

# MMGCN: Multimodal Fusion via Deep Graph Convolution Network for Emotion Recognition in Conversation

Jingwen Hu, Yuchen Liu, Jinming Zhao, Qin Jin\*

School of Information, Renmin University of China

{hujingwen\_benja, Liuyuchen\_Alfred, zhaojinming, qjin}@ruc.edu.cn

## Abstract

Emotion recognition in conversation (ERC) is a crucial component in affective dialogue systems, which helps the system understand users' emotions and generate empathetic responses. However, most works focus on modeling speaker and contextual information primarily on the textual modality or simply leveraging multimodal information through feature concatenation. In order to explore a more effective way of utilizing both multimodal and long-distance contextual information, we propose a new model based on multimodal fused graph convolutional network, MMGCN, in this work. MMGCN can not only make use of multimodal dependencies effectively, but also leverage speaker information to model inter-speaker and intra-speaker dependency. We evaluate our proposed model on two public benchmark datasets, IEMOCAP and MELD, and the results prove the effectiveness of MMGCN, which outperforms other SOTA methods by a significant margin under the multimodal conversation setting.

## 1 Introduction

Emotion is an important part of human daily communication. Emotion Recognition in Conversation (ERC) aims to automatically identify and track the emotional status of speakers during a dialogue. It has attracted increasing attention from researchers in the field of natural language processing and multimodal processing. ERC has a wide range of potential applications such as assisting conversation analysis for legal trials and e-health services etc. It is also a key component for building natural human-computer interactions that can produce emotional responses in a dialogue.

The fast growing availability of conversational data on social media is one of the factors that boost

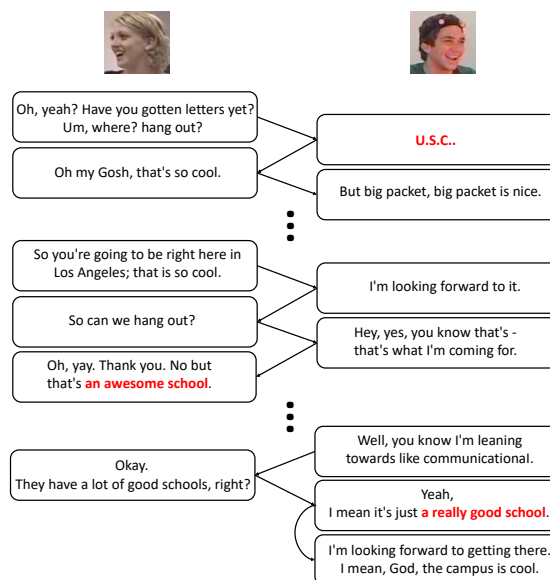


Figure 1: Illustration of an example conversation in the IEMOCAP dataset

the research focus on emotion recognition in conversation. Different from traditional emotion recognition on isolated utterances, emotion recognition in conversation requires context modeling of individual utterances. The context can be attributed to the preceding utterances, temporality in conversation turns, or speaker related information etc. Different models have been proposed to capture the contextual information in previous works, including the LSTM-based model (Poria et al., 2017), the conversational memory network (CMN) model (Hazarika et al., 2018b), interactive conversational memory network (ICON) model (Hazarika et al., 2018a), and DialogueRNN model (Majumder et al., 2019) etc. In the example conversation as shown in Figure 1, the two speakers are chatting in the context of the male speaker being admitted to USC. In this chatting scene, they change topics a few

\*Corresponding Author

times, such as the female speaker inviting the male speaker out to play and so on. But they keep coming back to the topic of USC, and then both of them express an excitement emotional status. It shows that long-distance contextual information is of great help to the prediction of speakers' emotions. However, previous models can not effectively capture both speaker and long-distance dialogue contextual information simultaneously in multi-speaker conversation scenarios. Ghosal et al. (Ghosal et al., 2019), therefore, first propose the DialogueGCN model which applies graph convolutional network (GCN) to capture long-distance contextual information in a conversation. DialogueGCN takes each utterance as a node and connects any nodes that are in the same window within a conversation. It can well model both the dialogue context and speaker information which leads to the state-of-the-art ERC performance. However, like most previous models, DialogueGCN only focuses on the textual modality of the conversation, ignoring effective combination of other modalities such as visual and acoustic modalities. Works that consider multimodal contextual information often conduct the simple feature concatenation type of multimodal fusion.

In order to effectively explore the multimodal information and at the same time capture long-distance contextual information, we propose a new multimodal fused graph convolutional network (MMGCN) model in this work. MMGCN constructs the fully connected graph in each modality, and builds edge connections between nodes corresponding to the same utterance across different modalities, so that contextual information across different modalities can interact. In addition, the speaker information is injected into MMGCN via speaker embedding. Furthermore, different from DialogueGCN, which is a non-spectral domain GCN and its many optimized matrices occupy too much computing resource, we encode the multimodal graph using spectral domain GCN and extend the GCN from a single layer to deep layers. To verify the effectiveness of the proposed model, we carry out experiments on two benchmark multimodal conversation datasets, IEMOCAP and MELD. MMGCN significantly outperforms other models on both datasets.

The rest of the paper is organized as follows: Section 2 discusses some related works; Section 3 introduces the proposed MMGCN model in details;

Section 4 and 5 present the experiment setups on two public benchmark datasets and the analysis of experiment results and ablation study; Finally, Section 6 draws some conclusions.

## 2 Related Work

### 2.1 Emotion Recognition in Conversation

With the fast development of social media, much more interaction data become available, including several open-sourced conversation datasets such as IEMOCAP (Busso et al., 2008), AVEC (Schuller et al., 2012), MELD (Poria et al., 2018), etc. ERC has attracted much research attention recently.

Many previous works focus on modeling contextual information due to its importance in ERC. Poria et al. (Poria et al., 2017) leverage a LSTM-based model to capture interaction history context. Hazarika et al. (Hazarika et al., 2018b,a) first pay attention to the importance of speaker information and exploit different memory networks to model different speakers. DialogueRNN (Majumder et al., 2019) leverage distinct GRUs to capture speakers' contextual information. DialogueGCN (Ghosal et al., 2019) construct the graph considering both speaker and conversation sequential information and achieve the state-of-the-art performance.

### 2.2 Multimodal Fusion

Most recent studies on ERC focus primarily on the textual modality. (Poria et al., 2017; Hazarika et al., 2018b,a) leverage multimodal information through concatenating features from three modalities without modeling the interaction between modalities. (Chen et al., 2017) conduct multimodal fusion at the word-level for emotion recognition of isolated utterances. (Sahay et al., 2018) consider contextual information and use relations in the emotion labels across utterances to predict the emotion (Zadeh et al., 2018) propose MFN to fuse information of multi-views, which aligns features from different modalities well. However, MFN neglects to model speaker information, which is significant to ERC as well. The state-of-the-art dialogueGCN model only considers the textual modality. In order to explore a more effective way of fusing multiple modalities and at the same time capturing contextual conversation information, we propose MMGCN which constructs a graph based on all three modalities.

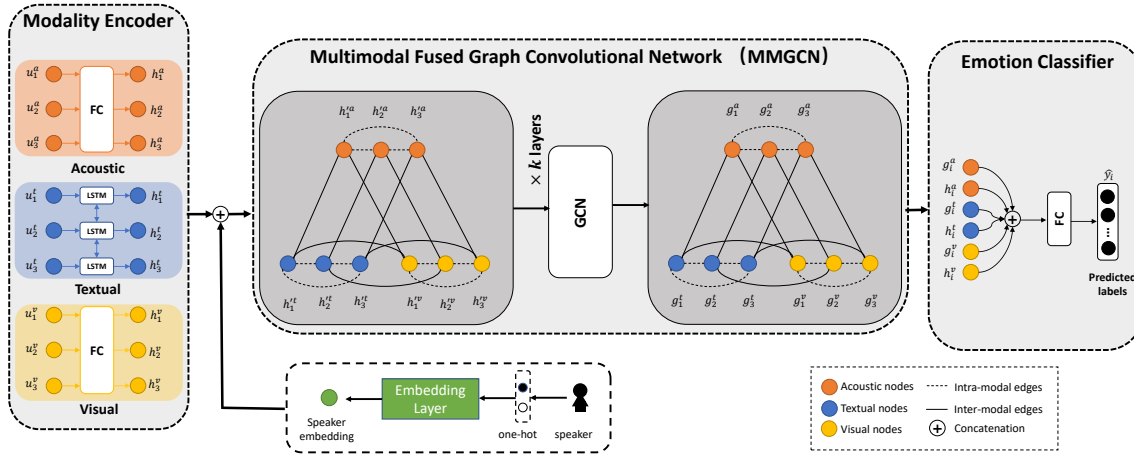


Figure 2: Framework illustration of the MMGCN based emotion recognition in conversation, which consists of three key components: Modality Encoder, Multimodal Graph Convolutional Network, Emotion Classifier.

### 2.3 Graph Convolutional Network

Graph convolutional networks have been widely used in the past few years for their ability to cope with non-Euclidean data. Mainstream GCN methods can be divided into spectral domain methods and non-spectral domain methods (Veličković et al., 2017). Spectral domain GCN methods (Zhang et al., 2019) are based on Laplace Spectral decomposition theory. They can only deal with undirected graphs. Non-spectral domain GCN methods (Veličković et al., 2017; Schlichtkrull et al., 2018; Li et al., 2015) can be applied to both directed and undirected graphs, but consuming larger computing resource. Recently, researchers have proposed methods to make spectral domain GCN deeper without over-smoothing (Li et al., 2019; Chen et al., 2020). In order to further improve MMGCN on ERC, we encode the multimodal graph using spectral domain GCN with deep layers.

## 3 Method

A dialogue can be defined as a sequence of utterances  $\{u_1, u_2, \dots, u_N\}$ , where  $N$  is the number of utterances. Each utterance involves three sources of utterance-aligned data corresponding to three modalities, including acoustic (a), visual (v) and textual (t) modalities, which can be represented as follows:

$$u_i = \{u_i^a, u_i^v, u_i^t\} \quad (1)$$

where  $u_i^a$ ,  $u_i^v$ ,  $u_i^t$  denote the raw feature representation of  $u_i$  from the acoustic, visual and textual

modality, respectively. The emotion recognition in conversation task aims to predict the emotional status label for each utterance  $u_i$  in the conversation based on the available information from all three modalities. Figure 2 illustrates the overall framework of our proposed emotion recognition in conversation system, which consists of three key modules: Modality Encoder, Multimodal Fused Graph Convolutional Network (MMGCN), and Emotion Classifier.

### 3.1 Modality Encoder

As we mentioned above, the dialog context information is important for predicting the emotion label of each utterance. Therefore, it is beneficial to encode the contextual information into the utterance feature representation. We generate the context-aware utterance feature encoding for each modality through the corresponding modality encoder. To be specific, we apply a bidirectional Long Short Term Memory (LSTM) network to encode the sequential textual context information for the textual modality. For the acoustic and visual modalities, we apply a fully connected network. The context-aware feature encoding for each utterance can be formulated as follows:

$$\begin{aligned} h_i^t &= [\overrightarrow{\text{LSTM}}(u_i^t, h_{i-1}^t), \overleftarrow{\text{LSTM}}(u_i^t, h_{i+1}^t)] \\ h_i^a &= W_e^a u_i^a + b_i^a \\ h_i^v &= W_e^v u_i^v + b_i^v \end{aligned} \quad (2)$$

where  $u_i^a$ ,  $u_i^v$ ,  $u_i^t$  are the context-independent raw feature representation of utterance  $i$  from the acoustic, visual and textual modalities, respectively. The

modality encoder outputs the context-aware raw feature encoding  $h_i^a$ ,  $h_i^v$ , and  $h_i^t$  accordingly.

### 3.2 Multimodal fused GCN (MMGCN)

In order to capture the utterance-level contextual dependencies across multiple modalities, we propose a Multimodal fused Graph Convolutional Network (MMGCN). We construct a spectral domain graph convolutional network to encode the multimodal contextual information inspired by (Li et al., 2019; Chen et al., 2020). We also stack more layers to construct a deep GCN. Furthermore, we add learned speaker-embeddings to encode the speaker-level contextual information.

#### 3.2.1 Speaker Embedding

As mentioned above, speaker information is important for ERC. In order to encode the speaker identity information, we add speaker embeddings to the features before constructing the graph. Assuming there are  $M$  parties in a dialogue, then the size of the speaker embedding is  $M$ . We show a two-speaker conversation case in Figure 2. The original speaker identity can be denoted with a one-hot vector  $s_i$  and the speaker embedding  $S_i$  is calculated as follows:

$$S_i = W_s s_i + b_i^s \quad (3)$$

The speaker embedding can then be leveraged to attach speaker information in the graph construction.

#### 3.2.2 Graph Construction

A dialogue with  $N$  utterances can be represented as an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  ( $|\mathcal{V}| = 3N$ ) denotes utterance nodes in three modalities and  $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$  is a set of relationships containing context, speaker and modality dependency. We construct the graph as follows:

**Nodes:** Each utterance is represented by three nodes  $v_i^a$ ,  $v_i^v$ ,  $v_i^t$  in a graph, initialized with  $h_i^a, h_i^v, h_i^t$ , which represent  $[h_i^a, S_i]$ ,  $[h_i^v, S_i]$ ,  $[h_i^t, S_i]$  respectively, corresponding to the three modalities. Thus, given a dialogue with  $N$  utterances, we construct a graph with  $3N$  nodes.

**Edges:** We assume that each utterance has certain connection to other utterances in the same dialogue. Therefore, any two nodes in the same modality in the same dialogue are connected in the graph. Furthermore, each node is connected with the nodes which correspond to the same utterance but from

different modalities. For example,  $v_i^a$  will be connected with  $v_i^v$  and  $v_i^t$  in the graph.

**Edge Weighting:** We assume that if two nodes have higher similarity, the information interaction between them is also more important, and the edge weight between them should be higher. In order to capture the similarities between node representations, following (Skianis et al., 2018), we use the angular similarity to represent the edge weight between two nodes.

There are two types of edges in the graph: 1) edges connecting nodes from the same modality, and 2) edges connecting nodes from different modalities. To differentiate them, we use different edge weighting strategies. For the first type of edges, the edge weight is computed as:

$$A_{ij} = 1 - \frac{\arccos(\text{sim}(n_i, n_j))}{\pi} \quad (4)$$

where  $n_i$  and  $n_j$  denote the feature representations of the  $i$ -th and  $j$ -th node in the graph. For the second type of edges, the edge weight is computed as:

$$A_{ij} = \gamma \left( 1 - \frac{\arccos(\text{sim}(n_i, n_j))}{\pi} \right) \quad (5)$$

where  $\gamma$  is a hyper parameter.

**Graph Learning:** Inspired by (Chen et al., 2020), we build a deep graph convolutional network based on the undirected graph formed following the above construction steps to further encode the contextual dependencies. To be specific, given the undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , let  $\tilde{\mathcal{P}}$  be the renormalized graph Laplacian matrix (Kipf and Welling, 2016) of  $\mathcal{G}$ :

$$\begin{aligned} \tilde{\mathcal{P}} &= \tilde{\mathcal{D}}^{-1/2} \tilde{\mathcal{A}} \tilde{\mathcal{D}}^{-1/2} \\ &= (\mathcal{D} + \mathcal{I})^{-1/2} (\mathcal{A} + \mathcal{I}) (\mathcal{D} + \mathcal{I})^{-1/2} \end{aligned} \quad (6)$$

where  $\mathcal{A}$  denotes the adjacency matrix,  $\mathcal{D}$  denotes the diagonal degree matrix of graph  $G$ , and  $\mathcal{I}$  denotes identity matrix. The iteration of GCN from different layers can be formulated as:

$$\mathcal{H}^{(l+1)} = \sigma(((1-\alpha)\tilde{\mathcal{P}}\mathcal{H}^{(l)} + \alpha\mathcal{H}^{(0)})((1-\beta^{(l)})\mathcal{I} + \beta^{(l)}\mathcal{W}^{(l)})) \quad (7)$$

where  $\alpha$  and  $\beta^{(l)}$  are two hyper parameters,  $\sigma$  denotes the activation function and  $\mathcal{W}^{(l)}$  is a learnable weight matrix. To ensure the decay of the weight matrix adaptively increases when stacking more layers, we set  $\beta^{(l)} = \log(\frac{\eta}{l} + 1)$ , where  $\eta$  is also a hyper parameter. A residual connection to the first layer  $\mathcal{H}^{(0)}$  is added to the representation  $\tilde{\mathcal{P}}\mathcal{H}^{(l)}$  and an identity mapping  $\mathcal{I}$  is added to the weight matrix  $\mathcal{W}^{(l)}$ . With such residual connection, we can make MMGCN deeper to further improve performance.



### 3.3 Emotion Classifier

As described in sec. 3.2.2, we initialize nodes with the combination of utterance feature and speaker embedding,  $h'_i$ .

$$h'_i = [h'_i{}^a, h'_i{}^v, h'_i{}^t]. \quad (8)$$

Let  $g_i^a$ ,  $g_i^v$  and  $g_i^t$  be the features of different modalities encoded by the GCN. The features corresponding to the same utterance are concatenated:

$$g_i = [g_i^a, g_i^v, g_i^t]. \quad (9)$$

We then can concatenate  $g_i$  and  $h'_i$  to generate the final feature representation for each utterance:

$$e_i = [h'_i, g_i], \quad (10)$$

$e_i$  is then fed into a MLP with fully connected layers to predict the emotion label  $\hat{y}_i$  for the utterance:

$$\begin{aligned} l_i &= \text{RELU}(W_l e_i + b_l) \\ P_i &= \text{Softmax}(W_{\text{softmax}} l_i + b_{\text{softmax}}) \\ \hat{y}_i &= \arg \min_k (P_i[k]) \end{aligned} \quad (11)$$

### 3.4 Training Objectives

We use categorical cross-entropy along with L2-regularization as the loss function during training:

$$\mathcal{L} = - \frac{1}{\sum_{s=1}^N c(s)} \sum_{i=1}^N \sum_{j=1}^{c(i)} \log P_{i,j}[y_{i,j}] + \lambda \|\theta\|_2 \quad (12)$$

where  $N$  is the number of dialogues,  $c(i)$  is the number of utterances in dialogue  $i$ ,  $P_{i,j}$  is the probability distribution of predicted emotion labels of utterance  $j$  in dialogue  $i$ ,  $y_{i,j}$  is the expected class label of utterance  $j$  in dialogue  $i$ ,  $\lambda$  is the L2-regularization weight, and  $\theta$  is the set of all trainable parameters. We use stochastic gradient descent based Adam (Kingma and Ba, 2014) optimizer to train our network. Hyper parameters are optimized using grid search.

## 4 Experiment Setups

### 4.1 Dataset

We evaluate our proposed MMGCN model on two benchmark datasets, IEMOCAP (Busso et al., 2008) and MELD (Poria et al., 2018). Both are multi-modal datasets with aligned acoustical, visual and textual information of each utterance in a conversation. Followed (Ghosal et al., 2019), we partition both datasets into train and test sets with roughly

Dataset	dialogues		utterances	
	train+val	test	train+val	test
IEMOCAP	120	31	5810	1623
MELD	1153	280	11098	2610

Table 1: Data distribution of IEMOCAP and MELD

8:2 ratio. Table 1 shows the distribution of train and test samples for both datasets.

**IEMOCAP:** The dataset contains 12 hours of videos of two-way conversations from ten unique speakers, where only the first eight speakers from session one to four are used in the training set. Each video contains a single dyadic dialogue, segmented into utterances. There are in total 7433 utterances and 151 dialogues. Each utterance in the dialogue is annotated with an emotion label from six classes, including happy, sad, neutral, angry, excited and frustrated.

**MELD:** Multi-modal Emotion Lines Dataset (MELD) is a multi-modal and multi-speaker conversation dataset. Compared to the Emotion Lines dataset (Chen et al., 2018), MELD has three modality-aligned conversation data with higher quality. There are in total 13708 utterances, 1433 conversations and 304 different speakers. Specifically, different from dyadic conversation datasets such as IEMOCAP, MELD has three or more speakers in a conversation. Each utterance in the dialogue is annotated with an emotion label from seven classes, including anger, disgust, fear, joy, neutral, sadness and surprise.

### 4.2 Utterance-level Raw Feature Extraction

The textual raw features are extracted using TextCNN following (Hazarika et al., 2018a). The acoustic raw features are extracted using the *OpenS-mile* toolkit with IS10 configuration (Schuller et al., 2011). The visual facial expression features are extracted using a DenseNet (Huang et al., 2015) pre-trained on the Facial Expression Recognition Plus (FER+) corpus (Barsoum et al., 2016).

### 4.3 Implementation Details

The hyperparameters are set as follows: the number of GCN layers are both 4 for IEMOCAP and MELD. The dropout is 0.4. The learning rate is 0.0003. The L2 regularization parameter is 0.00003.  $\alpha$ ,  $\eta$  and  $\gamma$  are set as 0.1, 0.5 and 0.7 respectively. Considering the class-imbalance in

	IEMOCAP							MELD
	Happy	Sad	Neutral	Angry	Excited	Frustrated	Average(w)	Average(w)
BC-LSTM	34.43	60.87	51.81	56.73	57.95	58.92	54.95	56.80
CMN	30.38	62.41	52.39	59.83	60.25	60.69	56.13	-
ICON	29.91	64.57	57.38	63.04	63.42	60.81	58.54	-
DialogueRNN	39.16	<b>81.69</b>	59.77	67.36	72.91	60.27	64.58	57.11
DialogueGCN	<b>47.1</b>	80.88	58.71	66.08	70.97	61.21	65.04	58.23
MMGCN	42.34	78.67	<b>61.73</b>	<b>69.00</b>	<b>74.33</b>	<b>62.32</b>	<b>66.22</b>	<b>58.65</b>

Table 2: ERC performance (F1-score) of different approaches on both IEMOCAP and MELD datasets under the multimodal setting, which means the input includes all the acoustic, visual, and textual modalities; bold font denotes the best performance. Average(w) means weighted average. (The result of CMN and ICON are deficient for suiting two-way conversations only)

MELD, we use focal loss when training MMGCN on MELD. In addition, we add layer normalization after the speaker embedding.

#### 4.4 Evaluation Metrics and Significance Test

Following previous works (Hazarika et al., 2018a; Majumder et al., 2019; Ghosal et al., 2019), we use weighted average f1-score as the evaluation metric. Paired t-test is performed to test the significance of performance improvement with a default significance level of 0.05.

#### 4.5 Compared Baselines

In order to verify the effectiveness of our model, We implement and compare the following models on emotion recognition in conversation.

**BC-LSTM (Poria et al., 2017):** it encodes contextual information through Bi-directional LSTM (Hochreiter and Schmidhuber, 1997) network. The context-aware features are then used for emotion classification. BC-LSTM ignores speaker information as it doesn’t attach any speaker-related information to their model.

**CMN (Hazarika et al., 2018b):** it leverages speaker-dependent GRUs to model utterance context combining dialogue history information. The utterance features with contextual information are subject to two distinct memory networks for both speakers. Due to the fixed number of Memory network blocks, CMN can only serve in dyadic conversation scenarios.

**ICON (Hazarika et al., 2018a):** it extends CMN to model distinct speakers respectively. Same with CMN, two speaker-dependent GRUs are leveraged. Besides, A global GRU is used to track the change of emotion status in the entire conversation and multi-layer memory networks are leveraged to model the global emotion status. Though ICON

improves the result of ERC, it still cannot adapt to a multi-speaker scenario.

**DialogueRNN (Majumder et al., 2019):** it models speakers and sequential information in dialogues through three different GRUs, which include Global GRU, Speaker GRU and Emotion GRU. Specifically, Global GRU models context information, while Speaker dependent GRU models the status of the certain speaker. The two modules update interactively. Emotion GRU detects emotion of utterances in conversation. Furthermore, in the multimodal setting, the concatenation of acoustical, visual, and textual features is used when the speaker talks, but only use visual features otherwise. However, DialogueRNN doesn’t improve much in multimodal settings.

**DialogueGCN (Ghosal et al., 2019):** it applies GCN to ERC, in which the generated features can integrate rich information. Specifically, utterance-level features encoded by bi-lstm are used to initialize the nodes of the graph, edges are constructed within a certain window. Utterances in the same dialogue but with long distance can be connected directly. Relation GCN(Schlichtkrull et al., 2018) and GNN(Morris et al., 2019), which are both non-spectral domain GCN models, are leveraged to encode the graph. However, DialogueGCN only focuses on the textual modality. In order to compare with our MMGCN under the multimodal setting, we extend DialogueGCN by simply concatenating features of three modalities.

## 5 Results and Discussions

We compare our proposed MMGCN with all the baseline models presented in section 4.5 on IEMOCAP and MELD datasets under the multimodal setting. In order to compare the results under the same experiment settings, we reimplement the models in

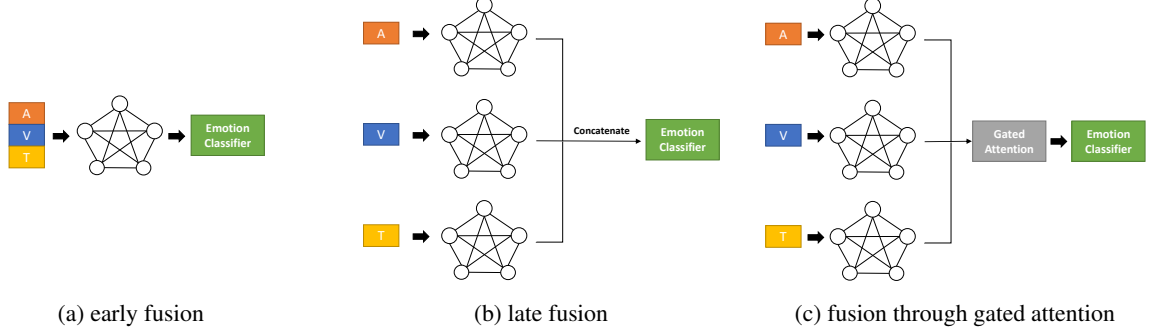


Figure 3: Illustration of the three types of multi-modal fusion methods

the following experiments.

### 5.1 Comparison with other models

Table 2 shows the performance comparison of MMGCN with other models on the two benchmark datasets under the multimodal setting. DialogueGCN was the best performing model when using only the textual modality. Under the multimodal setting, DialogueGCN which is fed with the concatenation of acoustic, visual and textual features achieves some slight improvement over the single textual modality. Our proposed MMGCN improves the F1-score performance over DialogueGCN under the multimodal setting by absolute 1.18% on IEMOCAP and 0.42% on MELD on average, and the improvement is significant with  $p$ -value  $< 0.05$ .

### 5.2 MMGCN under various modality setting

Table 3 shows the performance comparison of MMGCN under different multimodal settings on both benchmark datasets. From Table 3 we can see that the best single modality performance is achieved on the textual modality and the worst is on the visual modality, which is consistent with previously reported findings. Adding acoustic and visual modalities can bring additional performance improvement over the textual modality.

### 5.3 Comparison with other fusion methods

To verify the effectiveness of MMGCN in multi-modal fusion, we compare it with other multimodal fusion methods, including early fusion, late fusion, fusion through gated attention and other representative fusion methods such as MFN(Zadeh et al., 2018) and MulT(Tsai et al., 2019). The first three fusion methods are illustrated in Figure 3. As for

modality	IEMOCAP	MELD
a	54.66	42.63
v	33.86	33.27
t	62.35	57.72
at	65.70	58.02
vt	62.89	57.92
avt	<b>66.22</b>	<b>58.65</b>

Table 3: ERC performance of MMGCN under different multimodal settings, which means the input contains different combination of the three modalities

early fusion, multimodal features are concatenated and fed into GCN directly. As for late fusion, features of different modalities are fed into different GCNs respectively and concatenated afterwards. As for fusion through gated attention, features are fed into different GCNs the same way as in late fusion, and then to a gated attention module. Specifically, the gated attention module can be formulated as follows:

$$r_i^{m_j} = \tanh(W_{m_j} \cdot h_i^{m_j}) \quad (13)$$

$$r_i^{m_k} = \tanh(W_{m_k} \cdot h_i^{m_k}) \quad (14)$$

$$z = \sigma(W_z \cdot h_i^{m_j}) \quad (15)$$

$$r_i^{(m_j, m_k)} = z * r_i^{m_j} + (1 - z) * r_i^{m_k} \quad (16)$$

$$e_i = [r_i^{(a,v)}, r_i^{(a,t)}, r_i^{(v,t)}] \quad (17)$$

where  $m_j$  and  $m_k$  could be any modality among  $\{a, v, t\}$ ,  $h_i^{m_j}$  and  $h_i^{m_k}$  represent the feature encoded by the corresponding modality encoder,  $e_i$  represents the final feature representation for the  $i^{th}$  utterance. Considering MFN and MulT are leveraged to fuse multimodal information sequentially, they are used to replace the Modality Encoder. The fused multimodal features are fed to the GCN module subsequently.

Table 4 shows that MMGCN with the graph-based multimodal fusion outperforms all other com-

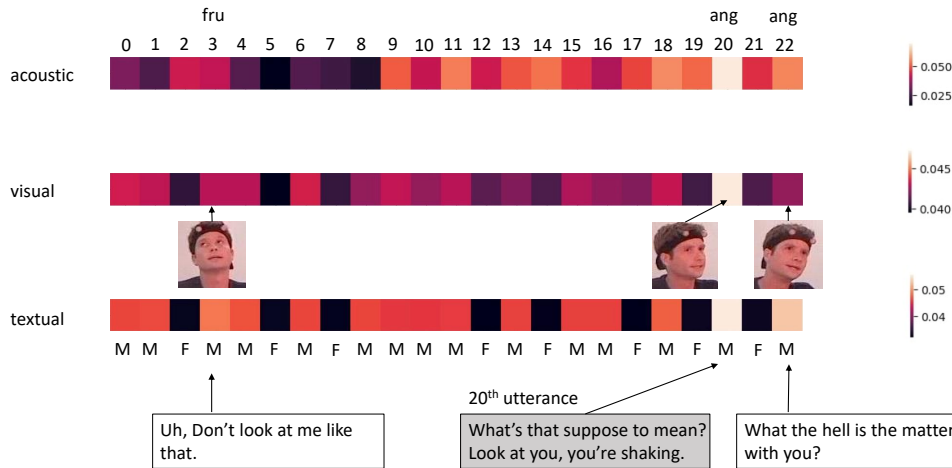


Figure 4: Visualization of the heatmap of the adjacent matrix for the 20<sup>th</sup> utterance in a conversation with three modalities. 'M' and 'F' refer to the male and female speakers respectively

	IEMOCAP	MELD
<i>DeepGCN<sub>early-fusion</sub></i>	64.46	57.94
<i>DeepGCN<sub>late-fusion</sub></i>	64.62	58.26
<i>DeepGCN<sub>gated-attention</sub></i>	64.45	58.18
<i>DeepGCN<sub>MFN</sub></i>	62.77	58.21
<i>DeepGCN<sub>MulT</sub></i>	62.37	57.93
<b>MMGCN</b>	<b>66.22</b>	<b>58.65</b>

Table 4: ERC performance comparison of MMGCN and other multimodal fusion methods

layers	IEMOCAP	MELD
1	66.12	58.40
2	66.17	58.38
4	<b>66.22</b>	<b>58.65</b>
8	66.10	58.54
16	66.06	58.38
32	66.10	58.42

Table 5: ERC performance comparison of MMGCN with different number of layers

pared multimodal fusion methods.

#### 5.4 MMGCN with different layers

We investigate the impact of the number of layers in MMGCN on the ERC performance in Table 5. The experiment results show that a different number of layers does affect the ERC recognition performance. Specifically, MMGCN achieves the best performance with 4 layers on both IEMOCAP and MELD.

	MMGCN	IEMOCAP	MELD
w/ spkr embedding		<b>66.22</b>	<b>58.65</b>
w/o spkr embedding		65.76	58.38

Table 6: Ablation study of the speaker embedding impact on ERC performance

#### 5.5 Impact of Speaker Embedding

Speaker Embedding can differentiate input features from different speakers. Previous works have reported that speaker information can help improve emotion recognition performance. We conduct the ablation study to verify the contribution of speaker embedding in MMGCN as shown in Table 6. As expected, dropping speaker embedding in MMGCN leads to performance degradation, which is significant by t-test with  $p < 0.05$ .

#### 5.6 Case Study

Fig 4 depicts a scene in which a man and a woman quarrel with each other over a female friend of the man who came to meet with him across 700 miles. They are frustrated or angry in most cases. At the beginning of the conversation, their emotion states are both neutral. Over time, they become emotional. They are both angry at the end of the conversation. The heatmaps of the adjacent matrix for the 20<sup>th</sup> utterance in the conversation from the three modalities demonstrate that different from simple sequential models, MMGCN pays attention not only to the close context, but also relate to the context in long-distance. For example, as shown



in the textual heatmap, MMGCN can successfully aggregate information from the most relevant utterances, even from long-distance utterances, for example the 3<sup>rd</sup> utterance.

## 6 Conclusion

In this paper, we propose a multimodal fused graph convolutional network (MMGCN) for multimodal emotion recognition in conversation (ERC). MMGCN provides a more effective way of utilizing both multimodal and long-distance contextual information. It constructs a graph that captures not only intra-speaker context dependency but also inter-modality dependency. With the residual connection, MMGCN can have deep layers to further improve recognition performance. We carry out experiments on two public benchmark datasets, IEMOCAP and MELD, and the experiment results prove the effectiveness of MMGCN, which outperforms other state-of-the-art methods by a significant margin under the multimodal conversation setting.

## 7 Acknowledgement

This work was supported by the National Key R&D Program of China under Grant No. 2020AAA0108600, National Natural Science Foundation of China (No. 62072462), National Natural Science Foundation of China (No. 61772535), and Beijing Natural Science Foundation (No. 4192028).

## References

- Emad Barsoum, Cha Zhang, Cristian Canton Ferrer, and Zhengyou Zhang. 2016. Training deep networks for facial expression recognition with crowd-sourced label distribution. New York, NY, USA. Association for Computing Machinery.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335.
- Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. 2020. Simple and deep graph convolutional networks. In *International Conference on Machine Learning*, pages 1725–1735. PMLR.
- Minghai Chen, Sen Wang, Paul Pu Liang, Tadas Baltrušaitis, Amir Zadeh, and Louis-Philippe Morency. 2017. Multimodal sentiment analysis with word-level fusion and reinforcement learning. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pages 163–171.
- Sheng-Yeh Chen, Chao-Chun Hsu, Chuan-Chun Kuo, Lun-Wei Ku, et al. 2018. Emotionlines: An emotion corpus of multi-party conversations. *arXiv preprint arXiv:1802.08379*.
- Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. 2019. Dialoguecn: A graph convolutional neural network for emotion recognition in conversation. *arXiv preprint arXiv:1908.11540*.
- Devamanyu Hazarika, Soujanya Poria, Rada Mihalcea, Erik Cambria, and Roger Zimmermann. 2018a. Icon: Interactive conversational memory network for multimodal emotion detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2594–2604.
- Devamanyu Hazarika, Soujanya Poria, Amir Zadeh, Erik Cambria, Louis-Philippe Morency, and Roger Zimmermann. 2018b. Conversational memory network for emotion recognition in dyadic dialogue videos. In *Proceedings of NAACL-HLT*, pages 2122–2132.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2015. Densely connected convolutional networks.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. 2019. Deepgcn: Can gcn go as deep as cnn? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9267–9276.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. 2015. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*.
- Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander Gelbukh, and Erik Cambria. 2019. Dialoguernn: An attentive rnn for emotion detection in conversations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6818–6825.
- Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In

- Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4602–4609.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. Context-dependent sentiment analysis in user-generated videos. In *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 873–883.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508*.
- Saurav Sahay, Shachi H Kumar, Rui Xia, Jonathan Huang, and Lama Nachman. 2018. Multimodal relational tensor network for sentiment and emotion classification. *arXiv preprint arXiv:1806.02923*.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.
- Björn Schuller, Anton Batliner, Stefan Steidl, and Dino Seppi. 2011. Recognising realistic emotions and affect in speech: State of the art and lessons learnt from the first challenge. *Speech Communication*, 53(9-10):1062–1087.
- Björn Schuller, Michel Valster, Florian Eyben, Roddy Cowie, and Maja Pantic. 2012. Avec 2012: the continuous audio/visual emotion challenge. In *Proceedings of the 14th ACM international conference on Multimodal interaction*, pages 449–456.
- Konstantinos Skianis, Fragkiskos Malliaros, and Michalis Vazirgiannis. 2018. Fusing document, collection and label graph-based representations with word embeddings for text classification. In *Proceedings of the Twelfth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-12)*, pages 49–58.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Memory fusion network for multi-view sequential learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Dong Zhang, Liangqing Wu, Changlong Sun, Shoushan Li, Qiaoming Zhu, and Guodong Zhou. 2019. Modeling both context-and speaker-sensitive dependence for emotion detection in multi-speaker conversations. In *IJCAI*, pages 5415–5421.