

Multimodal Contrastive Learning via Uni-Modal Coding and Cross-Modal Prediction for Multimodal Sentiment Analysis

Ronghao Lin and Haifeng Hu*

Sun Yat-sen University, China

linrh7@mail2.sysu.edu.cn, huhaiif@mail.sysu.edu.cn

Abstract

Multimodal representation learning is a challenging task in which previous work mostly focus on either uni-modality pre-training or cross-modality fusion. In fact, we regard modeling multimodal representation as building a skyscraper, where laying stable foundation and designing the main structure are equally essential. The former is like encoding robust uni-modal representation while the later is like integrating interactive information among different modalities, both of which are critical to learning an effective multimodal representation. Recently, contrastive learning has been successfully applied in representation learning, which can be utilized as the pillar of the skyscraper and benefit the model to extract the most important features contained in the multimodal data. In this paper, we propose a novel framework named MultiModal Contrastive Learning (MMCL) for multimodal representation to capture intra- and inter-modality dynamics simultaneously. Specifically, we devise uni-modal contrastive coding with an efficient uni-modal feature augmentation strategy to filter inherent noise contained in acoustic and visual modality and acquire more robust uni-modality representations. Besides, a pseudo siamese network is presented to predict representation across different modalities, which successfully captures cross-modal dynamics. Moreover, we design two contrastive learning tasks, instance- and sentiment-based contrastive learning, to promote the process of prediction and learn more interactive information related to sentiment. Extensive experiments conducted on two public datasets demonstrate that our method surpasses the state-of-the-art methods.

1 Introduction

With the surge of user-generated videos, Multimodal Sentiment Analysis (MSA) have become a hot research field, which aims to infer people's

sentiment based on multimodal data including text, audio and video (Zadeh et al., 2017; Tsai et al., 2019a, 2020; Poria et al., 2020). To successfully understand human behaviours and interpret human intents, it is necessary to attain an effective and powerful multimodal representation for the model. However, two major challenges in learning such multimodal representation exist: the accurate extraction of uni-modal features and the heterogeneities across different modalities bring difficulty of modeling cross-modal interaction.

To acquire powerful uni-modal features, Devlin et al. (2019) presents a large-scale language model named BERT for textual modality and Wu et al. (2022) introduce an audio representation learning method for audio modality by distilling from Radford et al. (2021) which targets at transferable model for visual modality. In MSA, previous methods (Yu et al., 2021; Han et al., 2021) mainly utilize BERT for textual modality while vague feature extractor such as COVAREP (Degottex et al., 2014) and Facet (iMotions 2017) for acoustic and visual modality, where the inherent noise contain in uni-modal features may still exist. To avoid uni-modal noise interfering downstream sentiment inference task, we design Uni-Modal Contrastive Coding (UMCC) which employs feature cutoff strategy inspired by (Shen et al., 2020) and generates augmentation features to construct contrastive learning task with origin uni-modal representation. As shown in Figure 1, we then obtain robust and efficient representations for acoustic and visual modalities.

To alleviate the impact of modality heterogeneity, previous MSA models propose various modalities fusion methods to learn cross-modality interaction information (Hazarika et al., 2020; Rahman et al., 2020). Modality translation is a popular method to explicitly translate source modality to the target one, which directly manipulates the commonalities across modalities (Tsai et al., 2019a; Wu et al., 2021; Zhao et al., 2021). However, due to the exist-

*Corresponding author.

tence of discrepancy modality-specific information and huge modality gap, it is undesirable and extremely difficult to project the representations from different modalities to the same one. Different with these explicit modality translation methods, we propose Cross-Modal Contrastive Prediction (CMCP) composed of a pseudo siamese predictive network and two designed contrastive learning tasks to predict cross-modal representation in an implicitly contrastive way. The predictive representation efficiently capture cross-modality dynamics and concurrently preserve modality-specific features for the modalities.

The novel contributions of our work can be summarized as follows:

- 1) We propose a framework named MultiModal Contrastive Learning (MMCL), consisting of Uni-Modal Contrastive Coding (UMCC) which mitigates the interference of modality inherent noise and learns robust uni-modal representations, and Cross-Modal Contrastive Prediction (CMCP) with a pseudo siamese predictive network which learn commonalities and interactive features across different modalities.
- 2) We design two contrastive learning tasks, instance- and sentiment-based contrastive learning, in order to improve the convergence of the predictive network and capture sentiment-related information contained in the multimodal data.
- 3) We conduct extensive experiments on two publicly available datasets, and gain superior results to the state-of-the-art MSA models.

2 Related Work

2.1 Multimodal Sentiment Analysis (MSA)

MSA focus on integrating textual, acoustic and visual modalities to comprehend varied human sentiment (Morency et al., 2011). Previous research mainly comprises of two steps: uni-modal representation learning and multimodal fusion. For uni-modal representation, Tsai et al. (2019b) factorizes them into two independent sets while Hazarika et al. (2020) projects them into two distinct subspaces. Large pre-trained Transformer-based language models such as BERT have shown great performance improvement on downstream NLP tasks (Devlin et al., 2019). However, for acoustic

and visual modalities, the features are extracted by CMU-MultimodalSDK with a vague description of feature and backbone selection(Zadeh et al., 2018c) in MSA task. We argue that without powerful pre-trained language tokenizer to extract features as textual modality does, the inherent noise of acoustic and visual features may disturb the inference of sentiment.

Different from direct processing the uni-modal features, we present uni-modal features coding to learn robust acoustic and visual representations. For multimodal fusion, Zadeh et al. (2017); Liu et al. (2018); Zadeh et al. (2018c) present early fusion at the feature level while Poria et al. (2017); Zadeh et al. (2018a); Yu et al. (2021) adopt late fusion at the decision level. However, the former methods limit capabilities in modeling cross-modal dynamics due to the inconsistent space of different modalities while the later methods suffer from neglecting modality-specific information with the absence of low-level feature process. To avoid the respective issues from the two methods, Tsai et al. (2019a,b); Hazarika et al. (2020) propose hybrid fusion which perform multimodal fusion at both input and output level. Guided by this thought, we construct a cross-modal predictive network to process representations from different modalities at early and late fusion stage, which effectively exploit the intra- and inter-modality dynamics in a prediction manner.

2.2 Contrastive Learning

The core idea of contrastive learning is to measure the similarities of sample pairs in the representation space (Hadsell et al., 2006), which is firstly adopted in the field of computer vision (He et al., 2020) and then extend to the field of nature language analysis (Gao et al., 2021). Previous work based on contrastive learning mostly only consider uni-modal data and utilize contrastive losses in a discrimination manner. Different with discrimination models, Oord et al. (2018) combines predicting future observations named predictive coding with a probabilistic contrastive loss called InfoNCE. Inspired by but diverse from this work, we apply predictive network with contrastive learning in multimodal feature to capture cross-modal dynamics and enhance the interaction among different modalities.

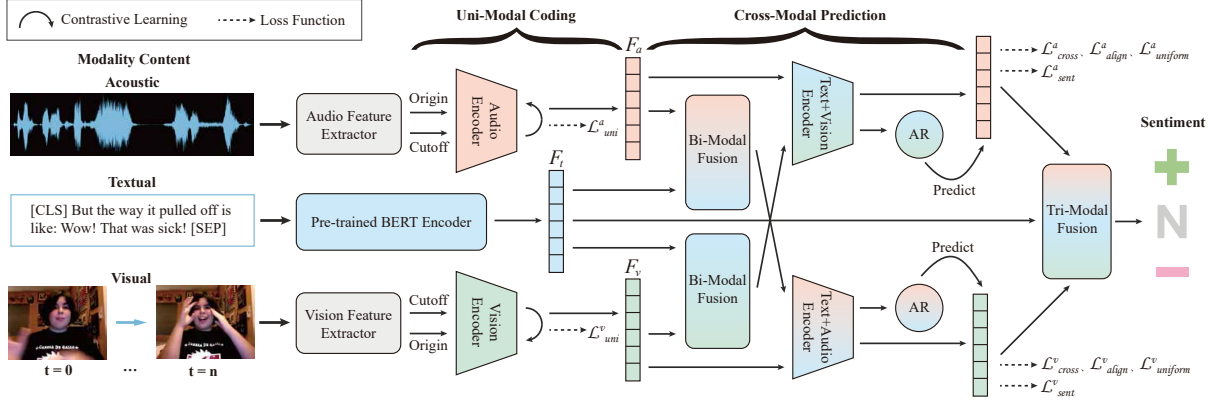


Figure 1: The overall architecture of our proposed MMCL framework.

3 Method

3.1 Problem Definition

In MSA task, the input is utterance consisting of three modalities: textual, acoustic and visual modality, where $m \in \{t, a, v\}$. The sequences of these three modalities are represented as triplet (T, A, V) , including $T \in \mathbb{R}^{N_t \times d_t}$, $A \in \mathbb{R}^{N_a \times d_a}$ and $V \in \mathbb{R}^{N_v \times d_v}$ where N_m denotes the sequence length of corresponding modality and d_m denotes the dimensionality. The goal of MSA task is to learn a mapping $f(T, A, V)$ to infer the sentiment score $\hat{y} \in \mathbb{R}$.

3.2 Overall Architecture

As shown in Figure 1, we firstly process raw input into sequential feature vectors with fixed feature extractor for audio and vision data while pre-trained BERT (Devlin et al., 2019) encoder for text. Then we utilize contrastive learning in both uni-modal coding and cross-modal prediction, which are the two key modules in our proposed model. The uni-modal coding drive the model to focus on informative features which then implicitly filter out inherent noise and produces robust and effective uni-modal representation for acoustic and visual modalities. The cross-modal prediction captures commonalities among different modalities and outputs predictive representation full of interaction dynamics. Lastly, we fuse predictive acoustic and visual representations with textual representation to derive the final multimodal representation which contains both modality-specific and cross-modal dynamics most related to sentiment.

3.3 Uni-Modal Contrastive Coding

For uni-modality, we encode the sequential triplet (T, A, V) into corresponding representa-

tions. Specifically, we use BERT (Devlin et al., 2019) to encode input sentences to obtain the hidden representations of textual modality. The embedding from the last Transformer layer’s output can be represented as:

$$F_t = BERT(T; \theta_t^{BERT}) \in \mathbb{R}^{L_t \times d_t} \quad (1)$$

To acquire more robust acoustic and visual representations, we design Uni-Modal Contrastive Coding (UMCC) for both modalities. Firstly, we encode audio and vision inputs by two uni-modal bi-directional LSTMs (Hochreiter and Schmidhuber, 1997) to capture temporal characteristic:

$$\begin{aligned} h_a &= bLSTM(A; \theta_a^{bLSTM}) \in \mathbb{R}^{L_a \times d_a} \\ h_v &= bLSTM(V; \theta_v^{bLSTM}) \in \mathbb{R}^{L_v \times d_v} \end{aligned} \quad (2)$$

To construct contrastive learning, we treat the encoded acoustic and visual representations as query samples q and get the corresponding positive key samples k^+ by feature augmentation strategy. In natural language understanding and generation task, Shen et al. (2020) introduces an efficient data augmentation approach named cutoff to erase part of the information within an input sentence and yield its restricted views during the fine-tuning stage. Inspired by this work, we utilize random feature cutoff strategy on acoustic and visual representations which randomly convert a certain proportion of embedding dimensions of every token within the sequence into a vector of zeros. As shown in Figure 2, we then generate augmented version of uni-modal representations, denoted as h_a^\dagger and h_v^\dagger .

In order to fuse with textual representations in the similar semantic space later, we design uni-modal Transformer models for acoustic and visual

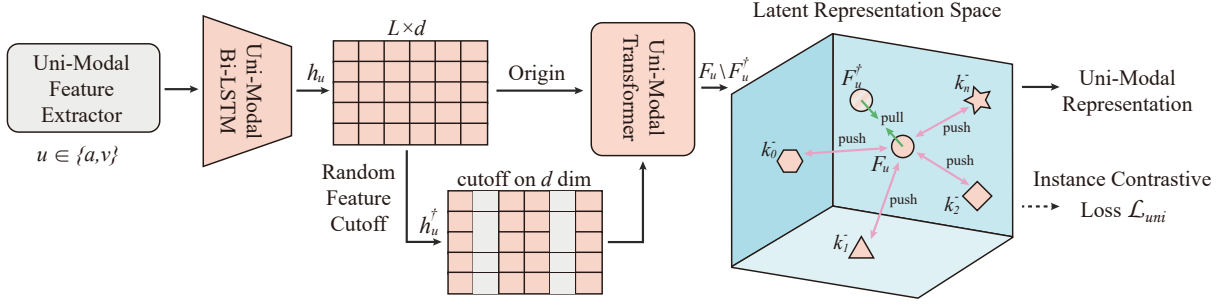


Figure 2: Illustration of Uni-Modal Contrastive Coding.

modalities, respectively. For $u \in \{a, v\}$, the query sample q and positive key sample k^+ are respectively denoted as F_u and F_u^+ :

$$\begin{aligned} F_u &= \text{Transformer}(h_u; \theta_u^{Tf}) \in \mathbb{R}^{L_u \times d_u} \\ F_u^+ &= \text{Transformer}(h_u^+; \theta_u^{Tf}) \in \mathbb{R}^{L_u \times d_u} \end{aligned} \quad (3)$$

Given a batch set $F_{uni} = \{F_u^0, F_u^1, \dots, F_u^{n-1}\}$, noted that there is a single positive key F_u^+ (as k^+) that each encoded query F_u^i (as q , $i \in [1, n]$) matches, while the other representations F_u^j ($j \in [0, n]$ and $j \neq i$) in the same batch are considered as negative key samples k^- , as the instance discrimination task (Wu et al., 2018; Ye et al., 2019) does. With the similarity measured by dot product, we present the uni-modal instance contrastive loss \mathcal{L}_{uni} in InfoNCE (Oord et al., 2018) form:

$$\begin{aligned} \mathcal{L}_{uni}^u &\triangleq -\log \frac{\exp(q \cdot k^+ / \tau)}{\sum_{i=1}^n \exp(q \cdot q^i / \tau)} \\ &= -\mathbb{E}_{F_{uni}} \left[\log \frac{\exp(F_u \cdot F_u^+ / \tau)}{\sum_{i=1}^n \exp(F_u \cdot F_u^i / \tau)} \right] \end{aligned} \quad (4)$$

where τ is a temperature hyper-parameter that controls the probability distribution over distinct instances (Hinton et al., 2015). Due to $u \in \{a, v\}$, the final uni-modal instance contrastive loss $\mathcal{L}_{uni} = \mathcal{L}_{uni}^a + \mathcal{L}_{uni}^v$.

Since each dimension of representation contains certain features, with a certain number of features erased entirely, the UMCC can impel uni-modal Transformer models to implicitly eliminate the inherent noise of uni-modal data and capture informative semantic information which are the most essential to predict the sentiment. Doing so, we can finally acquire a robust modality-specific representation for acoustic and visual modalities.

3.4 Cross-Modal Contrastive Prediction

To further learn inter-modality dynamics and focus on the commonalities related to sentiment among

different modalities, we present Cross-Modal Contrastive Prediction (CMCP) as shown in Figure 3. Specifically, we construct a pseudo siamese predictive network to utilize two modalities to predict another one. Through the network we can attain the original representation and the corresponding predictive representation. For intra- and inter-embedding of these two representations, we design two contrastive learning tasks based on data instances and sentiment labels, respectively. The experiment in Section 4.3 indicate the effectiveness of the designed contrastive learning tasks.

3.4.1 Pseudo Siamese Predictive Network

The proposed pseudo siamese predictive network consists of a bi-modal fusion network, pseudo siamese convolutional neural networks (CNNs), an auto-regressive model and a cross-modal predictor.

As shown in Figure 3, we take acoustic and textual modalities to predict visual modality as an example. Firstly, on account of various feature dimensions of different modalities, we utilize a linear layer to project three representations to the same dimension.

Secondly, we concatenate acoustic and textual representations and feed them into the bi-modal fusion network, which is a multilayer perceptron (MLP). After that, we can obtain the bi-modal representation denoted as F_{ta} , where $F_{ta} = MLP(\{F_t; F_a\})$. To mine out the commonalities among three modalities, the bi-modal representation F_{ta} are utilized to predict the visual modality representation F_v subsequently.

Thirdly, compressing high-dimension embedding into a more compact latent representation space can make conditional predictions easier to the predictive network (Oord et al., 2018). To project F_{ta} and F_v into the same compact space, we design two CNNs encoders in pseudo siamese structure (Bromley et al., 1993). As shown in Figure 3, the

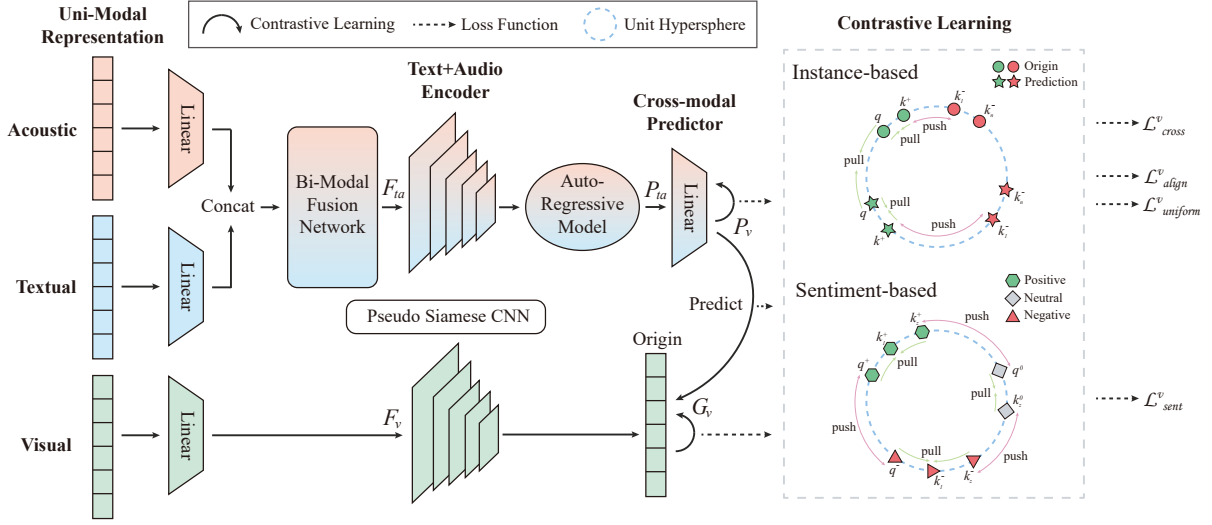


Figure 3: Illustration of Cross-Modal Contrastive Prediction along with two contrastive learning tasks.

pseudo siamese CNNs means that two CNNs have the same architecture setting but with unshared parameters, which can model the inputs in the same way while with different neuron weights according to their own characteristics. In this compact latent space, visual representation is denoted as G_v , which is the prediction target subsequently.

After mapping the bi-modal representations F_{ta} by CNN encoder, an auto-regressive model is applied to summarize and produce a context latent representation P_{ta} , where $P_{ta} = sLSTM(F_{ta})$. We regard P_{ta} as an inter-modality dynamics container which covers common features shared by bi-modalities and visual modality.

At last, we use a linear projection layer as a cross-modal predictor and get the predictive visual representation P_v , where $P_v = Linear(P_{ta})$. For $u \in \{a, v\}$, Now the problem has become how to make the prediction P_u more consistent with the target G_u while retaining the most useful information for downstream tasks.

3.4.2 Instance-based Contrastive Learning

To maximize the similarity between the prediction representation P_u and target representation G_u , we propose the Instance-based Contrastive Learning (ICL) to force the predictive network learning effective features in the original representation.

In details, we firstly perform L2-normalization on both P_u and G_u to restrict the predictive space to the unit hypersphere (Wang and Isola, 2020), as shown in Figure 3 with $u = v$ as an example. Considering the instability of prediction training

at the beginning, we then successively take the combination of original and prediction representation as query and key, meaning that $\{query, key\}$ can be $\{origin, origin\}$, $\{predict, predict\}$ and $\{origin, predict\}$. The reason of this setting will be further discussed with experiment evidence in Section 5.1. In ICL, each query q has a corresponding key as k^+ while the other representations in the same batch are seen as k^- . Similar with uni-modal instance contrastive loss \mathcal{L}_{uni} , the cross-modal instance contrastive loss \mathcal{L}_{cross} is presented as:

$$\mathcal{L}_{cross} \triangleq - \mathbb{E}_{F_{cross}} \left[\log \frac{\exp(F_c \cdot F_c^+ / \tau)}{\sum_{i=1}^n \exp(F_c \cdot F_c^i / \tau)} \right] \quad (5)$$

where $F_c \setminus F_c^+ \in \{P_u, G_u\}$, $u \in \{a, v\}$ and $F_{cross} = \{F_c^1, \dots, F_c^n\}$.

Due to the discrepancy modality-specific information contained in different modalities, there exists huge modality gap across the representations of modalities. Intuitively, contrastive learning task on cross-modal prediction is far more difficult than on uni-modal coding. Inspired by (Wang and Isola, 2020) which identifies two key properties related to contrastive loss, we introduce two more contrastive loss for better optimization of the predictive cross-modal instance-based contrastive learning:

- Cross-modal alignment loss \mathcal{L}_{align} to map query sample q and positive key sample k^+ from different modalities to nearby features and thus be mostly invariant to the modality gap, straightforwardly defined with the ex-

pected distance among two samples:

$$\mathcal{L}_{align} \triangleq \mathbb{E}[\|P_u - G_u\|_2^\lambda], \lambda > 0 \quad (6)$$

- Cross-modal uniformity loss $\mathcal{L}_{uniform}$ to distribute predictive and target representations roughly uniformly on the unit hypersphere and preserve as much modality-specific information as possible, defined as the logarithm of the average pairwise Gaussian potential (Cohn and Kumar, 2007):

$$\mathcal{L}_{uniform} \triangleq \log \mathbb{E} \left[e^{-\kappa \|P_u - G_u\|_2^2} \right], \kappa > 0 \quad (7)$$

3.4.3 Sentiment-based Contrastive Learning

For the purpose of concentrating on sentiment related features contained in multimodal representations, we construct Sentiment-based Contrastive Learning (SCL) to help the modal learn more discriminative representation for different polarities of sentiments.

Firstly, according to the sentiment labels, we divide the representations into three sentiment classes $\{positive(+), neutral(0), negative(-)\}$. Then, we treat the representations from the same sentiment class as positive pairs, represented as $\{q^s, k_1^s, \dots, k_z^s\}$ where $s \in \{+, 0, -\}$ and $z \leq n$. The other representations from different sentiment classes in the same batch are treated as negative samples. At last, the sentiment-based contrastive learning loss \mathcal{L}_{sent} can be given by:

$$\mathcal{L}_{sent} \triangleq - \mathbb{E}_{F_{sent}} \left[\log \frac{\sum_{j=1}^z \exp(q_s \cdot k_j^s / \tau)}{\sum_{i=1}^n \exp(q_s \cdot k_i^s / \tau)} \right] \quad (8)$$

where F_{sent} denotes original and corresponding predictive representations divided by sentiment classes. Here we consider both original and predictive representations which is conducive to guide the predictive network to preform hard sample mining when the sentiments across different modalities are variant.

By means of the sentiment-based contrastive learning, we extend contrastive learning approach to a fully-supervised setting as (Khosla et al., 2020) does and effectively leverages sentiment label information for better downstream sentiment classification performance.

3.5 Total Training Loss

At last, we fuse the predictive acoustic and visual representations along with the textual representation to attain the final multimodal representation,

represented as $F_M = MLP(\{F_t; P_a; P_v\})$, and predict the final sentiment score \hat{y} . Along with the truth sentiment label y , we can have the regression task loss \mathcal{L}_{reg} as:

$$\mathcal{L}_{reg} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (9)$$

where n is the number of training samples. Combined with the designed contrastive losses, the total loss for training is formulated as:

$$\begin{aligned} \mathcal{L}_{total} = & \mathcal{L}_{reg} + \mu \mathcal{L}_{uni} + \eta \mathcal{L}_{sent} \\ & + \alpha \mathcal{L}_{cross} + \beta \mathcal{L}_{align} + \gamma \mathcal{L}_{uniform} \end{aligned} \quad (10)$$

where μ, η, α, β and γ are weighted hyperparameters that adjust the impact of various loss functions.

4 Experiments

We conduct extensive experiments on two public datasets which offer both token-aligned and -unaligned data for multimodal sentiment analysis. The details about the datasets, evaluation metrics and baseline methods are provided subsequently.

4.1 Datasets and Evaluation Metrics

CMU-MOSI (Zadeh et al., 2016) is a popular benchmark dataset collected from YouTube in MSA research, including 93 monologues where speakers make their comment on a specific topics. The dataset consists of 2,199 opinion video segments with a total of 26,295 words in the utterances and is annotated with sentiment intensity label ranged from -3 (strongly negative) to +3 (strongly positive).

CMU-MOSEI (Zadeh et al., 2018b) is a large dataset of multimodal sentiment analysis and emotion recognition, containing 23,454 YouTube monologues video segments covering 250 distinct topics from 1,000 distinct speakers. The utterances in the dataset are randomly chosen on various movie review topics, annotated with sentiment scores between -3 and +3 and 6 different emotion classes.

Both of the datasets are split into train, validation and test sets as Han et al. (2021) does.

We use public evaluation metrics of classification and regression to demonstrate the performance of our proposed framework and further compare with baselines: seven-class classification accuracy (Acc7) indicating the correct sentiment label predictions in the range of [-3, +3], binary classification

Models	CMU-MOSI					CMU-MOSEI				
	Acc7↑	Acc2↑	F1↑	MAE↓	Corr↑	Acc7↑	Acc2↑	F1↑	MAE↓	Corr↑
TFN*	33.7	78.3/80.2	78.2/80.1	0.925	0.662	52.2	81.0/82.6	81.1/82.3	0.570	0.716
LMF*	32.7	77.5/80.1	77.3/80.0	0.931	0.670	52.0	81.3/83.7	81.6/83.8	0.568	0.727
MFN	34.2	77.9/80.0	77.8/80.0	0.951	0.665	51.1	81.8/84.0	81.9/83.9	0.575	0.720
MFM	33.3	77.7/80.0	77.7/80.1	0.948	0.664	50.8	80.3/83.4	80.7/83.4	0.580	0.722
MuT	35.0	79.0/80.5	79.0/80.5	0.918	0.685	52.1	81.3/84.0	81.6/83.9	0.564	0.732
MISA	43.5	81.8/83.5	81.7/83.5	0.752	0.784	52.2	81.6/84.3	82.0/84.3	0.550	0.758
MAG-BERT	45.1	82.4/84.6	82.2/84.6	0.730	0.789	52.8	81.9/85.1	82.3/85.1	0.558	0.761
Self-MM	45.8	82.7/84.9	82.6/84.8	0.731	0.785	53.0	82.6/85.2	82.8/85.2	0.540	0.763
MMIM	45.0	83.0/85.1	82.9/85.0	0.738	0.781	53.1	81.9/85.1	82.3/85.0	0.547	0.752
MMCL(ours)	46.5	84.0/86.3	83.8/86.2	0.705	0.797	53.6	84.8/85.9	84.8/85.7	0.537	0.765

Table 1: Performance Comparison between MMCL and baselines on CMU-MOSI and CMU-MOSEI datasets. The multimodal data are token-aligned except for unaligned models with *.

(Acc2) and F1-score using two calculation settings marked as segmentation symbol ‘-/-’ where left represents non-negative/negative (has-0) and right denotes positive/negative (non-0); mean absolute error (MAE) computing the average absolute difference between predicted and truth labels, Pearson correlation (Corr) measuring the degree of prediction skew.

4.2 Baselines

The mentioned baselines in the experiment are introduced in detail in the following.

TFN (Zadeh et al., 2017) Tensor Fusion Network introduces a multi-dimensional tensor by calculating the outer-product among different modalities to capture uni-modal, bi-modal and tri-modal interactions.

LMF (Liu et al., 2018) Low-rank Multimodal Fusion decomposes stacked high-order multimodal tensors into low-rank weight tensors to reduce computational complexity and perform efficient fusion.

MFN (Zadeh et al., 2018a) Memory Fusion Network separately leverages LSTM to encodes information from each modality and utilizes a delta-memory attention network with a multi-view gated memory to explicitly accounts for the cross-view interaction.

MFM (Tsai et al., 2019b) Multimodal Factorization Model presents jointly optimize multimodal discriminative factors and modality-specific generative factors to reconstruct missing modalities and interpret the interactions that influence multimodal learning.

MuT (Tsai et al., 2019a) Multimodal Trans-

former extends three sets of Transformers with directional pairwise cross-modal attention which latently adapts streams from one modality to another.

MISA (Hazarika et al., 2020) Modality-Invariant and -Specific Representations projects modalities to two distinct subspaces, modality-invariant and -specific subspace, to provide a holistic view of multimodal data.

MAG-BERT (Rahman et al., 2020) Multimodal Adaption Gate for BERT applies multimodal adaption gate to generate a shift for the internal representation of pre-trained Transformer models.

Self-MM (Yu et al., 2021) Self-Supervised Multi-Task Learning automatically generates uni-modal labels which are weight-adjusted by multimodal labels to learn consistency and difference across modalities.

MMIM (Han et al., 2021) Multimodal Mutual Information Maximization maintain task-related information through maximizing the mutual information in uni-modal input pairs and between multimodal fusion output and uni-modal input.

4.3 Results and Ablation Study

In accord with baselines, our proposed model is run five times under the same setting of hyperparameters and present the average performance in Table 1. Significantly, MMCL outperforms SOTA in all metrics whether on CMU-MOSI or on CMU-MOSEI. These outcomes preliminarily demonstrate the effectiveness of our method in MSA task.

To further explore the contributions of proposed uni-modal coding and cross-modal predictive net-

Description	Acc7↑	Acc2↑	F1↑	MAE↓	Corr↑
(1) UMCC					
rp Transformer	51.0	81.9/84.7	82.2/84.6	0.567	0.741
w/o \mathcal{L}_{uni}	52.4	82.6/85.5	82.9/85.3	0.554	0.752
(2) CMCP					
w/o ICL	52.0	81.7/84.7	82.0/84.6	0.550	0.755
w/o \mathcal{L}_{cross}	53.0	83.8/85.2	83.7/84.9	0.543	0.761
w/o \mathcal{L}_{align}	52.9	83.9/85.0	83.8/84.6	0.554	0.756
w/o $\mathcal{L}_{uniform}$	52.3	84.4/85.5	84.4/85.2	0.553	0.757
w/o SCL	52.2	82.2/84.9	82.6/84.8	0.546	0.756
(3) No Contrast					
MMCL	53.6	84.8/85.9	84.8/85.7	0.537	0.765

Table 2: Ablation study of MMCL on CMU-MOSEI dataset. Note that "rp Transformer" denotes to replace uni-modal Transformers with linear layers.

work with corresponding contrastive loss functions in MMCL, we carry out a series of ablation experiments on CMU-MOSEI. As shown in Table 2, for UMCC, we firstly substitute uni-modal Transformers with linear layers and discover the great performance degradation of the proposed model, which means that the powerful Transformer encoders play a crucial part in uni-modal coding. Then we remove the uni-modal instance contrastive loss \mathcal{L}_{uni} and note that although Acc2 (non-0) and F1 (non-0) drop slightly, the model have a sharp decline on other metrics. We consider that \mathcal{L}_{uni} mainly focus on the inherent noise contained in uni-modal data. Without feature augmentation with \mathcal{L}_{uni} , the model is easier to be interrupted by hard samples such as the nearly neutral samples with sentiment labels around 0. For CMCP, we eliminate the loss terms in ICL individually and the results indicates that all the loss functions are useful for the sentiment prediction, which can be further observed in Section 5.1. Moreover, ablating either ICL or SCL hurt the model performance significantly. This demonstrates the effectiveness of two designed contrastive learning tasks to the cross-modal predictive network. Finally, we spot a clear drop in all metrics when taking off all contrastive losses and just training the model with just regression task loss \mathcal{L}_{reg} . The results of ablation study imply that each module and loss function are necessary and essential for MMCL to achieve the best performance.

5 Further Analysis

5.1 Losses Tracing

To better understand how every contrastive loss works, we visualize the variation of all losses during training in Figure 4. At the beginning of train-

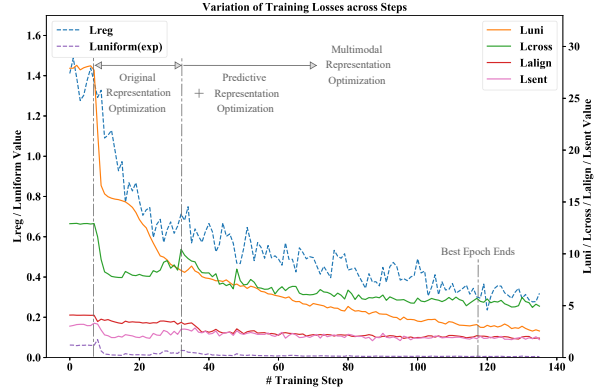


Figure 4: Visualization of losses changing as training proceeds on CMU-MOSI. The values for plotting are the average losses in a constant interval of every 5 steps and $\mathcal{L}_{uniform}$ is exponentiated for plotting purposes.

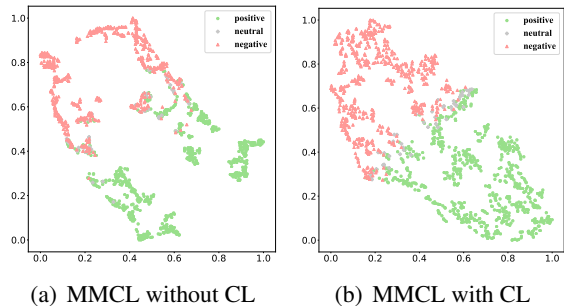


Figure 5: T-SNE (van der Maaten and Hinton, 2008) visualization of multimodal representation in the embedding space on the training set of CMU-MOSI.

ing, MMCL tends to optimize the original representation for each modality since all losses decrease almost concurrently. Then when other losses flatten out, \mathcal{L}_{uni} continue to decline which means UMCC starts to filter out the inherent noises of hard samples. Next, we observe that \mathcal{L}_{cross} and $\mathcal{L}_{uniform}$ increased slightly where MMCL attempt to optimize the predictive representation, which is due to the extreme difficulty for the predictive network to inference across different modalities. Succeeding in reducing the modality gap and promoting the network to converge, the designed contrastive tasks ICL and SCL make the losses drop again with a smaller slope. Finally, the training of MMCL ends at the best epochs when the task loss on the validation set reaches the minimum. In the optimization process of multimodal representation, the designed losses are productive to the model in an alternating manner.

Unaligned SOTA Models	CMU-MOSI					CMU-MOSEI				
	Acc7↑	Acc2↑	F1↑	MAE↓	Corr↑	Acc7↑	Acc2↑	F1↑	MAE↓	Corr↑
MuT	36.5	78.8/80.2	78.7/80.2	0.899	0.687	52.5	81.0/84.0	81.4/84.0	0.561	0.737
MISA	43.6	81.6/83.7	81.5/83.6	0.752	0.779	52.0	81.1/84.3	81.6/84.4	0.554	0.753
Self-MM	45.9	82.6/84.9	82.5/84.8	0.722	0.789	53.1	82.2/84.9	82.4/85.0	0.540	0.760
MMIM	45.5	82.6/85.0	82.5/84.9	0.726	0.786	53.2	82.2/85.1	82.6/85.0	0.544	0.756
MMCL (ours)	46.8	83.8/85.8	83.7/85.7	0.692	0.800	53.4	84.5/85.7	84.6/85.6	0.538	0.765

Table 3: Performance Comparison between MMCL and baselines on CMU-MOSI and CMU-MOSEI datasets in a token-unaligned manner with token-unaligned multimodal data.

5.2 Representation Visualization

Figure 5 displays the visualization of fusion multimodal representation F_M calculated by MMCL with contrastive learning losses or not. Without contrastive learning, the representations in the same batch are clustered too tight while far away from each other in the different batches, which means that the model is inclined to overfit the training dataset. Meanwhile, the hard samples with nearly neutral sentiment are indistinguishable for the model due to the absence of \mathcal{L}_{uni} and \mathcal{L}_{sent} . After introducing designed contrastive learning, MMCL is more capable of mining the commonalities between different instances with the same sentiment polarity, while retaining the specific features contained in different instances. In the interim, the hard samples are more distinguishable for corresponding sentiment class which also proves the effectiveness of the designed contrastive learning tasks on representation learning.

5.3 Token-Unaligned Exploration

Additionally, we explore the performance of MMCL framework on token-unaligned data as shown in Table 3. All models are run five times the same as the experiments on token-aligned data. Basically, the performance of models trained on token-unaligned data are similar to the ones trained on token-aligned data with a slight drop on some metrics. Nevertheless, compared with unaligned state-of-the-art models, MMCL still achieves the best result on all metrics, which manifests the significant power of the proposed framework.

6 Conclusion

In this paper, we present a novel framework named MultiModal Contrastive Learning (MMCL) which efficiently adopts contrastive learning in multi-

modal sentiment analysis. MMCL consists of Uni-Modal Contrastive Coding (UMCC) which learns robust uni-modal representations by focusing on informative features and reducing the interference of inherent noise for acoustic and visual modalities, and Cross-Modal Contrastive Prediction (CMCP) which construct a powerful pseudo siamese predictive network to learn commonalities and interactive dynamics across different modalities. Besides, we propose two effective contrastive learning tasks, instance- and sentiment-based contrastive learning, to help the predictive network maintain modality-specific dynamics and learn sentiment-related information in multimodal data, which improve the performance of the model significantly. The extensive experiments conducted on CMU-MOSI and CMU-MOSEI demonstrate the superiority of our model and the efficacy of contrastive learning. Moreover, the further visualization of losses and representation space provide comprehensive insight into our model. We believe our work can boost the creativity in contrastive learning and multimodal representation learning in the future.

Limitations

Below we describe the limitations of the proposed model in our view and suggest directions for future work. Firstly, there are still a lot of data augmentation strategy can be further examined in the process of encoding uni-modal representation. Then, with the proposed powerful predictive network, we can extend the predictive method into dealing with missing modality problem which is a more challenging research emphasis. At last, the training process of proposed model is still a supervised learning process. Since contrastive learning is increasingly adopted in semi- and self-supervised learning which burst out with great potential in

computer vision, we attempt to remove the supervised labels and explore its capability in the field of multimodal in the future work.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62076262.

References

- Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. [Signature verification using a " siamese" time delay neural network](#). *Advances in neural information processing systems*, 6.
- Henry Cohn and Abhinav Kumar. 2007. [Universally optimal distribution of points on spheres](#). *Journal of the American Mathematical Society*, 20(1):99–148.
- 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)*, 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.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- R. Hadsell, S. Chopra, and Y. LeCun. 2006. [Dimensionality reduction by learning an invariant mapping](#). In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742.
- Wei Han, Hui Chen, and Soujanya Poria. 2021. [Improving multimodal fusion with hierarchical mutual information maximization for multimodal sentiment analysis](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9180–9192, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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*, pages 1122–1131.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. [Momentum contrast for unsupervised visual representation learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#). *arXiv preprint arXiv:1503.02531*, 2.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- iMotions 2017. Facial expression analysis. [Online]. <https://imotions.com/>.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. [Supervised contrastive learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, AmirAli Bagher Zadeh, and Louis-Philippe Morency. 2018. [Efficient low-rank multimodal fusion with modality-specific factors](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2247–2256, Melbourne, Australia. Association for Computational Linguistics.
- Louis-Philippe Morency, Rada Mihalcea, and Payal Doshi. 2011. [Towards multimodal sentiment analysis: Harvesting opinions from the web](#). In *Proceedings of the 13th International Conference on Multimodal Interfaces, ICMI '11*, page 169–176, New York, NY, USA. Association for Computing Machinery.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. [Representation learning with contrastive predictive coding](#). *arXiv e-prints*, pages arXiv–1807.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. [Context-dependent sentiment analysis in user-generated videos](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 873–883, Vancouver, Canada. Association for Computational Linguistics.

- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, and Rada Mihalcea. 2020. [Beneath the tip of the iceberg: Current challenges and new directions in sentiment analysis research](#). *IEEE Transactions on Affective Computing*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, AmirAli Bagher Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. 2020. [Integrating multimodal information in large pretrained transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2359–2369, Online. Association for Computational Linguistics.
- Dinghan Shen, Mingzhi Zheng, Yelong Shen, Yanru Qu, and Weizhu Chen. 2020. [A simple but tough-to-beat data augmentation approach for natural language understanding and generation](#). *arXiv preprint arXiv:2009.13818*.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019a. [Multimodal transformer for unaligned multimodal language sequences](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6558–6569, Florence, Italy. Association for Computational Linguistics.
- Yao-Hung Hubert Tsai, Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019b. [Learning factorized multimodal representations](#). In *International Conference on Learning Representations*.
- Yao-Hung Hubert Tsai, Martin Ma, Muqiao Yang, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. [Multimodal routing: Improving local and global interpretability of multimodal language analysis](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1823–1833, Online. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Tongzhou Wang and Phillip Isola. 2020. [Understanding contrastive representation learning through alignment and uniformity on the hypersphere](#). In *International Conference on Machine Learning*, pages 9929–9939. PMLR.
- Ho-Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello. 2022. [Wav2clip: Learning robust audio representations from clip](#). In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4563–4567. IEEE.
- Yang Wu, Zijie Lin, Yanyan Zhao, Bing Qin, and Lidan Zhu. 2021. [A text-centered shared-private framework via cross-modal prediction for multimodal sentiment analysis](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4730–4738, Online. Association for Computational Linguistics.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. [Unsupervised feature learning via non-parametric instance discrimination](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742.
- Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. 2019. [Unsupervised embedding learning via invariant and spreading instance feature](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6210–6219.
- Wenmeng Yu, Hua Xu, Ziqi Yuan, and Jiele Wu. 2021. [Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10790–10797.
- Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2017. [Tensor fusion network for multimodal sentiment analysis](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1103–1114, Copenhagen, Denmark. Association for Computational Linguistics.
- Amir Zadeh, Paul Pu Liang, Navonil Majumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018a. [Memory fusion network for multi-view sequential learning](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Amir Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018b. [Multimodal language analysis in the wild: CMU-MOSEI dataset and interpretable dynamic fusion graph](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2236–2246, Melbourne, Australia. Association for Computational Linguistics.
- Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, and Louis-Philippe Morency. 2018c. [Multi-attention recurrent network for human communication comprehension](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. 2016. [Mosi: multimodal cor-](#)

pus of sentiment intensity and subjectivity analysis in online opinion videos. *arXiv preprint arXiv:1606.06259*.

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)*, pages 2608–2618, Online. Association for Computational Linguistics.

A Appendix

A.1 Implementation Details

To fairly compare the effectiveness of baselines and our proposed frameworks, we reproduce the baselines based on the same pre-trained BERT language model as text features encoder and ran hyper-parameter grid search for the best results. All models are trained on a single GTX 1080Ti GPU. As shown in Table 1, noted that except for unaligned models TFN and LMF, the multimodal data are token-aligned in the pre-processing stage.

We perform grid-search over finite sets of options for hyper-parameters, including τ in $\{0.05, 0.1, 0.7\}$, μ in $\{0.6, 0.7, 0.8\}$, η in $\{0.8, 0.9, 1.0\}$, α in $\{0.8, 0.9, 1.0\}$, β in $\{0.6, 0.75, 0.9\}$ and γ in $\{0.05, 0.1, 0.2\}$ and dropout values in $\{0.1, 0.6\}$, to select the best model. For optimization, we apply Adam (Kingma and Ba, 2014) as the optimizer with a linear warmup learning rate scheduler, using the highest learning rate $5e-5$ for BERT finetuning and $5e-3$ on CMU-MOSI and $1e-3$ on CMU-MOSEI for other parameters. The batch size on both datasets is set as 32. In UMCC, the uni-modal Transformers have a stack of 3 identical layers with 8 parallel attention heads for both acoustic and visual modalities. The random cutoff ratio is set as 0.2. In CMCP, the pseudo siamese CNNs have five convolutional layers with strides $[5, 4, 2, 2, 2]$, filter-sizes $[10, 8, 4, 4, 4]$ and 256 hidden units with *Tanh* activation layers while the sLSTM as auto-regressive model has 128 dimensional hidden states.

A.2 Explanation of Eq.6-7

Eq.6 aims at aligning the query-positive samples pairs in the contrastive learning due to the fact that the feature distance of samples pairs should be minimized if they are from query-positive samples pairs, which are the prediction and target representations in CMCP. Here $\lambda \in \{1, 2\}$ illustrates the

form of loss function as L1 and L2 loss, respectively.

Eq. 7 is proposed according to the thought that the feature distribution of the predictive and target representations need to preserve maximal modality-specific information, which is the uniform distribution on the unit hypersphere. For this purpose, the Gaussian potential kernel (Cohn and Kumar, 2007), known as the Radial Basis Function kernel K are considered as:

$$K \triangleq e^{-\kappa \|P-G\|_2^2}, \quad \kappa > 0 \quad (11)$$

Then in Eq. 7, the uniformity is further defined as the logarithm of the average pairwise Gaussian potential, which is properly tied with the uniform distribution on the unit hypersphere.

A.3 Ablation Study on Multiple Modalities

Modality	Acc7 \uparrow	Acc2 \uparrow	F1 \uparrow	MAE \downarrow	Corr \uparrow
<i>A</i>	41.6	70.0/62.8	60.5/51.1	0.835	0.116
<i>V</i>	40.9	70.9/62.8	59.2/48.9	0.841	0.109
<i>T</i>	52.6	80.7/85.3	81.4/85.4	0.552	0.761
<i>A + V</i>	41.4	71.0/62.9	59.0/48.5	0.837	0.148
<i>T + A</i>	53.3	82.4/85.5	82.8/85.4	0.540	0.762
<i>T + V</i>	53.6	83.6/ 85.9	83.9/ 85.9	0.538	0.763
<i>T + A + V</i>	53.6	84.8/85.9	84.8/85.7	0.537	0.765

Table 4: Ablation study of MMCL on multiple modalities on CMU-MOSEI dataset. Note that *T*, *A*, *V* represent textual, acoustic and vision modality, respectively.

To evaluate the contribution of each modality, we conduct ablation study of MMCL on multiple modalities as shown in Table 4. For better observing the effect in performance without changing the architecture of the proposed model, we remove modality by setting the corresponding representation to all zero vector. Firstly, we conclude that the multimodal representation learned by MMCL provides the best performance, indicating that the designed network and multimodal contrastive learning tasks can effectively capture complementary features across different modalities. In addition, the performance dramatically drops when the textual modality is removed which do not happen when removing the other modalities. The results show that the textual modality plays a significant dominant part in MSA task due to inherently better data quality and powerful pre-trained model. Although UMCC can reduce the interference of uni-modal

noise and focus more on the task-related features for acoustic and visual modalities, the randomly initialized uni-modal feature extractor still lack of comparable information extraction capability with BERT.

A.4 Discussion on Varying Hyperparameters

As shown in Table 5, to further explore the relationship of different losses, we conduct ablation experiments of MMCL on CMU-MOSEI dataset under different loss hyperparameters setting. We can observe that there are no significant performance drop in Table 5, which concludes that the performance of MMCL is independent from the balance hyperparameters setting. Due to the different losses contribute to different modules of MMCL, the variation of hyperparameters can even help the model achieve better accuracy in several metrics.

Varying Hyperparameters					
Value	Acc7 \uparrow	Acc2 \uparrow	F1 \uparrow	MAE \downarrow	Corr \uparrow
$\mu = 0.4$	54.1	83.2/85.5	83.5/85.4	0.536	0.768
$\mu = 0.5$	53.5	83.1/ 86.1	83.5/ 86.0	0.537	0.760
$\mu = 0.6$	53.6	83.6/85.9	83.8/85.7	0.536	0.763
$\mu = 0.7$	53.6	84.8/85.9	84.8/85.6	0.537	0.765
$\mu = 0.8$	54.1	83.0/85.6	83.3/85.5	0.533	0.764
$\mu = 0.9$	52.7	83.8/85.6	84.0/85.5	0.543	0.764
$\eta = 0.4$	53.7	83.0/85.7	83.3/85.6	0.532	0.765
$\eta = 0.6$	52.4	84.5/84.9	84.4/84.6	0.544	0.760
$\eta = 0.8$	54.1	84.1/85.1	84.2/84.9	0.536	0.763
$\eta = 1.0$	53.6	84.8/85.9	84.8/85.6	0.537	0.765
$\eta = 1.2$	53.0	83.8/85.8	84.0/ 85.7	0.541	0.766
$\eta = 1.4$	53.6	84.1/85.5	84.1/85.2	0.538	0.760
$\alpha = 0.4$	53.6	82.4/85.3	82.7/85.2	0.532	0.768
$\alpha = 0.6$	53.8	82.8/85.5	83.1/85.4	0.539	0.759
$\alpha = 0.8$	53.3	84.4/85.4	84.4/85.2	0.542	0.765
$\alpha = 1.0$	53.6	84.8/85.9	84.8/85.6	0.537	0.765
$\alpha = 1.2$	53.8	83.7/85.6	83.8/85.4	0.536	0.762
$\alpha = 1.4$	52.3	84.8/86.1	84.8/85.9	0.546	0.765

Table 5: Ablation study of MMCL on varying the loss hyperparameters μ, η, α on CMU-MOSEI dataset, which aims at reweighting the contribution of $\mathcal{L}_{uni}, \mathcal{L}_{sent}, \mathcal{L}_{cross}$ in Eq. 10, respectively.