.com/WooyooHL/Noise-robust\\_MER](https://github.com/WooyooHL/Noise-robust_MER).

\*Equal contributions.

†Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MRAC '24, November 1, 2024, Melbourne, VIC, Australia.

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

ACM ISBN 979-8-4007-1203-6/24/11

<https://doi.org/10.1145/3689092.3689411>

## CCS Concepts

• **Computing methodologies** → *Computer vision*; • **Human-centered computing** → *HCI design and evaluation methods*.

## Keywords

Multimodal emotion recognition, incomplete modalities

### ACM Reference Format:

Qi Fan, Haolin Zuo, Rui Liu, Zheng Lian, and Guanglai Gao. 2024. Learning Noise-Robust Joint Representation for Multimodal Emotion Recognition under Incomplete Data Scenarios. In *Proceedings of the 2nd International Workshop on Multimodal and Responsible Affective Computing (MRAC '24)*, November 1, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3689092.3689411>




Figure 1: The incompleteness of multimodal data in realistic scenarios poses a significant challenge to multimodal emotion recognition.

## 1 Introduction

Multimodal emotion recognition (MER) aims to take multimodal signals, including text, audio, and visual, as input to predict the emotion category [41]. As a cutting-edge technology, it is widely used in various scenarios, including virtual intelligent assistants, robot customer service, and other applications, offering tailored and empathetic user experiences by responding appropriately to the emotional cues of users. In the research of MER, remarkable performance depends heavily on the complete multimodal data and robust joint representation learning [15]. However, in realistic

scenarios, data often becomes incomplete due to two main cases: either data is **absolutely absence** caused by sensor malfunction [39], or it is **partially incompleteness** attributed to diminished network bandwidth or various forms of noise interference [25, 40], etc. We collectively call these data noisy or incomplete data, which present huge challenges for multimodal emotion recognition [39].

In recent years, researchers have tried to enhance the robustness of MER by simulating incomplete data during training. Previous researchers proposed two methods to simulate incomplete data according to the usual intuition: 1) set feature vectors of missing data to zero: in some situations, zero is a value that signifies “no information” or “lack of information” [39, 41], 2) discard data randomly with a predefined probability: dropped data can be regarded as completely missing [5, 23]. On this basis, MER work in the face of incomplete data focuses on two main areas: 1) missing data completion [4, 9, 23, 35] and 2) multimodal joint representation learning with available data [13, 32]. For example, Cai et al. [4] proposed to use adversarial learning to generate missing modality images. Zeng et al. [38] proposed an ensemble learning method to use several models to solve the missing problem jointly. Zhao et al. [39] proposed a Missing Modality Imagination Network (MMIN) using AutoEncoder and cycle consistency construction to learn joint representations while predicting missing modalities. Zuo et al. [41] proposed introducing the modality invariant feature into MMIN to learn robust multimodal joint representations. Liu et al. [26] combined contrastive learning and invariant features to imagine the feature of the missing modality. The above work lays a solid foundation for MER work under incomplete data. Note that the framework that combines both *missing data completion* and *multimodal joint representation learning* has become the mainstream scheme in the research of this area.

However, the above approach faces two main issues, including 1) the method of simulating incomplete data is unreasonable and impractical: traditional data setup methods set parts of the feature to zero vectors or discard some data directly at a set percentage. These methods neither reliably simulate real-world scenarios nor reserve the availability of incomplete data. Moreover, assuming the data contains some type or level of noise, is not it better to try to mine useful information from it rather than just discard it? and 2) the structure is redundant: traditional model structure learns a multimodal joint representation in the process of completing the missing data [41], and such a process tends to cause the errors generated in the first step to have a negative impact on the second step. Whether robust multimodal joint feature representations can be learned directly from noisy data?

To answer the above two questions, we propose a novel Noise-robust MER model, termed NMER. Specifically, we design a noise scheduler at the embedding level. It creates noisy training and testing data by adding various types and intensities of noise to the embedding, thus simulating the influence of diverse incomplete situations in realistic scenarios. It is worth mentioning that we explored an incomplete condition of three new modalities through our noise scheduler, which was impossible to achieve in previous works. Then, we present a Variational AutoEncoder (VAE) [22, 34] based multimodal joint representation learning network to reconstruct robust multimodal joint representations from the noisy data. In this way, we simulate incomplete data in realistic scenarios,

make full use of the valuable information of the existing noisy data, and then leverage the powerful generative capabilities of VAE to reconstruct robust multimodal joint representations from noisy data and achieve multimodal emotion recognition. We conduct experiments on the benchmark dataset IEMOCAP [3] and CMU-MOSEI [24], which are widely used. Our code has been released in the supplementary material.

The main contributions of this work are as follows:

- We propose a noise-robust MER model, termed NMER, to generate robust multimodal joint representations under noisy incomplete data and proceed with multimodal emotion recognition.
- We explore a new method to simulate realistic data under various incomplete conditions utilizing a noise scheduler while enabling a new condition with the incompleteness of all three modalities.
- Experimental results under various noise types and intensity conditions show that our NMER outperforms most baselines and demonstrates robustness in the face of incomplete data.

## 2 NMER: Methodology

### 2.1 Overall Architecture

The overall architecture of NMER is illustrated in Figure 2, which consists of 1) *Noise Scheduler*; 2) *VAE-based Network*; and 3) *Classifier*. Specifically, we first employ our *Noise Scheduler* to add configurable noise to original embeddings to get incomplete data and send them into the *VAE-based Network*. Then, the *VAE-based Network* seeks to extract the useful features to reconstruct the robust multimodal joint representation. At last, the joint representation will be fed into *Classifier* to predict the final emotional result.

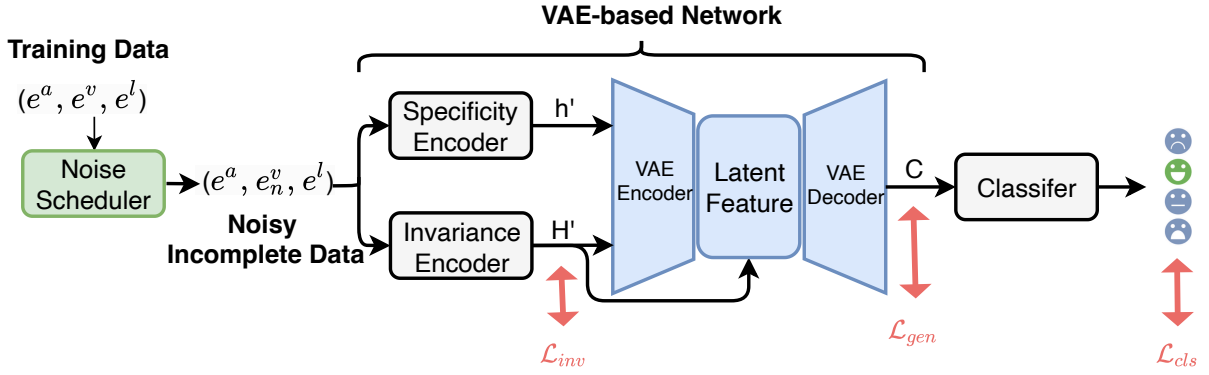
### 2.2 Noise Scheduler

#### 2.2.1 Rationale and Noise Selection.

Inspired by the previous work [17], we can obtain a blurred picture by gradually adding Gaussian noise to the picture. Similarly, we can obtain a “noisy embedding” by adding noise to the embedding [30, 33], which enhances the robustness of the model while simulating the modality incomplete condition. We examine two common noise types as examples: Gaussian noise, as described by C. F. Gauss in [11], and impulse noise, following the research of Kim et al. [20]

Gaussian noise is extensively employed to simulate errors and disturbances occurring in natural and technical processes due to its distinct statistical characteristics, namely a well-defined mean  $\mu$  and standard deviation  $\sigma$  that conform to a normal distribution. The selection of this type of noise is grounded in a crucial observation: deviations influenced by various factors tend to approximate a normal distribution in most natural and technological systems. Consequently, the use of Gaussian noise in practical applications not only maintains statistical rigor but also ensures that models exhibit greater robustness and adaptability when confronted with the complexities of real-world data [2, 14].

Impulse noise is another typical noise form that often appears in signal processing. Due to its ability to simulate brief and intense disturbances, it is widely used to model sudden anomalies in data transmission or sensor inputs. Additionally, impulse noise can also be employed to enhance the capability of the model to recognize and process extreme data points [27, 29].



**Figure 2: The structure of our NMER model, including noise scheduler, VAE-based network, and the classifier.  $\mathcal{L}_{inv}$  stands the invariant loss,  $\mathcal{L}_{gen}$  is the generation loss,  $\mathcal{L}_{cls}$  refers to the classification loss.**

The integration of noise at the embedding level is a pivotal aspect of simulating realistic scenarios. Previous works have proved its effectiveness and robustness facing real situations. For the incomplete modality problem, we can choose various techniques for corrupting original data, such as text (masking, reversing word order), audio (adding reverberation or noise), and video (introducing frame loss or blur). However, quantifying the specific impact of such corruption on each modality is quite difficult. For instance, determining the equivalent effect of adding a certain decibel level of noise to audio, or masking corresponding words, if a frame in a video is completely missing, is particularly difficult. Therefore, we introduce noise on the embedding to strike a balance between realism and analytical feasibility.

Based on the reasons above, we select those two types of noise on the embedding level, bring quantitative loss of information through control of the noise level, and simulate the negative impact on information from noise in reality. To be clear, our noise scheduler is designed to separate the processes of generating and adding noise. Users have the freedom to select or customize any type of noise, provided it can be created as a tensor using a programming language. This design ensures that the choice of noise type does not impact the addition process, offering users significant flexibility to adapt to various data processing needs.

### 2.2.2 Noisy Embedding Construction.

Following the diffusion process in [17], we construct noisy embedding through a similar process. Given complete multimodal data samples, encompassing acoustic, visual, and lexical components (denoted as  $a$ ,  $v$ , and  $l$  respectively), the embeddings are noted as  $E^a$ ,  $E^v$ , and  $E^l$ . We use  $E^a$  as an example to illustrate the process of making incomplete data.

For constructing the noisy embedding using Gaussian noise, we build a noise schedule across time steps represented by  $\beta_s$ :

$$\beta_s = (\beta_1, \dots, \beta_T) \quad (1)$$

It initializes a sequence of values starting from a small positive number  $\beta_{start}$  (denoting a low noise level) and progressively increasing towards a larger number  $\beta_{end}$  (denoting a high noise level) after  $T$  steps. The total number of steps  $T$  influences the gradual

transition of data toward a noise-dominated state, where  $T$  should be an integer greater than zero. The  $\beta_{start}$  and  $\beta_{end}$  generally refer to previous work [17]. The embedding is noted as  $E_0^a$  at time step 0. We then use Algorithm 1 to generate the noisy embedding  $E_T^a$ .

---

#### Algorithm 1: Process for Generating Noisy Embedding $E_T^a$

---

**Input:** Original embedding  $E_0^a$ , number of steps  $T$ , noise schedule parameters  $\beta_1, \beta_2, \dots, \beta_T$

**for**  $t = 1$  **to**  $T$  **do**

    Sample noise vector  $\epsilon_t \sim \mathcal{N}(0, 1)$ ;

    Update  $E_t^a \leftarrow \sqrt{1 - \beta_t} E_{t-1}^a + \sqrt{\beta_t} \epsilon_t$ ;

**end**

**Output:** Noisy embedding  $E_T^a$

---

However, this method is computationally intensive. To optimize this process for directly obtaining  $E_T^a$  from  $E_0^a$  without iterating through each intermediate step from  $E_1$  to  $E_{T-1}$ , we precompute the cumulative product of  $1 - \beta$  terms, denoted as  $\bar{\alpha}_t$ :

$$\bar{\alpha}_T = \prod_{t=1}^T (1 - \beta_t) \quad (2)$$

It significantly reduces computational complexity, which reflects the total variance retained in the data up to step  $T$ . Utilizing  $\bar{\alpha}_t$ , the noisy embedding  $E_T^a$  is obtained from  $E_0^a$  with a single update:

$$E_T^a = \sqrt{\bar{\alpha}_T} E_0^a + \sqrt{1 - \bar{\alpha}_T} \epsilon \quad (3)$$

where  $\epsilon$  is a freshly sampled noise vector from a Gaussian distribution,  $\epsilon \sim \mathcal{N}(0, 1)$ . This method enables an efficient bypass of the sequential update steps, directly synthesizing the noisy vector at any time step  $T$  that we desired.

For impulse noise, the sampled noise vector  $\epsilon$  consists only of random values 1s and -1s. An appearance frequency  $p$  is employed to adjust its values, setting  $1 - p$  percent of them to zero and leaving the remaining  $p$  percent unchanged. This adjustment introduces noise to  $p$  percent of data. Then we still use Equation 3 to get  $E_T^a$ .

### 2.2.3 Noise Intensity Calculation.

In the noisy embedding construction process, the  $\beta_{\text{start}}$ ,  $\beta_{\text{end}}$ , and  $T$  control the intensity of noise jointly. Therefore, we better quantify the noise intensity by calculating the signal-to-noise ratio (SNR) between the original and noisy data, in other words, the intensity ratio between noise and raw data.

In Equation 3, we regard  $\sqrt{\bar{\alpha}_T}$  and  $\sqrt{1 - \bar{\alpha}_T}$  as the power of the original data part and the noise part, separately. According to Equation 2 and 5, it's easy to know that when  $\sigma^2$  is invariant, the SNR value shares the opposite trend with  $T$  growing. Hence we can obtain a higher level of noise by increasing the value of  $T$ .

$$\text{SNR} = \frac{P_{\text{original}}}{P_{\text{noise}}} = \frac{\bar{\alpha}_T \sigma^2}{1 - \bar{\alpha}_T} \quad (4)$$

where the  $\sigma_2$  is the variance of the original data. More intuitively, SNR is always represented in decibels(dB):

$$\text{SNR}_{dB} = 10 \log_{10} \left( \frac{P_{\text{original}}}{P_{\text{noise}}} \right) = 10 \log_{10} \left( \frac{\bar{\alpha}_T \sigma^2}{1 - \bar{\alpha}_T} \right) \quad (5)$$

According to Equation 2 and 5, it's easy to know that when  $\sigma^2$  is invariant, the SNR value shares the opposite trend with  $T$  growing. Hence we can obtain a higher level of noise by increasing the value of  $T$ . Due to the variance differences across different datasets and modalities, we select a distinct T-value for each modality to maintain a uniform noise level. For instance, within the CMU-MOSEI dataset, we adjust the  $T$ -value for  $a$ ,  $v$ , and  $l$  modalities to [140, 60, 5] respectively to introduce noise at -10dB. By doing so, we are able to impose a similar level of noise on different datasets and modalities, which simulate diverse levels of information loss.

## 2.3 VAE-based Network

The VAE-based Network includes Specificity and Invariance Encoders [41] and the VAE Module. The Specificity Encoders, employing LSTM [18] and TextCNN [21] structures, are tasked with extracting modality-specific emotion features  $\mathbf{h}'$  from each modality, which is achieved by mapping the embeddings of different modalities into distinct vector spaces. Concurrently, the Invariance Encoder, utilizing a linear structure, extracts modality-invariant emotion features  $\mathbf{H}'$  across various modalities by mapping these embeddings into a unified vector space [26, 41].

Upon sending the concatenated features  $\mathbf{h}'$  and  $\mathbf{H}'$  into the VAE model, the VAE Encoder compresses and maps them, yielding mean and variance parameters within the latent space. This step effectively translocates the feature representation from the original data space to a probabilistic distribution in the latent space. A stochastic sampling process, facilitated by the reparameterization trick [22], then generates samples that reflect this latent space distribution. These samples embody the latent representations of the input features.

Subsequently, these latent variables pass through the decoder network. The decoder reconstructs the feature, remapping it to the original data space, thereby producing a reconstructed multimodal joint representation denoted as  $C$ . Notably, the invariant feature  $\mathbf{H}'$  plays a pivotal role during decoding, guiding the model to focus on the common emotional features in the multimodal data.

Noisy features are transformed into normal ones throughout the compression, sampling, and reconstruction phases, which provide the denoising effect. The final output, the multimodal joint representation  $C$ , is then fed into the Classifier to derive the result.

## 2.4 Loss Functions

As shown in Fig. 2, the total loss  $\mathcal{L}$  for NMER includes three parts:  $\mathcal{L} = \lambda_1 \mathcal{L}_{\text{gen}} + \lambda_2 \mathcal{L}_{\text{inv}} + \lambda_3 \mathcal{L}_{\text{cls}}$ , where  $\lambda_s$  are the balance factors.

Generation loss  $\mathcal{L}_{\text{gen}}$  aims to calculate the distance between the generation result  $C$  (from incomplete data) and the multimodal joint representation  $\hat{C}$  (from complete data).

Note that the  $\mathcal{L}_{\text{gen}}$  consists of two items, that are  $\mathcal{L}_{\text{kl}}$  in Equation 6 and  $\mathcal{L}_{\text{mse}}$  in Equation 7.  $\mathcal{L}_{\text{kl}}$  aims to make the hidden variables generated by the encoder conform to the standard normal distribution, while  $\mathcal{L}_{\text{mse}}$  seeks to make the generated multi-modal joint representation more similar to the target that extracted from complete data,

$$\mathcal{L}_{\text{kl}} = -\frac{1}{2} \left( \log \sigma^2 - \sigma^2 - \mu^2 + 1 \right) \quad (6)$$

$$\mathcal{L}_{\text{mse}} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 \quad (7)$$

where  $\sigma$  is the variance of the distribution of the latent vector while the  $\mu$  is the mean;  $N$  is the total number of the real values,  $y_i$  is the real value, and  $\hat{y}_i$  is the predicted value.

The invariant loss  $\mathcal{L}_{\text{inv}}$  shares the same spirit as in [41]. It adopts the MSE loss style to reduce the distance between the modality-invariant feature  $\mathbf{H}'$  during training (under incomplete conditions) and the real modality-invariant feature  $\mathbf{H}$ .  $\mathcal{L}_{\text{cls}}$  adopts the Cross-Entropy loss function to measure and minimize the disparities between the predicted and actual emotion category labels.

## 3 Experiments and Results

### 3.1 Data Setup

We perform experiments on the Interactive Emotional Dyadic Motion Capture (IEMOCAP) [3] and CMU Multimodal Opinion Sentiment and Emotion Intensity (CMU-MOSEI) dataset [37], which are both widely used in the research of MER.

**Table 1: Emotion distribution and division of two datasets.**

| Dataset   | Label    | Samples | Train Set | Test Set           |
|-----------|----------|---------|-----------|--------------------|
| IEMOCAP   | Happy    | 1636    | 4446      | 3342 (557 samples) |
|           | Angry    | 1103    |           |                    |
|           | Sad      | 1084    |           |                    |
|           | Neutral  | 1708    |           |                    |
| CMU-MOSEI | Positive | 14842   | 16265     | 4643               |
|           | Negative | 6066    |           |                    |

In the IEMOCAP dataset, training and testing conditions are consistently following the research [39]. For the training set, each

**Table 2: The results of the comparative studies on the IEMOCAP and CMU-MOSEI dataset, which employs two types of noise. The note “-10dB(Avg)” means the average performance across six common incomplete conditions under the noise intensity -10dB. “WA” stands for the weighted accuracy and “UA” refers to the unweighted accuracy. Bold values imply the best accuracy on that dataset.**

| Dataset                 | System                  | -10dB (Avg)    |               |               |               | -20dB (Avg)    |               |               |               | -30dB (Avg)    |               |               |               | -40dB (Avg)    |               |               |               |
|-------------------------|-------------------------|----------------|---------------|---------------|---------------|----------------|---------------|---------------|---------------|----------------|---------------|---------------|---------------|----------------|---------------|---------------|---------------|
|                         |                         | Gaussian noise |               | Impulse noise |               | Gaussian noise |               | Impulse noise |               | Gaussian noise |               | Impulse noise |               | Gaussian noise |               | Impulse noise |               |
|                         |                         | WA             | UA            | WA            | UA            | WA             | UA            | WA            | UA            | WA             | UA            | WA            | UA            | WA             | UA            | WA            | UA            |
| IEMOCAP                 | MEN                     | 0.7120         | 0.7175        | 0.7448        | 0.7453        | 0.6747         | 0.6714        | 0.7026        | 0.7133        | 0.6370         | 0.6349        | 0.6758        | 0.6850        | 0.6047         | 0.6004        | 0.6348        | 0.6364        |
|                         | MCTN                    | 0.7215         | 0.7431        | 0.7400        | 0.7531        | 0.6989         | 0.7003        | 0.7270        | 0.7314        | 0.6778         | 0.6712        | 0.7112        | 0.7157        | 0.6674         | 0.6790        | 0.6978        | 0.7041        |
|                         | MMIN                    | 0.7551         | 0.7640        | 0.7750        | 0.7870        | 0.7271         | 0.7394        | 0.7521        | 0.7671        | 0.7155         | 0.7209        | 0.7397        | 0.7516        | 0.7003         | 0.7067        | 0.7226        | 0.7352        |
|                         | IF-MMIN                 | 0.7543         | 0.7655        | 0.7738        | 0.7858        | <b>0.7345</b>  | <b>0.7493</b> | 0.7589        | 0.7719        | 0.7184         | 0.7225        | <b>0.7382</b> | 0.7509        | 0.7048         | 0.7155        | 0.7343        | 0.7452        |
|                         | <b>Ours</b>             | <b>0.7598</b>  | <b>0.7675</b> | <b>0.7773</b> | <b>0.7872</b> | 0.7307         | 0.7406        | <b>0.7651</b> | <b>0.7790</b> | <b>0.7197</b>  | <b>0.7284</b> | 0.7304        | <b>0.7532</b> | <b>0.7085</b>  | <b>0.7176</b> | <b>0.7391</b> | <b>0.7503</b> |
|                         | w/o VAE                 | 0.7382         | 0.7275        | 0.7509        | 0.7631        | 0.7004         | 0.7132        | 0.7467        | 0.7500        | 0.6850         | 0.6882        | 0.7136        | 0.7280        | 0.6793         | 0.6801        | 0.7133        | 0.7238        |
|                         | w/o $\mathcal{L}_{inv}$ | 0.7401         | 0.7338        | 0.7630        | 0.7712        | 0.7121         | 0.7293        | 0.7489        | 0.7549        | 0.7010         | 0.6942        | 0.7202        | 0.7276        | 0.6843         | 0.6805        | 0.7158        | 0.7208        |
|                         | CMU-MOSEI               | MEN            | 0.7285        | 0.6593        | 0.7341        | 0.6638         | 0.6966        | 0.6380        | 0.7010        | 0.6495         | 0.6789        | 0.6099        | 0.6539        | 0.6260         | 0.6528        | 0.5472        | 0.6383        |
| MCTN                    |                         | 0.7330         | 0.6656        | 0.7455        | 0.6827        | 0.7254         | 0.6451        | 0.7309        | 0.6678        | 0.7231         | 0.6228        | 0.7180        | 0.6454        | 0.7103         | 0.6107        | 0.7007        | 0.6280        |
| MMIN                    |                         | 0.7438         | <b>0.6810</b> | 0.7713        | 0.7035        | 0.7440         | <b>0.6676</b> | 0.7580        | 0.6802        | 0.7414         | 0.6545        | 0.7331        | 0.6540        | 0.7428         | 0.6464        | 0.7277        | 0.6391        |
| IF-MMIN                 |                         | 0.7566         | 0.6750        | 0.7701        | 0.7026        | 0.7431         | 0.6609        | 0.7522        | 0.6703        | 0.7389         | 0.6501        | 0.735         | 0.6558        | 0.7432         | 0.6502        | 0.7393        | 0.6498        |
| <b>Ours</b>             |                         | <b>0.7596</b>  | 0.6760        | <b>0.7744</b> | <b>0.7105</b> | <b>0.7454</b>  | 0.6617        | <b>0.7588</b> | <b>0.6831</b> | <b>0.7487</b>  | <b>0.6549</b> | <b>0.7361</b> | <b>0.6572</b> | <b>0.7482</b>  | <b>0.6543</b> | <b>0.7380</b> | <b>0.6530</b> |
| w/o VAE                 |                         | 0.7439         | 0.6610        | 0.7621        | 0.6800        | 0.7328         | 0.6305        | 0.7483        | 0.6622        | 0.7289         | 0.6300        | 0.7151        | 0.6445        | 0.7018         | 0.6204        | 0.6971        | 0.6210        |
| w/o $\mathcal{L}_{inv}$ |                         | 0.7521         | 0.6732        | 0.7709        | 0.7034        | 0.7387         | 0.6713        | 0.7511        | 0.6732        | 0.7261         | 0.6659        | 0.7272        | 0.6419        | 0.7230         | 0.6336        | 0.7045        | 0.6223        |

sample is subjected to a randomly selected incomplete condition. In contrast, the test set consists of 557 unique samples. Because of the small number of test samples, each sample will be evaluated under six predefined incomplete conditions, resulting in a total of 3342 test samples. In the CMU-MOSEI dataset, both train and test samples are randomly influenced by one incomplete condition, the same as the train set in the IEMOCAP dataset.

The only difference in data with previous works is that incomplete samples are created with our noise scheduler. Besides, the training and testing of the incomplete condition ( $E_T^a, E_T^v, E_T^l$ ) will be carried out individually to show the impact on the data in detail.

### 3.2 New Incomplete Condition

We introduce one novel incomplete condition ( $E_T^a, E_T^v, E_T^l$ ) utilizing the noise scheduler, which represents the addition of noise to all three modalities. This condition was not previously considered in existing research. In prior methodologies, once one modality is regarded as *incomplete* or *missing*, they will drop the data of this modality or set it to zero vectors, rendering the analysis of all three modalities’ incompleteness as both impractical and meaningless. But in real-world scenarios, this kind of incompleteness is not uncommon, and the incompleteness level also varies. The controllable intensity of noise allows us to report results under this extreme condition.

### 3.3 Experimental Setup

On the IEMOCAP dataset, we follow research [39, 41] to extract the original embeddings  $E^a$ ,  $E^v$ , and  $E^l$ . The audio, visual, and lexical embeddings are 130-dim *OpenSMILE* features with the configuration of *IS13\_ComParE* [10], *Denseface* [19] embeddings extracted by a pre-trained DenseNet model of 342 dimensions, and 1024-dim *BERT* [8] word embeddings, respectively.

On the CMU-MOSEI dataset, we employ the feature extraction method from the work of Liang et al. [24] Audio features are 74 dimensions extracted using *COVAREP* [6], while visual embeddings are 35 dimensions derived from the pool5 layer of an ImageNet [7]-trained *ResNet-152* [16] model for each video frame, which underwent preprocessing steps including resizing, center cropping, and normalizing. Facial expression features are obtained using the *OpenFace* [10] tool. Lexical features are represented using 300-dimensional *GloVe* word vectors [31]. We adjust the balance factors  $\lambda_s$  to 1, 10, and 1 for scaling the losses accordingly.

The hidden size of the LSTM structure is set to 128. The TextCNN contains 3 convolution blocks with kernel sizes of 3, 4, 5 and an output size of 128. The output size of the Invariance Encoder is also set to 128. The VAE Module includes a Transformer Encoder of 5 layers, 768 dimensions, and 16 heads as the encoder while Linear layers with dimensions of {128, 256, 384} as the decoder. The classifier contains three linear layers of size {384, 128, 4}. For the noise intensity, we conduct experiments on  $SNR_{dB}$  of [-10dB,

$-20dB$ ,  $-30dB$ ,  $-40dB$ ]. For impulse noise, the appearance frequency  $p$  is set at 0.3, with other parameters remaining unchanged. All experiments are run on an NVIDIA A100 80GB GPU.

We utilize the AdamW [28] as the optimizer and use the Lambda LR [36] to dynamically update the learning rate. The initial learning rate is 0.0001. The batch size is 128 and the dropout rate is 0.5. We run experiments with 10-fold cross-validation, where each fold contains 60 epochs, and report the result on the test set. Each result is run three times and averaged to reduce the effect of random initialization of parameters. We employ the same evaluation metrics as those used in previous works [39], [41], *Weighted Accuracy* (WA) [1] and *Unweighted Accuracy* (UA) [12], to assess various systems.

### 3.4 Comparison and Ablation Study

In our study, we benchmark our NMER model against four advanced MER baselines to establish its relative performance.

1) **Modality Encoder Network (MEN)** [39]: This model serves as the **complete-modality baseline**. MEN is trained under complete modality conditions and tested on incomplete modality conditions.

2) **MCTN** [32]: MCTN uses translation-based method with cycle consistency loss to learn joint representations between every two modalities in multimodal data, which is used as a popular method.

3) **MMIN** [39]: This model employs a cascade residual AutoEncoder coupled with cycle consistency construction to learn joint representations, particularly for predicting missing modalities.

4) **IF-MMIN** [41]: An enhancement of MMIN. IF-MMIN integrates the modality invariant feature to learn robust joint representations and is recognized as a state-of-the-art incomplete modalities multimodal emotion recognition system. MCTN, MMIN, and IF-MMIN are categorized as **incomplete-modality baselines**, with training and testing both under incomplete modality conditions.

Our experimental design not only evaluates the capability of the multimodal emotion recognition systems but also demonstrates the fine-grained control offered by our noise scheduler. We first conduct experiments on traditional six incomplete conditions, where the noise intensity progressively increases from  $-10dB$  to  $-40dB$ . Besides, we design the experiment on the brand new incompleteness condition ( $E_T^a, E_T^v, E_T^l$ ), to explore and evaluate the potential of models under a more complex noise environment. Moreover, we launched a simple benchmark test on the IEMOCAP dataset that directly sends unimodality data with various intensities of noise to the classifier and reports the result. This test aims to establish a benchmark for the subsequent incomplete modality multimodal emotion recognition technology, that is, how much it has improved on the basis of a single modality data and its various types of incompleteness.

Additionally, ablation studies are performed to ascertain the contribution of the VAE model and the importance of the invariant features to the overall model performance. In the study, labeled as “w/o VAE”, we remove the VAE model. Instead, we directly concatenate the extracted emotion-specific feature  $h'$  and the emotion-invariant feature  $H'$ , then feed this representation into the classifier. This study is designed to highlight the significant role of the VAE model in reconstructing the multimodal joint representation from noisy features. Another ablation study that removed the invariant loss, labeled as “w/o  $\mathcal{L}_{inv}$ ”, aims to demonstrate the invariant features can help guide the VAE model to generate more exact representations.

## 3.5 Results

### 3.5.1 Results About Models.

The average results of the comparative studies, across six incomplete conditions, four noise intensities and two noise types, are shown in Table 2, including  $\{a\}$ ,  $\{v\}$ ,  $\{l\}$ ,  $\{a, v\}$ ,  $\{a, l\}$ ,  $\{v, l\}$ ,<sup>1</sup> with Gaussian and impulse noise at four levels of intensities.

1) The NMER model outperforms most baseline models across various noise intensities and types. For instance, under  $-10dB$  Gaussian noise in the IEMOCAP dataset, NMER’s WA is 0.7598 compared to 0.7120 for MEN, 0.7215 for MCTN, 0.7551 for MMIN, and 0.7543 for IF-MMIN. UA results show a similar pattern.

2) There is a notable difference in accuracy (mainly in high-level noise conditions, for example  $-40dB$ , about 6% - 10%) between the complete-modality baseline (MEN) and the incomplete-modality baselines (MCTN, MMIN, and IF-MMIN) as shown in Table 2, which indicates the importance of using incomplete training data.

3) The results of our model on both datasets under six incomplete conditions are detailed in Table 3 (taking Gaussian noise as an example), the WA values for single clean modality conditions  $\{a\}$  and  $\{v\}$  drop by about 9% with the noise intensity increases from  $-10dB$  to  $-40dB$ , whereas in two clean modalities conditions  $\{a, v\}$ ,  $\{v, l\}$ , and  $\{v, l\}$ , the decrease is slower, ranging from about 2% - 4%. When modality  $a$  or  $v$  is noise-influenced, the performance decreases are roughly equivalent. This indicates that the resistance of the model to noise worsens significantly when only one modality is noise-free. However, the model shows relative robustness when the two modalities are noise-free. Notably, modality  $l$  plays a critical role in maintaining performance even when other modalities are compromised by noise. The WA value decreased about 4% when only modality  $l$  is clean, and about 2% when two modalities including  $l$  are noise-free (condition  $\{v, l\}$  and  $\{v, l\}$ ).

4) The WA and UA results on the CMU-MOSEI dataset, as shown in Table 2, exhibit different trends and absolute values compared to those on the IEMOCAP dataset, primarily due to the imbalanced label distribution (as shown in Table 1). In such scenarios, classifiers are inclined to categorize samples into the more samples class because this strategy statistically enhances their accuracy performance. Because of the way the WA value is calculated, if the classification performance of the larger number samples class is good, then even if the noise increases, the weighted accuracy (WA) may not decrease significantly as long as the prediction for that class remains accurate. Because that class contributes more to the overall accuracy, the decline in the accuracy of fewer sample class is offset. These reasons result in the different variation trends of WA and UA values.

### 3.5.2 Results About the Noise Scheduler.

1) It is evident that the performance of all models deteriorates as noise intensity increases, which aligns with the expectation that increasing noise levels obscure the original data’s information. For instance, in Table 2, the WA of the IF-MMIN system on the IEMOCAP dataset under Gaussian noise declines from 0.7543 at  $-10dB$  to 0.7048 at  $-40dB$ . These results highlight the effectiveness of our noise scheduler, which allows for precise control over noise intensities and the simulation of various noise conditions.

<sup>1</sup>{·} means the clean modality. The noisy modalities are omitted.

**Table 3: The detailed results of our NMER model under the six incomplete conditions utilizing Gaussian noise. “a” means that the modality  $a$  is clean, and the other two modalities ( $v$  and  $l$ ) are noise-influenced, “Avg” indicates the average result of six conditions.**

| Dataset   | Intensity | a      |        | v      |        | l      |        | a, v   |        | a, l   |        | v, l   |        | Avg    |        |
|-----------|-----------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
|           |           | WA     | UA     | WA     | UA     | WA     | UA     | WA     | UA     | WA     | UA     | WA     | UA     | WA     | UA     |
| IEMOCAP   | -10dB     | 0.7275 | 0.7349 | 0.7440 | 0.7503 | 0.7750 | 0.7820 | 0.7432 | 0.7508 | 0.7813 | 0.7917 | 0.7877 | 0.7952 | 0.7598 | 0.7675 |
|           | -20dB     | 0.6752 | 0.6873 | 0.6911 | 0.6949 | 0.7623 | 0.7735 | 0.7068 | 0.7178 | 0.7720 | 0.7846 | 0.7771 | 0.7864 | 0.7307 | 0.7406 |
|           | -30dB     | 0.6574 | 0.6626 | 0.6715 | 0.6688 | 0.7513 | 0.7592 | 0.7021 | 0.7071 | 0.7649 | 0.7778 | 0.7712 | 0.7771 | 0.7197 | 0.7254 |
|           | -40dB     | 0.6347 | 0.6484 | 0.6529 | 0.6511 | 0.7348 | 0.7484 | 0.6978 | 0.7024 | 0.7632 | 0.7777 | 0.7672 | 0.7773 | 0.7085 | 0.7176 |
| CMU-MOSEI | -10dB     | 0.7274 | 0.5905 | 0.7262 | 0.6204 | 0.7787 | 0.7362 | 0.7366 | 0.6187 | 0.8030 | 0.7507 | 0.7856 | 0.7383 | 0.7596 | 0.6760 |
|           | -20dB     | 0.7030 | 0.5644 | 0.6924 | 0.5885 | 0.7783 | 0.7317 | 0.7084 | 0.5965 | 0.7982 | 0.7474 | 0.7922 | 0.7395 | 0.7454 | 0.6617 |
|           | -30dB     | 0.7099 | 0.5497 | 0.6959 | 0.5690 | 0.7837 | 0.7394 | 0.7094 | 0.5859 | 0.8012 | 0.7453 | 0.7924 | 0.7380 | 0.7487 | 0.6549 |
|           | -40dB     | 0.7120 | 0.5579 | 0.6872 | 0.5694 | 0.7813 | 0.7263 | 0.6973 | 0.5844 | 0.8096 | 0.7456 | 0.8017 | 0.7412 | 0.7482 | 0.6543 |

**Table 4: The WA and UA declining on uni-modality testing.**

| Intensity | a      |        | v      |        | l      |        |
|-----------|--------|--------|--------|--------|--------|--------|
|           | WA     | UA     | WA     | UA     | WA     | UA     |
| 0dB       | 0.6693 | 0.6776 | 0.5723 | 0.5560 | 0.6484 | 0.6592 |
| -10dB     | 0.5633 | 0.5588 | 0.4252 | 0.4107 | 0.5421 | 0.5415 |
| -20dB     | 0.4769 | 0.4879 | 0.3941 | 0.3780 | 0.4665 | 0.4644 |
| -30dB     | 0.4483 | 0.4651 | 0.3624 | 0.3443 | 0.3867 | 0.3878 |
| -40dB     | 0.4161 | 0.4182 | 0.3502 | 0.3363 | 0.3795 | 0.3333 |

2) From the accuracy declining and the noise intensity increasing process in Table 2, it becomes crucial to recognize the shortcomings of previous studies that used zero vectors or drops data to simulate noise. These methods lack rationality and precise control over noise intensity and fail to enable meaningful comparisons across different noise types and intensities. In contrast, the experiments demonstrate the utility of our method in managing noises, achieving more comprehensive evaluations of the models.

3) In Table 4, we test the uni-modality result on the IEMOCAP dataset using Gaussian noise. The results in the first row show that different modalities have different quantities of information and difficulties of recognition. The performance of different modalities at the same noise level shows that different data types (audio, video, lexical) have different sensitivity to ambient noise. For example, the video ( $v$ ) modality has a greater performance degradation at -10dB (about 15%) than the audio ( $a$ ) and lexical ( $l$ ) modalities (about 10%). With the incremental noise intensities (from 0dB to -40dB), the accuracy of each modality has its own downtrend. This downward trend highlights the difficulty of effective emotion recognition in high-noise environments, while also providing data support for the

importance of multimodal fusion strategies. This table emphasizes the necessity and superiority of multimodal joint representation learning at the same time. When one mode is seriously disturbed, other modes may still be able to maintain good recognition performance, thus improving the overall emotion recognition accuracy.

### 3.5.3 Results About New Incomplete Condition.

The result of the new incomplete condition is listed in Table 5.

1) The introduction of scenarios with full-modality noise interference in multimodal emotion recognition presents a new challenge, reflecting complex real-world environments where multiple information sources can simultaneously experience quality degradation. This expands the boundaries of current research. Traditionally, models could rely on at least one clean modality. However, this new setup eliminates such possibilities, demanding greater robustness and denoising capability from the models themselves. With noise affecting all modalities, the accuracy significantly decreases. For instance, in Table 5, on the IEMOCAP dataset, the WA value of the MMIN model falls by about 4% from -10dB to -20dB noise levels and further declines by approximately 13% at -40dB. It indicates that the traditional multimodal fusion strategy may not be able to effectively deal with the simultaneous modal degradation, exposing the limitations of existing techniques against cross-modal interference.

2) There is a noticeable acceleration in performance degradation as noise intensity increases, which may be caused by the non-linear nature of emotion data and the complex response of models to varying noise levels. This information loss complicates the interaction between different modalities. Notably, all models demonstrate a similar trend under this condition. Future research should thus concentrate on studying how noise variations affect data and exploring more strategies to mitigate the impact of diverse noise.

### 3.5.4 Ablation Results.

The result of the last two rows in Table 2 shows the crucial role of the VAE model and the invariant features’ guidance. For example, in the IEMOCAP dataset, under Gaussian noise, the WA of two

**Table 5: The results of various systems utilizing Gaussian noise and adapting the new noise condition, the Impulse noise condition has a similar trend.**

| Dataset   | System      | -10dB         |               | -20dB         |               | -30dB         |               | -40dB         |               |
|-----------|-------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
|           |             | WA            | UA            | WA            | UA            | WA            | UA            | WA            | UA            |
| IEMOCAP   | MEN         | 0.6496        | 0.6355        | 0.5868        | 0.5545        | 0.5425        | 0.5233        | 0.4767        | 0.4631        |
|           | MCTN        | 0.6874        | 0.6930        | 0.6295        | 0.6250        | 0.5954        | 0.6032        | 0.5469        | 0.5503        |
|           | MMIN        | 0.7340        | 0.7453        | 0.6893        | 0.7011        | 0.6498        | 0.6566        | 0.6024        | 0.5997        |
|           | IF-MMIN     | <b>0.7450</b> | <b>0.7546</b> | 0.6940        | 0.6985        | 0.6472        | 0.6531        | 0.6053        | 0.5911        |
|           | <b>Ours</b> | 0.7436        | 0.7528        | <b>0.6963</b> | <b>0.7026</b> | <b>0.6510</b> | <b>0.6577</b> | <b>0.6103</b> | <b>0.6136</b> |
| CMU-MOSEI | MEN         | 0.7591        | 0.6471        | 0.7319        | 0.6061        | 0.7017        | 0.5707        | 0.6825        | 0.5560        |
|           | MCTN        | 0.7534        | 0.6635        | 0.7346        | 0.6285        | 0.7100        | 0.5844        | 0.6880        | 0.5613        |
|           | MMIN        | 0.7566        | 0.6855        | 0.7436        | 0.6459        | 0.7301        | 0.6081        | 0.7204        | 0.5625        |
|           | IF-MMIN     | <b>0.7680</b> | <b>0.6893</b> | <b>0.7588</b> | 0.6467        | 0.7327        | 0.6170        | 0.7240        | 0.5706        |
|           | <b>Ours</b> | 0.7651        | 0.6842        | 0.7570        | <b>0.6513</b> | <b>0.7341</b> | <b>0.6181</b> | <b>0.7266</b> | <b>0.5735</b> |

ablation studies achieves 0.7382 and 0.7401 on the noise intensity -10dB whereas the NMER achieves 0.7598. The outcomes clearly indicate the inclusion of the contributions that these parts provided to the overall performance of our NMER model.

## 4 Conclusion

This work proposed a Noise-robust Multimodal Emotion Recognition model (NMER) that effectively mitigates the impact of incomplete data and reconstructs the robust multimodal joint representations from incomplete data. Experimental results show that Our noise scheduler can effectively create different types and intensities of noisy data to simulate various noise corruptions. Our NMER model achieves robust performance across various incomplete situations. Notably, there exist huge challenges in the new incomplete condition, which proposes more requirements to the MER systems. Our research is only the first step in addressing this problem.

## 5 Acknowledgements

The research by Rui Liu was funded by the Young Scientists Fund of the National Natural Science Foundation of China (No. 62206136), the Guangdong Provincial Key Laboratory of Human Digital Twin (No. 2022B1212010004), and the “Inner Mongolia Science and Technology Achievement Transfer and Transformation Demonstration Zone, University Collaborative Innovation Base, and University Entrepreneurship Training Base” construction Project (Supercomputing Power Project) (21300-231510). The work is also supported by the following fundings: National Natural Science Foundation of China (62066033), the Outstanding Youth Foundation of the Natural Science Foundation of Inner Mongolia (2022JQ05), the Science and Technology Program of Inner Mongolia Autonomous Region (2021GG0158), the Hohhot Collaborative Innovation Project for Universities and Institutes, the Youth Science and Technology Talent

Cultivation Project of Inner Mongolia University (21221505), the fund of Supporting the Reform and Development of Local Universities (Disciplinary Construction) and the special research project of First-class Discipline of Inner Mongolia A. R. of China under Grant (YLXKZX-ND-036). College of Computer Science, Inner Mongolia University, National & Local Joint Engineering Research Center of Intelligent Information Processing Technology for Mongolian, and Inner Mongolia Key Laboratory of Multilingual Artificial Intelligence Technology provided equipment support to this work.

## References

- [1] Ishwar Baidari and Nagaraj Honnikoll. 2020. Accuracy weighted diversity-based online boosting. *Expert Syst. Appl.* 160 (2020), 113723. <https://doi.org/10.1016/j.eswa.2020.113723>
- [2] Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern recognition and machine learning*. Vol. 4. Springer.
- [3] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42, 4 (2008), 335–359.
- [4] Lei Cai, Zhengyang Wang, Hongyang Gao, Dinggang Shen, and Shuiwang Ji. 2018. Deep adversarial learning for multi-modality missing data completion. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1158–1166.
- [5] Jiayi Chen and Aidong Zhang. 2020. Hgmf: heterogeneous graph-based fusion for multimodal data with incompleteness. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 1295–1305.
- [6] Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer. 2014. COVAREP—A collaborative voice analysis repository for speech technologies. In *2014 IEEE international conference on acoustics, speech and signal processing (icassp)*. IEEE, 960–964.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT (1)*. Association for Computational Linguistics, 4171–4186.
- [9] Changde Du, Changying Du, Hao Wang, Jinpeng Li, Wei-Long Zheng, Bao-Liang Lu, and Huiguang He. 2018. Semi-supervised deep generative modelling of incomplete multi-modality emotional data. In *Proceedings of the 26th ACM international conference on Multimedia*. 108–116.



- [10] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [11] Carl Friedrich Gauss. 1877. *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*. Vol. 7. FA Perthes.
- [12] Shruti Gupta, Md. Shah Fahad, and Akshay Deepak. 2020. Pitch-synchronous single frequency filtering spectrogram for speech emotion recognition. *Multim. Tools Appl.* 79, 31-32 (2020), 23347–23365. <https://doi.org/10.1007/s11042-020-09068-1>
- [13] 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 Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5861–5865. <https://doi.org/10.1109/ICASSP.2019.8782465>
- [14] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*. Vol. 2. Springer.
- [15] Devamanyu Hazarika, Roger Zimmermann, and Soujanya Poria. 2020. Misa: Modality-invariant and-specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM international conference on multimedia*. 1122–1131.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- [18] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (1997), 1735–1780.
- [19] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*. 2261–2269. <https://doi.org/10.1109/CVPR.2017.243>
- [20] Seong Rag Kim and Adam Efron. 1995. Adaptive robust impulse noise filtering. *IEEE Transactions on Signal Processing* 43, 8 (1995), 1855–1866.
- [21] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *EMNLP, ACL*, 1746–1751.
- [22] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- [23] Zheng Lian, Lan Chen, Licai Sun, Bin Liu, and Jianhua Tao. 2023. GCNet: graph completion network for incomplete multimodal learning in conversation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023).
- [24] Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Yufan Chen, Peter Wu, Michelle A Lee, Yuke Zhu, et al. 2021. MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- [25] Rui Liu, Berrak Sisman, Björn W. Schuller, Guanglai Gao, and Haizhou Li. 2022. Accurate Emotion Strength Assessment for Seen and Unseen Speech Based on Data-Driven Deep Learning. In *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*. ISCA, 5493–5497.
- [26] Rui Liu, Haolin Zuo, Zheng Lian, Bjorn W Schuller, and Haizhou Li. 2024. Contrastive Learning based Modality-Invariant Feature Acquisition for Robust Multimodal Emotion Recognition with Missing Modalities. *IEEE Transactions on Affective Computing* (2024).
- [27] Sicong Liu, Liang Xiao, Lianfen Huang, and Xianbin Wang. 2019. Impulsive noise recovery and elimination: A sparse machine learning based approach. *IEEE Transactions on Vehicular Technology* 68, 3 (2019), 2306–2315.
- [28] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg6RiCqY7>
- [29] Mehdi Mafi, Harold Martin, Mercedes Cabrerizo, Jean Andrian, Armando Barreto, and Malek Adjouadi. 2019. A comprehensive survey on impulse and Gaussian denoising filters for digital images. *Signal Processing* 157 (2019), 236–260.
- [30] Takeru Miyato, Andrew M. Dai, and Ian Goodfellow. 2017. Adversarial Training Methods for Semi-Supervised Text Classification. In *International Conference on Learning Representations*. [https://openreview.net/forum?id=r1X3g2\\_xl](https://openreview.net/forum?id=r1X3g2_xl)
- [31] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [32] 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6892–6899.
- [33] Jocelyn Sietsma and Robert JF Dow. 1991. Creating artificial neural networks that generalize. *Neural networks* 4, 1 (1991), 67–79.
- [34] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* 28 (2015).
- [35] Qiuling Suo, Weida Zhong, Fenglong Ma, Ye Yuan, Jing Gao, and Aidong Zhang. 2019. Metric Learning on Healthcare Data with Incomplete Modalities.. In *IJCAI*. 3534–3540.
- [36] Neo Wu, Bradley Green, Xue Ben, and Shawn O'Banion. 2020. Deep transformer models for time series forecasting: The influenza prevalence case. *arXiv preprint arXiv:2001.08317* (2020).
- [37] Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, and Louis-Philippe Morency. 2018. Multi-attention recurrent network for human communication comprehension. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [38] Jiandian Zeng, Jiantao Zhou, and Tianyi Liu. 2022. Mitigating Inconsistencies in Multimodal Sentiment Analysis under Uncertain Missing Modalities. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 2924–2934.
- [39] Jinming Zhao, Ruichen Li, and Qin Jin. 2021. Missing modality imagination network for emotion recognition with uncertain missing modalities. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 2608–2618.
- [40] Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li. 2021. Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 920–924.
- [41] 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 *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.