

Curriculum Learning Meets Directed Acyclic Graph for Multimodal Emotion Recognition

Cam-Van Thi Nguyen, Cao-Bach Nguyen, Quang-Thuy Ha, Duc-Trong Le

VNU University of Engineering and Technology, Hanoi, Vietnam

{vanntc, 19020218, thuyhq, trongld}@vnu.edu.vn

Abstract

Emotion recognition in conversation (ERC) is a crucial task in natural language processing and affective computing. This paper proposes MultiDAG+CL, a novel approach for Multimodal Emotion Recognition in Conversation (ERC) that employs Directed Acyclic Graph (DAG) to integrate textual, acoustic, and visual features within a unified framework. The model is enhanced by Curriculum Learning (CL) to address challenges related to emotional shifts and data imbalance. Curriculum learning facilitates the learning process by gradually presenting training samples in a meaningful order, thereby improving the model's performance in handling emotional variations and data imbalance. Experimental results on the IEMOCAP and MELD datasets demonstrate that the MultiDAG+CL models outperform baseline models. We release the code for MultiDAG+CL and experiments: <https://github.com/vanntc711/MultiDAG-CL>.

Keywords: Multimodal Emotion Recognition, Curriculum Learning, Directed Acyclic Graph

1. Introduction

Online social networks' growing popularity has sparked interest in capturing emotions in conversations. Emotion Recognition in Conversation (ERC) has emerged as a critical task in various domains such as chatbots (Ghosh et al., 2017), healthcare (Li et al., 2019), and social media analysis (Polzin and Waibel, 2000). In the field of ERC, researches can be broadly categorized into unimodal and multimodal approaches. Unimodal approaches usually focus on using text as the main modality for emotion recognition. Several models have been proposed in the past to tackle unimodal ERC task. DialogueRNN (Majumder et al., 2019) introduces a recurrent network to track speaker states and context during the conversation. DialogueGCN (Ghosal et al., 2019) utilizes graph structures to combine contextual dependencies.

Multimodal Emotion Recognition in Conversation (Multimodal ERC) classifies emotions in conversation turns using text, audio, and visual cues. By incorporating multiple modalities, it provides a comprehensive representation of emotional expressions, including tone of voice, facial expressions, and body language, resulting in improved accuracy and robustness in emotion recognition compared to traditional unimodal ERC approaches. Several models have been proposed to address the task of multimodal ERC. The MFN (Zadeh et al., 2018) synchronizes multimodal sequences using a multi-view gated memory. ICON (Hazarika et al., 2018) provides conversational features from modalities through multi-hop memories. The bc-LSTM (Poria et al., 2017) leverages an utterance-level LSTM to capture multimodal features. MMGCN (Hu et al., 2021) uses a graph-based fusion module to cap-

ture intra- and inter-modality contextual features. CTNet (Lian et al., 2021) utilizes a transformer-based structure to model interactions among multimodal features. CORECT (Nguyen et al., 2023) leverages relational temporal GNNs with cross-modality interaction support, effectively capturing conversation-level interactions and utterance-level temporal relations.

A Directed Acyclic Graph (DAG) is a directed graph without any directed cycles, comprising vertices and edges, where each edge is directed from one vertex to another, ensuring no closed loops. Building upon this concept, Yu et al. (2019) introduced Directed Acyclic Graph Neural Network (DAG-GNN). Additionally, Shen et al. (2021) presented DAG-ERC, a model combining graph-based and recurrence-based neural architectures to capture information flow in long-distance conversations. However, DAG-ERC's focus has been primarily on unimodal text data, with limited exploration in other modalities. Curriculum Learning (CL), inspired by human learning, progressively introduces more complex concepts starting from a simple initial state. It establishes a sequence of curricula where the best curriculum with the simplest examples is used to train the classifier in each learning round (Bengio et al., 2009; Soviany et al., 2022). CL incorporates two key factors: a *difficulty measurer* to assess the difficulty level of training examples, and a *training scheduler* to determine the order of example presentation during training. The difficulty measurer assesses the difficulty level of training examples, while the training scheduler determines the order in which examples are presented to the model during training. For the ERC task, Yang et al. (2022) proposes a hybrid CL framework specifically for the textual modality only.

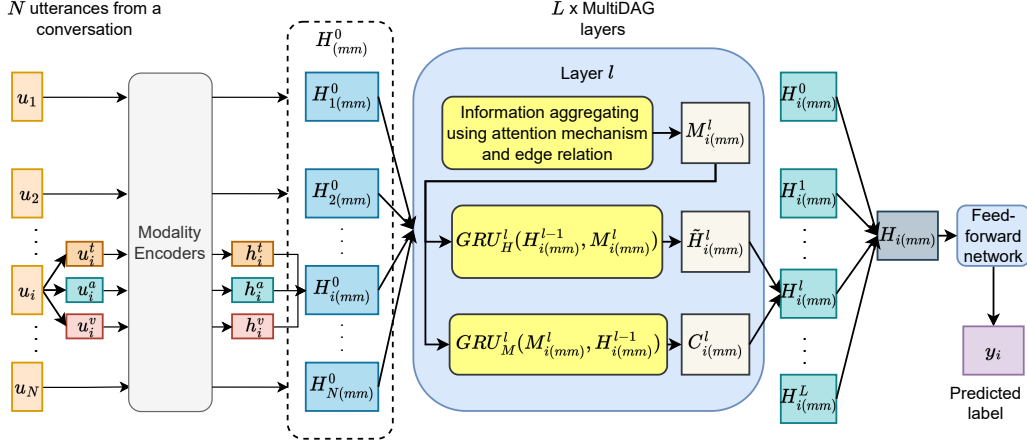


Figure 1: Overall structure of MultiDAG.

In this paper, we propose **MultiDAG+CL**, a multimodal model inspired by DAG-ERC (Shen et al., 2021), designed to overcome the limitations of text-based approaches. It integrates multimodal features using DAG-GNN, enabling a comprehensive understanding of emotions in conversations. Leveraging Curriculum Learning, our model, **MultiDAG+CL**, addresses emotional shift issues and imbalanced data, significantly enhancing ERC model performance on IEMOCAP and MELD datasets. Notably, we are the first to integrate multimodal ERC models with Curriculum Learning strategy.

2. Methodology

Consider a conversation C having utterances $\{u_1, u_2, \dots, u_N\}$ where N is the number of utterances. An utterance is a coherent piece of information conveyed by a single participant p_m at a specific moment, where $m \geq 2$. The task of Emotion Recognition in Conversation (ERC) is to predict emotion label of each utterance u_i with predefined emotion label set $E = \{y_1, y_2, \dots, y_r\}$. Following the multimodal approach, we represent an utterance in terms of three different modalities: audio (a), visual (v), and textual (l). The raw feature representation of utterance u_i is $u_i = \{u_i^a, u_i^v, u_i^l\}$.

MultiDAG+CL consists of two core components: *MultiDAG* and *Curriculum Learning-CL*. The *MultiDAG* component represents the model that combines multimodal features without CL integration. The *-CL* component is where Curriculum Learning is incorporated to enhance model performance.

2.1. Multimodal ERC with Directed Acyclic Graph - MultiDAG

2.1.1. Modality Encoder

We use modality-specific encoders to generate context-aware utterance feature encoding. For the

textual modality, a bidirectional LSTM network captures sequential textual context information, while a Fully Connected Network is used for the acoustic and visual modalities as follows:

$$h_i^a = Enc_A(u_i^a); h_i^v = Enc_V(u_i^v); h_i^l = Enc_L(u_i^l) \quad (1)$$

where Enc_A, Enc_V, Enc_L are modality encoder for audio, visual, textual modalities, respectively. These encoders generate the context-aware raw feature encodings h_i^a, h_i^v, h_i^l accordingly. The multimodal feature vector for an utterance $u_{i(mm)}$ corresponding to available modalities is:

$$H_{i(mm)}^0 = h_i^a \oplus h_i^v \oplus h_i^l \quad (2)$$

2.1.2. MultiDAG Construction

Each utterance in a conversation receives information exclusively from past utterances. This one-way information flow is effectively represented by a Directed Acyclic Graph (DAG), where information moves from predecessors to successors. This characteristic allows the DAG to gather information for a query utterance not only from neighboring utterances but also from more distant ones. Following the multimodal representation input, we initialize the Directed Acyclic Graph Gated Neural Network (DAG-GNN) (Yu et al., 2019). The integration of both remote and local information is executed in a manner analogous to the approach undertaken in DAG-ERC by Shen et al. (2021). The comprehensive architecture of MultiDAG is visually represented in Figure 1.

At each layer l of the MultiDAG, the hidden state of the utterances is continuously computed from the first utterance to the last utterance. For each utterance $u_{i(mm)}$, the attention weight between $u_{i(mm)}$ and the preceding nodes is calculated by using the hidden state of $u_{i(mm)}$ at layer $l-1$ to attend to the hidden states of the nodes at layer l :

$$a_{ij}^l = S_{j \in \mathcal{N}_{i(mm)}}(W_\alpha^l [H_j^l \| H_{i(mm)}^{l-1}]) \quad (3)$$

Here, S denotes the Softmax function; $N_{i(mm)}$ represents the set of preceding nodes leading to u_i , W_a^l is a trainable weight matrix, $H_{i(mm)}^{l-1}$ is the hidden state of $u_{i(mm)}$ at layer $l-1$, and \parallel denotes concatenation.

The attention weight is further utilized in combination with edge relationships to aggregate information.

$$M_{i(mm)}^l = \sum_{j \in N_{i(mm)}} a_{ij(mm)} W_{r_{ij}}^l H_{j(mm)}^l \quad (4)$$

where $W_{r_{ij}}^l \in \{W_0^l, W_1^l\}$ are trainable parameters. The 0/1 value represents the edge relationship, distinguishing different or same speakers.

The aggregated information $M_{i(mm)}^l$ interacts with the previous layer's hidden state of $u_{i(mm)}$, $H_{i(mm)}^{l-1}$, through a GRU to generate the final hidden state $\tilde{H}_{i(mm)}^l$ at the current layer:

$$\tilde{H}_{i(mm)}^l = GRU_H^l(H_{i(mm)}^{l-1}, M_{i(mm)}^l) \quad (5)$$

where $H_{i(mm)}^{l-1}$, $M_{i(mm)}^l$, and $\tilde{H}_{i(mm)}^l$ represent the input, hidden state, and output of the GRU network, respectively. This step is the node information unit. Another GRU serves as the context information unit, modeling the flow of information from the historical context through a layer. In this unit, the roles of $H_{i(mm)}^{l-1}$ and $M_{i(mm)}^l$ in the GRU are exchanged, where $H_{i(mm)}^{l-1}$ controls the propagation of $M_{i(mm)}^l$:

$$C_{i(mm)}^l = GRU_M^l(M_{i(mm)}^l, H_{i(mm)}^{l-1}) \quad (6)$$

The hidden states of u_i from all layers are concatenated together to create final representation:

$$H_{i(mm)} = \parallel_{l=0}^L (\tilde{H}_{i(mm)}^l + C_{i(mm)}^l) \quad (7)$$

This representation is then passed through a Feed-Forward Network to perform emotion prediction. The objective function used to train the model is the cross-entropy loss function.

2.2. Curriculum Learning - CL

We design a **Difficulty Measure Function (DMF)** based on the frequency of emotional shift in conversations, and simultaneously construct a **Training Scheduler** to implement the training process according to the predefined learning curriculum.

2.2.1. Difficulty Measure Function (DMF)

When designing the difficulty measurement function for a conversation, it is essential to determine what makes a conversation easier or more difficult than others. Taking inspiration from [Yang et al. \(2022\)](#), we constructed a function to calculate the difficulty of a conversation based on the frequency

Algorithm 1 CL Training with DMF

Input: \mathcal{D} - training dataset; M - training model
 k - number of buckets in baby step scheduler
 DLF - difficulty measure function
 t - number of epochs; n - number of utterances
 e - the emotion label of the utterances
 $p(u_i)$ - the speaker's corresponding utterance u_i

S - Set containing the emotion sequence of speakers; $S[p][i]$ represents the emotion in the i -th utterance of speaker p

Output: M^* - the optimal model

$S = \emptyset, N_{es} = 0$

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

$S[p][i] \leftarrow S[p][i] \cup \{e[i]\}$

end for

$N_{sp} = \text{length}(S)$

for $p \in S$ **do**

for $i = 1$ to $\text{length}(S[p]) - 1$ **do**

if $S[p][i] \neq S[p][i+1]$ **then**

$N_{shift} \leftarrow N_{shift} + 1$

end if

end for

end for

$DLF = \frac{N_{shift} + N_{sp}}{n + N_{sp}}$

$\mathcal{D}' = \text{sort}(\mathcal{D}, DLF)$

$\mathcal{D}' = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^k\}$ where $DLF(d_a) < DLF(d_b), d_a \in \mathcal{D}^i, d_b \in \mathcal{D}^j, \forall i < j$

$\mathcal{D}^{train} = \emptyset$

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

if $i \leq k$ **then**

$\mathcal{D}^{train} = \mathcal{D}^{train} \cup \mathcal{D}^i$

end if

TRAIN(M, \mathcal{D}^{train})

end for

return M^*

of emotional shift. Here, an emotional shift is defined as occurring when the emotion expressed in two consecutive utterances by the same speaker is different. Specifically, $e(u_i) \neq e(u_k), p(u_i) = p(u_k), \nexists j : i < j < k, p(u_i) = p(u_j) = p(u_k)$. Here, $e(u_i)$ and $e(u_k)$ is the emotions of two consecutive utterances u_i and u_k , respectively. The more frequent the emotional shift occur in a conversation, the more it is considered difficult. Therefore, the difficulty of i -th conversation c_i is as follows:

$$DLF(c_i) = \frac{N_{shift}(c_i) + N_{sp}(c_i)}{N_u(c_i) + N_{sp}(c_i)} \quad (8)$$

where $N_{shift}(c_i)$ and $N_u(c_i)$ represent the number of emotional shift in conversation c_i and the total number of utterances in c_i , respectively. $N_{sp}(c_i)$ is the number of speakers appearing in conversation c_i and acts as a smoothing factor. The algorithm for calculating the difficulty of the conversation is fully described in the Algorithm 1.

Model	IEMOCAP							MELD			
	Happy	Sad	Neutral	Angry	Excited	Frustrated	Acc. (%)	w-F1 (%)	Acc. (%)	w-F1 (%)	
bc-LSTM (Poria et al., 2017)	33.82	78.76	56.75	64.35	60.25	60.75	60.51	60.42	59.62	57.29	
MFN (Zadeh et al., 2018)	48.19	73.41	56.28	63.04	64.11	61.82	61.24	61.60	60.80	57.80	
ICON (Hazarika et al., 2018)	32.80	74.40	60.60	68.20	68.40	66.20	64.00	63.50	58.20	56.30	
DialogueRNN (Majumder et al., 2019)	32.20	80.26	57.89	62.82	73.87	59.76	63.52	62.89	60.31	57.66	
DialogueGCN (Ghosal et al., 2019)	51.57	80.48	57.69	53.95	72.81	57.33	63.22	62.89	58.62	56.36	
DAG-ERC (Shen et al., 2021)	47.59	79.83	69.36	66.67	66.79	68.66	67.53	68.03	61.04	63.66	
MMGCN (Hu et al., 2021)	45.14	77.16	64.36	<u>68.82</u>	<u>74.71</u>	61.40	66.36	66.26	60.42	58.31	
CTNet (Lian et al., 2021)	51.3	79.9	65.8	67.2	78.7	58.8	68.0	67.5	62.0	60.5	
DAG-ERC+HCL (Yang et al., 2022)	-	-	-	-	-	-	-	<u>68.73</u>	-	<u>63.89</u>	
COGMEN (Joshi et al., 2022)	-	-	-	-	-	-	68.2	67.6	-	-	
MultiDAG (Ours)	<u>49.65</u>	79.83	66.40	67.59	71.78	<u>67.90</u>	<u>68.30</u>	68.45	<u>64.29</u>	63.87	
MultiDAG+CL (Ours)	45.26	81.40	69.53	70.33	71.61	66.94	69.11	69.08	64.41	64.00	

Table 1: Performance of approaches on IEMOCAP and MELD datasets. **Bold** indicates the highest performance, and underlining denotes the second-highest. “-” represents missing values due to unavailability in original papers.

2.2.2. Training Scheduler

The training scheduler is used to organize and schedule the training process by arranging conversations. Specifically, the dataset \mathcal{D} is divided into multiple different bins $\{\mathcal{D}_1, \dots, \mathcal{D}_k\}$, where conversation with similar difficulty are grouped into the same bin. The training process starts with the easiest bin. After training for a certain number of epochs, the next bin is mixed into the current training dataset. Finally, once all bins have been mixed and used, additional epochs of training are performed.

3. Experimental Setup

3.1. Datasets and Baselines

We evaluate our approach on the following two ERC datasets: IEMOCAP (Busso et al., 2008) and MELD (Poria et al., 2019). The detailed statistics of the datasets are reported in Table 2. For the data processing, we use the same split as the work in (Hu et al., 2021). We compare our method against several state-of-the-art baselines, including unimodal and multimodal learning approaches. (Due to the space limit, they are brief described in Section 1). The evaluation metrics used are Accuracy (Acc.) and weighted average F1-score (w-F1).

Datasets	Conversation			Utterances			Avg. utterances
	Train	Valid	Test	Train	Valid	Test	
IEMOCAP	120		31	5810	1623		66.8
MELD	1038	114	280	9989	1109	2610	9.57

Table 2: Dataset statistics

3.2. Implementation Details

We perform hyperparameter tuning for our proposed model on each dataset using hold-out validation with separate validation sets. For the IEMOCAP dataset, the hyperparameter configuration

includes a learning rate of 0.0005, a dropout rate of 0.4, 30 epochs of training, and 4 layers of MultiDAG+CL. For the MELD dataset, the hyperparameter configuration for the MultiDAG+CL model is as follows: a learning rate of 0.00001, a dropout rate of 0.1, 60 epochs of training, and 2 layers of Multi-DAG.

4. Results and Analysis

4.1. Comparison with Baselines

We conducted a comprehensive comparison of our proposed approach with SOTA multimodal ERC methods, and the results are summarized in Table 1. Due to space constraints, we only report Acc. and w-F1 for the MELD dataset. Our approach, *MultiDAG+CL*, which combines the *MultiDAG* model with a curriculum learning strategy, achieves SOTA performance on both the IEMOCAP and MELD datasets. *MultiDAG+CL* outperforms previous SOTAs by 1.05% (DAG-ERC on IEMOCAP) and 0.34% (DAG-ERC on MELD), respectively. Specifically, our models achieve improvements in individual emotion recognition tasks in most cases, especially for the *Sad*, *Neutral* and *Angry* emotions. In the meantime, we find *Happy*, *Sad*, and *Angry* emotions can be confused with the *Neutral* emotion in some cases (as shown in Fig. 2). Such phenomenon is related to imbalanced class distribution.

4.2. Effect of Modality

Table 3 compares the performance of MultiDAG and MultiDAG+CL under various multimodal settings on both benchmark datasets. In IEMOCAP, the textual modality performs best among the unimodal settings, while the visual modality shows the lowest results due to noise from factors like camera position and environmental conditions. In bimodal settings, the combination of textual and acoustic modalities performs the best, while the combina-

tion of visual and acoustic modalities yields the lowest result. Similar observations are made in the MELD dataset.

Modality	MultiDAG		MultiDAG+CL	
	IEMOCAP	MELD	IEMOCAP	MELD
T	68.17	63.66	67.12	63.47
A	49.37	40.27	50.58	40.17
V	33.79	31.27	36.69	31.27
T + A	68.42	63.61	68.45	63.56
T + V	67.56	63.69	67.40	63.62
A + V	52.40	40.54	51.86	39.99
T + V + A	68.45	63.87	69.08	64.00

Table 3: Results of MultiDAG and MultiDAG+CL under different modality settings. T, A, V represent the text, audio, visual modality, respectively.

4.3. Effect of Curriculum Learning

The MultiDAG+CL model demonstrates notable performance improvement by incorporating curriculum learning for both the IEMOCAP and MELD datasets. The effectiveness of curriculum learning relies on factors like the difficulty measure design and training strategy, including the number of buckets in the training set. We perform experiments to select the optimal number of buckets in the CL training scheduler. The results shown in the Table 4, indicate that for the IEMOCAP dataset, the optimal number of buckets is 5, while for the MELD dataset, it is 12. These findings suggest that the CL strategy is effective in improving the performance of the MultiDAG model on both datasets, with the specific number of buckets tailored to each dataset’s representations. In summary, our proposed MultiDAG+CL model with curriculum learning, significantly contribute to the achieved results.


IEMOCAP		MELD	
Number of buckets	w-F1	Number of buckets	w-F1
4	68.05	5	63.94
5	69.08	8	63.83
7	68.84	10	63.89
10	68.38	12	64.00
15	68.36	14	63.96

Table 4: Results of MultiDAG+CL for different number of buckets in CL training scheduler.

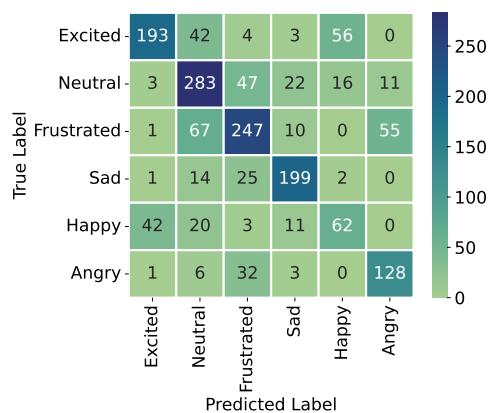
4.4. Performance for Emotion-shift

From the confusion matrices of the MultiDAG and MultiDAG+CL models (Figure 2), it can be observed that the prediction accuracy for the “Happy”, “Neutral”, “Sad”, and “Angry” labels is improved when CL is incorporated into the model. Particularly, the misclassification rate of the “Neutral” label as “Disgust” decreases significantly from 19.3% in the MultiDAG model to only 12.3% in MultiDAG+CL.

However, the prediction accuracy for the “Disgust” and “Happy” labels decreases.



(a) MultiDAG



(b) MultiDAG+CL

Figure 2: The confusion matrices on the IEMOCAP.

5. Conclusion

This paper proposes MultiDAG+CL, a novel approach for Multimodal Emotion Recognition in Conversation that leverages Directed Acyclic Graphs to integrate textual, acoustic, and visual features within a unified framework. The incorporation of Curriculum Learning (CL) addresses challenges related to emotional shifts and data imbalance, enhancing the model’s performance. Through extensive experiments, we evaluate the performance of both the proposed MultiDAG+CL models. Future work includes exploring alternative training schedulers for Curriculum Learning and incorporating a learning-based approach to model emotion label similarity.

Acknowledgements

Cam-Van Thi Nguyen was funded by the Master, PhD Scholarship Programme of Vingroup Innovation Foundation (VINIF), code VINIF.2023.TS147.

References

- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- 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):335–359.
- Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. 2019. DialogueGCN: A graph convolutional neural network for emotion recognition in conversation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 154–164, Hong Kong, China. Association for Computational Linguistics.
- Sayan Ghosh, Mathieu Chollet, Eugene Laksana, Louis-Philippe Morency, and Stefan Scherer. 2017. Affect-LM: A neural language model for customizable affective text generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 634–642, Vancouver, Canada. Association for Computational Linguistics.
- Devamanyu Hazarika, Soujanya Poria, Rada Mihalcea, Erik Cambria, and Roger Zimmermann. 2018. 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, Brussels, Belgium. Association for Computational Linguistics.
- Jingwen Hu, Yuchen Liu, Jinming Zhao, and Qin Jin. 2021. MMGCN: Multimodal fusion via deep graph convolution network for emotion recognition in conversation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5666–5675, Online. Association for Computational Linguistics.
- Abhinav Joshi, Ashwani Bhat, Ayush Jain, Atin Singh, and Ashutosh Modi. 2022. COGMEN: COntextualized GNN based multimodal emotion recognition. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4148–4164, Seattle, United States. Association for Computational Linguistics.
- Runnan Li, Zhiyong Wu, Jia Jia, Yaohua Bu, Sheng Zhao, and Helen Meng. 2019. Towards discriminative representation learning for speech emotion recognition. In *IJCAI*, pages 5060–5066.
- Zheng Lian, Bin Liu, and Jianhua Tao. 2021. Ctnet: Conversational transformer network for emotion recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:985–1000.
- 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.
- Cam Van Thi Nguyen, Tuan Mai, Son The, Dang Kieu, and Duc-Trong Le. 2023. Conversation understanding using relational temporal graph neural networks with auxiliary cross-modality interaction. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15154–15167, Singapore. Association for Computational Linguistics.
- Thomas S Polzin and Alexander Waibel. 2000. Emotion-sensitive human-computer interfaces. In *ISCA tutorial and research workshop (ITRW) on speech and emotion*.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. Context-dependent sentiment analysis in user-generated videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 873–883, Vancouver, Canada. Association for Computational Linguistics.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. MELD: A multimodal multi-party dataset for emotion recognition in conversations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536, Florence, Italy. Association for Computational Linguistics.
- Weizhou Shen, Siyue Wu, Yunyi Yang, and Xiaojun Quan. 2021. Directed acyclic graph network

- for conversational emotion recognition. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1551–1560, Online. Association for Computational Linguistics.
- Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. 2022. Curriculum learning: A survey. *International Journal of Computer Vision*, 130(6):1526–1565.
- Lin Yang, Yi Shen, Yue Mao, and Longjun Cai. 2022. Hybrid curriculum learning for emotion recognition in conversation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11595–11603.
- Yue Yu, Jie Chen, Tian Gao, and Mo Yu. 2019. DAG-GNN: DAG structure learning with graph neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7154–7163. PMLR.
- 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.