Curriculum Learning Meets Weakly Supervised Multimodal Correlation Learning

Sijie Mai and **Ya Sun** and **Haifeng Hu**

School of Electronics and Information Technology, Sun Yat-sen University

`{maisj,suny278}@mail2.sysu.edu.cn, huhaif@mail.sysu.edu.cn`

Abstract

In the field of multimodal sentiment analysis (MSA), a few studies have leveraged the inherent modality correlation information stored in samples for self-supervised learning. However, they feed the training pairs in a random order without consideration of difficulty. Without human annotation, the generated training pairs of self-supervised learning often contain noise. If noisy or hard pairs are used for training at the easy stage, the model might be stuck in bad local optimum. In this paper, we inject curriculum learning into weakly supervised modality correlation learning. The weakly supervised correlation learning leverages the label information to generate scores for negative pairs to learn a more discriminative embedding space, where negative pairs are defined as two unimodal embeddings from different samples. To assist the correlation learning, we feed the training pairs to the model according to difficulty by the proposed curriculum learning, which consists of elaborately designed scoring and feeding functions. The scoring function computes the difficulty of pairs using pre-trained and current correlation predictors, where the pairs with large losses are defined as hard pairs. Notably, the hardest pairs are discarded in our algorithm, which are assumed as noisy pairs. Moreover, the feeding function takes the difference of correlation losses as feedback to determine the feeding actions ('stay', 'step back', or 'step forward'). The proposed method reaches state-of-the-art performance on MSA.

1 Introduction

With the rapid development of social media, multimodal data have been widely-used to perform many downstream tasks, including multimodal sentiment analysis (Rahman et al., 2020), multi-omics integrative analysis(Lee and Schaar, 2021), human action recognition(Kay et al., 2017), etc. Multimodal sentiment analysis, aiming to mine humans' sentiments and opinions from language, audio, and visual sequences, has attracted significant attention.
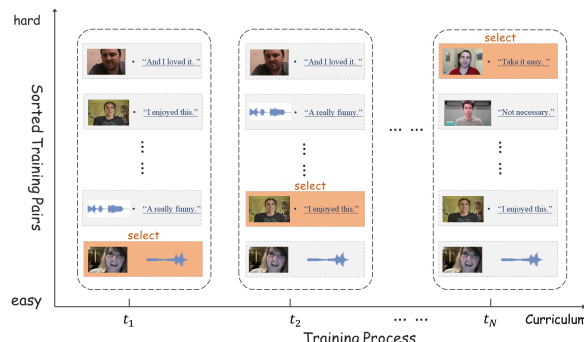

Figure 1: **Illustration of the curriculum learning.**

Manual annotation of large-scale multimodal datasets is challenging and time-consuming (Amrani et al., 2021). Many models suffer from overfitting due to limited training samples and the large amount of parameters introduced by complex fusion algorithms(Zadeh et al., 2017; Tsai et al., 2019) or large pre-trained transformers(Devlin et al., 2019; Yang et al., 2019), which degrades the generalizability. Recently, to leverage the correlation information between modalities inherently stored in each sample, a few contrastive learning frameworks are proposed (Mai et al., 2021; Udandarao et al., 2020; Tian et al., 2020; Chuang et al., 2022), which generate abundant training pairs of unimodal embeddings across samples to train a robust model. These methods learn the similarity between modalities, such that unimodal embeddings from the same sample are pulled closer while those from different samples or classes are pushed away. However, they feed the generated pairs to the model in a random order without consideration of their difficulty, which might cause inaccurate learning. If hard pairs are used for training at early stage where the model is not well trained, the optimization may be stuck in a bad local optimum (Liu et al., 2021). More severely, without human supervision, the training data of unsupervised or weakly supervised tasks often contain noise. For instance, two modalities from the same sample might

convey contradictory semantic meanings. On the other hand, two unimodal pairs can share semantic similarity even if they come from different samples. The noisy pairs are often the hardest pairs that cannot be well predicted by the model, which should be identified to prevent the feature extraction capability from being damaged.

In contrast, the education of humans is highly organized, which is based on a curriculum that introduces different concepts at different times according to difficulty, exploiting previously learned concepts to ease the learning of new abstractions. Inspired by the mode of human education, we focus on selecting the generated pairs for training dynamically (see Fig. 1), and follow the idea of curriculum learning which enables training data to be fed into the model in a certain order(Hacohen and Weinshall, 2019; Bengio et al., 2009; Soviany et al., 2022). By choosing which data to present and in which order to present them, we can guide the learning and find a better local minimum for the model.

Specifically, we first construct a weakly supervised modality correlation learning task to capture the shared information between modalities, where a positive pair and a negative pair consists of unimodal representations from the same sample and from different samples respectively. The task is said to be weakly supervised as it incorporates label information to score the pairs to learn a more discriminative embedding space for downstream tasks (while the score is not always correct). Then, we inject curriculum learning into correlation learning which consists of difficulty scoring and pair feeding functions. The scoring function determines the hardness of training pairs based on pre-trained and current correlation predictors, and sorts the training pairs based on hardness. The feeding function determines the pace by which the training pairs are fed to the model. Different from traditional methods(Hacohen and Weinshall, 2019; Liu et al., 2021), our feeding function has three candidate actions ('stay', 'step back', or 'step forward'), which enables a more refined and accurate feeding strategy. It takes the difference of current and previous correlation losses as feedback to determine the feeding actions. The proposed curriculum learning enables the model to be learned gradually from easy pairs to hard pairs which effectively prevents the model from falling into bad local optimum. To address the issue of noisy pairs, we dynamically identify and discard the hardest pairs during training to minimize the negative effect of noisy data.

In brief, the contributions are listed as follows:

- We construct a Weakly Supervised Correlation Learning framework and innovatively equip the framework with Curriculum Learning (WSCL-CL) that enables the model to progressively learn from easy pairs to hard pairs. In this way, we can prevent the model from falling into bad local optimum.

- In curriculum learning, difficulty scoring function is introduced to sort the training pairs based on pre-trained and current predictors, and feeding function is elaborately designed to determine the pace by which the training pairs are fed to the model. Moreover, we address the noisy pair problem during training which is not considered in previous methods(Udandarao et al., 2020; Mai et al., 2021).

- The proposed method achieves state-of-the-art performance on multimodal sentiment analysis. We show that the introduced curriculum learning can improve the performance of the weakly supervised learning.

2 Related Work

With the availability of multimodal data, MSA draws increasing attention for its capability to interpret human language and mine sentiments and opinions. Many works focus on designing fusion strategies to obtain informative multimodal embedding (Poria et al., 2017a). Two simple and explicit fusion methods are early fusion (i.e., feature-level fusion) (Wollmer et al., 2013; Poria et al., 2016, 2017b) and late fusion (i.e., decision-level fusion)(Nojavanasghari et al., 2016; Kampman et al., 2018; Zhao et al., 2021). Recently, modern fusion methods are proposed to explore intra-/inter-modal interactions. Specifically, tensor-based methods (Liang et al., 2019; Zadeh et al., 2017) draw increasing attention because they can learn joint embedding with high expressive power. Graph-based fusion methods(Mai et al., 2020; Zadeh et al., 2018c; Behmanesh et al., 2021) learn interactions on the graph across time series and modalities. To highlight important information across modalities, many algorithms apply cross-modal attention mechanisms (Zadeh et al., 2018b; Tsai et al., 2019; Shenoy and Sardana, 2020). To improve the

interpretablility of multimodal fusion, quantum-inspired fusion (Li et al., 2021) and capsule-based fusion (Tsai et al., 2020; Wu et al., 2021) are proposed. Moreover, a few methods use KL-divergence, canonical correlation analysis and so on to regularize the learning of unimodal distributions and better match different modalities (Sun et al., 2020; Dumpala et al., 2019; Shankar, 2022; Sun et al., 2022). With the success of BERT(Devlin et al., 2019), there has been a trend to fine-tune pre-trained transformers using multimodal data(Yang et al., 2020; Rahman et al., 2020; Hazarika et al., 2020). However, these methods do not fully leverage the modality correlation information inherently stored in each sample, and may easily suffer from overfitting and have poor generalizability especially when the dataset is small(Mai et al., 2022).

More recently, self-supervised learning on multimodal data has attracted significant research attention, where the self-supervised learning tasks typically include modality matching (especially image-text pair) and masked reconstruction (Gan et al., 2020; Lu et al., 2019; Tan and Bansal, 2019; Mai et al., 2022). Another popular strategy of self-supervised learning is to apply contrastive learning to learn representations by pushing unimodal representations describing the same sample closer, and pushing those of different samples apart (Ruan et al., 2022; Udandarao et al., 2020; Pielawski et al., 2020; Chuang et al., 2022). For example, HyCon (Mai et al., 2021) explores intra-modal and inter-modal dynamics across samples, which defines two unimodal representations from the same class as positive pairs. However, these methods do not consider feeding the training pairs into the model based on difficulty, which might cause the model to be stuck in the bad local optimum. In contrast, we perform weakly supervised modality correlation learning to learn a more discriminative embedding space for downstream task, and leverage curriculum learning to enable a more smooth and accurate learning. Moreover, we also seek to identify noisy modality pairs for minimizing the negative effect of noisy data. Curriculum learning has been applied in various tasks and shows promising results (Hacohen and Weinshall, 2019; Wang et al., 2021; Liu et al., 2021; Gui et al., 2017; Yang et al., 2022; Bengio et al., 2009). To the best of our knowledge, we are the first to introduce curriculum learning into weakly supervised/unsupervised multimodal learning algorithms.

3 Algorithm

3.1 Notations and Task Definitions

Our downstream task is multimodal sentiment analysis (MSA). The input of MSA is an utterance (Olson, 1977), i.e., a segment of a video bounded by a sentence. Each utterance has three modalities, i.e., acoustic (a), visual (v), and language (l). The input sequences of acoustic, visual, and language modalities are denoted as $\boldsymbol{U_a} \in \mathbb{R}^{T_a \times d_a}$, $\boldsymbol{U_v} \in \mathbb{R}^{T_v \times d_v}$, and $\boldsymbol{U_l} \in \mathbb{R}^{T_l \times d_l}$ respectively, where T_m and d_m are the sequence length and feature dimensionality respectively ($m \in \{a, v, l\}$). MSA aims to predict the sentiment score based on three unimodal sequences. After being processed by unimodal learning network F^m, the input sequence $\boldsymbol{U_m} \in \mathbb{R}^{T_m \times d_m}$ becomes the unimodal representation $\boldsymbol{x_m} \in \mathbb{R}^d$. For conciseness, the structures of F^m are illustrated in the Appendix. The pipeline of our algorithm is illustrated in Fig. 2. We first introduce the modality correlation learning task, and then elaborate on our curriculum learning.

3.2 Modality Correlation Learning Task

Formally, let $\boldsymbol{x}^j_{m_1}$, $\boldsymbol{x}^j_{m_2}$, ..., $\boldsymbol{x}^j_{m_k}$ denote the representations of k modalities describing the same sample j ($k = 3$ in the case of MSA), where $j \in \{1, .., n\}$ and n is the batch size. Correlation learning aims to enable unimodal networks to extract high-level semantic representations and capture the shared information across modalities within each sample(Mai et al., 2022). We take modality m_i and modality m_j as an example to illustrate the pipeline. Firstly, to construct the negative pairs, we sample N unimodal representations from modality m_i and modality m_j respectively, where N is the number of negative pairs. Then we concatenate the representations of each two modalities to generate negative bimodal representations:

$$\boldsymbol{x}^o_{m_i m_j} = \boldsymbol{x}^{o_1}_{m_i} \oplus \boldsymbol{x}^{o_2}_{m_j}, \quad o_1 \neq o_2 \quad (1)$$

where $\oplus$ denotes concatenation operation, $\boldsymbol{x}^o_{m_i m_j} \in \mathbb{R}^{2 \cdot d}$ ($o \in \{1, 2, .., N\}$) is a negative bimodal representation whose constructed modalities come from different samples. N is set to $\beta \cdot n$, where β is the negative sampling factor. By sampling abundant modality pairs for training, the negative effect of small dataset can be mitigated. Similarly, we construct positive bimodal representations as:

$$\boldsymbol{x}^o_{m_i m_j} = \boldsymbol{x}^{o_1}_{m_i} \oplus \boldsymbol{x}^{o_2}_{m_j}, \quad o_1 = o_2, \ o \in \{1, 2, .., n\} \quad (2)$$

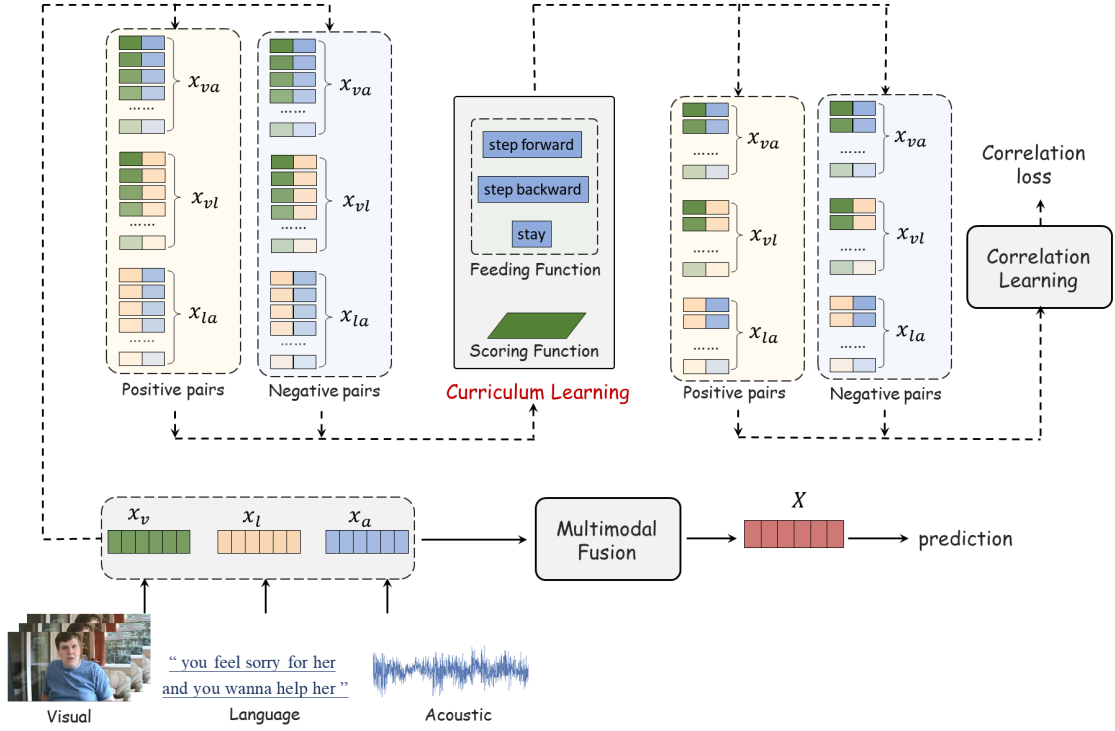To determine the score of the generated pairs,

Figure 2: **The diagram of the proposed WSCL-CL.**

we first calculate the distance between the labels of the associated samples:

$$d^o = \sqrt{(\bar{l}_{o_1} - \bar{l}_{o_2})^2} \quad (3)$$

The final correlation score is defined as:

$$y^o = \begin{cases} \frac{1}{d^o + \gamma}, & o_1 \neq o_2 \\ 1, & o_1 = o_2 \end{cases} \quad (4)$$

where $\bar{l}_{o_1}$ and $\bar{l}_{o_2}$ are the true labels of the sample o_1 and o_2 respectively, which are continuous values that indicate the intensity of sentiment. γ is a hyperparameter that determines the maximum score of negative pairs ($\gamma > 1$). We assign a score within zero and one to each modality pair. If two unimodal representations are from the same object, y^o is equal to 1. If they come from different objects but share the same label, we have $y^o = \frac{1}{\gamma}$. And if they have different labels, we have $y^o < \frac{1}{\gamma}$. Consequently, the more semantic similarity the two modalities share, the larger y^o. Label information of each object is considered here to learn a more discriminative feature space for downstream task (that is why we call it weakly supervised).

A correlation predictor is then leveraged to predict the correlation score of each training pair, and then the correlation loss is calculated based on the predicted score:

$$s^o = CP(\boldsymbol{x}^o_{m_i m_j}; \boldsymbol{W}_{cp}) = \boldsymbol{W}_{cp}\boldsymbol{x}^o_{m_i m_j} \quad (5)$$

$$\ell_{m_i m_j} = \frac{1}{2N_p}\sum_{o=1}^{N_p}|s^o - y^o| + \frac{1}{2N_n}\sum_{o=1}^{N_n}|\hat{s}^o - \hat{y}^o| \quad (6)$$

$$\ell_c = \frac{1}{C_k^2}\sum_{i}^{i<j}\sum_{j=2}^{k}\ell_{m_i m_j}, \quad i, j \in \{1, 2, .., k\} \quad (7)$$

where C_k^2 is the number of bimodal correlation losses, ℓ_c is the overall correlation loss, CP denotes the correlation predictor, $\boldsymbol{W}_{cp} \in \mathbb{R}^{1 \times 2d}$ is the parameter matrix of the correlation predictor, s^o denotes the predictive score of the o^{th} positive pair, and $\hat{s}^o$ is the score of negative pair (here we add the hat to identify positive and negative pairs). Note that N_p and N_n are the number of selected positive and negative pairs respectively (determined by the curriculum learning). In Eq. 6, mean absolute error (MAE) is used to calculate the learning loss. By minimizing the correlation loss of positive pairs, the model can learn to identify whether two modalities are describing the same object, encouraging the model to capture shared information across various modalities within each object. Via minimizing the correlation loss of negative pairs, the model can learn to discover the distinguishable information of each sample with respect to other samples. Instead of directly using dot product or distance metric to calculate the similarity of unimodal embeddings as in contrastive learning (Mai et al., 2021; Udandarao et al., 2020; Ruan et al., 2022; Weinberger and Saul,

2009), correlation learning utilizes a learnable predictor which can explore complex correspondence relationship between modalities(Mai et al., 2022).

3.3 Curriculum Learning

The curriculum learning determines how to choose the training pairs presented to the model and in which order to present them(Bengio et al., 2009). In this way, we can guide the learning of modality correlation learning task and find a better local minimum for the model. The core and the main innovation of our curriculum learning lies in the design of difficulty scoring function and pair feeding schedule, which are introduced below. Notably, the design of the curriculum learning is independent of concrete learning tasks. Other than the correlation learning task, the proposed curriculum learning is applicable to any tasks as long as the self-/weakly-supervised learning loss is given.

3.3.1 Difficulty Scoring Function

Scoring Function is designed to determine the difficulty of the training pairs. In this paper, we propose to score the pairs by leveraging the pre-trained and current predictors. Compared to inner product or Euclidean distance that cannot reveal high-level correlation, our method is more target-oriented.

To identify the hardness of the pairs, we first pre-train a correlation predictor CP_{pre} on the same dataset. At the warm-up training stage, we only use the pre-trained CP_{pre} to determine the difficulty. For a batch of training pairs $\{\boldsymbol{x}_{m_i m_j}^o | o = 1, 2, ...\}$, we feed them into CP_{pre}, which outputs the correlation scores $\{s_{m_i m_j}^o | o = 1, 2, ...\}$. We then compute the correlation loss of each training pair $\{l_{m_i m_j}^{o'} | o = 1, 2, ...\}$, and finally sort the pairs by their loss values (which are defined as the difficulties of the pairs). However, as the training strategies of CP_{pre} and current CP are different, the CP_{pre} might not serve as an excellent difficulty scoring function. Therefore, after the warm-up training stage where the current CP becomes discriminative, we incorporate the predictive loss of CP to generate difficulty score:

$$s_d^o = l_{m_i m_j}^{o'} + \lambda \cdot l_{m_i m_j}^o, \quad o \in \{1, 2, ..\} \quad (8)$$

where $l_{m_i m_j}^o$ is the correlation loss produced by current CP, s_d^o denotes the difficulty score after warm-up training stage, and λ is set to 0.8 in our experiment. The sorted training pairs are then sent into the feeding function.

3.3.2 Pair Feeding Function

The feeding function is designed to determine by which the training data are fed to the model. In this paper, we elaborately design the feeding function according to the feedback of the correlation loss to adaptively change the feeding strategy.

Different from previous methods(Hacohen and Weinshall, 2019; Bengio et al., 2009; Liu et al., 2021), our feeding function has three actions to be taken at each iteration, namely, 'step forward', 'step backward', and 'stay'. Firstly, we divide the sorted training pairs ($\{\boldsymbol{x}_{m_i m_j}^o\}$ for short) into c partitions based on difficulty. The feeding function is responsible for choosing one partition for training at each iteration, where the pairs from the same partition are assumed to have approximately the same difficulty. Notably, for positive and negative pairs, the value of c can be different to ensure a more flexible choosing mechanism. Generally, we follow the idea that if the model learns fast from the tasks of a particular difficulty (i.e., the correlation loss decreases intensely), we should sample the tasks of the same difficulty at the next iteration (i.e., 'stay') to enable the model to be learned fast. In contrast, if the model learns little from the tasks of a particular difficulty for a certain number of iterations, we should 'step forward' and train the model with more challenging tasks. Finally, if the performance drops considerably, we should 'step backward' to let the model learn more from easy pairs and become more expressive before feeding harder pairs. The choosing procedure is shown in Algorithm 1. The feeding function is run for $2 \cdot C_k^2$ times to select positive/negative unimodal pairs.

As shown in Algorithm 1, we introduce the counting factor $count$ and patience factor p to enable a more stable training. Only when $count$ is larger or equal to p which implies that the model learns little from the tasks of current difficulty and becomes steady during training, can we feed more difficult tasks to the model. In this way, we can prevent the correlation loss from oscillating. Notably, as presented in Line 5 of Algorithm 1, when the model is already trained by the hardest partition, we randomly sample other pairs for training, which prevents the catastrophic forgetting of previous knowledge.

Moreover, as elaborated in Introduction section, without human supervision, the generated training data of unsupervised or weakly supervised tasks often contain noise. To minimize the negative effect

Algorithm 1 Pair Feeding Function

input: Patience factor p, counting factor $count$, forward factor f, backward factor b, number of partitions c, training pairs $\{\boldsymbol{x}^o_{m_i m_j}\}$, previous choosing index c_i, previous correlation loss $\ell^{pre}_{m_i m_j}$

output: A partition of training pairs $\{\boldsymbol{x}^{o'}_{m_i m_j}\}$

1. Compute current correlation loss $\ell^{now}_{m_i m_j}$ using all training pairs in $\{\boldsymbol{x}^o_{m_i m_j}\}$.
2. Divide $\{\boldsymbol{x}^o_{m_i m_j}\}$ into c partitions.
3. If $(\ell^{now}_{m_i m_j} - \ell^{pre}_{m_i m_j})/\ell^{pre}_{m_i m_j} > b$:
    $count = 0$
    if $c_i > 1$:
        $c_i = c_i - 1$ %step backward
   elseif $(\ell^{pre}_{m_i m_j} - \ell^{now}_{m_i m_j})/\ell^{pre}_{m_i m_j} > f$:
    $count = 0$ %stay
   elseif $count >= p$:
    $count = 0$
    if $c_i < c$:
        $c_i = c_i + 1$ %step forward
   else:
    $count = count + 1$ %stay.
4. Choose the $(c_i)^{th}$ partition $\{\boldsymbol{x}^{o'}_{m_i m_j}\}$.
5. if $c_i == c$:
    Randomly sample other training pairs and put them into $\{\boldsymbol{x}^{o'}_{m_i m_j}\}$.
6. $\ell^{pre}_{m_i m_j} \longleftarrow \ell^{now}_{m_i m_j}$.

---

of noisy data, we assume the hardest pairs (with correlation loss greater than 95% of all pairs) as noisy and discard them during training.

3.4 Sentiment Prediction

Our algorithm is independent of the concrete fusion method. In practice, we apply a deep neural network on the concatenated unimodal representations to learn high-level interactions between modalities:

$$\boldsymbol{X} = DNN(\boldsymbol{x}_a \oplus \boldsymbol{x}_v \oplus \boldsymbol{x}_l; \theta_{DNN}) \quad (9)$$

where DNN consists of several fully connected layers. We then learn a multimodal predictor that is composed of fully connected layers to generate the sentiment prediction based on $\boldsymbol{X}$:

$$\tilde{l} = MP(\boldsymbol{X}; \theta_{MP}) \quad (10)$$

where $\tilde{l}$ is the predicted label, and MP is the multimodal predictor. We use mean square error (MSE) to compute the task loss:

$$\mathcal{L} = (\bar{l} - \tilde{l})^2 \quad (11)$$

where $\bar{l}$ is the ground truth label, and $\mathcal{L}$ is the MSE for prediction. The total loss is weighted sum of the task loss and correlation loss:

$$\mathcal{L} \longleftarrow \mathcal{L} + \alpha \cdot \ell_c \quad (12)$$

where α is the weight of the correlation loss ℓ_c.

4 Experiments

In this section, we conduct extensive experiments to evaluate the effectiveness of the proposed WSCL-CL. **For the lack of space, more details about the introduction of baselines and the feature extraction details are placed in the Appendix.**

4.1 Datasets

1) **CMU-MOSI**(Zadeh et al., 2016) is a widely-used dataset for multimodal sentiment analysis from online websites, which contains a collection of more than 2k video clips. Each video clip is annotated with sentiment on a [-3,3] Likert scale, where +3 denotes the strongest positive sentiment and -3 the strongest negative. To be consistent with prior works (Sun et al., 2022; Rahman et al., 2020), we use 1,281 utterances for training, 229 utterances for validation, and 685 utterances for testing.

2) **CMU-MOSEI**(Zadeh et al., 2018c) is a large dataset of multimodal sentiment analysis. The dataset consists of more than 20k utterances from more than 1,000 YouTube speakers, covering 250 distinct topics. All the utterances are randomly chosen from various topics and monologue videos. In our algorithm, we use the sentiment label of CMU-MOSEI to perform sentiment label, where the range of the sentiment label is the same as that of the CMU-MOSI. Following previous works (Sun et al., 2022; Rahman et al., 2020; Mai et al., 2021), we use 16,265 utterances as training set, 1,869 utterances as validation set, and 4,643 utterances as testing set.

4.1.1 Evaluation Metrics

We adopt the following metrics to evaluate the performance of the model: 1) Acc7: 7-way accuracy, sentiment score classification; 2) Acc2: binary accuracy, positive or negative; 3) F1 score; 4) MAE: mean absolute error and 5) Corr: the correlation between the model's prediction and human evaluation. Notably, to compute the 7-way accuracy, we round up the model prediction to an integer from -3 to 3. Note that following (Rahman et al., 2020),

Table 1: Hyperparameters of WSCL-CL.

	CMU-MOSI	CMU-MOSEI
Batch Size	64	48
Learning Rate	2e-5	1e-5
Weight α	1	1
Training Epochs	20	50
Dimensionality d	100	100
Negative Sampling Factor β (pre-training)	4	4
Negative Sampling Factor β	30	30
Scale Factor γ	1.4	1.4
patience Factor p	6	5
forward Factor f	0.1	0.1
backward Factor b	0.15	0.15
c for negative pairs	10	10
c for positive pairs	2	8

Table 2: **The comparison with baselines on CMU-MOSI.** The best results are highlighted in bold, and the second best results are marked with underlines. The results of ICCN (Sun et al., 2020) are copied from the original paper since the codes are unavailable.

	Acc7 (↑)	Acc2 (↑)	F1 (↑)	MAE (↓)	Corr (↑)
TFN-BERT (Zadeh et al., 2017)	44.7	82.6	82.6	0.761	0.789
LMF-BERT (Liu et al., 2018)	45.1	84.0	84.0	0.742	0.785
MFN-BERT (Zadeh et al., 2018a)	44.1	83.5	83.5	0.759	0.786
MULT-BERT (Tsai et al., 2019)	41.5	83.7	83.7	0.767	**0.799**
GFN-BERT (Mai et al., 2020)	47.0	84.3	84.3	0.736	0.790
ICCN-BERT (Sun et al., 2020)	39.0	83.0	83.0	0.860	0.710
MAG-BERT (Rahman et al., 2020)	42.9	83.5	83.5	0.790	0.769
TFR-Net (Yuan et al., 2021)	42.6	84.0	83.9	0.787	0.788
AMML (Sun et al., 2022)	46.3	84.9	84.8	0.723	0.792
HyCon(Mai et al., 2021)	46.6	85.2	85.1	0.713	0.790
WSCL-CL	**47.5**	**86.3**	**86.2**	**0.712**	0.798

when calculating Acc2, F1 score, corr, and MAE, we do not use the neutral utterances.

4.2 Experimental Settings

We develop our model with the PyTorch framework on GTX2080Ti on CUDA 10.1 and torch version 1.4.0. The proposed model is trained using Adam (Kingma and Ba, 2015) optimizer. The detailed hyperparameter setting can be referred to Table 1. Notably, the negative sampling factor β is set to be larger than that in the pre-training stage, but we only sample a part of the generated pairs for training via the pair feeding function.

4.3 Comparison with Baselines

In this section, we compare our proposed method with baselines on the task of multimodal sentiment analysis. As shown in Table 2 and 3, the contrastive learning based algorithm HyCon performs better than other baselines and sets up a high baseline. Nevertheless, WSCL-CL still outperforms HyCon and obtains the best performance on the majority of the evaluation metrics. Specifically, on CMU-MOSI dataset, WSCL-CL outperforms HyCon by 0.9% on Acc7, 1.1% on Acc2 and F1 score. The improvement is remarkable considering that the humans do not always perform well on

Table 3: **The comparison with baselines on CMU-MOSEI dataset.**

	Acc7 (↑)	Acc2 (↑)	F1 (↑)	MAE (↓)	Corr (↑)
TFN-BERT (Zadeh et al., 2017)	51.8	84.5	84.5	0.622	0.781
LMF-BERT (Liu et al., 2018)	51.2	84.2	84.3	0.612	0.779
MFN-BERT (Zadeh et al., 2018a)	52.6	84.8	84.8	0.607	0.771
MULT-BERT (Tsai et al., 2019)	50.7	84.7	84.6	0.625	0.775
GFN-BERT (Mai et al., 2020)	51.8	85.0	85.0	0.611	0.774
ICCN-BERT (Sun et al., 2020)	51.6	84.2	84.2	**0.565**	0.713
MAG-BERT (Rahman et al., 2020)	51.9	85.0	85.0	0.602	0.778
TFR-Net (Yuan et al., 2021)	51.7	85.2	85.1	0.606	0.781
AMML (Sun et al., 2022)	52.4	85.3	85.2	0.614	0.776
HyCon(Mai et al., 2021)	52.8	85.4	85.6	0.601	0.776
WSCL-CL	**53.3**	**86.1**	**86.0**	0.577	**0.794**

multimodal sentiment analysis(Zadeh et al., 2017). On CMU-MOSEI dataset, WSCL-CL yields 0.7% improvement on Acc2 and 1.2% improvement on F1 score compared to the current state-of-the-art HyCon(Mai et al., 2021). Our method also reaches state-of-the-art performance on Acc7 and Corr. Compared to HyCon(Mai et al., 2021), WSCL-CL allows the modality-specific information to be preserved in an elegant way, addresses the noisy pair problem, and considers to feed the training pairs according to difficulty to find a better local optimum for the model. These results demonstrate the effectiveness of our proposed model, indicating the importance of learning the correlation between modalities and injecting curriculum learning into weakly-supervised learning task.

4.4 Ablation Study

In this section, we conduct extensive ablation studies to evaluate the effectiveness of the components in our WSCL-CL:

1) Correlation Learning: In the case of 'W/O Correlation Learning' (see Table 4), we remove the correlation learning task and the accompanied curriculum learning. As shown in Table 4, the performance drops dramatically by about 3.5 points in Acc7 and 2.5 points in Acc2 and F1 score, demonstrating the importance of correlation learning to better correlate unimodal representations in a more discriminative embedding space;

2) Curriculum Learning: In the case of 'W/O Curriculum Learning', we remove the curriculum learning but retain the correlation learning. The results suggest that the curriculum learning brings over 1% improvement to the multimodal systems on Acc7, Acc2 and F1 score. The performance on other evaluation metrics is also improved. These results reveal that selecting pairs for training dynamically based on difficulty is of great significance to correlation learning, which can find a better local minimum for the parameters of the model;

Table 4: **Ablation studies on the CMU-MOSI dataset.**

	Acc7	Acc2	F1	MAE	Corr
W/O Correlation Learning	43.8	83.8	83.8	0.756	0.784
W/O Curriculum Learning	46.1	84.9	84.8	0.715	0.797
W/O Patience Factor	47.0	85.7	85.6	0.728	0.793
W/O 'Backward'	45.4	85.5	85.3	0.714	0.797
W/O Discarding Hardest Pairs	46.7	85.8	85.7	0.718	0.793
W/O Random Sampling	47.3	86.1	86.1	**0.710**	**0.798**
WSCL-CL	**47.5**	**86.3**	**86.2**	0.712	**0.798**

3) Patience Factor: In the case of 'W/O Patience Factor', the feeding function selects harder pairs immediately if no significant gain is obtained on the easier pairs. The results on all the evaluation metrics decrease considerably. We argue that this is because without the patience factor, the learning of modality correlation suffers from oscillating and a worse local optimum is reached;

4) Backward Action: In the case of 'W/O Backward', we remove the backward action in the feeding function such that the feeding function cannot select easier pairs after harder pairs are used for training. We observe that Acc7 and Acc2 drop by about 2% and 1% respectively, demonstrating the importance of the design of backward action in learning a more expressive model by 'looking back' at easier training data;

5) Discarding Hardest Pairs: The performance on all the evaluation metrics declines when the hardest pairs are used for training. The results suggest that the hardest pairs are harmful to the performance of the model which are often the noisy pairs and needed to be addressed during training;

6) Random Sampling: In the case of 'W/O Random Sampling', we only use the hardest training pairs if the algorithm has already selected the hardest partition of $\{x^o_{m_i m_j}\}$ for training. The performance of the model slightly drops, which might be due to the catastrophic forgetting of previous knowledge. Although random sampling does not improve the performance by a large margin, it remains a good practice considering that it is simple and introduces no additional parameters.

4.5 Discussion on Difficulty Scoring Function

We analyze the performance of different kinds of difficulty scoring functions in this section. The compared methods include the widely-used inner product, Euclidean distance, current correlation predictor alone, and pre-trained correlation predictor alone. As shown in Table 5, combining the performance on all the evaluation metrics, the combination of the pre-trained CP_{pre} and current CP

Table 5: **Discussion on Difficulty Scoring Function.** In the case of 'Euclidean distance', we define hard positive (negative) pairs as the pairs with long (short) distance, and 'Inner Product' vice verse.

	Acc7	Acc2	F1	MAE	Corr
Euclidean distance	47.5	85.5	85.4	**0.703**	**0.804**
Inner Product	46.9	85.7	85.6	0.710	0.795
Pre-trained CP_{pre}	47.2	85.7	85.6	0.716	0.797
Current CP	**47.6**	85.5	85.4	0.719	0.792
$CP_{pre} + CP$	47.5	**86.3**	**86.2**	0.712	0.798

gives the best result, which demonstrates the effectiveness of the proposed method in identifying the difficulty of training pairs.

In addition, all the scoring functions reach remarkable results, indicating that the curriculum learning is widely applicable given reasonable difficulty scoring function and pair feeding function. Euclidean distance achieves the best performance on MAE and Corr, but our method still slightly outperforms it considering that our improvement on Acc2 and F1 score is remarkable. We argue that this is because Euclidean distance scores the pairs by the similarity between two unimodal embeddings which cannot reveal the high-level correlation between different modalities. In contrast, our method utilizes the correlation loss as difficulty score, which is more target-oriented.

Notably, using pre-train CP_{pre} as difficulty scoring function is time-consuming considering that we have to pre-train the model. Actually, although the current CP might not be discriminative at the start of training, it ends up performing well. The other difficulty scoring functions also reach competitive performance. For the sake of time complexity, one can freely choose the time-efficient difficulty scoring functions listed in Table 5.

4.6 Visualization of Feeding Function

In this section, we visualize the values of c_i for different positive and negative training pairs on CMU-MOSEI dataset. For CMU-MOSEI, the number of partitions of training pairs is set to 10 and 8 for negative and positive pairs, respectively. As shown in Fig. 3, we observe that the learning of positive pairs is much faster than negative pairs. Especially for the positive language-visual (L+V) and acoustic-visual (A+V) pairs, their choosing index c_i reaches the largest value at the early stage of the training, and almost remains unchanged as training deepens, which indicates that the correlation between these modalities is easy to learn. For positive language-

acoustic (L+A) pairs, the choosing index c_i also reaches the largest value at the early stage, but it suffers from oscillating as training deepens, suggesting that the correlation between language and acoustic modalities is hard to learn. We argue that this is because the literal meaning of spoken language and the tone does not always coordinated (e.g., the 'tone' modality is often negative but the spoken language is positive literally if sarcastic tone was used). For negative pairs, the trend of the changes with respect to c_i is opposite to that of the positive pairs, that is, they generally suffer from large oscillation. These results suggest that identifying the modalities belonging to the same sample is easy, while identifying modalities from different samples is much more difficult. It indicates that discovering the distinguish information of each sample and learning the correspondence relationship across different samples are difficult, and points out a possible way to improve our correlation learning: paying more attention to correctly locate different samples in the embedding space.
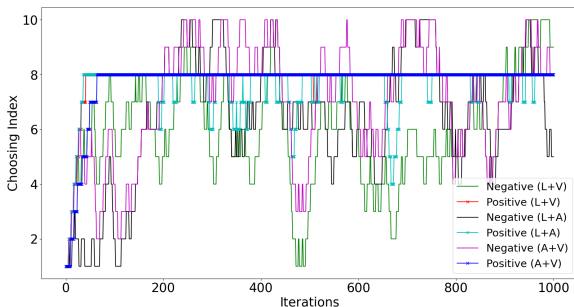

Figure 3: **Visualization of the value of the choosing index c_i for different training pairs.**

4.7 Analysis on Model Complexity

In this section, we evaluate the time complexity and space complexity of WSCL-CL. For the training time, all the compared baselines are tested under the same GPU device and experimental setting (batch size is consistent across all methods). Quantitatively, as shown in Table 6, combining the time of pre-training and training, the time consumption of WSCL-CL with pre-trained CP_{pre} is about 524 seconds on CMU-MOSI dataset. Without pre-trained CP_{pre}, the time consumption of WSCL-CL (which also achieves competitive performance) is 385 seconds. In comparison, the running time of HyCon, MULT-BERT and MFN-BERT is 440, 378, and 455 seconds, respectively. For space complexity, considering that we use a simple fusion method

and the additional trainable parameters brought by WSCL-CL are only the parameters in the correlation predictor, the space complexity is acceptable compared to other baselines using complex fusion structures. Specifically, the number of parameters of WSCL-CL is 109,917,955. For comparison, the number of parameters of HyCon, MULT-BERT and MFN-BERT is 109,209,149, 110,838,399, and 111,256,325, respectively. Overall, the model complexity of WSCL-CL is acceptable.

Table 6: **The comparison of model complexity on CMU-MOSI dataset.**

	The number of parameters	Running Time
MAG-BERT (Rahman et al., 2020)	110,853,121	293
TFN-BERT (Zadeh et al., 2017)	161,409,399	275
GFN-BERT (Mai et al., 2020)	109,232,350	274
MFN-BERT (Zadeh et al., 2018a)	111,256,325	455
MULT-BERT (Tsai et al., 2019)	110,838,399	378
HyCon (Mai et al., 2021)	109,209,149	440
WSCL-CL (W/O CP_{pre})	109,917,955	385
WSCL-CL	109,917,955	524

5 Conclusion

We introduce the curriculum learning and inject it into a weakly supervised modality correlation learning task. Via curriculum learning, we can feed the model with training pairs based on their difficulties, such that the model does not fall into a bad local optimum. The main novelty of this paper is the design of the curriculum learning, in which the difficulty scoring and pair feeding functions are proposed. The scoring function computes the difficulty of pairs using pre-trained and current correlation predictors, and the feeding function takes the difference in correlation losses as feedback to determine the feeding action. The proposed method achieves state-of-the-art performance on multimodal sentiment analysis.

6 Acknowledgement

Limitations

The main limitation of this paper is that using the learned correlation predictor might still not be a perfect scoring function. This is because the learning loss of the learned predictor does not approximate zero, and thus the learned predictor might wrongly classify a training pair into 'easy' or 'hard' pair. Moreover, the training of the pre-trained correlation predictor increases the time complexity.

References

Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. 2021. Noise estimation using density estimation for self-supervised multimodal learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6644–6652.

Maysam Behmanesh, Peyman Adibi, Mohammad Saeed Ehsani, and Jocelyn Chanussot. 2021. Geometric multimodal deep learning with multi-scaled graph wavelet convolutional network. *arXiv preprint arXiv:2111.13361*.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.

Ching-Yao Chuang, R Devon Hjelm, Xin Wang, Vibhav Vineet, Neel Joshi, Antonio Torralba, Stefanie Jegelka, and Yale Song. 2022. Robust contrastive learning against noisy views. *arXiv preprint arXiv:2201.04309*.

Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer. 2014. Covarep: A collaborative voice analysis repository for speech technologies. In *ICASSP*, pages 960–964.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Sri Harsha Dumpala, Imran Sheikh, Rupayan Chakraborty, and Sunil Kumar Kopparapu. 2019. Audio-visual fusion for sentiment classification using cross-modal autoencoder. In *32nd conference on neural information processing systems (NIPS 2018)*, pages 1–4.

Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. 2020. Large-scale adversarial training for vision-and-language representation learning. *Advances in Neural Information Processing Systems*, 33:6616–6628.

Liangke Gui, Tadas Baltrušaitis, and Louis-Philippe Morency. 2017. Curriculum learning for facial expression recognition. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, pages 505–511. IEEE.

Guy Hacohen and Daphna Weinshall. 2019. On the power of curriculum learning in training deep networks. In *International Conference on Machine Learning*, pages 2535–2544. PMLR.

Devamanyu Hazarika, R. 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*.

Onno Kampman, Elham J. Barezi, Dario Bertero, and Pascale Fung. 2018. Investigating audio, visual, and text fusion methods for end-to-end automatic personality prediction. In *ACL short paper*.

Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.

Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representations (ICLR)*.

Changhee Lee and Mihaela Schaar. 2021. A variational information bottleneck approach to multi-omics data integration. In *International Conference on Artificial Intelligence and Statistics*, pages 1513–1521. PMLR.

Qiuchi Li, Dimitris Gkoumas, Christina Lioma, and Massimo Melucci. 2021. Quantum-inspired multimodal fusion for video sentiment analysis. *Information Fusion*, 65:58 – 71.

Paul Pu Liang, Zhun Liu, Yao-Hung Hubert Tsai, Qibin Zhao, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2019. Learning representations from imperfect time series data via tensor rank regularization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1569–1576, Florence, Italy. Association for Computational Linguistics.

Fenglin Liu, Shen Ge, and Xian Wu. 2021. Competence-based multimodal curriculum learning for medical report generation. 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)*, pages 3001–3012.

Zhun Liu, Ying Shen, Paul Pu Liang, Amir Zadeh, and Louis Philippe Morency. 2018. Efficient low-rank multimodal fusion with modality-specific factors. In *ACL*, pages 2247–2256.

Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.

Sijie Mai, Haifeng Hu, and Songlong Xing. 2020. Modality to modality translation: An adversarial representation learning and graph fusion network for multimodal fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 164–172.

Sijie Mai, Ya Sun, Ying Zeng, and Haifeng Hu. 2022. Excavating multimodal correlation for representation learning. *Information Fusion*.

Sijie Mai, Ying Zeng, Shuangjia Zheng, and Haifeng Hu. 2021. Hybrid contrastive learning of tri-modal representation for multimodal sentiment analysis. *arXiv preprint arXiv:2109.01797*.

Behnaz Nojavanasghari, Deepak Gopinath, Jayanth Koushik, Tadas Baltrušaitis, and Louis-Philippe Morency. 2016. Deep multimodal fusion for persuasiveness prediction. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pages 284–288.

David Olson. 1977. From utterance to text: The bias of language in speech and writing. *Harvard Educational Review*, 47(3):257–281.

Nicolas Pielawski, Elisabeth Wetzer, Johan Öfverstedt, Jiahao Lu, Carolina Wählby, Joakim Lindblad, and Natasa Sladoje. 2020. Comir: Contrastive multimodal image representation for registration. *Advances in neural information processing systems*, 33:18433–18444.

Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. 2017a. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 37:98–125.

Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis Philippe Morency. 2017b. Context-dependent sentiment analysis in user-generated videos. In *ACL*, pages 873–883.

Soujanya Poria, Iti Chaturvedi, Erik Cambria, and Amir Hussain. 2016. Convolutional mkl based multimodal emotion recognition and sentiment analysis. In *Proceedings of IEEE International Conference on Data Mining (ICDM)*, pages 439–448.

Wasifur Rahman, M. Hasan, Sangwu Lee, Amir Zadeh, Chengfeng Mao, Louis-Philippe Morency, and E. Hoque. 2020. Integrating multimodal information in large pretrained transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2020, pages 2359–2369.

Yue Ruan, Han-Hung Lee, Ke Zhang, and Angel X Chang. 2022. Tricolo: Trimodal contrastive loss for fine-grained text to shape retrieval. *arXiv preprint arXiv:2201.07366*.

Shiv Shankar. 2022. Multimodal fusion via cortical network inspired losses. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1167–1178.

Aman Shenoy and Ashish Sardana. 2020. Multilogue-net: A context aware rnn for multi-modal emotion detection and sentiment analysis in conversation. *arXiv preprint arXiv:2002.08267*.

Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. 2022. Curriculum learning: A survey. *International Journal of Computer Vision*, pages 1–40.

Ya Sun, Sijie Mai, and Haifeng Hu. 2022. Learning to learn better unimodal representations via adaptive multimodal meta-learning. *IEEE Transactions on Affective Computing*.

Zhongkai Sun, Prathusha Sarma, William Sethares, and Yingyu Liang. 2020. Learning relationships between text, audio, and video via deep canonical correlation for multimodal language analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8992–8999.

Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.

Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020. Contrastive multiview coding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer.

Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6558–6569, Florence, Italy. Association for Computational Linguistics.

Yao-Hung Hubert Tsai, Martin Q. Ma, Muqiao Yang, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. Multimodal routing: Improving local and global interpretability of multimodal language analysis. *arXiv preprint arXiv:2001.08735*.

Vishaal Udandarao, Abhishek Maiti, Deepak Srivatsav, S. R. Vyalla, Yifang Yin, and R. Shah. 2020. Cobra: Contrastive bi-modal representation algorithm. *ArXiv*, abs/2005.03687.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

Xin Wang, Yudong Chen, and Wenwu Zhu. 2021. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. 2019. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7216–7223.

Kilian Q Weinberger and Lawrence K Saul. 2009. Distance metric learning for large margin nearest neighbor classification. *Journal of machine learning research*, 10(2).

Martin Wollmer, Felix Weninger, Tobias Knaup, Bjorn Schuller, Congkai Sun, Kenji Sagae, and Louis Philippe Morency. 2013. Youtube movie reviews: Sentiment analysis in an audio-visual context. *IEEE Intelligent Systems*, 28(3):46–53.

Jianfeng Wu, Sijie Mai, and Haifeng Hu. 2021. Graph capsule aggregation for unaligned multimodal sequences. In *Proceedings of the 2021 International Conference on Multimodal Interaction*, pages 521–529.

Kaicheng Yang, Hua Xu, and Kai Gao. 2020. Cm-bert: Cross-modal bert for text-audio sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 521–528.

Lin Yang, Yi Shen, Yue Mao, and Longjun Cai. 2022. Hybrid curriculum learning for emotion recognition in conversation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11595–11603.

Z. Yang, Zihang Dai, Yiming Yang, J. Carbonell, R. Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*.

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 International Conference on Multimedia*, pages 4400–4407.

Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis Philippe Morency. 2017. Tensor fusion network for multimodal sentiment analysis. In *EMNLP*, pages 1114–1125.

Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, and Louis Philippe Morency. 2018a. Memory fusion network for multiview sequential learning. In *AAAI*, pages 5634–5641.

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 *AAAI*, pages 5642–5649.

Amir Zadeh, Paul Pu Liang, Jonathan Vanbriesen, Soujanya Poria, Edmund Tong, Erik Cambria, Minghai Chen, and Louis Philippe Morency. 2018c. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *ACL*, pages 2236–2246.

Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis Philippe Morency. 2016. Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6):82–88.

Yuxuan Zhao, Xinyan Cao, Jinlong Lin, Dunshan Yu, and Xixin Cao. 2021. Multimodal affective states recognition based on multiscale cnns and biologically inspired decision fusion model. *IEEE Transactions on Affective Computing*, pages 1–1.

A Appendix

A.1 Unimodal Networks F_m

Due to the excellent capability of Transformer(Vaswani et al., 2017) to model long sequences, we follow HyCon(Mai et al., 2021) to use Transformer-based(Vaswani et al., 2017) model to extract high-level representations from unimodal sequences. We apply a regular Transformer encoder to extract the acoustic and visual features. Following the state-of-the-art algorithms (Hazarika et al., 2020; Rahman et al., 2020; Mai et al., 2021), BERT (Devlin et al., 2019) is used to extract the high-level language representation.

Specifically, the procedures of the BERT network are shown as below:

$$\hat{\boldsymbol{X}}_l = \text{BERT}(\boldsymbol{U}_l)$$
$$\boldsymbol{X}_l = \text{Conv 1D}\left(\hat{\boldsymbol{X}}_l, K_l\right) \in \mathbb{R}^{T_l \times d} \quad (13)$$

where Conv 1D denotes the temporal convolution with kernel size K_l set to 3, which is used for mapping the output dimensionality of BERT to the shared dimensionality d. The final language representation $\boldsymbol{x}_l$ is defined as representation of the last time step of $\boldsymbol{X}_l$, for the reason that it incorporates the information from prior time steps during processing. For the Transformer encoder network, the procedures are presented as follows:

$$\hat{\boldsymbol{X}}_m = \text{Conv 1D}\left(\boldsymbol{U}_m, K_m\right) \in \mathbb{R}^{T_m \times d}$$
$$\boldsymbol{X}_m = \text{Transformer}(\hat{\boldsymbol{X}}_m) \in \mathbb{R}^{T_m \times d} \quad (14)$$

Here we also take the representation in the last time step of $\boldsymbol{X}_m$ as the final representation $\boldsymbol{x}_m$. The generated unimodal representation $\boldsymbol{x}_m$ is used for fusion and multimodal association learning.

A.2 EXPERIMENT

A.2.1 Feature extraction Details

Facet[1] is used for extracting a set of visual features that are composed of facial action units, facial landmarks, head pose, etc. COVAREP (Degottex et al., 2014) is used for extracting acoustic features, which consist of 12 Mel-frequency cepstral coefficients, pitch tracking, speech polarity, glottal closure instants, spectral envelope, etc. Following the state-of-the-art methods (Rahman et al., 2020; Hazarika et al., 2020; Yang et al., 2020; Mai et al., 2021), BERT (Devlin et al., 2019) is used to extract the high-level textual representation. All the

baselines use the same feature set in our experiment. For CMU-MOSEI dataset, the input feature dimensionality of language, acoustic, and visual modality are 768, 74, and 35, respectively. For CMU-MOSI, the input feature dimensionality of language, acoustic, and visual modality are 768, 74, and 47, respectively.

2) Training details: We develop our model with the PyTorch framework on GTX2080Ti with CUDA 10.1 and torch version of 1.4.0. Our proposed model is trained using Adam (Kingma and Ba, 2015) optimizer.

A.2.2 Baselines

1) **Memory Fusion Network** (**MFN**) (Zadeh et al., 2018a), which designs delta-attention module and multi-view gated memory network to explore cross-modal interactions;

2) **Multimodal Transformer** (**MULT**) (Tsai et al., 2019), which learns multimodal representation by translating the source modality into the target modality using cross-modal Transformer (Vaswani et al., 2017);

3) **Graph Fusion Network** (**GFN**)(Mai et al., 2020), which designs a graph neural network to explore unimodal, bimodal, and trimodal interactions;

4) **Tensor Fusion Network** (**TFN**) (Zadeh et al., 2017), which applies outer product over unimodal representations to jointly learn unimodal, bimodal, and trimodal dynamics;

5) **Low-rank Modality Fusion** (**LMF**) (Liu et al., 2018), which leverages low-rank weight tensors to deal with the high complexity problem of tensor fusion;

6) **Interaction Canonical Correlation Network** (**ICCN**) (Sun et al., 2020), which fuses language features with visual and acoustic features respectively to obtain two bimodal representations, and then applies a canonical correlation analysis (CCA) network on the bimodal representations to generate multimodal representation;

7) **Multimodal Adaption Gate BERT** (**MAG-BERT**) (Rahman et al., 2020), which proposes an attachment called multimodal adaptation gate derived from RAVEN(Wang et al., 2019) that enables large pre-trained transformers to accept multimodal data during fine-tuning;

8) **Transformer-based Feature Reconstruction Network** (**TFR-Net**) (Yuan et al., 2021), which proposes a feature reconstruction network to improve the robustness of multimodal model for the random missing in modality sequences;

9) **Hybrid Contrastive Learning** (**HyCon**) (Mai et al., 2021), which explores inter-class relationships and intra-/inter-modal interactions using intra-/inter-modal contrastive learning and semi-contrastive learning;

10) **Adaptive Multimodal Meta-Learning** (**AMML**) (Sun et al., 2022), which applies meta-learning to learn better unimodal representation for multimodal fusion.

---

[1]iMotions 2017. https://imotions.com/