

Self-MI: Efficient Multimodal Fusion via Self-Supervised Multi-Task Learning with Auxiliary Mutual Information Maximization

Cam-Van Thi Nguyen, Ngoc-Hoa Thi Nguyen,
Duc-Trong Le, Quang-Thuy Ha

Faculty of Information Technology
VNU - University of Engineering and Technology, Hanoi, Vietnam

Abstract

Multimodal representation learning poses significant challenges in capturing informative and distinct features from multiple modalities. Existing methods often struggle to exploit the unique characteristics of each modality due to unified multimodal annotations. In this study, we propose Self-MI in the self-supervised learning fashion, which also leverage Contrastive Predictive Coding (CPC) as an auxiliary technique to maximize the Mutual Information (MI) between unimodal input pairs and the multimodal fusion result with unimodal inputs. Moreover, we design a label generation module, ULG_{MI} for short, that enables us to create meaningful and informative labels for each modality in a self-supervised manner. By maximizing the Mutual Information, we encourage better alignment between the multimodal fusion and the individual modalities, facilitating improved multimodal fusion. Extensive experiments on three benchmark datasets including CMU-MOSI, CMU-MOSEI, and SIMS, demonstrate the effectiveness of Self-MI in enhancing the multimodal fusion task.

1 Introduction

Multimodal sentiment analysis (MSA) has emerged as a prominent research area, garnering significant attention in recent years (Bagher Zadeh et al., 2018; Poria et al., 2023). It demonstrates remarkable enhancements compared to the traditional unimodal analysis. Real-life multimodal data typically involves three distinct channels: visual, acoustic, and textual. Effective representation and information fusion from these diverse sources are essential components in MSA.

In multimodal learning, Baltrušaitis et al. (2018) highlight five fundamental challenges including alignment, translation, representation, fusion, and co-learning. Among these challenges, representation learning is a significant and challenging task. The goal is to obtain effective representations that

capture both the commonalities and distinctions across different modalities. An effective representation should include two key aspects: consistency and differentiation (Poria et al., 2020). The representations should be consistent across modalities to enable seamless integration and analysis of information from various sources. Equally important, it should also capture the unique characteristics of each modality to facilitate a comprehensive understanding of the data.

Yu et al. (2021) point out the shortcomings of current methods in capturing distinctive information due to their reliance on unified multimodal annotation. However, incorporating additional unimodal annotations can be burdensome in terms of time and human resource. As a solution for this challenge, Yu et al. (2021) introduce a novel approach named Self-MM. The main idea is to incorporate a measure based on the distance between modality representations and the class centroids, which correlates positively with the model’s output. In order to generate unimodal labels, the authors implement a self-supervised learning strategy, of which the absolute distance calculation between two modality representations in different spaces is intractable. Dealing with this problem, Han et al. (2021) propose MultiModal InfoMax (MMIM), which maximize the shared information in multimodal fusion via enhancing two types of mutual information: between unimodal representations and between fusion embeddings and unimodal representations.

Building upon the insights from MMIM (Han et al., 2021) and addressing the limitations of Self-MM (Yu et al., 2021), this paper presents a novel approach named Self-MI that computes the similarity between modality representations using the maximize mutual information technique, while also incorporating self-supervised learning for the MSA task. Additionally, designing a multitask learning model using the hard parameter sharing strategy

allows us to gain a comprehensive understanding of how unimodal tasks impact multimodal tasks. To be clear, our contributions can be summarized as the following:

- We propose an unified model called Self-MI, which is based on Self-supervised learning and Multi-task Learning with the auxiliary of Mutual Information (MI).
- We design a module named Mutual Information Maximization for Unimodal Labels Generation (ULG_{MI}) that utilizes the Contrastive Predictive Coding (CPC) method to estimate the correlation between multimodal representations.
- We perform extensive experiments on three benchmark datasets including CMU-MOSI, CMU-MOSEI, and CMU-MOSEI. Our model results in superior or comparable results compared to state-of-the-art models.

2 Related Works

In this section, we literately review several classes of previous research works related to the multimodal sentiment analysis task and our proposed model Self-MI.

2.1 Multimodal Sentiment Analysis

Multimodal Sentiment Analysis (MSA) is a comprehensive approach that combines verbal and non-verbal features to perform user sentiment analysis. This field focuses on extracting emotions, interpretations, and feelings by analyzing various sources such as language, facial expressions, speech, music, and movements (Kaur and Kautish, 2022). The field of multimodal sentiment analysis commonly utilizes various public datasets such as CMU-MOSI (Zadeh et al., 2016), CMU-MOSEI (Bagher Zadeh et al., 2018), CH-SIMS (Yu et al., 2020), IEMO-CAP (Busso et al., 2008), and several others.

MSA research could be divided into four groups: 1) *early multimodal fusion methods* like Tensor Fusion Network TFN (Zadeh et al., 2017), Low-rank Multimodal Fusion LMF (Liu et al., 2018), and Multimodal Factorization Model MFM (Tsai et al., 2019b), and 2) *the methods that fuse multimodality through modeling modality interaction*, such as multimodal Transformer MulT (Tsai et al., 2019a) and modal-temporal attention graph MTAG (Yang et al., 2021) and 3) *the methods focusing on the consistency and the difference of modality*, in which

MISA (Hazarika et al., 2020), Self-MM (Yu et al., 2021), and MMIM (Han et al., 2021). Our work focuses on the third group of methods, in line with previous studies (Yu et al., 2021; Han et al., 2021). We also use these models as the baseline models in this study. We provide detailed descriptions of these models in the corresponding Section 4.

2.2 Self-Supervised Multimodal Learning

Multimodal learning, which seeks to comprehend and analyze information from various modalities, has made significant advancements in the supervised regime in recent years. However, the reliance on paired data and costly human annotations hinders the scalability of models. On the other hand, with the abundance of unannotated data available in the wild, self-supervised learning has emerged as an appealing strategy to address the annotation bottleneck (Zong et al., 2023). The objective functions for training self-supervised multimodal algorithms can be categorized into three main types: instance discrimination, clustering, and masked prediction. In this study, our model belongs to the category of instance discrimination, specifically contrastive learning. By leveraging the ULGM module in Self-MM (Yu et al., 2021), we re-design this module and name it ULG_{MI} (Mutual Information Maximization for Unimodal Labels Generation). The purpose of this module is to generate uni-modal supervision values using multimodal annotations and modality representations. In contrast to Self-MM (Yu et al., 2021), we replace the Relative Distance Value, which evaluates the relative distance from the modality representation to the positive center and the negative center, with the computation of Mutual Information Maximization, utilizing the auxiliary of Contrastive Predictive Coding (CPC) (Oord et al., 2018a).

2.3 Multi-task Learning

Multi-task learning aims to enhance the generalization performance of multiple related tasks by leveraging the knowledge from different tasks (Zhang and Yang, 2021). In contrast to single-task learning, multi-task learning faces two primary challenges during the training stage. The first challenge is how to effectively share network parameters, which can be achieved through hard-sharing or soft-sharing methods. The second challenge is to balance the learning process across different tasks. Multi-task learning has been widely applied in multimodal sentiment analysis (MSA) (Akhtar et al., 2019).

In this work, we propose the incorporation of unimodal subtasks to assist in the process of learning modality-specific representations.

2.4 Mutual Information

In probability theory and information theory, Mutual Information (MI) is a measure of the dependence between two random variables. It quantifies how much knowing the value of one variable reduces uncertainty about the value of the other variable. In the context of machine learning, maximizing mutual information can help to learn informative representations and improve the performance of tasks such as classification and generation.

$$I(X; Y) = \mathbb{E}_{p(x,y)} \left[\log \frac{p(x,y)}{p(x)p(y)} \right] \quad (1)$$

Alemi et al. (2016) (Alemi et al., 2016) were the pioneers in integrating optimization techniques related to mutual information into deep learning models. Since then, numerous studies (Amjad and Geiger, 2019; He et al., 2020) have investigated and demonstrated the advantages of maximizing mutual information principles in various contexts. However, estimating mutual information directly in high-dimensional spaces is often considered impractical or challenging.

3 Methodology

In this section, we present the overall architecture of Self-MI as well as its mathematical formulation for the multimodal sentiment analysis task.

3.1 Task setup

A MSA model takes unimodal raw sequences $X_m \in \mathbb{R}^{l_m \times d_m}$ from the same video fragment as input, where l_m is the sequence length, and d_m is the dimension of the representation vector for modality m . Here, we consider three modalities denoted as $m \in \{t, a, v\}$, representing text, visual, and acoustic modalities respectively. The main objective of Self-MI is to integrate information from these input vectors, and create a unified representation to accurately predict the sentiment intensity reflected via the truth value y .

The overall architecture of Self-MI is illustrated in Figure 1. The model comprises one multimodal task and three separate unimodal subtasks. Between the multimodal task and the unimodal tasks, we exploit the MI (Mutual Information) maximization with Contrastive Predictive Coding (CPC) as the fusion module in our model.

3.2 Multimodal Representation

The multimodal sentiment analysis task is tackled as a classification model, which consists of three key components: (1) **The feature representation module**, (2) **The feature fusion module**, and (3) **The classification module**.

For text modality, the sequence is denoted as $T = \{w_C, w_0, \dots, w_i, \dots, w_S\}$, where w_C and w_S represent the special tokens [CLS] and [SEP] respectively. The pre-trained 12-layers BERT model is utilized to extract the global text representation and the local text representation. Empirically, the first-word vector in the last layer is identified as the optimal representation for the entire sentence denoted as X_t :

$$X^t = BERT(T; \theta_t^{bert}) \quad (2)$$

where $X^t = \{x_C^t, x_0^t, \dots, h_i^t, \dots, h_S^t\}$, $X^t \in \mathbb{R}^{l^t \times d^t}$; l^t is the max length of the text sequence; and d^t is the dimension of the text representation.

For audio and visual modality, we use a pre-trained ToolKits to extract the initial vector features, i.e., A and V. Inspired by previous works (Hazarika et al., 2020; Yu et al., 2021), a single directional Long Short-Term Memory (sLSTM) model (Hochreiter and Schmidhuber, 1997) is employed to capture temporal characteristics. The final step adopts the end-state hidden vectors as the comprehensive sequence representations:

$$X^a = sLSTM(A; \theta_a^{lstm}) \in \mathbb{R}^{d^a} \quad (3)$$

$$X^v = sLSTM(V; \theta_v^{lstm}) \in \mathbb{R}^{d^v} \quad (4)$$

where $X^a = \{x_0^a, \dots, h_j^a, \dots, h_{l^a}^a\}$, $X^a \in \mathbb{R}^{l^a \times d^a}$, $X^v = \{x_0^v, \dots, h_k^v, \dots, h_{l^v}^v\}$, $X^v \in \mathbb{R}^{l^v \times d^v}$; l^a and l^v are the sequence length of audio and vision, respectively.

Subsequently, all uni-modal representations are concatenated and projected into a lower-dimensional space \mathbb{R}^{d_m} .

$$X^m = [X^t, X^a, X^v] \quad (5)$$

$$Z^m = ReLU(W_{l_1}^{mT} X^m + b_{l_1}^m) \quad (6)$$

where $W_{l_1}^{mT} \in \mathbb{R}^{(d_t+d_a+d_v) \times d_m}$ and $ReLU$ is the relu activation function.

Finally, the fusion representation Z_m is utilized to generate the multimodal sentiment prediction as:

$$\hat{y}_m = W_{l_2}^{mT} Z^m + b_{l_2}^m \quad (7)$$

where $W_{l_2}^{mT} \in \mathbb{R}^{d_m \times 1}$.




Figure 1: The overall architecture of our propose model Self-MI

3.3 Unimodal Task

The three unimodal tasks leverage the same modality representations as the multimodal task. To reduce the dimensionality gap between different modalities, we project the representations into a new feature space. Then, linear regression is used to obtain the unimodal results as follows:

$$Z^s = \text{ReLU}(W_{l_1}^s X^s + b_{l_1}^s) \quad (8)$$

$$\hat{y}_s = W_{l_2}^s Z^s + b_{l_2}^s \quad (9)$$

where $s \in \{t, a, v\}$. In order to guide the unimodal task's training process, we design a Unimodal Label Generation with Mutual Information Module (ULG_{MI}) to get unimodal labels. Section 3.4 delves into the specifics of the ULG_{MI} :

$$y_s = ULG_{MI}(y_m, Z^m, Z^s) \quad (10)$$

Ultimately, the multimodal task and three unimodal tasks are jointly learned using human-annotated multimodal labels, i.e., m-labels, and auto-generated unimodal labels, i.e., u-labels. It should be emphasized that these unimodal tasks only exist during the training stage. Consequently, y^m is used as the final output.

3.4 Mutual Information Maximization for Unimodal Labels Generation - ULG_{MI}

Since there exists a generation path from X^s to Z^m , we expect an opposite path to construct X^s with $s \in \{t, a, v\}$. The objective is to create unimodal supervised data representations based on human-labeled multimodal data and modality-specific representations.

Unimodal supervised data representations have a strong correlation with multimodal data labels, hence, the Contrastive Predictive Coding (CPC) method, a widely used unsupervised approach for high-dimensional data, is applied to estimate the correlation between the multimodal representation. (Oord et al., 2018b) use CPC to measure the mutual information between contextual information and future elements across a time horizon. In other words, CPC seeks to capture the ‘‘slow features’’ that extend over multiple time steps, ensuring their retention in the encoding process. CPC uses a contrastive loss function based on estimated contrastive noise, called InfoNCE (Normalized Mutual Information-based Contrastive Estimation). The InfoNCE loss function is designed to measure the similarity between pairs of data samples by calculating the mutual information between them. Specifically, InfoNCE leverages the concepts of en-

tropy and mutual information from information theory to compute the distance between pairs of data samples. Therefore, when using InfoNCE, we are essentially estimating the shared information and maximizing it using gradient descent. Inspired by this idea, Self-MI implements Z^m to reversely predict representation across modalities so that more modality-invariant information can be passed to Z^m .

As depicted in Figure 1, after extracting unimodal representations and multimodal representations, we seek to maximize the correlation between the two modality’s representations via CPC. The CPC score function is defined as follows:

$$\overline{G_\phi(Z^m)} = \frac{G_\phi(Z^m)}{\|G_\phi(Z^m)\|^2} \quad (11)$$

$$\overline{Z^s} = \frac{Z^s}{\|Z^s\|^2} \quad (12)$$

We utilizes Euclidean normalization to compute the unit-length vectors. Specifically, G_ϕ represents a simple neural network with parameters ϕ .

$$s(Z^m, Z^s) = \exp(\overline{Z^s}(\overline{G_\phi(Z^m)})^T) \quad (13)$$

$$\mathcal{L}_N(Z^m, Z^s) = \mathbb{E}_F[\log \frac{s(Z^m, Z^s)}{\sum_j^N s(Z^m, Z^s)}] \quad (14)$$

Here, \mathcal{L}_N representing the CPC loss function between two vectors Z^m and Z^s , where $s \in \{t, a, v\}$, and N being the number of samples in a batch, the sample index is the sum of elements within the same batch, excluding the current element. From this, we can derive the following relationship:

$$\mathcal{L}_{CPC} = \mathcal{L}_N^{mt} + \mathcal{L}_N^{ma} + \mathcal{L}_N^{mv} \quad (15)$$

The CPC output will be passed through a regression multilayer perceptron (MLP) to generate the unimodal label y^s . Along with the human-generated multimodal label y_m , we can define the loss function as follows:

$$\mathcal{L}_{task} = MAE(y_s, y_m) \quad (16)$$

The unimodal labels update policy is shown in Algorithm 1.

Algorithm 1 ULG_{MI}

Input: unimodal input X_t, X_a, X_v , m-labels y_m

Output: u-label $y_t^{(i)}, y_a^{(i)}, y_v^{(i)}$ where i is the number of training epochs

Initialize model parameters $M(\theta; x)$

Initialize u-labels $y_t^{(1)} = y_m, y_a^{(1)} = y_m, y_v^{(1)} = y_m$

Initialize global representations $F_t^g = 0, F_a^g = 0, F_v^g = 0, F_m^g = 0$

for $n \in [1, end]$ **do**

for mini-batch in dataLoader **do**

 Compute mini-batch modality representations Z^t, Z^a, Z^v, Z^m

 Compute loss \mathcal{L} using Equation 16, 17

 Compute parameters gradient $\frac{\partial \mathcal{L}}{\partial \theta}$

 Update model parameters: $\theta = \theta - n \frac{\partial \mathcal{L}}{\partial \theta}$

if $n \neq 1$ **then**

 Compute CPC loss \mathcal{L}_{CPC} using Eq. 14 and Eq.15.

 Compute y_t, y_a, y_v from mutual information learning from CPC model + MLP.

end if

 Update global representations F_s^g using Z^s , where $s \in \{m, t, a, v\}$

end for

end for

3.5 Optimization Objectives

Finally, we utilize L1Loss as a fundamental optimization method. For unimodal tasks, we employ the discrepancy between the unimodal label \mathbf{u} and the multimodal label \mathbf{m} as the weighting factor for the loss function. This indicates that the system should pay more attention to samples with significant differences.

$$L = \frac{1}{N} \sum_i^N (|y_m^i - y_m^i| + \sum_s^{\{t,a,v\}} W_s^i \cdot |y_s^i - y_s^i|) \quad (17)$$

where N representing the number of samples used for training, $W_s^i = \tanh(|y_s - y_m|)$ denotes the weight of the i^{th} sample supporting task s .

4 Experimental Settings

In this section, we present the details of extensive experiments, which investigate the advantages of

Table 1: Dataset statistics in CMU-MOSI, CMU-MOSEI, and SIMS.

Dataset	Train	Valid	Test	Total
MOSI	1,284	229	686	2,199
SIMS	1,368	456	457	2,281

Self-MI compared to state-of-the-art baselines in term of evaluation metrics.

4.1 Dataset

In this study, we utilize two publicly available multimodal sentiment analysis datasets: CMU-MOSI (Zadeh et al., 2016), and SIMS (Yu et al., 2020). Table 1 provides an overview of the basic statistics for each dataset.

CMU-MOSI: The CMU-MOSI dataset (Zadeh et al., 2016), is widely recognized as one of the prominent benchmark datasets for Multimodal Sentiment Analysis (MSA). Each sample in the dataset is annotated by human annotators with a sentiment score ranging from -3 (strongly negative sentiment) to 3 (strongly positive sentiment).

MOSEI: The CMU-MOSEI dataset (Bagher Zadeh et al., 2018) provides a larger number of utterances, increased sample variety, speakers, and topics compared to CMU-MOSI. Similarly to MOSI, annotators in the CMU-MOSEI dataset label each sample with a sentiment score ranging from -3 (strongly negative) to 3 (strongly positive).

SIMS: The SIMS dataset (Yu et al., 2020) is a unique benchmark dataset for Chinese Multimodal Sentiment Analysis that offers fine-grained annotations of modality. Each sample in the dataset is annotated by human annotators with a sentiment score ranging from -1 (strongly negative sentiment) to 1 (strongly positive sentiment).

4.2 Baselines

To thoroughly evaluate the performance of our model Self-MI, we conduct a comprehensive comparison with state-of-the-art baselines for the MSA task as follows:

- **Tensor Fusion Network (Zadeh et al., 2017):** TFN disentangles unimodal data into tensors through a threefold Cartesian product and computes the outer product of these tensors for fusion.
- **Low-rank Multimodal Fusion (Liu et al., 2018):** LMF decomposes stacked high-order

Table 2: The details of Hyperparameter Setting

Hyper-parameters	CMU-MOSI	MOSEI	SIMS
batch size	32	64	32
learning rate			
textual modality (BERT)	5e-4	1e-3	5e-3
BERT embedding size	768	768	768
text dropout	0.0	0.1	0.1
learning rate			
audio modality	0.001	0.001	0.005
audio embedding size	16	16	16
audio dropout	0.1	0.1	0.0
learning rate			
visual modality	1e-3	1e-3	0.005
visual embedding size	16	64	16
visual dropout	0.1	0.1	0.0
fusion dropout	0.1	0.1	0.0
sentiment score range	-3 to 3	-3 to 3	-1 to 1

tensors into multiple low-rank factors and efficiently performs fusion based on these factors.

- **Multimodal Transformer (Tsai et al., 2019a):** MulT constructs an architecture with separate unimodal and crossmodal transformer networks and completes the fusion process using attention mechanisms.
- **Modal-temporal Attention Graph (Yang et al., 2021):** MTAG is an interpretable graph attention network model capable of both fusion and alignment.
- **MISA (Hazarika et al., 2020):** MISA projects features into separate two spaces with specific constraints and performs fusion on these features.
- **Self-MM (Yu et al., 2021):** Self-MM assigns each modality a unimodal training task with automatically generated labels, aiming to adjust the gradient backpropagation.
- **MMIM (Han et al., 2021):** MMIM improves Multimodal Fusion with Hierarchical Mutual Information Maximization.

4.3 Experimental Settings

The hyperparameters used in the experiments are described in Table 2.

4.4 Evaluation Metrics

Followed by (Han et al., 2021), we consider various evaluation metrics namely mean absolute error (MAE), Pearson correlation (Corr), binary classification accuracy (Acc₂), F1 score computed for positive/negative and non-negative/negative classification.

Table 3: Experimental results on CMU-MOSI and CMU-MOSEI dataset. For F1_score and Acc_2, the “/” sign separates the value of “negative/non-negative” or “negative/positive” performance. The _ sign is no information. Equally important, the sign ↓ indicates “smaller is better” while ↑ signifies “higher is better”. The bold style stands for the best-performing model whilst the underline marks the second best.

Model	CMU-MOSI				MOSEI			
	MAE ↓	Corr ↑	Acc_2 ↑	F1_score ↑	MAE ↓	Corr ↑	Acc_2 ↑	F1_score ↑
TFN	0.901	0.698	_/80.8	_/80.7	0.593	0.677	-/82.5	-/82.1
LMF	0.917	0.695	_/82.5	_/82.4	0.623	0.700	-/82.0	-/82.1
MFM	0.877	0.706	-/81.7	-/81.6	0.568	0.703	-/84.4	-/84.3
MTAG	0.866	0.722	-/82.3	-/82.1	-	-	-	-
MuT	0.861	0.711	81.50/84.10	80.60/83.90	0.580	0.713	-/82.5	-/82.3
MISA	0.804	0.764	80.79/82.10	80.77/82.03	0.568	0.717	82.59/84.23	82.67/83.97
Self-MM	<u>0.718</u>	<u>0.791</u>	82.56/84.48	82.46/84.44	0.530	0.765	82.81/85.17	82.53/85.30
MMIM	0.72	0.78	<u>82.8/84.76</u>	<u>82.63/84.66</u>	0.526	<u>0.772</u>	82.24/85.97	82.66/ 85.94
Ours (Self-MI)	0.699	0.80	83.09/85.37	82.95/85.3	<u>0.527</u>	0.773	83.86/85.45	83.90/85.37

5 Results & Discussion

In this section, we provide a detailed analysis and discussion on the experimental results, which are obtained from Self-MI and baselines, on three datasets: CMU-MOSI, CMU-MOSEI, and SIMS.

5.1 Comparison With Baselines

Table 3 shows the comparative results on CMU-MOSI and CMU-MOSEI datasets. Generally, Self-MI significantly outperforms baselines across all evaluation metrics on the CMU-MOSI dataset. For the CMU-MOSEI dataset, Self-MI outperforms the SOTA models on Corr, Acc_2, and F1 score, while being comparable in MAE. These results provide initial evidence of the effectiveness of our approach in MSA tasks.

As the SIMS dataset consists of unaligned data, we compare Self-MI with TFN, LMF, and Self-MM methods. The experimental results are presented in Table 4. It is clear that Self-MI has a sustainable performance over all baselines.

Table 4: Experimental results on SIMS dataset. The sign ↓ indicates “smaller is better” while ↑ signifies “higher is better”. The bold style stands for the best-performing model whilst the underline marks the second best.

Model	MAE ↓	Corr ↑	Acc_2 ↑	F1_score ↑
TFN	0.428	0.605	79.86	80.15
LMF	0.431	0.600	79.37	78.65
Self-MM	0.415	0.608	78.21	78.2
Ours (Self-MI)	0.402	0.6138	79.65	79.82

5.2 Ablation Study

5.2.1 Effectiveness of unimodal subtasks

Clearly, Self-MI contributes significantly to the overall performance of the model when examining

the effectiveness of combining various unimodal tasks. The experimental results under different ablation settings are presented in Table 5. In comparing to the single-task model with the incorporation of unimodal subtasks, we observe a notable improvement in performance. Interestingly, the combinations “M, T” and “M, T, V” achieve comparable or even better results than the combination “M, T, A, V”. Additionally, the subtasks “T” and “A” seem to have a more positive impact on the overall performance compared to the subtask “V”. These findings offer valuable insights into the effectiveness of Self-MI in leveraging different unimodal tasks to enhance multimodal learning.

5.2.2 Effectiveness of CPC loss

To highlight the advantages of the CPC loss functions in Self-MI, we conduct a series of ablation experiments on the CMU-MOSI dataset. The results are listed in Table 6, where we eliminated one of these loss terms (\mathcal{L}_N^{ma} , \mathcal{L}_N^{mv} , \mathcal{L}_N^{mt}) from the total CPC loss. We observe the similar phenomena on other datasets.

From the initial CPC loss function, there exists a performance degradation after removing a portion of the CPC loss. Specifically, when using \mathcal{L}_{CPC} , Self-MI shows a stable performance with a slightly higher MAE score compared to using only \mathcal{L}_N^{mt} or \mathcal{L}_N^{ma} , with insignificant differences of 0.0042 and 0.0075, respectively. Similarly, in the Acc_2 metric, when using \mathcal{L}_{CPC} , the performance was slightly lower than when using only \mathcal{L}_N^{mt} or \mathcal{L}_N^{ma} , with insignificant differences of 0.3 and 0.15, respectively. This demonstrates the effectiveness of maximizing mutual information as proposed in the Self-MI model.

Table 5: Perspective analysis results from multimodal data with different tasks on CMU-MOSI dataset. M, T, A, V are multimodal, text, audio and visual tasks, respectively. For F1_score and Acc_2, the “/” sign separates the value of “negative/non-negative” or “negative/positive” performance. Equally important, the sign ↓ indicates “smaller is better” while ↑ signifies “higher is better”. The bold style stands for the best-performing model.

Task	MAE ↓	Corr ↑	Acc_2 ↑	F1_score ↑
M	0.7155	0.796	81.67/82.13	81.55/83.08
M,T	0.7217	0.7911	81.34/83.23	81.32/83.27
M,A	0.6955	0.8027	82.67/82.98	83.59/85.96
M,V	0.6908	0.8051	83.53/85.12	83.42/85.8
M,T,A	0.7215	0.801	83.61/85.15	84.52/86.21
M,T,V	0.7172	0.8012	84.49/85.93	84.51/85.01
M,V,A	0.7121	0.7962	82.36/83.84	82.35/83.88
M,T,V,A	0.6999	0.8014	83.09/85.37	82.95/85.3

Table 6: Results of CPC on CMU-MOSI. Symbol t, v, a, m represent text, images, audio and video. For F1_score and Acc_2, the “/” sign separates the value of “negative/non-negative” or “negative/positive” performance. Equally important, the sign ↓ indicates “smaller is better” while ↑ signifies “higher is better”. The bold style stands for the best-performing model.

Loss	MAE ↓	Corr ↑	Acc_2 ↑	F1_score ↑
w/o \mathcal{L}_N^{mt}	0.6957	0.8016	83.53/85.67	83.41/85.62
w/o \mathcal{L}_N^{ma}	0.6924	0.8019	83.24/85.52	83.05/85.42
w/o \mathcal{L}_N^{mv}	0.7014	0.7982	81.78/83.69	81.4/83.71
\mathcal{L}_{CPC} $\mathcal{L}_N^{mt}, \mathcal{L}_N^{ma}, \mathcal{L}_N^{mv}$	0.6999	0.8014	83.09/85.37	82.95/85.30

6 Conclusion

In this paper, we introduce Self-MI, a model designed to enhance efficient multimodal representation. Our method includes a label generation module based on self-supervised learning, which enables us to obtain independent unimodal supervisions. Moreover, we maximize the Mutual Information (MI) between unimodal input pairs and between the multimodal fusion result and unimodal inputs, utilizing the auxiliary of Contrastive Predictive Coding (CPC). This comprehensive approach allows for better representation learning and more effective multimodal fusion in MSA task. We thoroughly evaluate the performance of our model on three benchmark datasets, and our comprehensive ablation study further validates the effectiveness of our proposed approach. We anticipate that our work will serve as a source of inspiration for advancing representation learning and multimodal sentiment analysis in future research.

7 Acknowledgement

Cam-Van Nguyen Thi was funded by the Master, PhD Scholarship Programme of Vingroup Innovation Foundation (VINIF), code VINIF.2022.TS143.

References

- Md Shad Akhtar, Dushyant Chauhan, Deepanway Ghosal, Soujanya Poria, Asif Ekbal, and Pushpak Bhattacharyya. 2019. [Multi-task learning for multimodal emotion recognition and sentiment analysis](#). 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 370–379, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Rana Ali Amjad and Bernhard C Geiger. 2019. Learning representations for neural network-based classification using the information bottleneck principle. *IEEE transactions on pattern analysis and machine intelligence*, 42(9):2225–2239.
- AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. [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.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443.

- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42:335–359.
- 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ramandeep Kaur and Sandeep Kautish. 2022. Multimodal sentiment analysis: A survey and comparison. *Research Anthology on Implementing Sentiment Analysis Across Multiple Disciplines*, pages 1846–1870.
- 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.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018a. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018b. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- 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*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, and Rada Mihalcea. 2023. Beneath the tip of the iceberg: Current challenges and new directions in sentiment analysis research. *IEEE Trans. Affect. Comput.*, 14(1):108–132.
- 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 *ICLR*.
- Jianing Yang, Yongxin Wang, Ruitao Yi, Yuying Zhu, Azaan Rehman, Amir Zadeh, Soujanya Poria, and Louis-Philippe Morency. 2021. [MTAG: Modality-temporal attention graph for unaligned human multimodal language sequences](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1009–1021, Online. Association for Computational Linguistics.
- Wenmeng Yu, Hua Xu, Fanyang Meng, Yilin Zhu, Yixiao Ma, Jiele Wu, Jiyun Zou, and Kaicheng Yang. 2020. [CH-SIMS: A Chinese multimodal sentiment analysis dataset with fine-grained annotation of modality](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3718–3727, Online. Association for Computational Linguistics.
- Wenmeng Yu, Hua Xu, Ziqi Yuan, and Jiele Wu. 2021. Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis. 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, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. 2016. [MOSI: multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos](#). *CoRR*, abs/1606.06259.
- Yu Zhang and Qiang Yang. 2021. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609.
- Yongshuo Zong, Oisín Mac Aodha, and Timothy Hospedales. 2023. Self-supervised multimodal learning: A survey. *arXiv preprint arXiv:2304.01008*.