

HyDiscGAN: A Hybrid Distributed cGAN for Audio-Visual Privacy Preservation in Multimodal Sentiment Analysis

Zhuojia Wu¹, Qi Zhang^{1,4}, Duoqian Miao^{1*}, Kun Yi², Wei Fan³ and Liang Hu^{1,4}

¹Tongji University, Shanghai, China

²Beijing Institute of Technology, Beijing, China

³University of Oxford, Oxford, United Kingdom

⁴DeepBlue Academy of Sciences, Shanghai, China

{wuzhuojia, zhangqi_cs, dqmiao, lianghu}@tongji.edu.cn, yikun@bit.edu.cn, wei.fan@wrh.ox.ac.uk

Abstract

Multimodal Sentiment Analysis (MSA) aims to identify speakers’ sentiment tendencies in multimodal video content, raising serious concerns about privacy risks associated with multimodal data, such as voiceprints and facial images. Recent distributed collaborative learning has been verified as an effective paradigm for privacy preservation in multimodal tasks. However, they often overlook the privacy distinctions among different modalities, struggling to strike a balance between performance and privacy preservation. Consequently, it poses an intriguing question of maximizing multimodal utilization to improve performance while simultaneously protecting necessary modalities. This paper forms the first attempt at modality-specified (i.e., audio and visual) privacy preservation in MSA tasks. We propose a novel **Hybrid Distributed cross-modality cGAN** framework (**HyDiscGAN**), which learns multimodality alignment to generate fake audio and visual features conditioned on shareable de-identified textual data. The objective is to leverage the fake features to approximate real audio and visual content to guarantee privacy preservation while effectively enhancing performance. Extensive experiments show that compared with the state-of-the-art MSA model, HyDiscGAN can achieve superior or competitive performance while preserving privacy.

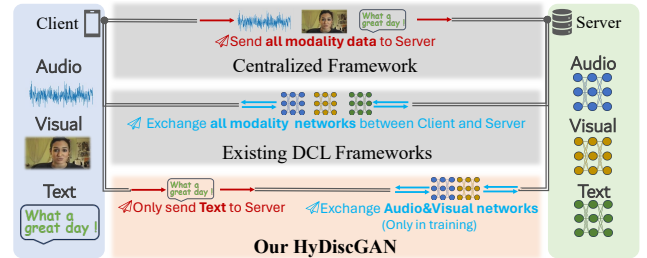
1 Introduction

With the growing prevalence of video content on social media, Multimodal Sentiment Analysis (MSA) is poised to provide new opportunities by leveraging multimodal data to enhance and go beyond traditional text-based sentiment analysis [Zhao *et al.*, 2023]. MSA aims to predict the speaker’s sentiment by utilizing extra information available in audio and visual content instead of only textual content. The audio/visual modality extracts facial emotions and vocal expressions, enabling a more comprehensive sentiment understanding in wide-ranging applications.

*Duoqian Miao is the corresponding author.



(a) Video data often contains personal information, whereas text content, when properly de-identified, easily maintains privacy.



(b) Differences among our MSA framework, the centralized MSA framework, and existing DCL-based MSA frameworks.

Figure 1: Illustration of privacy preservation and MSA frameworks.

Notably, social video data contains massive private information, including Personally Identifiable Information (PII) and biometric data like face images and voiceprints [Regulation, 2016]. Unfortunately, the misuse of personal information often causes a string of public security incidents, raising widespread concerns regarding personal privacy and security within society [Nguyen *et al.*, 2021; Yang *et al.*, 2020]. Upon closer examination of video data, we uncover a crucial yet often overlooked fact: *different modalities carry varying requirements for privacy*, as depicted in Figure 1(a). For instance, legislative efforts aimed at preserving privacy data have emphasized the privacy of personal audio or visual data over textual data [Regulation, 2016]. In addition, techniques such as introducing noises or blurring faces (e.g., differential privacy [Dwork, 2006]) to de-identify audio and visual data can significantly impede the recognition of sentiment cues. In contrast, de-identifying textual data, such as removing sensitive words, can effectively protect privacy without altering the primary semantics [Wang *et al.*, 2023]. The observations encourage us to contemplate an intriguing question: *how to protect the privacy of specific modalities (i.e., audio and visual) when building MSA models?*

Existing MSA approaches typically adopt a centralized

paradigm where multimodal data is collected from personal devices and stored centrally for training, achieving excellent performance yet posing considerable challenges and risks in preserving personal privacy, as shown in Figure 1(b). Instead, increasing efforts have been made to apply Distributed Collaborative Learning (DCL) to multimodal tasks [Yu *et al.*, 2022; Chen and Zhang, 2022; Bao *et al.*, 2023]. Existing DCL frameworks [Kairouz *et al.*, 2021], such as federated learning [Nguyen *et al.*, 2021] and split learning [Thapa *et al.*, 2022], have gained prominence, offering privacy preservation by avoiding centralized data hosting and access. They rely on exchanging multimodal networks between a central server and clients that hold unshareable data for model training and testing, however, they struggle to strike a balance between performance and privacy preservation. Additionally, these endeavors primarily focus on scenarios where all multimodal content is isolated on separate clients, which does not align with the goal of modality-specific privacy preservation in practice. These insights prompt us to explore strategies for *maximizing data utilization to improve performance while simultaneously protecting necessary modalities*.

An intuitive solution to modality-specific privacy preservation is combining centralized and DCL frameworks to create a hybrid distributed learning paradigm, guaranteeing data utilization and protection, respectively. Accordingly, we can initiate a primary idea: train shareable modality data (text) centrally while private modality data (audio and visual) distributively. However, in the setting, where we have separate copies of shareable modality and private modality data on the server and clients respectively, we face a dilemma when it comes to performing model inference. On one hand, performing inference on the server requires access to the data representations of the private modality, which can increase communication costs and pose risks to privacy [Thapa *et al.*, 2022]. On the other hand, performing inference on the client-side necessitates each client to train the entire MSA model to guarantee effective multimodal fusion, needing more client computational resources. Note that it is practically impossible for personal devices (clients) such as smartphones or laptops to have sufficient computing power to accommodate widespread large-scale MSA models. As a result, the hybrid distributed mode inevitably poses two primary challenges stemming from the misaligned treatment between modalities: 1) *achieving effective multimodal alignment*, and 2) *ensuring efficient collaborative communication*.

In light of the above discussion, we propose a novel hybrid distributed collaborative learning framework based on a cross-modality conditional Generative Adversarial Network (cGAN), termed HyDiscGAN. Specifically, we build an audio generator and a visual generator to generate fake features of private audio and visual data respectively in an autoregressive manner. The generators are placed in the server to approximate real features in the clients. On one hand, the generated features are sent to the corresponding audio and visual discriminators in the clients which are regulated by two customized contrastive losses and a cGAN loss. Generators and discriminators are based on Transformers to cater to the sequential audio and visual data. On the other hand, the features are fed into Transformer layers, followed by the gated atten-

tion unit to fuse multimodal features of text, visual, and audio, to perform the downstream sentiment analysis. Note that HyDiscGAN is trained in two stages: 1) the cross-modality cGAN is pre-trained to guarantee *effective multimodal alignment*, where global generators and local discriminators are distributively optimized in an alternating manner; 2) the MSA components are trained along with fine-tuning the generators under keeping the discriminators frozen. Its learning process simulates guessing audio and visual (semantic) features conditioned on text inputs, which is inspired by the empirical observation that individuals can envision the tone and facial expressions associated with a piece of text when it is narrated. Consequently, HyDiscGAN does not require any client-side computation during inference, reducing large collaborative costs to boost *efficient communication*.

Our key contributions can be summarized as follows:

- We propose a novel hybrid DCL framework HyDiscGAN for audio-visual privacy preservation in MSA. To the best of our knowledge, this forms the first endeavor to address modality-specified privacy preservation.
- We customize a cross-modality cGAN to achieve effective multimodal alignment and efficient collaborative communication in HyDiscGAN.
- Experiments on two MSA benchmarks show HyDiscGAN achieves desirable performance competent to SOTA baselines while preserving audio-visual privacy.

2 Related Work

2.1 Multimodal Sentiment Analysis

The current approaches to MSA can be broadly categorized: representation-based and fusion-based methods [Yu *et al.*, 2023; Lao *et al.*, 2024]. Representation-based methods aim to acquire effective representations for each modality, facilitating subsequent fusion processes. One perspective argues that effective sentiment representations should encompass both modality-invariant and modality-variant features [Hazrika *et al.*, 2020; Yu *et al.*, 2021; Lin and Hu, 2022]. Another viewpoint suggests that, in multimodal data, the textual modality predominates, seeking to enhance text representation by integrating textual and non-textual modality information [Wang *et al.*, 2019; Yang *et al.*, 2021; Guo *et al.*, 2022; Su *et al.*, 2023]. Recently, Yang *et al.* [Yang *et al.*, 2023] further improved multimodal informative representation by incorporating contrastive learning and contrastive feature decomposition alongside representation learning. Regarding fusion-based methods, early research categorized them into early-fusion and late-fusion [Yu *et al.*, 2023]. Early-fusion emphasizes learning dependencies in multimodal sequence data [Zadeh *et al.*, 2018a], while late-fusion initially learns independent unimodal representations, integrating them later for sentiment inference [Zadeh *et al.*, 2017]. Recently, Zhao *et al.* [Zhao *et al.*, 2023] achieved SOTA performance in MSA by acquiring more enriched multimodal representations through data augmentation strategies on limited datasets.

2.2 Distributed Collaborative Learning

DCL has gained significant attention in recent years due to its data protection capabilities [Kairouz *et al.*, 2021]. The two

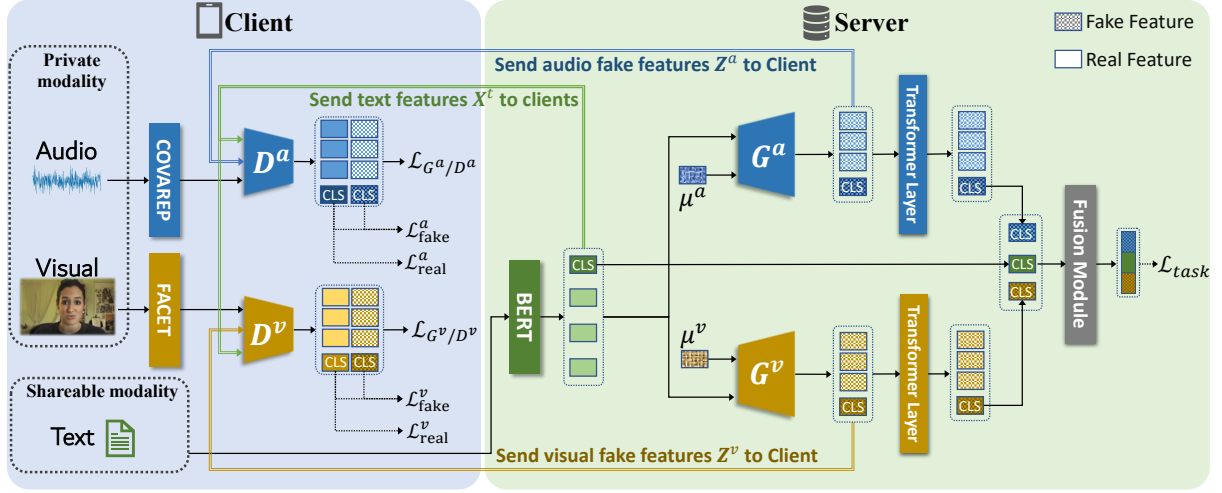


Figure 2: The overall process of HyDiscGAN generates “sufficiently realistic” fake features for private modalities (audio and visual) through hybrid DCL between the server and clients and subsequently fuses them with the real features of the shareable modality (text) for MSA tasks.

most popular frameworks are federated learning [McMahan *et al.*, 2017] and split learning [Gupta and Raskar, 2018]. For federated learning, the FedAvg was initially proposed by Brendan *et al.* [McMahan *et al.*, 2017]. In FedAvg, a complete model is trained on each local client holding data, and the locally updated models are then sent to the server for aggregation, resulting in a global model. Subsequent researchers have made further improvements, such as introducing penalty terms to address non-convex problems [Li *et al.*, 2020] and incorporating momentum mechanisms to enhance its convergence speed and performance [Hsu *et al.*, 2019; Zhu *et al.*, 2020]. The advantages of federated learning lie in the parallelization of computations across multiple clients, while its drawback is its inapplicability to scenarios where client resources are limited. On the contrary, split learning [Gupta and Raskar, 2018] divides a model, such as a deep neural network, into multiple parts, and then performs computations on different devices. Thapa *et al.* [Thapa *et al.*, 2022] proposed the integration of federated learning and split learning, introducing federated split learning, which eliminates the inherent limitations of both frameworks.

2.3 Generative Adversarial Networks

Generative Adversarial Networks (GANs) were initially proposed by Goodfellow *et al.* [Goodfellow *et al.*, 2014]. Subsequently, researchers have proposed various improvements and variations [Radford *et al.*, 2015; Almahairi *et al.*, 2018; He *et al.*, 2024]. Importantly, conditional Generative Adversarial Network (cGAN) [Mirza and Osindero, 2014] introduces conditional information during the training process, enabling the generator to produce samples related to the given conditions. GANs were initially applied in computer vision to generate realistic images in a self-supervised manner and later spread rapidly to other fields [Goodfellow *et al.*, 2014], for natural language processing, like text generation [Zhang *et al.*, 2017], adversarial training [Zhang *et al.*, 2016], and data augmentation [Zhao *et al.*, 2023].

3 Methodology

3.1 Problem Statement

MSA is formulated as a binary/multi-classification or regression task for predicting sentiment labels. In contrast to all previous centralized models, our HyDiscGAN is implemented in a more realistic and secure scenario, encompassing a central server and numerous personal clients. Each client C holds N_C video clips as training or test samples. According to the Introduction, each sample comprises shareable modality data, i.e., text (t), as well as two private modality data, namely, audio (a) and visual (v). The raw data and extracted true features of private modalities are securely maintained on their personal clients throughout the entire process.

For each sample, we obtain its real feature embedding sequences $X^m = [x_1^m, x_2^m, \dots, x_{L^m}^m] \in \mathbb{R}^{L^m \times d^m}$ from three modality data using BERT [Kenton and Toutanova, 2019], COVAREP [Degottex *et al.*, 2014], and FACET [De la Torre and Cohn, 2011], respectively. Here, L^m denotes the length of the sequence, and d^m is the feature dimension. $m \in \{t, *\}$ and $* \in \{a, v\}$ denote the set of private modalities. Following BERT, we introduce $\langle \text{CLS} \rangle$ tag features $x_{\langle \text{CLS} \rangle}^*$ at the end of the audio and visual feature sequences to represent the comprehensive semantics of the sequence. $x_{\langle \text{CLS} \rangle}^*$ is initialized through the average pooling of all features in the sequence.

It is worth noting that our primary motivation is to generate fake feature sequences $Z^* = [z_1^*, z_2^*, \dots, z_{L^*}^*, z_{\langle \text{CLS} \rangle}^*] \in \mathbb{R}^{(L^*+1) \times d^*}$, which approximate the real features X^* extracted from the raw audio and visual data, rather than the raw data itself. This not only reduces sentiment-irrelevant redundant computations but also applies gradient truncation to prevent adversaries from reconstructing the raw data from gradients [Thapa *et al.*, 2022]. Subsequently, Z^* and X^t are utilized for the training or testing of MSA models.

3.2 Hybrid Distributed Collaborative Learning

The training pipeline of HyDiscGAN for MSA is shown in Figure 2. Specifically, its training consists of two steps:

Algorithm 1 Training Cross-Modality cGAN

Input: Multiple training clients; training epoch T ; audio and visual generators G^* ; audio and visual global discriminators D^*

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

S training clients are randomly selected;

▷ Send global discriminators D^* to each client

for each client C in $[S]$ **in parallel do**

▷ Send textual data t to the central server

Server Executes:

$X^t = \text{BERT}(t)$;

$Z^* = G^*(\mu_C^*, X^t; \theta_{G^*}), * \in \{a, v\}$;

▷ Send features X^t and Z^* to client C

Client Executes:

$X^a = \text{COVAREP}(a), X^v = \text{FACET}(v)$;

$\mathcal{L}_{D^*}, \mathcal{L}_{\text{real}}^* = D^*([X^* || Z^*], X^t; \theta_{D^*}), * \in \{a, v\}$;

Update local discriminators D^* based on $(1 - \lambda_D)\mathcal{L}_{D^*} + \lambda_D\mathcal{L}_{\text{real}}^*$;

$\mathcal{L}_{G^*}, \mathcal{L}_{\text{fake}}^* = D^*(Z^*, X^t; \theta_{D^*}), * \in \{a, v\}$;

▷ Send losses $\mathcal{L}_{G^*}, \mathcal{L}_{\text{fake}}^*$, and updated local discriminators D^* to the central server

end for

Server Executes:

Update generators G^* based on $(1 - \lambda_G)\mathcal{L}_{G^*} + \lambda_G\mathcal{L}_{\text{fake}}^*$;

Update global discriminators D^* by averaging the local discriminator parameters received from S clients

end for

Output: Updated audio and visual generators G^*

(1) **Training Cross-Modality cGAN** involves hybrid distributed collaborative learning among numerous clients and a central server to ensure effective multimodal alignment of generators on the server. As shown in Algorithm 1, at the beginning of each training epoch, the central server receives textual data from a group of clients. Guided by the text semantics, the generators G^* on the server generate fake features for the audio and visual modalities, which are then transmitted to their respective clients. Each client computes losses $\mathcal{L}_{D^*}, \mathcal{L}_{G^*}, \mathcal{L}_{\text{real}}^*$, and $\mathcal{L}_{\text{fake}}^*$ using its own real features and the received fake features, and then sends generator’s losses and the local discriminator’s parameters back to the server. The central server updates generators G^* based on the received losses and updates discriminators D^* by averaging the local parameters received from multiple clients.

(2) **Training MSA Component** begins with specific Transformer Layers to further encode the fake features from the generators G^* corresponding to the private modalities. Subsequently, the Fusion Module is employed to combine these private modality fake features with the real features of the shareable modality, which are then used for computing the MSA task loss $\mathcal{L}_{\text{task}}$. During this stage, the generators G^* are fine-tuned as part of the MSA Component, while the global discriminators remain frozen.

3.3 Cross-Modality cGAN

cGAN [Mirza and Osindero, 2014] is a variant of the Generative Adversarial Network [Goodfellow *et al.*, 2014] designed to enable targeted sample generation based on given conditions. It comprises a generator and a discriminator, where the generator produces samples satisfying specific conditions, while the discriminator is used to determine whether the input

samples are generated by the generator or are real.

In our conception, we aim to generate the fake features of private modalities that semantically align with corresponding shareable modality features. To achieve this, we use text information, i.e., the feature sequence X^t encoded by BERT, as the conditional input. Additionally, since private modality features are sequential data, to maintain the contextual correlation in generated fake feature sequences Z^* , we adopt an autoregressive manner to generate features at each temporal position. The generation process can be formalized as:

$$z_i^* = G^*(z_{0:i-1}^*, X^t; \theta_{G^*}) \quad (1)$$

where θ_{G^*} is the set of trainable parameters. Especially, $z_0^* = \mu^*$ and $\mu^* \sim \mathcal{N}(0, 1)$ is a random feature vector sampled from a Gaussian distribution.

Transformer Layer

In our framework, apart from the classifier and Fusion Module, all other components utilize the Transformer [Vaswani *et al.*, 2017] as the backbone, with distinctions solely in input and output. Transformer is an efficient neural architecture for modeling sequential data. The core computation is the Scaled Dot-Product Attention, which is defined as:

$$\text{Att}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_K}}\right)V \quad (2)$$

where Q, K , and V are obtained by linear mapping of the input feature sequence. Furthermore, the ‘‘Multi-head’’ operation [Vaswani *et al.*, 2017] is used to jointly focus on different parts of the input sequence on multiple subspaces, enhancing the ability to capture information.

Transformer-based Autoregressive Generator

Inspired by the neural machine translation model [Vaswani *et al.*, 2017], we construct a Transformer-based Autoregressive Generator. It is a simple variant of the basic Transformer Layer. Specifically, it comprises two attention structures serving different purposes: (1) Intra-modality Multi-Head Attention is employed to capture contextual relationships within the unimodal feature sequence, and corresponding $Q_{\text{ra}}^*, K_{\text{ra}}^*$, and V_{ra}^* are all derived from the mapping of $Z_{0:i-1}^*$, i.e., $[Q_{\text{ra}}^*, K_{\text{ra}}^*, V_{\text{ra}}^*] = [z_{0:i-1}^* W_{Q_{\text{ra}}}^*, z_{0:i-1}^* W_{K_{\text{ra}}}^*, z_{0:i-1}^* W_{V_{\text{ra}}}^*]$, where $W_{Q_{\text{ra}}}^*, W_{K_{\text{ra}}}^*$, and $W_{V_{\text{ra}}}^*$ are parameter matrices. (2) Inter-modality Multi-Head Attention is used to capture cross-modality alignment information attentive to the shareable modality feature sequence X^t . Hence, $Q_{\text{er}}^*, K_{\text{er}}^*$, and V_{er}^* are calculated as $[Q_{\text{er}}^*, K_{\text{er}}^*, V_{\text{er}}^*] = [z_{0:i-1}^* W_{Q_{\text{er}}}^*, X^t W_{K_{\text{er}}}^*, X^t W_{V_{\text{er}}}^*]$, where $W_{Q_{\text{er}}}^*, W_{K_{\text{er}}}^*$, and $W_{V_{\text{er}}}^*$ are parameter matrices.

Transformer-based Discriminator

For discriminators D^* , their structure is identical to generators G^* , except for the exclusion of autoregressive iteration steps. The inputs includes fake feature sequences Z^* generated by G^* and corresponding real feature sequences X^* . An additional binary classifier is added to the output layer to discriminate between the generated and real features. The discriminator plays a crucial role by providing feedback to the generator to enhance its ability to generate ‘‘sufficiently realistic’’ fake features.

3.4 MSA Component

We further introduce two Transformer Layers for learning deep semantic representations of non-textual modality features. Specifically, the generated fake audio and visual feature sequences Z^* are encoded through corresponding Transformer Layers before being fed to the Fusion Module.

Fusion Module

This module is used to fuse $\langle \text{CLS} \rangle$ tag features of different modalities and regulate the influence of each modality feature in the final sentiment prediction via the gated attention unit [Dhingra *et al.*, 2017]. The operation of the gated attention unit is formulated for each modality as follows:

$$h_{\text{output}}^m = \text{GAtt}(h_{\text{input}}^m; \theta_{\text{GAtt}}^m) \odot h_{\text{input}}^m \quad (3)$$

where the gated attention function GAtt is a fully connected linear layer with sigmoid activation, and its output dimension is equal to the input dimension. θ_{GAtt} is the set of trainable parameters. The symbol \odot denotes the Hadamard product [Horn, 1990]. Specifically, $h_{\text{input}} = x_{\langle \text{CLS} \rangle}^t$ for the textual modality and $h_{\text{input}} = z_{\langle \text{CLS} \rangle}^*$ for the audio and visual modalities. Finally, the tensor $h_{\text{final}} = [h_{\text{output}}^v; h_{\text{output}}^a; h_{\text{output}}^t]$ is formed by concatenating features from the three modalities, which is then utilized for sentiment prediction.

3.5 Learning Objectives

Our framework contains three learning objectives: cGAN Losses, customized Contrastive Losses, and MSA task Loss.

cGAN Losses

For a training sample that has private modality real feature sequences X^* and generated fake feature sequences Z^* , cGAN losses are defined as:

$$\mathcal{L}_{G^*} = \frac{1}{L^* + 1} \sum_{i=1}^{L^*+1} [\log(1 - D^*(G^*(z_{0:i-1}^*, X^t)))] \quad (4)$$

$$\mathcal{L}_{D^*} = \frac{1}{L^* + 1} \sum_{i=1}^{L^*+1} [\log(1 - D^*(x_{0:i-1}^*, X^t)) + \log D^*(G^*(z_{0:i-1}^*, X^t))] \quad (5)$$

where \mathcal{L}_{G^*} and \mathcal{L}_{D^*} represent the losses of the generator and discriminator, respectively. L^* are lengths of feature sequences. Specifically, the $\langle \text{CLS} \rangle$ tag feature is also utilized for computation. In the training of Cross-modality cGAN, \mathcal{L}_{G^*} and \mathcal{L}_{D^*} are alternately minimized.

Contrastive Losses

We design two sample separation loss terms based on NT-Xent contrastive loss [Chen *et al.*, 2020], which are used to further regularize the learning process for both the discriminator and the generator. Specifically, for a sample in training client C , its real and fake $\langle \text{CLS} \rangle$ tag features are $x_{\langle \text{CLS} \rangle}^*$ and $z_{\langle \text{CLS} \rangle}^*$, respectively. (1) the Real-Real contrastive loss $\mathcal{L}_{\text{real}}^*$ is employed to regulate the discriminator:

$$\mathcal{L}_{\text{real}}^* = -\log \frac{e^{\text{sim}(x_{\langle \text{CLS} \rangle}^*, x_{\langle \text{CLS} \rangle}^{*+})/\tau}}{\sum_{\{x_{\langle \text{CLS} \rangle}^{*-}\} \in C} e^{\text{sim}(x_{\langle \text{CLS} \rangle}^*, x_{\langle \text{CLS} \rangle}^{*-})/\tau}} \quad (6)$$

Dataset	Train		Valid		Test	
	#S	#Sp	#S	#Sp	#S	#Sp
MOSI	1,284	52	229	10	686	31
MOSEI	16,326	150	1,871	50	4,659	100

Table 1: Statistics of the MOSI and MOSEI Datasets. #S represents the number of video clips, i.e., the number of samples. #Sp indicates the number of distinct speakers, i.e., the number of clients.

where sim is Cosine similarity function and τ is the temperature parameter. $\{x_{\langle \text{CLS} \rangle}^{*-}\} \in C$ denotes the feature set of samples in client C with a different sentiment polarity from the sample corresponding to $x_{\langle \text{CLS} \rangle}^*$. Conversely, $x_{\langle \text{CLS} \rangle}^{*+}$ is the feature of a sample with the same sentiment polarity, randomly sampled from client C .

(2) the Real-Fake contrastive loss $\mathcal{L}_{\text{fake}}^*$ is introduced to regulate the generator:

$$\mathcal{L}_{\text{fake}}^* = -\log \frac{e^{\text{sim}(z_{\langle \text{CLS} \rangle}^*, x_{\langle \text{CLS} \rangle}^*)/\tau}}{\sum_{\{z_{\langle \text{CLS} \rangle}^{*\text{other}}\} \in C} e^{\text{sim}(z_{\langle \text{CLS} \rangle}^*, z_{\langle \text{CLS} \rangle}^{*\text{other}})/\tau}} \quad (7)$$

where $\{z_{\langle \text{CLS} \rangle}^{*\text{other}}\} \in C$ is the feature set of samples in client C , excluding the sample corresponding to $z_{\langle \text{CLS} \rangle}^*$.

MSA Loss

Let y and \hat{y} denote the true and predicted sentiment labels of a sample, respectively. The MSA task loss $\mathcal{L}_{\text{task}}$ is defined:

$$\mathcal{L}_{\text{task}} = \begin{cases} \frac{1}{N_B} \sum_{n=1}^{N_B} y_n \cdot \log \hat{y}_n & \text{for classification} \\ \frac{1}{N_B} \sum_{n=1}^{N_B} (y_n - \hat{y}_n)^2 & \text{for regression} \end{cases} \quad (8)$$

where N_B is the batch size. \hat{y} is obtained through classification or regression predictions on h_{final} .

4 Experiments

4.1 Datasets and Distributed Settings

Two popular MSA benchmark datasets, MOSI [Zadeh *et al.*, 2016] and MOSEI [Zadeh *et al.*, 2018b], are utilized to evaluate the performance of our HyDiscGAN. Statistical details are shown in Table 1. Our results, averaged over five random seed runs, encompass both classification (Acc-2, F1-Score, and Acc-7) and regression (Mean Absolute Error (MAE) and Correlation Coefficient (Corr)) tasks.

4.2 Baselines

To validate the performance of the features generated by HyDiscGAN, we conducted a comparison with several advanced and SOTA MSA models [Zhao *et al.*, 2023; Yang *et al.*, 2023]. These baseline models can be broadly categorized based on their backbone networks: (1) LSTM-based models, denoted as TFN [Zadeh *et al.*, 2017], LMF [Liu *et al.*, 2018], MFN [Zadeh *et al.*, 2018a], MISA [Hazarika *et al.*, 2020], and Self-MM [Yu *et al.*, 2021]; (2) Transformer-based models, denoted as MulT [Tsai *et al.*, 2019], TMMDA [Zhao *et al.*, 2023], and ConFEDE [Yang *et al.*, 2023]. Additionally, there is a model MTAG [Yang *et al.*, 2021] based on GNN. Note that only our HyDiscGAN utilizes generated private modality fake features for MSA, while all other baseline models neglect the preservation of the speaker’s privacy.

Model	MOSI					MOSEI				
	Acc-2 \uparrow	F1-Score \uparrow	Acc-7 \uparrow	MAE \downarrow	Corr \uparrow	Acc-2 \uparrow	F1-Score \uparrow	Acc-7 \uparrow	MAE \downarrow	Corr \uparrow
(G) TFN [Zadeh <i>et al.</i> , 2017]	- / 80.8	- / 80.7	34.9	0.901	0.698	- / 82.5	- / 82.1	51.6	0.593	0.700
(G) LMF [Liu <i>et al.</i> , 2018]	- / 82.4	- / 82.4	33.2	0.917	0.695	78.5 / 81.9	79.0 / 81.7	51.6	0.573	0.714
(G) MFN [Zadeh <i>et al.</i> , 2018a]	77.4 / -	77.3 / -	34.1	0.965	0.632	79.0 / 82.9	79.6 / 82.9	51.3	0.573	0.718
(G) MulT [Tsai <i>et al.</i> , 2019]	- / 83.0	- / 82.8	40.0	0.871	0.698	- / 82.5	- / 82.3	52.8	0.580	0.703
(B) MISA [Hazarika <i>et al.</i> , 2020]	81.8 / 83.4	81.7 / 83.6	42.3	0.783	0.761	<u>83.6</u> / 85.5	83.8 / 85.3	52.2	0.555	0.756
(B) MTAG [Yang <i>et al.</i> , 2021]	- / 82.3	- / 82.1	-	0.866	0.722	-	-	-	-	-
(B) Self-MM [Yu <i>et al.</i> , 2021]	83.4 / 85.5	83.4 / 85.4	46.7	0.708	0.796	83.8 / 85.2	83.8 / 84.9	53.9	<u>0.531</u>	<u>0.765</u>
(B) TMMDA [Zhao <i>et al.</i> , 2023]	- / 86.9	- / 86.9	-	0.703	0.801	-	-	-	-	-
(B) ConFEDE [Yang <i>et al.</i> , 2023]	84.2 / 85.5	84.1 / 85.5	42.3	0.742	0.784	81.7 / <u>85.8</u>	<u>82.2</u> / <u>85.8</u>	54.9	0.522	0.780
(B) HyDiscGAN (ours)	<u>84.1</u> / <u>86.7</u>	<u>83.7</u> / <u>86.3</u>	<u>43.2</u>	0.749	0.782	81.9 / 86.3	82.1 / 86.2	54.4	0.533	0.761

Table 2: Predicted results of different MSA models on MOSI and MOSEI datasets. “ \uparrow ” indicates that larger values represent better results and “ \downarrow ” signifies the opposite. (G) and (B) represent using Glove and BERT as text feature extractors, respectively. In Acc-2 and F1 score columns, the number on the left side of “/” corresponds to “negative/non-negative” and the number on the right side corresponds to “negative/positive”. Bold values represent optimal performance and underlined values indicate suboptimal performance.

Term	ConFEDE	-FL	-SL	-SFL	HyDiscGAN
Privacy preservation	\times	\checkmark	\checkmark	\checkmark	\checkmark
Distributed computing	\times	\checkmark	\checkmark	\checkmark	\checkmark
Generative capacity	\times	\times	\times	\times	\checkmark
No computations on testing clients	\checkmark	\times	\times	\times	\checkmark
Client-side training	-	Parallel	Sequential	Parallel	Parallel
Scale of parameters (per client)	-	109.5M	23.9M	23.9M	77.8K
Scale of communication parameters (one epoch)	-	109.5M \times 2S	23.9M \times 2S	23.9M \times 2S	77.8K \times 2S

Table 3: Comparison of key attributes and training costs (on the MOSI dataset) of foundational DCL frameworks, including Federated Learning (-FL), Split Learning (-SL), Federated Split Learning (-SFL), and our HyDiscGAN. S is the count of training clients in one epoch.

Model	Acc-2 \uparrow	F1-Score \uparrow	Acc-7 \uparrow	MAE \downarrow	Corr \uparrow
ConFEDE	84.2 / 85.5	84.1 / 85.5	42.3	0.742	0.784
-FL	81.4 / 81.7	81.3 / 81.5	40.7	0.803	0.721
-SL	83.5 / 84.2	83.1 / 83.9	41.6	0.765	0.767
-SFL	82.8 / 83.2	82.7 / 83.0	41.3	0.811	0.734
HyDiscGAN	84.1 / 86.7	83.7 / 86.3	43.2	0.749	0.782

Table 4: Predicted results of different DCL frameworks on the MOSI dataset, including Federated Learning (-FL), Split Learning (-SL), Federated Split Learning (-SFL), and our HyDiscGAN.

To assess the MSA performance and communication costs of HyDiscGAN in distributed training, we deployed ConFEDE across three widely used DCL frameworks: Federated Learning (-FL) [McMahan *et al.*, 2017], Split Learning (-SL) [Gupta and Raskar, 2018], and Federated Split Learning (-SFL) [Thapa *et al.*, 2022]. ConFEDE and its three variants are implemented using the provided codes from the authors. In -FL, each client independently trains a complete ConFEDE model using its local data, which is then aggregated at the central server. In -SL and -SFL, we adhere to the minimum split principle, aiming to perform as much computation as possible on the central server.

4.3 Performance Analysis

Comparison with Advanced MSA Models

Table 2 presents the comparative results between our proposed HyDiscGAN and other MSA models on MOSI and MOSEI datasets. In detail, HyDiscGAN achieves suboptimal performance across all classification metrics on the MOSI dataset. In the fundamental binary sentiment classification metrics (Acc-2 and F1-Score), the results of HyDiscGAN are, on average, only 0.325% lower than the SOTA performance.

Note that in the “negative/positive” binary classification, HyDiscGAN ranks second, closely trailing behind the SOTA model TMMDA which incorporates data augmentation techniques. Moreover, on the MOSEI dataset, HyDiscGAN significantly outperforms ConFEDE and achieves SOTA performance in this task, exhibiting an average improvement of 0.45% higher than the suboptimal model ConFEDE. Furthermore, in other metrics, HyDiscGAN also demonstrates competitiveness. This indicates that the private modality fake features generated by HyDiscGAN contain high-quality sentiment cues, comparable to real features.

HyDiscGAN does not exhibit the same level of performance in regression tasks, i.e., MAE and Corr, as in classification tasks. One possible reason is that the Real-Real contrastive loss can only separate samples with different sentiment polarities, lacking regularization for samples with the same sentiment polarity but differing only in intensity during the feature generation process.

Comparison with Existing DCL Frameworks

Table 3 presents a comparison of key attributes and training costs between our HyDiscGAN and three existing DCL frameworks that deploy the SOTA MSA model ConFEDE. Overall, the training costs on the client side are significantly reduced with HyDiscGAN (reduced by 99.93% compared to -FL, and 99.68% compared to -SL and -SFL). Simultaneously, HyDiscGAN possesses the capability to generate fake features for private modalities, resulting in zero costs on the client side during testing. In contrast, other DCL frameworks require the same costs during testing as in training, making HyDiscGAN more suitable for testing scenarios with completely limited resources.

The results of comparative MSA methods are outlined in

Variant	Acc-2 \uparrow	F1-Score \uparrow	MAE \downarrow	Corr \uparrow
Real feature (Only Audio) [†]	58.2	57.0	1.150	0.144
Fake feature (Only Audio)	65.2	61.6	1.147	0.162
Real feature (Only Visual) [†]	57.4	57.0	1.160	0.143
Fake feature (Only Visual)	65.3	65.1	1.139	0.168
cGAN loss (Only)	85.3	84.9	0.751	0.778
w/o \mathcal{L}_{real}^*	85.4	85.2	0.752	0.774
w/o \mathcal{L}_{fake}^*	86.0	85.7	0.750	0.779
HyDiscGAN	86.7	86.3	0.749	0.782

Table 5: Ablation results of real/fake features for private modalities (audio-visual). \uparrow indicates the results from the baseline TMMDA. “w/o” denotes “without”.

Table 4. While HyDiscGAN may not surpass the centralized model ConFEDE on specific evaluation metrics, it demonstrates notable superiority over ConFEDE across all metrics when ConFEDE is deployed in existing DCL frameworks. The influence arises from the label distribution skew [Zhang *et al.*, 2022] in client data within existing DCL frameworks for ConFEDE. When applied to MSA tasks, HyDiscGAN follows the two-stage training approach and employs the hybrid DCL strategy exclusively during the stage of learning the private modality real feature distribution on the client. This stage involves self-supervised learning and is therefore not influenced by the distribution of sentiment labels.

4.4 Ablation Study

Effects of Generated Fake Features

Table 5 upper section displays the performance of predicting sentiment tendencies using only visual or audio modality features. One observation is that in both modalities, the fake features we generated show significant performance improvements across all metrics compared to real features. This is attributed to the Cross-Modality cGAN we constructed, which generates non-textual modality features from text features. As the textual modality contains more sentiment cues, the generated features carry more sentiment information.

Effects of Customized Contrastive Losses

Table 5 lower section demonstrates the impact of two customized contrastive losses \mathcal{L}_{real}^* and \mathcal{L}_{fake}^* , developed by us, on the performance of MSA tasks. We observed a performance enhancement with both losses, particularly with \mathcal{L}_{real}^* . This is because \mathcal{L}_{real}^* is based on a regularization term applied to real features with different sentiment polarities. It promotes the aggregation of samples with the same polarity in the feature space while encouraging the separation of samples with different polarities, leading to a more distinct representation of sentiment information.

4.5 Convergence Analysis

When training Cross-modality cGAN, there is a mutual game between the generator and discriminator, which may lead to training instability [Radford *et al.*, 2015]. We present the convergence curves of losses during the training of Cross-Modality cGAN in HyDiscGAN, as shown in Figure 3. It can be observed that the losses of generators and discriminators eventually converge to low values. This indicates that HyDiscGAN is capable of generating “sufficiently realistic” fake features.

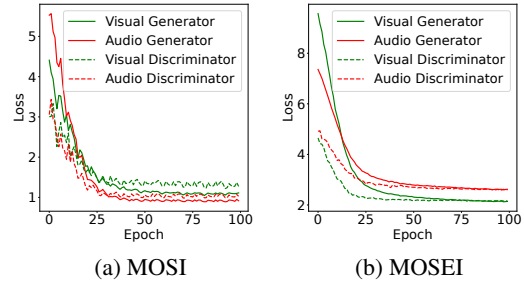


Figure 3: Convergence visualization of training the cross-modality cGAN in HyDiscGAN on MOSI and MOSEI datasets, respectively.

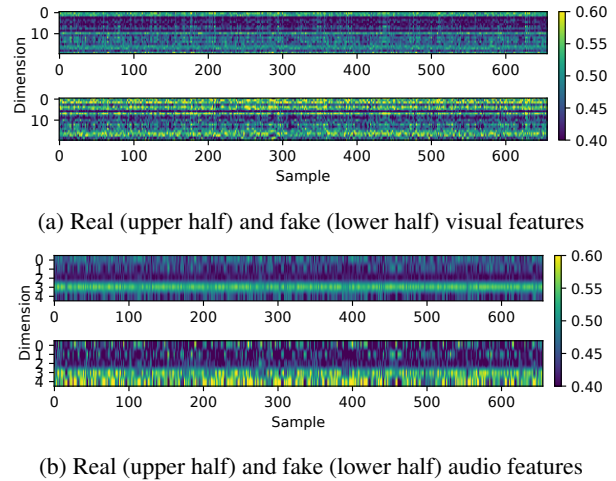


Figure 4: Visualization of the gated attention weights in the Fusion Module for audio-visual features on the test set of MOSI. Brighter regions imply higher unimodal information flow through the gates.

4.6 Visualization

To further validate the effectiveness of generated fake features for MSA tasks, we qualitatively visualize the differences in their contributions to the final sentiment prediction compared to real features. As shown in Figure 4, the information in audio and visual fake features generated by HyDiscGAN is more retained in h_{final} , indicating their broader involvement in sentiment prediction and underscoring their effectiveness.

5 Conclusion

In this paper, we propose a novel hybrid DCL framework HyDiscGAN for audio-visual privacy preservation in MSA. HyDiscGAN conducts training through direct communication between the server and clients, aiming to avoid constructing centralized datasets that expose personal privacy. Compared to other DCL frameworks, HyDiscGAN achieves a better balance between performance and privacy preservation. Additionally, it demonstrates significantly superior training efficiency on the client side, making it more suitable for scenarios with limited client resources. Extensive experiments verify that, while preserving privacy, HyDiscGAN competes comparably with the SOTA models in MSA tasks.

Acknowledgments

This work was supported partially by the National Key Research and Development Program of China under Grant No.2022YFB3104700, and the National Science Foundation of China under Grants No.62376198, 62276190, 62076182, and 62006172.

References

- [Almahairi *et al.*, 2018] Amjad Almahairi, Sai Rajeshwar, Alessandro Sordani, Philip Bachman, and Aaron Courville. Augmented cycleGAN: Learning many-to-many mappings from unpaired data. In *ICML*, pages 195–204. ACM Press, 2018.
- [Bao *et al.*, 2023] Guanyin Bao, Qi Zhang, Duoqian Miao, Zixuan Gong, and Liang Hu. Multimodal federated learning with missing modality via prototype mask and contrast. *arXiv preprint arXiv:2312.13508*, 2023.
- [Chen and Zhang, 2022] Jiayi Chen and Aidong Zhang. Fedmsplit: Correlation-adaptive federated multi-task learning across multimodal split networks. In *SIGKDD*, pages 87–96. ACM Press, 2022.
- [Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. ACM Press, 2020.
- [De la Torre and Cohn, 2011] Fernando De la Torre and Jeffrey F Cohn. Facial expression analysis. *Visual analysis of humans: Looking at people*, pages 377–409, 2011.
- [Degottex *et al.*, 2014] Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer. Covarep—a collaborative voice analysis repository for speech technologies. In *ICASSP*, pages 960–964. IEEE, 2014.
- [Dhingra *et al.*, 2017] Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. Gated-attention readers for text comprehension. In *ACL*, pages 1832–1846. ACL Press, 2017.
- [Dwork, 2006] Cynthia Dwork. Differential privacy. In *International Colloquium on Automata, Languages, and Programming*, pages 1–12. Springer, 2006.
- [Goodfellow *et al.*, 2014] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680. MIT Press, 2014.
- [Guo *et al.*, 2022] Jiwei Guo, Jiajia Tang, Weichen Dai, Yu Ding, and Wanzeng Kong. Dynamically adjust word representations using unaligned multimodal information. In *ACM MM*, pages 3394–3402. ACM Press, 2022.
- [Gupta and Raskar, 2018] Otkrist Gupta and Ramesh Raskar. Distributed learning of deep neural network over multiple agents. *Journal of Network and Computer Applications*, 116:1–8, 2018.
- [Hazarika *et al.*, 2020] Devamanyu Hazarika, Roger Zimmermann, and Soujanya Poria. Misa: Modality-invariant and-specific representations for multimodal sentiment analysis. In *ACM MM*, pages 1122–1131. ACM Press, 2020.
- [He *et al.*, 2024] H. He, Q. Zhang, S. Wang, K. Yi, Z. Niu, and L. Cao. Learning informative representation for fairness-aware multivariate time-series forecasting: A group-based perspective. *IEEE Transactions on Knowledge Data Engineering*, 01:1–13, oct 2024.
- [Horn, 1990] Roger A Horn. The hadamard product. In *Matrix Theory and Applications*, pages 87–169. American Mathematical Society, 1990.
- [Hsu *et al.*, 2019] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.
- [Kairouz *et al.*, 2021] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- [Kenton and Toutanova, 2019] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186. ACL Press, 2019.
- [Lao *et al.*, 2024] An Lao, Qi Zhang, Chongyang Shi, Longbing Cao, Kun Yi, Liang Hu, and Duoqian Miao. Frequency spectrum is more effective for multimodal representation and fusion: A multimodal spectrum rumor detector. In *AAAI*, pages 18426–18434. AAAI Press, 2024.
- [Li *et al.*, 2020] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020.
- [Lin and Hu, 2022] Ronghao Lin and Haifeng Hu. Multimodal contrastive learning via uni-modal coding and cross-modal prediction for multimodal sentiment analysis. In *EMNLP*, pages 511–523. ACL Press, 2022.
- [Liu *et al.*, 2018] Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, AmirAli Bagher Zadeh, and Louis-Philippe Morency. Efficient low-rank multimodal fusion with modality-specific factors. In *ACL*, pages 2247–2256. ACL Press, 2018.
- [McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, pages 1273–1282. JMLR, 2017.
- [Mirza and Osindero, 2014] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

- [Nguyen *et al.*, 2021] Dinh C Nguyen, Ming Ding, Pubudu N Pathirana, Aruna Seneviratne, Jun Li, and H Vincent Poor. Federated learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 23(3):1622–1658, 2021.
- [Radford *et al.*, 2015] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [Regulation, 2016] Protection Regulation. Regulation (eu) 2016/679 of the european parliament and of the council. *Regulation (eu)*, 679:2016, 2016.
- [Su *et al.*, 2023] Xiangrui Su, Qi Zhang, Chongyang Shi, Jiachang Liu, and Liang Hu. Syntax tree constrained graph network for visual question answering. In *ICONIP (5)*, volume 14451 of *Lecture Notes in Computer Science*, pages 122–136. Springer, 2023.
- [Thapa *et al.*, 2022] Chandra Thapa, Pathum Chamikara Mahawaga Arachchige, Seyit Camtepe, and Lichao Sun. Splitfed: When federated learning meets split learning. In *AAAI*, pages 36–8:8485–8493. AAAI Press, 2022.
- [Tsai *et al.*, 2019] YH Tsai, S Bai, JZ Kolter, LP Morency, R Salakhutdinov, et al. Multimodal transformer for unaligned multimodal language sequences. In *ACL*, volume 2019, pages 6558–6569. ACL Press, 2019.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 6000–6010. MIT Press, 2017.
- [Wang *et al.*, 2019] Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. In *AAAI*, pages 33–01:7216–7223. AAAI Press, 2019.
- [Wang *et al.*, 2023] Yinggui Wang, Wei Huang, and Le Yang. Privacy-preserving end-to-end spoken language understanding. In *IJCAI*, pages 5224–5232. Morgan Kaufmann, 2023.
- [Yang *et al.*, 2020] Liu Yang, Ben Tan, Vincent W Zheng, Kai Chen, and Qiang Yang. Federated recommendation systems. *Federated Learning: Privacy and Incentive*, pages 225–239, 2020.
- [Yang *et al.*, 2021] Jianing Yang, Yongxin Wang, Ruitao Yi, Yuying Zhu, Azaan Rehman, Amir Zadeh, Soujanya Poria, and Louis-Philippe Morency. Mtag: Modal-temporal attention graph for unaligned human multimodal language sequences. In *NAACL*, pages 1009–1021. ACL Press, 2021.
- [Yang *et al.*, 2023] Jiuding Yang, Yakun Yu, Di Niu, Weidong Guo, and Yu Xu. Confede: Contrastive feature decomposition for multimodal sentiment analysis. In *ACL*, pages 7617–7630. ACL Press, 2023.
- [Yu *et al.*, 2021] Wenmeng Yu, Hua Xu, Ziqi Yuan, and Jiele Wu. Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis. In *AAAI*, volume 35(12), pages 10790–10797. AAAI Press, 2021.
- [Yu *et al.*, 2022] Qiying Yu, Yang Liu, Yimu Wang, Ke Xu, and Jingjing Liu. Multimodal federated learning via contrastive representation ensemble. In *ICLR*, 2022.
- [Yu *et al.*, 2023] Yakun Yu, Mingjun Zhao, Shi-ang Qi, Feiran Sun, Baoxun Wang, Weidong Guo, Xiaoli Wang, Lei Yang, and Di Niu. Conki: Contrastive knowledge injection for multimodal sentiment analysis. In *ACL*, pages 13610–13624. ACL Press, 2023.
- [Zadeh *et al.*, 2016] Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6):82–88, 2016.
- [Zadeh *et al.*, 2017] Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Tensor fusion network for multimodal sentiment analysis. In *EMNLP*, pages 1103–1114. ACL Press, 2017.
- [Zadeh *et al.*, 2018a] Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Memory fusion network for multi-view sequential learning. In *AAAI*, pages 5634–5641. AAAI Press, 2018.
- [Zadeh *et al.*, 2018b] AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal language analysis in the wild: Cmmosei dataset and interpretable dynamic fusion graph. In *ACL*, pages 2236–2246. ACL Press, 2018.
- [Zhang *et al.*, 2016] Yizhe Zhang, Zhe Gan, and Lawrence Carin. Generating text via adversarial training. In *NeurIPS workshop on Adversarial Training*, volume 21, pages 21–32. MIT Press, 2016.
- [Zhang *et al.*, 2017] Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, and Lawrence Carin. Adversarial feature matching for text generation. In *ICML*, pages 4006–4015. ACM Press, 2017.
- [Zhang *et al.*, 2022] Jie Zhang, Zhiqi Li, Bo Li, Jianghe Xu, Shuang Wu, Shouhong Ding, and Chao Wu. Federated learning with label distribution skew via logits calibration. In *ICML*, pages 26311–26329. ACM Press, 2022.
- [Zhao *et al.*, 2023] Xianbing Zhao, Yixin Chen, Sicen Liu, Xuan Zang, Yang Xiang, and Buzhou Tang. Tmmda: A new token mixup multimodal data augmentation for multimodal sentiment analysis. In *WWW*, pages 1714–1722. ACM Press, 2023.
- [Zhu *et al.*, 2020] Chengzhang Zhu, Qi Zhang, Longbing Cao, and Arman Abrahamyan. Mix2vec: Unsupervised mixed data representation. In *DSAA*, pages 118–127. IEEE, 2020.