

# UniS-MMC: Multimodal Classification via Unimodality-supervised Multimodal Contrastive Learning

Heqing Zou, Meng Shen, Chen Chen, Yuchen Hu, Deepu Rajan, Eng Siong Chng  
Nanyang Technological University, Singapore

{heqing001, meng005, chen1436, yuchen005}@e.ntu.edu.sg, {asdrajan, aseschnng}@ntu.edu.sg

## Abstract

Multimodal learning aims to imitate human beings to acquire complementary information from multiple modalities for various downstream tasks. However, traditional aggregation-based multimodal fusion methods ignore the inter-modality relationship, treat each modality equally, suffer sensor noise, and thus reduce multimodal learning performance. In this work, we propose a novel multimodal contrastive method to explore more reliable multimodal representations under the weak supervision of unimodal predicting. Specifically, we first capture task-related unimodal representations and the unimodal predictions from the introduced unimodal predicting task. Then the unimodal representations are aligned with the more effective one by the designed multimodal contrastive method under the supervision of the unimodal predictions. Experimental results with fused features on two image-text classification benchmarks UPMC-Food-101 and N24News show that our proposed **Unimodality-Supervised MultiModal Contrastive (UniS-MMC)** learning method outperforms current state-of-the-art multimodal methods. The detailed ablation study and analysis further demonstrate the advantage of our proposed method.

## 1 Introduction

Social media has emerged as an important avenue for communication. The content is often multimodal, e.g., via text, speech, audio, and videos. Multimodal tasks that employ multiple data sources include image-text classification and emotion recognition, which could be used for specific applications in daily life, such as web search (Chang et al., 2022), guide robot (Moon and Seo, 2019). Hence, there is a need for an effective representation strategy for multimodal content. A common way is to fuse unimodal representations. Despite the recent progress in obtaining effective

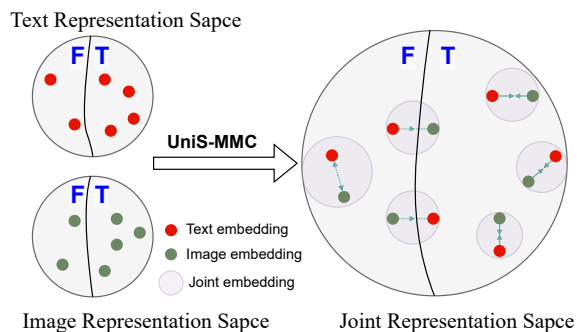


Figure 1: Unimodal representation of a single modality can be either effective or not. The effectiveness of different unimodal representations from the same sample also varies. To empower the interaction between modalities, our proposed method aligns the unimodal representation to the effective modality sample-wise and makes full use of the effective unimodal representation under the supervision of the unimodal prediction (F and T represent correct and incorrect predictions, respectively).

unimodal representations from large pre-trained models (Devlin et al., 2019; Liu et al., 2019; Dosovitskiy et al., 2021), fusing for developing more trustworthy and complementary multimodal representations remains a challenging problem in the multimodal learning area.

To solve the multimodal fusion problem, researchers propose aggregation-based fusion methods to combine unimodal representations. These methods include aggregating unimodal features (Castellano et al., 2008; Nagrani et al., 2021), aggregating unimodal decisions (Ramirez et al., 2011; Tian et al., 2020a), and aggregating both (Wu et al., 2022) of them. However, these aggregation-based methods ignore the relation between modalities that affects the performance of multimodal tasks (Udandarao et al., 2020). To solve this issue, the alignment-based fusion methods are introduced to strengthen the inter-modality relationship by aligning the embeddings among different modalities. Existing alignment-based methods can be

divided into two categories: architecture-based and contrastive-based. The architecture-based methods introduce a specific module for mapping features to the same space (Wang et al., 2016) or design an adaption module before minimizing the spatial distance between source and auxiliary modal distributions (Song et al., 2020). On the other hand, the contrastive-based methods efficiently align different modality representations through the contrastive learning on paired modalities (Liu et al., 2021b; Zolfaghari et al., 2021; Mai et al., 2022).

The unsupervised multimodal contrastive methods directly regard the modality pairs from the same samples as positive pairs and those modality pairs from different samples as negative pairs to pull together the unimodal representations of paired modalities and pull apart the unimodal representations of unpaired modalities in the embedding space. (Tian et al., 2020b; Akbari et al., 2021; Zolfaghari et al., 2021; Liu et al., 2021b; Zhang et al., 2021a; Taleb et al., 2022). Supervised multimodal contrastive methods are proposed to treat sample pairs with the same label as positive pairs and sample pairs with a different label as negative pairs in the mini-batch (Zhang et al., 2021b; Pinitas et al., 2022). In this way, the unimodal representations with the same semantics will be clustered.

Despite their effectiveness in learning the correspondence among modalities, these contrastive-based multimodal learning methods still meet with problems with the sensor noise in the in-the-wild datasets (Mittal et al., 2020). The current methods always treat each modality equally and ignore the difference of the role for different modalities, The final decisions will be negatively affected by those samples with inefficient unimodal representations and thus can not provide trustworthy multimodal representations. In this work, we aim to learn trustworthy multimodal representations by aligning unimodal representations towards the effective modality, considering modality effectiveness in addition to strengthening relationships between modalities. The modality effectiveness is decided by the unimodal prediction and the contrastive learning is under the weak supervision information from the unimodal prediction. As shown in Figure 1, the unimodal representations will be aligned towards those with correct unimodal predictions. In summary, our contributions are:

- To facilitate the inter-modality relationship for multimodal classification, we combine the

aggregation-based and alignment-based fusion methods to create a joint representation.

- We propose UniS-MMC to efficiently align the representation to the effective modality under weak supervision of unimodal prediction to address the issue of different contributions from the modalities.
- Extensive experiments on two image-text classification benchmarks, UPMC-Food-101 (Wang et al., 2015) and N24News (Wang et al., 2022) demonstrate the effectiveness of our proposed method.

## 2 Related Work

In this section, we will introduce the related work on contrastive learning and multimodal learning.

### 2.1 Contrastive Learning

Contrastive learning (Hadsell et al., 2006; Oord et al., 2018; Qin and Joty, 2022) captures distinguishable representations by drawing positive pairs closer and pushing negative pairs farther contrastively. In addition to the above single-modality representation learning, contrastive methods for multiple modalities are also widely explored. The common methods (Radford et al., 2021; Jia et al., 2021; Kamath et al., 2021; Li et al., 2021; Zhang et al., 2022; Taleb et al., 2022; Chen et al., 2022) leverage the cross-modal contrastive matching to align two different modalities and learn the inter-modality correspondence. Except the inter-modality contrastive, Visual-Semantic Contrastive (Yuan et al., 2021), XMC-GAN (Zhang et al., 2021a) and CrossPoint (Afham et al., 2022) also introduce the intra-modality contrastive for representation learning. Besides, CrossCLR (Zolfaghari et al., 2021) removes the highly related samples from negative samples to avoid the bias of false negatives. GMC (Poklukar et al., 2022) builds the contrastive learning process between the modality-specific representations and the global representations of all modalities instead of the cross-modal representations.

### 2.2 Multimodal Learning

Multimodal learning is expected to build models based on multiple modalities and to improve the general performance from the joint representation (Ngiam et al., 2011; Baltrušaitis et al., 2018; Gao et al., 2020). The fusion operation among multiple modalities is one of the key topics in multimodal

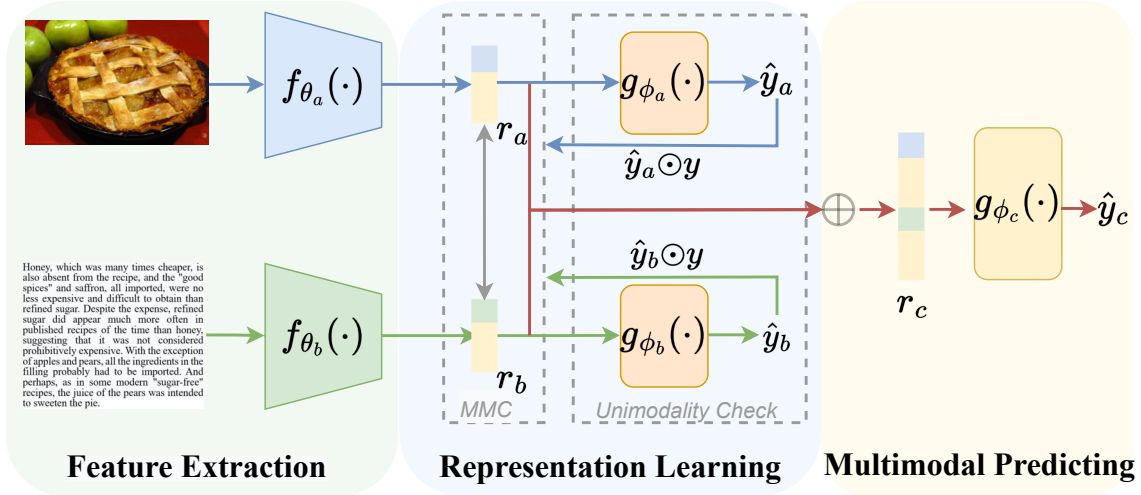


Figure 2: The framework for our proposed UniS-MMC.

learning to help the modalities complement each other (Wang, 2021). Multimodal fusion methods are generally categorized into two types: alignment-based and aggregation-based fusion (Baltrušaitis et al., 2018). Alignment-based fusion (Gretton et al., 2012; Song et al., 2020) aligns multimodal features by increasing the modal similarity to capture the modality-invariant features. Aggregation-based methods choose to create the joint multimodal representations by combining the participating unimodal features (early-fusion, Kalfaoglu et al. (2020); Nagrani et al. (2021); Zou et al. (2022)), unimodal decisions (late-fusion, Tian et al. (2020a); Huang et al. (2022)) and both (hybrid-fusion, Wu et al. (2022)). In addition to these joint-representation generating methods, some works further propose to evaluate the attended modalities and features before fusing. M3ER (Mittal et al., 2020) conducts a modality check step to finding those modalities with small correlation and Multimodal Dynamics (Han et al., 2022) evaluates both the feature- and modality-level informativeness during extracting unimodal representations.

### 3 Methodology

In this section, we present our method called UniS-MMC for multimodal fusion.

#### 3.1 Notation

Suppose we have the training data set  $\mathcal{D} = \{\{x_m^n\}_{m=1}^M, y^n\}_{n=1}^N$  that contains  $N$  samples  $\mathcal{X} = \{x_m^n \in \mathbb{R}^{d_m}\}_{m=1}^M$  of  $M$  modalities and  $N$  corresponding labels  $\mathcal{Y} = \{y^n\}_{n=1}^N$  from  $K$  categories. As shown in Figure 2, the unimodal representations

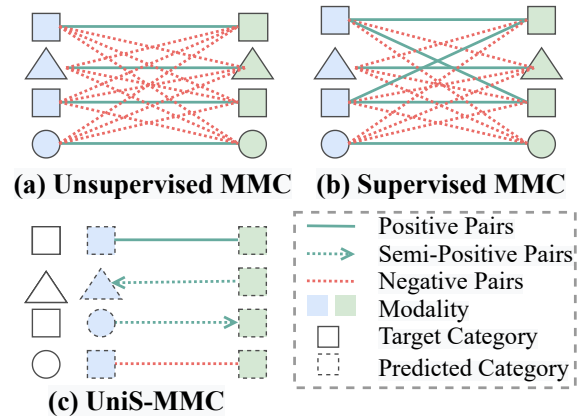


Figure 3: The relationship comparison between two modalities in training mini-batch of (a) unsupervised MMC, (b) supervised MMC and (c) UniS-MMC.

of modality  $a$  and  $b$  are extracted from the respective encoders  $f_{\theta_a}$  and  $f_{\theta_b}$ . Following the parameter sharing method in the multi-task learning (Pilault et al., 2021; Bhattacharjee et al., 2022), the representations are shared directly between unimodal prediction tasks and the following multimodal prediction task. With weak supervision information produced from the respective unimodal classifier  $g_{\phi_a}$  and  $g_{\phi_b}$ , the final prediction is finally learned based on the updated multimodal representations  $r_c$  and the multimodal classifier  $g_{\phi_c}$ .

#### 3.2 Unimodality-supervised Multimodal Contrastive Learning

First, the unimodal representations are extracted from the raw data of each modality by the pre-trained encoders. We introduce the uni-modality check step to generate the weak supervision for

checking the effectiveness of each unimodal representation. Then we illustrate how we design the unimodality-supervised multimodal contrastive learning method among multiple modalities to learn the multimodal representations.

### 3.2.1 Modality Encoder

Given multimodal training data  $\{\mathbf{x}_m\}_{m=1}^M$ , the raw unimodal data of modality  $m$  are firstly processed with respective encoders to obtain the hidden representations. We denote the learned hidden representation  $f_{\theta_m}(\mathbf{x}_m)$  of modality  $m$  as  $\mathbf{r}_m$ . We use the pretrained ViT (Dosovitskiy et al., 2021) as the feature encoder for images in both UPMC Food-101 and N24News datasets. We use only the pretrained BERT (Devlin et al., 2019) as the feature encoder for the textual description in these datasets. Besides, we also try the pretrained RoBERTa (Liu et al., 2019) for text sources in N24News.

### 3.2.2 Unimodality Check

**Unimodal prediction.** Different from the common aggregation-based multimodal learning methods which only use the unimodal learned representations for fusion, our method also use the unimodal representations as inputs to the unimodal predicting tasks. The classification module can be regarded as a probabilistic model:  $g_\phi : \mathcal{R} \rightarrow \mathcal{P}$ , which maps the hidden representation to a predictive distribution  $\mathbf{p}(\mathbf{y} | \mathbf{r})$ . For a unimodal predicting task, the predictive distribution is only based on the output of the unimodal classifier. The learning objective of the unimodal predicting task is to minimize each unimodal prediction loss:

$$\mathcal{L}_{uni} = - \sum_{m=1}^M \sum_{k=1}^K y^k \log p_m^k, \quad (1)$$

where  $y^k$  is the  $k$ -th element category label and  $[p_m^1; p_m^2; \dots; p_m^K] = \mathbf{p}_m(\mathbf{y} | \mathbf{r}_m)$  is the softmax output of unimodal classifiers on modality  $m$ .

**Unimodality effectiveness.** The above unimodal prediction results are used to check the supervised information for deciding the effectiveness of each modality. The unimodal representation with correct prediction is regarded as the effective representation for providing the information to the target label. Alternately, the unimodal representation with the wrong prediction is regarded as an ineffective representation.

### 3.2.3 Multimodal Contrastive Learning

We aim to reduce the multimodal prediction bias caused by treating modalities equally for each sam-

Table 1: Contrastive settings.

Uni-Prediction	Modality $a$	Modality $b$	Category
0	True	True	Positive
1	True	False	Semi-positive
2	False	True	
3	False	False	Negative

ple. This is done by learning to align unimodal representations towards the effective modalities sample by sample. We regulate each unimodal representation with the targets based on the multi-task-based multimodal learning framework. As shown in Figure 3 c), we propose a new multimodal contrastive method to encourage modalities with both correct unimodal predictions to share a stronger correspondence. For those samples with both wrong predictions, we encourage their unimodal representations to be more different from each other to get more complementary multimodal representations. It helps to a higher possibility of correct multimodal prediction. For those samples with mutually exclusive predictions, we encourage these unimodal representations to learn from each other under the supervision of unimodal predictions by aligning the ineffective modality with the effective one.

When considering two specific modalities  $m_a$  and  $m_b$  of  $n$ -th sample, we generate two unimodal hidden representations  $r_a^n$  and  $r_b^n$  from respective unimodal encoders. From the above unimodal predicting step, we also obtain the unimodal prediction results,  $\hat{y}_a^n$  and  $\hat{y}_b^n$ . As the summarization in Table 1, we define the following positive pair, negative pair and semi-positive pair:

**Positive pair.** If both the paired unimodal predictions are correct, we define these unimodal representation pairs are positive pairs, namely  $\mathbb{P}$ , where  $\mathbb{P} = \{n | \{\hat{y}_a^n \equiv y^n \text{ and } \hat{y}_b^n \equiv y^n\}_{n=1}^N\}$  in the mini-batch  $\mathbb{B}$ .

**Negative pair.** If both the paired unimodal predictions are wrong, we define these unimodal representation pairs are negative pairs, namely  $\mathbb{N}$ , where  $\mathbb{N} = \{n | \{\hat{y}_a^n \neq y^n \text{ and } \hat{y}_b^n \neq y^n\}_{n=1}^N\}$  in the mini-batch  $\mathbb{B}$ .

**Semi-Positive pair.** If the predictions of the paired unimodal representations are mutually exclusive, one correct and another wrong, we define these unimodal representation pairs are semi-positive pairs, namely  $\mathbb{S}$ , where  $\mathbb{S} = \{\{n | \{\hat{y}_a^n \equiv y^n \text{ and } \hat{y}_b^n \neq y^n\}_{n=1}^N\} \cup \{n | \hat{y}_a^n \neq y^n$

and  $\hat{y}_b^n \equiv y^n \setminus \{y^n\}_{n=1}^N$  in the mini-batch.

We further propose the multimodal contrastive loss for two modalities as follows:

$$\mathcal{L}_{b-mmcc} = -\log \frac{\sum_{n \in \mathbb{P}, \mathbb{S}} (\exp(\cos(r_a^n, r_b^n)/\tau))}{\sum_{n \in \mathbb{B}} (\exp(\cos(r_a^n, r_b^n)/\tau))}, \quad (2)$$

where  $\cos(r_a^n, r_b^n) = \frac{r_a^n \cdot r_b^n}{\|r_a^n\| \|r_b^n\|}$  is the cosine similarity between paired unimodal representations  $r_a^n$  and  $r_b^n$  for sample  $n$ ,  $\tau$  is the temperature coefficient. The similarity of positive pairs and semi-positive pairs is optimized towards a higher value while the similarity of negative pairs is optimized towards a smaller value. The difference between positive and semi-positive pairs is that the unimodal representations updated towards each other in positive pairs while only the unimodal representations of the wrong unimodal prediction updated towards the correct one in semi-positive pairs. We detach the modality feature with correct predictions from the computation graph when aligning with low-quality modality features for semi-positive pairs, which is inspired by GAN models (Arjovsky et al., 2017; Zhu et al., 2017) where the generator output is detached when updating the discriminator only.

Multimodal problems often encounter situations with more than two modalities. For more than two modalities, the multimodal contrastive loss for  $M$  modalities ( $M > 2$ ) can be computed by:

$$\mathcal{L}_{mmc} = \sum_{i=1}^M \sum_{j>i}^M \mathcal{L}_{b-mmcc}(m_i, m_j), \quad (3)$$

### 3.3 Fusion and Total Learning Objective

**Multimodal prediction.** When fusing all unimodal representations with concatenation, we get the fused multimodal representations  $r_c = r_1 \oplus r_2 \oplus \dots \oplus r_m$ . Similarly, the multimodal predictive distribution is the output of the multimodal classifier with inputs of the fused representations. For the multimodal prediction task, the target is to minimize the multimodal prediction loss:

$$\mathcal{L}_{multi} = -\sum_{k=1}^K y^k \log p_k^k, \quad (4)$$

where  $y^k$  is the  $k$ -th element category label and  $[p_k^1; p_k^2; \dots; p_k^K] = \mathbf{p}_c(\mathbf{y} | \mathbf{r}_c)$  is the softmax output of multimodal classifier.

**Total learning objective.** The overall optimization objective for our proposed UniS-MMC is:

$$\mathcal{L}_{UniS-MMC} = \mathcal{L}_{uni} + \mathcal{L}_{multi} + \lambda \mathcal{L}_{mmc}, \quad (5)$$

where  $\lambda$  is a loss coefficient for balancing the predicting loss and the multimodal contrastive loss.

## 4 Experiments

### 4.1 Experimental Setup

**Dataset and metric.** We evaluate our method on two publicly available image-text classification datasets UPMC-Food-101 and N24News. **UPMC-Food-101**<sup>1</sup> is a multimodal classification dataset that contains textual recipe descriptions and the corresponding images for 101 kinds of food. We get this dataset from their project website and split 5000 samples from the default training set as the validation set. **N24News**<sup>2</sup> is a news classification dataset with four text types (*Heading*, *Caption*, *Abstract* and *Body*) and images. In order to supplement the long text data of the FOOD101 dataset, we choose the first three text sources from N24News in our work. We use classification accuracy (Acc) as evaluation metrics for UPMC-Food-101 and N24News. The detailed dataset information can be seen in Appendix A.1.

**Implementation.** For the image-text dataset UPMC Food-101, we use pretrained BERT (Devlin et al., 2019) as a text encoder and pretrained vision transformer (ViT) (Dosovitskiy et al., 2021) as an image encoder. For N24News, we utilize two different pretrained language models, BERT and RoBERTa (Liu et al., 2019) as text encoders and also the same vision transformer as an image encoder. All classifiers of these two image-text classification datasets are three fully-connected layers with a ReLU activation function.

The default reported results on image-text datasets are obtained with BERT-base (or RoBERTa-base) and ViT-base in this paper. The performance is presented with the average and standard deviation of three runs on Food101 and N24News. The codes is available on GitHub<sup>3</sup>. The detailed settings of the hyper-parameter are summarized in Appendix A.2.

### 4.2 Baseline Models

The used baseline models are as follows:

- **MMBT** (Kiela et al., 2019) jointly finetunes pretrained text and image encoders by projecting image embeddings to text token space on BERT-like architecture.

<sup>1</sup>UPMC-Food-101: <https://visiir.isir.upmc.fr/>

<sup>2</sup>N24News: <https://github.com/billywzh717/N24News>

<sup>3</sup><https://github.com/Vincent-ZHQ/UniS-MMC>

Table 2: Comparison of multimodal classification performance on **a)** Food101 and **b)** N24News.

a) Model	Fusion		Backbone		Acc	b) Model	Fusion		Backbone		Multimodal		
	AGG	ALI	Image	Text			AGG	ALI	Image	Text	Headline	Caption	Abstract
MMBT	Early	✗	ResNet-152	BERT	92.1 $\pm$ 0.1	N24News	Early	✗	ViT	RoBERTa	79.41	77.45	83.33
HUSE	Early	✓	Graph-RISE	BERT	92.3	AggMM	Early	✗	ViT	BERT	78.6 $\pm$ 1.1	76.8 $\pm$ 0.2	80.8 $\pm$ 0.2
ViLT	Early	✓	ViT	BERT	92.0	UnSupMMC	Early	✓	ViT	BERT	79.3 $\pm$ 0.8	76.9 $\pm$ 0.3	81.9 $\pm$ 0.3
CMA-CLIP	Early	✓	ViT	Transformer	93.1	SupMMC	Early	✓	ViT	BERT	79.6 $\pm$ 0.5	77.3 $\pm$ 0.2	81.7 $\pm$ 0.8
ME	Early	✗	DenseNet	BERT	94.6	UniS-MMC	Early	✓	ViT	BERT	<b>80.2<math>\pm</math>0.1</b>	<b>77.5<math>\pm</math>0.3</b>	<b>83.2<math>\pm</math>0.4</b>
AggMM	Early	✗	ViT	BERT	93.7 $\pm$ 0.2	AggMM	Early	✗	ViT	RoBERTa	78.9 $\pm$ 0.3	77.9 $\pm$ 0.3	83.5 $\pm$ 0.2
UnSupMMC	Early	✓	ViT	BERT	94.1 $\pm$ 0.7	UnSupMMC	Early	✓	ViT	RoBERTa	79.9 $\pm$ 0.2	78.0 $\pm$ 0.1	83.7 $\pm$ 0.3
SupMMC	Early	✓	ViT	BERT	94.2 $\pm$ 0.2	SupMMC	Early	✓	ViT	RoBERTa	79.9 $\pm$ 0.4	77.9 $\pm$ 0.2	84.0 $\pm$ 0.2
UniS-MMC	Early	✓	ViT	BERT	<b>94.7<math>\pm</math>0.1</b>	UniS-MMC	Early	✓	ViT	RoBERTa	<b>80.3<math>\pm</math>0.1</b>	<b>78.1<math>\pm</math>0.2</b>	<b>84.2<math>\pm</math>0.1</b>

- **HUSE** (Narayana et al., 2019) creates a joint representation space by learning the cross-modal representation with semantic information.
- **ViLT** (Kim et al., 2021; Liang et al., 2022) introduces a BERT-like multimodal transformer architecture on vision-and-language data.
- **CMA-CLIP** (Liu et al., 2021a) finetunes the CLIP (Radford et al., 2021) with newly designed two types of cross-modality attention module.
- **ME** (Liang et al., 2022) is the state-of-the-art method on Food101, which performs cross-modal feature transformation to leverage cross-modal information.
- **N24News** (Wang et al., 2022) train both the unimodal and multimodal predicting task to capture the modality-invariant representations.
- **AggMM** finetunes the pretrained text and image encoders and concatenates the unimodal representations for the multimodal recognition task.
- **SupMMC** and **UnSupMMC** finetune the pretrained text and image encoders and then utilize the supervised and unsupervised multimodal contrastive method to align unimodal representations before creating joint embeddings, respectively.

### 4.3 Performance Comparison

#### Final classification performance comparison.

The final image-text classification performance on Food101 and N24News is presented in Table 2. We have the following findings from the experimental results: (i) focusing on the implemented methods, contrastive-based methods with naive alignment could get an improvement over the implemented aggregation-based methods; (ii) the implemented contrastive-based methods outperform many of the recent novel multimodal methods; (iii) the proposed UniS-MMC has a large improvement com-

pared with both the implemented contrastive-based baseline models and the recent start-of-art multimodal methods on Food101 and produces the best results on every kind of text source on N24News with the same encoders.

**T-sne visualization comparison with baseline models.** We visualize the representation distribution of the proposed uni-modality supervised multimodal contrastive method and compare it with the naive aggregation-based method and the typical unsupervised and supervised contrastive method.

As shown in Figure 4, unimodal representations are summarized and mapped into the same feature space. The previous typical contrastive methods, such as unsupervised and supervised contrastive methods will mix up different unimodal representations from different categories when bringing the representation of different modalities that share the same semantics closer. For example, the representations of two modalities from the same category are clustered well in Figure 4 (b) and (c) (green circle and orange circle). However, these contrastive-based methods can also bring two problems. One is that they map the unimodal embeddings into the same embedding space will lose the complementary information from different modalities. Another is that they heavily mix the representations from the specific class with other categories, such as the clusters (orange circle). As a comparison, our proposed method preserves the complementary multimodal information by maintaining the two parts of the distribution from two modalities (red line) well (Figure 4 (d) from the aggregation-based methods (Figure 4 (a)) in addition to a better cluster of unimodal representations.

We further summarized the visualization of the final multimodal representation in Figure 5. Comparing Figure 5 (a) and Figure 5 (d), the proposed

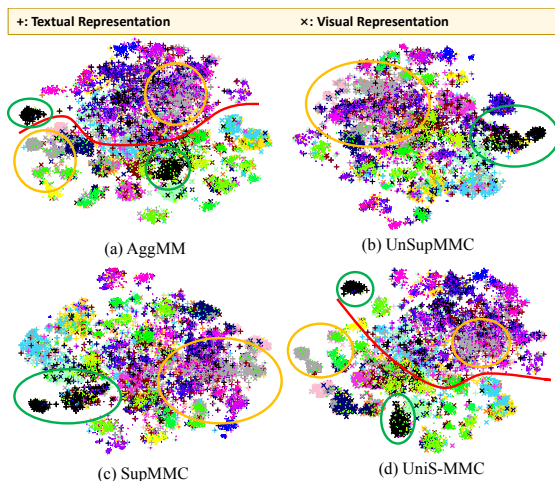


Figure 4: Unimodal representation distribution of the first 10 categories of the N24News test set across different methods: (a) aggregation-based method, (b) unsupervised multimodal method, (c) supervised contrastive method and (d) unimodality-supervised method.

UniS-MMC can create better class clusters, such as the **green circle**. Comparing Figure 5 (b), (c) and (d) (**orange circle**), the classification clusters are not separated by other classes in the proposed methods. It is different from the other two typical contrastive-based methods. Generally, our proposed method not only helps the unimodal representation learning process and gets better sub-clusters for each modality but also improves the classification boundary of the final multimodal representation.

#### 4.4 Analysis

**Classification with Different Combinations of Input Modalities.** We first perform an ablation study of classification on N24News with different input modalities. Table 3 provides the classification performance of unimodal learning with image-only, text-only, and traditional multimodal learning with the concatenation of visual and textual features and our proposed UniS-MMC. The text modality is encoded with two different encoders, RoBERTa or BERT. By comparing the models with different language encoders, we find that the feature encoder can significantly affect the multimodal performance, and the RoBERTa-based model usually performs better than the BERT-based model. This is because the multimodal classification task is influenced by each learned unimodal representation. Besides, all the multimodal networks perform better than unimodal networks. It reflects that multiple modalities will help make accurate decisions. Moreover, our proposed UniS-

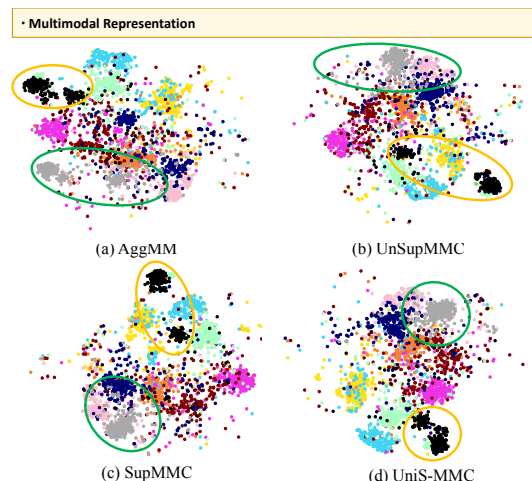


Figure 5: Multimodal representation distribution of the first 10 categories of the N24News test set across different methods: (a) aggregation-based method, (b) unsupervised multimodal method, (c) supervised contrastive method and (d) unimodality-supervised method.

MMC achieves 0.6% to 2.4% improvement over the aggregation-based baseline model with BERT and 0.3% to 1.4% improvement with RoBERTa.

**Ablation study on N24News.** We conduct the ablation study to analyze the contribution of the different components of the proposed UniS-MMC on N24News. AggMM is the baseline model of the aggregation-based method that combines the unimodal representation directly. The ablation works on three text source headline, caption and abstract with both BERT-based and RoBERTa-based models. Specifically,  $L_{uni}$  is the introduced unimodal prediction task,  $C_{Semi}$  and  $C_{Neg}$  are semi-positive pair and negative pair setting, respectively.

Table 4 presents the multimodal classification results of the above ablation stud with different participating components.  $L_{uni}$  and the setting of  $C_{Semi}$  align the unimodal representation towards the targets, with the former achieved by mapping different unimodal representations to the same target space and the latter achieved by feature distribution aligning. They can both provide a significant improvement over the baseline model.  $C_{Neg}$  further improve the performance by getting a larger combination of multimodal representation with more complementary information for those samples that are difficult to classify.

**Analysis on the learning process.** To further explore the role of our proposed UniS-MMC in aligning the unimodal representation towards the targets, we summarise the unimodal predicting re-

Table 3: Comparison to unimodal learning and the baseline model on N24News.

Dataset	Text	Image-only	BERT-based			RoBERTa-based		
			Text-only	AggMM	UniS-MMC	Text-only	AggMM	UniS-MMC
N24News	Headline		72.1 $\pm$ 0.2	78.6 $\pm$ 1.1	<b>80.2</b> $\pm$ 0.1 $\uparrow$ 1.6	71.8 $\pm$ 0.2	78.9 $\pm$ 0.3	<b>80.3</b> $\pm$ 0.1 $\uparrow$ 1.4
	Caption	54.1 $\pm$ 0.2	72.7 $\pm$ 0.3	76.8 $\pm$ 0.2	<b>77.5</b> $\pm$ 0.3 $\uparrow$ 0.7	72.9 $\pm$ 0.4	77.9 $\pm$ 0.3	<b>78.1</b> $\pm$ 0.2 $\uparrow$ 0.3
	Abstract		78.3 $\pm$ 0.3	80.8 $\pm$ 0.2	<b>83.2</b> $\pm$ 0.4 $\uparrow$ 2.4	79.7 $\pm$ 0.2	83.5 $\pm$ 0.2	<b>84.2</b> $\pm$ 0.1 $\uparrow$ 0.7

Table 4: Ablation study on N24News.

Method	Headline		Caption		Abstract	
	BERT	RoBERTa	BERT	RoBERTa	BERT	RoBERTa
AggMM	78.6 $\pm$ 1.1	78.9 $\pm$ 0.3	76.8 $\pm$ 0.2	77.9 $\pm$ 0.3	80.8 $\pm$ 0.2	83.5 $\pm$ 0.2
+ <i>L<sub>uni</sub></i>	79.4 $\pm$ 0.4	79.4 $\pm$ 0.3	77.3 $\pm$ 0.2	77.9 $\pm$ 0.1	82.5 $\pm$ 0.3	84.1 $\pm$ 0.2
+ <i>C<sub>Semi</sub></i>	80.1 $\pm$ 0.1	80.0 $\pm$ 0.3	77.3 $\pm$ 0.2	78.0 $\pm$ 0.3	82.7 $\pm$ 0.4	84.2 $\pm$ 0.3
+ <i>C<sub>Neg</sub></i>	<b>80.2</b> $\pm$ 0.1	<b>80.3</b> $\pm$ 0.1	<b>77.5</b> $\pm$ 0.3	<b>78.1</b> $\pm$ 0.2	<b>83.2</b> $\pm$ 0.4	<b>84.2</b> $\pm$ 0.1

sults of the validation set during the training process in Figure 6. Ideally, different participating modalities for the same semantic should be very similar and give the same answer for the same sample. However, in practical problems, the unimodal predictions are not usually the same as the actual noise. In our proposed method, the proportion of both wrong unimodal predicting is higher and the proportion of both correct unimodal predicting is lower when removing our setting of semi-positive pair and negative pair. It means that UniS-MMC could align the unimodal representations for the targets better and get more trustworthy unimodal representations.

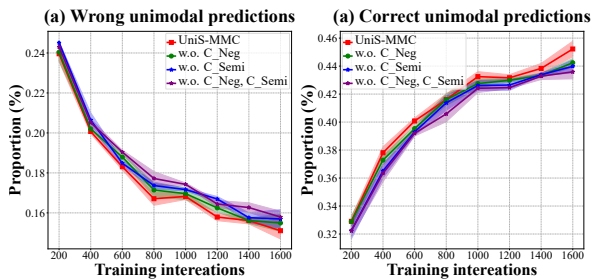


Figure 6: As the training progresses, the change of the proportion of both wrong (left), both correct (right) unimodal predictions of the validation set (N24News): the complete method (UniS-MMC), remove negative pair (w.o. C\_Neg), remove semi-positive pair (w.o. C\_Semi) and remove both (w.o. C\_Neg, C\_Semi).

**Analysis on the Final Multimodal Decision.** Compared with the proposed UniS-MMC, MT-MML is the method that jointly trains the unimodal and multimodal predicting task, without applying the proposed multimodal contrastive loss. We summarize unimodal performance on MT-MML and

UniS-MMC and present unimodal predictions in Figure 6. The unimodal prediction consistency here is represented by the consistency of the unimodal prediction for each sample. When focusing on the classification details of each modality pair, we find that the proposed UniS-MMC gives a larger proportion of samples with both correct predictions and a smaller proportion of samples with both wrong decisions and opposite unimodal decisions compared with MT-MML.




Figure 7: Consistency comparison of unimodal prediction between MT-MML and the UniS-MMC.

## 5 Conclusion

In this work, we propose the Unimodality-Supervised Multimodal Contrastive (UNniS-MMC), a novel method for multimodal fusion to reduce the multimodal decision bias caused by inconsistent unimodal information. Based on the introduced multi-task-based multimodal learning framework, we capture the task-related unimodal representations and evaluate their potential influence on the final decision with the unimodal predictions. Then we contrastively align the unimodal representation towards the relatively reliable modality under the weak supervision of unimodal predictions. This novel contrastive-based alignment method helps to capture more trustworthy multimodal representations. The experiments on four public multimodal classification datasets demonstrate the effectiveness of our proposed method.



## Limitations

Unlike the traditional multimodal contrastive loss focusing more on building the direct link between paired modalities, our proposed UniS-MMC aims to leverage inter-modality relationships and potential effectiveness among modalities to create more trustworthy and complementary multimodal representations. It means that UniS-MMC is not applied to all multimodal problems. It can achieve competitive performance in tasks that rely on the quantity of the joint representation, such as the multimodal classification task. It is not suitable for tasks that rely purely on correspondence between modalities, such as the cross-modal retrieval task.

## Acknowledgements

The computational work for this article was partially performed on resources of the National Supercomputing Centre, Singapore (<https://www.nsc.sg>)

## References

- Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. 2022. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9902–9912.
- Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. 2021. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34:24206–24221.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR.
- 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.
- Deblina Bhattacharjee, Tong Zhang, Sabine Süssstrunk, and Mathieu Salzmann. 2022. Mult: An end-to-end multitask learning transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12031–12041.
- Ginevra Castellano, Loic Kessous, and George Caridakis. 2008. Emotion recognition through multiple modalities: face, body gesture, speech. In *Affect and emotion in human-computer interaction*, pages 92–103. Springer.
- Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. 2022. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16495–16504.
- Chen Chen, Nana Hou, Yuchen Hu, Heqing Zou, Xiaofeng Qi, and Eng Siong Chng. 2022. Interactive audio-text representation for automated audio captioning with contrastive learning. *arXiv preprint arXiv:2203.15526*.
- 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.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- Jing Gao, Peng Li, Zhikui Chen, and Jianing Zhang. 2020. A survey on deep learning for multimodal data fusion. *Neural Computation*, 32(5):829–864.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773.
- Raia Hadsell, Sumit Chopra, and Yann 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. IEEE.
- Zongbo Han, Fan Yang, Junzhou Huang, Changqing Zhang, and Jianhua Yao. 2022. Multimodal dynamics: Dynamical fusion for trustworthy multimodal classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20707–20717.
- Yu Huang, Junyang Lin, Chang Zhou, Hongxia Yang, and Longbo Huang. 2022. Modality competition: What makes joint training of multi-modal network fail in deep learning?(provably). *arXiv preprint arXiv:2203.12221*.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR.

- M Kalfaoglu, Sinan Kalkan, and A Aydin Alatan. 2020. Late temporal modeling in 3d cnn architectures with bert for action recognition. In *European Conference on Computer Vision*, pages 731–747. Springer.
- Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. 2021. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790.
- Douwe Kiela, Suvrat Bhooshan, Hamed Firooz, Ethan Perez, and Davide Testuggine. 2019. Supervised multimodal bitransformers for classifying images and text. *arXiv preprint arXiv:1909.02950*.
- Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Tao Liang, Guosheng Lin, Mingyang Wan, Tianrui Li, Guojun Ma, and Fengmao Lv. 2022. Expanding large pre-trained unimodal models with multimodal information injection for image-text multimodal classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15492–15501.
- Huidong Liu, Shaoyuan Xu, Jinmiao Fu, Yang Liu, Ning Xie, Chien-Chih Wang, Bryan Wang, and Yi Sun. 2021a. Cma-clip: Cross-modality attention clip for image-text classification. *arXiv preprint arXiv:2112.03562*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yunze Liu, Qingnan Fan, Shanghang Zhang, Hao Dong, Thomas Funkhouser, and Li Yi. 2021b. Contrastive multimodal fusion with tupleinfonce. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 754–763.
- Sijie Mai, Ying Zeng, Shuangjia Zheng, and Haifeng Hu. 2022. Hybrid contrastive learning of tri-modal representation for multimodal sentiment analysis. *IEEE Transactions on Affective Computing*.
- Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. 2020. M3er: Multiplicative multimodal emotion recognition using facial, textual, and speech cues. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 1359–1367.
- Hee-Seung Moon and Jiwon Seo. 2019. Observation of human response to a robotic guide using a variational autoencoder. In *2019 Third IEEE International Conference on Robotic Computing (IRC)*, pages 258–261. IEEE.
- Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. 2021. Attention bottlenecks for multimodal fusion. *Advances in Neural Information Processing Systems*, 34:14200–14213.
- Pradyumna Narayana, Aniket Pednekar, Abishek Krishnamoorthy, Kazoo Sone, and Sugato Basu. 2019. Huse: Hierarchical universal semantic embeddings. *arXiv preprint arXiv:1911.05978*.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *ICML*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Jonathan Pilault, Amine Elhattami, and Christopher J. Pal. 2021. **Conditionally adaptive multi-task learning: Improving transfer learning in NLP using fewer parameters & less data**. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Kosmas Pinitas, Konstantinos Makantasis, Antonios Liapis, and Georgios N Yannakakis. 2022. Supervised contrastive learning for affect modelling. In *Proceedings of the 2022 International Conference on Multimodal Interaction*, pages 531–539.
- Petra Poklukar, Miguel Vasco, Hang Yin, Francisco S. Melo, Ana Paiva, and Danica Kragic. 2022. Geometric multimodal contrastive representation learning. In *International Conference on Machine Learning*.
- Chengwei Qin and Shafiq Joty. 2022. **Continual few-shot relation learning via embedding space regularization and data augmentation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2776–2789, Dublin, Ireland. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.

- Geovany A Ramirez, Tadas Baltrušaitis, and Louis-Philippe Morency. 2011. Modeling latent discriminative dynamic of multi-dimensional affective signals. In *International Conference on Affective Computing and Intelligent Interaction*, pages 396–406. Springer.
- Sijie Song, Jiaying Liu, Yanghao Li, and Zongming Guo. 2020. Modality compensation network: Cross-modal adaptation for action recognition. *IEEE Transactions on Image Processing*, 29:3957–3969.
- Aiham Taleb, Matthias Kirchler, Remo Monti, and Christoph Lippert. 2022. Contig: Self-supervised multimodal contrastive learning for medical imaging with genetics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20908–20921.
- Junjiao Tian, Wesley Cheung, Nathaniel Glaser, Yen-Cheng Liu, and Zsolt Kira. 2020a. Uno: Uncertainty-aware noisy-or multimodal fusion for unanticipated input degradation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5716–5723. IEEE.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020b. Contrastive multiview coding. In *European conference on computer vision*, pages 776–794. Springer.
- Vishaal Udandarao, Abhishek Maiti, Deepak Srivatsav, Suryatej Reddy Vyalla, Yifang Yin, and Rajiv Ratn Shah. 2020. Cobra: Contrastive bi-modal representation algorithm. *arXiv preprint arXiv:2005.03687*.
- Jinghua Wang, Zhenhua Wang, Dacheng Tao, Simon See, and Gang Wang. 2016. Learning common and specific features for rgb-d semantic segmentation with deconvolutional networks. In *European Conference on Computer Vision*, pages 664–679. Springer.
- Xin Wang, Devinder Kumar, Nicolas Thome, Matthieu Cord, and Frederic Precioso. 2015. Recipe recognition with large multimodal food dataset. In *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–6. IEEE.
- Yang Wang. 2021. Survey on deep multi-modal data analytics: Collaboration, rivalry, and fusion. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 17(1s):1–25.
- Zhen Wang, Xu Shan, Xiangxie Zhang, and Jie Yang. 2022. [N24News: A new dataset for multimodal news classification](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6768–6775, Marseille, France. European Language Resources Association.
- Nan Wu, Stanislaw Jastrzebski, Kyunghyun Cho, and Krzysztof J Geras. 2022. Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks. In *International Conference on Machine Learning*, pages 24043–24055. PMLR.
- Xin Yuan, Zhe Lin, Jason Kuen, Jianming Zhang, Yilin Wang, Michael Maire, Ajinkya Kale, and Baldo Faieta. 2021. Multimodal contrastive training for visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6995–7004.
- Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. 2021a. Cross-modal contrastive learning for text-to-image generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 833–842.
- Miaoran Zhang, Marius Mosbach, David Ifeoluwa Adelani, Michael A Hedderich, and Dietrich Klakow. 2022. Mcse: Multimodal contrastive learning of sentence embeddings. *arXiv preprint arXiv:2204.10931*.
- Wenjia Zhang, Lin Gui, and Yulan He. 2021b. Supervised contrastive learning for multimodal unreliable news detection in covid-19 pandemic. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 3637–3641.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232.
- Mohammadreza Zolfaghari, Yi Zhu, Peter Gehler, and Thomas Brox. 2021. Crossclr: Cross-modal contrastive learning for multi-modal video representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1450–1459.
- Heqing Zou, Yuke Si, Chen Chen, Deepu Rajan, and Eng Siong Chng. 2022. Speech emotion recognition with co-attention based multi-level acoustic information. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7367–7371. IEEE.

## A Appendix

### A.1 Datasets Usage Instructions

To make a fair comparison with the previous works, we adopt the following default setting of the split method, as shown in Table 5. Since the UPMC-Food101 dataset does not provide the validation set, we split 5000 samples out of the training set and use them as the validation set.

Table 5: Datasets information and the split results

Dataset	Modalities	#Category	#Train	#Valid	#Test
UPMC-Food-101	image, text	101	60085	5000	21683
N24News	image, text	24	48988	6123	6124

## A.2 Experimental Settings

The model is trained on NVIDIA V100-SXM2-16GB and NVIDIA A100-PCIE-40GB. The corresponding Pytorch version, CUDA version and CUDNN version are 1.8.0, 11.1 and 8005 respectively. We utilize Adam as the optimizer and use ReduceLROnPlateau to update the learning rate. We use Adam (Kingma and Ba, 2015) as the model optimizer. The temperature coefficient for contrastive learning is set as 0.07 and the loss coefficient in this paper is set as 0.1 to keep loss values in the same order of magnitude. The code is attached and will be available on GitHub. Some key settings of the model implementation are listed as followings:

Table 6: Detailed setting of the hyper-parameter for UPMC-Food-101, BRCA and ROSMAP

Item	UPMC-Food-101	N24News
Batch gradient	128	128
Batch size	32	32
Learning rate (m)	2e-5	1e-4
Dropout (m)	0	0
Weight decay	1e-4	1e-4

## A.3 Learning with a Single Modality

We show the unimodal classification results from different unimodal backbones on text-image datasets in the following Table 7.

Table 7: Unimodal classification performance with different backbones on Food101 and N24News.

Source	Backbone	Food101	N24News
<b>Image</b>	ViT	73.1 $\pm$ 0.2	54.1 $\pm$ 0.2
<b>Text</b>	BERT	86.8 $\pm$ 0.2	-
<i>Heading</i>	BERT	-	72.1 $\pm$ 0.2
	RoBERTa	-	71.8 $\pm$ 0.2
<i>Caption</i>	BERT	-	72.7 $\pm$ 0.3
	RoBERTa	-	72.9 $\pm$ 0.4
<i>Abstract</i>	BERT	-	78.3 $\pm$ 0.3
	RoBERTa	-	79.7 $\pm$ 0.2

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*Section 6 Limitations*
- A2. Did you discuss any potential risks of your work?  
*Not applicable. Our paper is foundational research that introduces a method of contrastive learning for the multimodal classification task. It is not tied to particular applications.*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*Abstract, Section 1 Introduction*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*No response.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*No response.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*No response.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*No response.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*No response.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*No response.*

### C Did you run computational experiments?

*Section 4 Experiments, Appendix A.3*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*Appendix A.2 Experimental Settings*

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*Appendix A.2 Experimental Settings*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*Section 4.1 Experimental Setup*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*These libraries are not necessary for our topic.*

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*No response.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*No response.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*No response.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*No response.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*No response.*