

# Dynamic fusion-aware graph convolutional neural network for multimodal emotion recognition in conversations

Tao Meng <sup>a</sup>, Weilun Tang <sup>a</sup>, Yuntao Shou <sup>a</sup>, Yilong Tan <sup>a</sup>, Jun Zhou <sup>a</sup>, Wei Ai <sup>a,\*</sup>, Keqin Li <sup>b</sup>

<sup>a</sup> College of Computer and Mathematics, Central South University of Forestry and Technology, 410004, Hunan, Changsha, China

<sup>b</sup> Department of Computer Science, State University of New York, New Paltz, New York, 12561, USA

## ARTICLE INFO

### Keywords:

Multimodal emotion recognition  
Dynamic fusion  
Graph convolutional neural network  
Prompt learning

## ABSTRACT

Multimodal emotion recognition in conversations (MERC) aims to identify and understand the emotions expressed by speakers during utterance interaction from multiple modalities (e.g., text, audio, images, etc.). Existing studies have shown that GCN can improve the performance of MERC by modeling dependencies between speakers. However, existing methods usually use fixed parameters to process multimodal features for different emotion types, ignoring the dynamics of fusion between different modalities, which forces the model to balance performance between multiple emotion categories, thus limiting the model's performance on some specific emotions. To this end, we propose a dynamic fusion-aware graph convolutional neural network (DF-GCN) for robust recognition of multimodal emotion features in conversations. Specifically, DF-GCN integrates ordinary differential equations into graph convolutional networks (GCNs) to capture the dynamic nature of emotional dependencies within utterance interaction networks and leverages the prompts generated by the global information vector (GIV) of the utterance to guide the dynamic fusion of multimodal features. This allows our model to dynamically change parameters when processing each utterance feature, so that different network parameters can be equipped for different emotion categories in the inference stage, thereby achieving more flexible emotion classification and enhancing the generalization ability of the model. Comprehensive experiments conducted on two public multimodal conversational datasets confirm that the proposed DF-GCN model delivers superior performance, benefiting significantly from the dynamic fusion mechanism introduced. To the best of our knowledge, this is the first framework that adaptively assigns different fusion weights to different emotion categories during inference, leading to more effective multimodal information integration. Our code is available at <https://github.com/yuntaoshou/DFGCN>.

## 1. Introduction

Multimodal emotion recognition in conversations (MERC) aims to exploit multimodal information (e.g., text, audio, video, etc.) in the dialogues to understand and identify the emotional state of the speaker [1–3]. MERC has become an important research topic that has attracted much attention due to its wide application in opinion mining [4,5], healthcare [6], psychological counseling [7], and building empathetic dialogue systems. MERC research not only helps us to deeply understand the emotional changes of users in multi-round dialogues but also improves the intelligence of dialogue systems, making them more capable of emotional perception and response. Unlike traditional non-conversational or unimodal emotion recognition (ER) methods, MERC focuses more on the fusion of semantic features from conversational context and multiple modalities [8]. The core challenge is how to ef-

fectively model and fuse the emotional dependencies between the context and multimodality in the conversation [9]. Specifically, MERC must fully capture the subtle changes in emotions in the conversation, understand the interaction patterns between multimodality, and reasonably fuse features to achieve more accurate emotion recognition [10].

The mainstream MERC methods mainly use a Transformer or GCN to model the emotional dependencies, and fuse the semantic features in the conversation [11–13]. For example, the Transformer-based self-distillation model (SDT) [1] adaptively learns weights related to different modalities through self-distillation technology and a hierarchical gated fusion strategy to perform feature fusion better. TBJE [14] leverages the Transformer architecture and a modular collaborative attention mechanism to effectively capture and integrate the key factors of emotion and sentiment expressed through multimodal language. DER-GCN [15] constructs multiple emotional relationship graphs through

\* Corresponding author.

E-mail address: [aiwei@hnu.edu.cn](mailto:aiwei@hnu.edu.cn) (W. Ai).

<https://doi.org/10.1016/j.knosys.2026.115846>

Received 16 September 2025; Received in revised form 7 March 2026; Accepted 20 March 2026

Available online 7 April 2026

0950-7051/© 2026 Elsevier B.V. All rights reserved, including those for text and data mining, AI training, and similar technologies.

events to capture the emotional dependencies between multimodalities in dialogue, and uses self-supervised learning and multi-information Transformer to optimize feature fusion. M3Net [16] leverages multi-frequency information within a graph to enhance the representation of multivariate relationships and capture complex relationships across modalities and context.

The above methods have achieved certain results in the MERC task through feature fusion. However, fixed parameters are often used in the inference stage to fuse the multimodal features across different emotion types. As shown in Fig. 1, although this unified parameter configuration simplifies the design of the model, the performance of the model on minority samples will drop significantly when faced with diverse emotion categories. In addition, due to the lack of specialized optimization for different emotion categories, the model has to balance performance between different categories, making it difficult to fully capture and represent the unique characteristics of each emotion. This not only affects the flexibility of the model but also limits its accuracy and sensitivity in certain emotion categories, making it difficult to achieve detailed distinction and recognition of complex emotions.

Therefore, MERC needs to consider the dynamics of semantic dependencies in the conversation context and adaptively equip different network parameters for the fusion of utterance features of different emotion categories as the starting point of model design. A key aspect of capturing this context is modeling the sequential relationships between utterances. While pre-trained language models like RoBERTa provide powerful utterance-level semantics, they lack built-in mechanisms for inter-utterance context. To this end, we incorporate a Bidirectional GRU (Bi-GRU) to efficiently capture these contextual dependencies in the conversation flow. Inspired by the above analysis, we propose a DF-GCN for robust recognition of multimodal emotion features in conversations. The core of DF-GCN is the design of two serial graph convolution blocks, namely, the Static Graph Convolution (SGCODE) block and the Dynamic Graph Convolution (DGCODE) with ordinary differential equations block. Specifically, we first construct an emotional interaction graph based on multimodal utterance features and input it into SGCODE to capture the dynamic emotional dependencies between multimodal utterances. Then, the multimodal utterance features with emotional dependencies are input into the Transformer to obtain the GIV, and the multi-layer perceptron is used to generate prompts to obtain the dynamic weight of the utterance. Finally, the dynamic weight and multimodal utterance features are input into DGCODE for dynamic fusion to complete the emotional recognition of the utterance. It is worth noting that in the inference stage of the model, the parameters of the DGCODE we constructed are dynamically changed so that different network parameters can be equipped for utterances of different emotion categories in the inference stage, to achieve more flexible emotion classification and enhance the generalization ability of the model. The main contributions of this paper are summarized as follows:

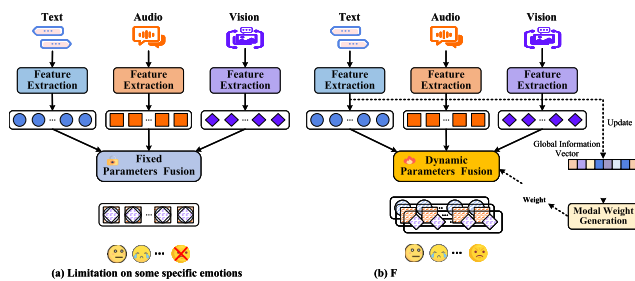


Fig. 1. Illustration of the difference between traditional static fusion and dynamic fusion processes. (a) Previous studies usually use fixed network parameters to fuse multimodal features. (b) Our proposed method generates modal parameters through GIV, thereby achieving dynamic multimodal feature fusion.

- We propose a DF-GCN for MERC. DF-GCN can adaptively equip utterances of different emotion categories with different network parameters to achieve more flexible emotion classification.
- We construct a novel prompt generation network to obtain GIV and dynamic fusion weights with rich semantics to guide the network to dynamically fuse utterance features with different emotion categories effectively.
- Extensive experiments were conducted on the IEMOCAP and MELD datasets, and the results verified the effectiveness of DF-GCN, especially in terms of WA and WF1, which significantly surpassed existing mainstream methods.

## 2. Related work

### 2.1. Multimodal emotion recognition

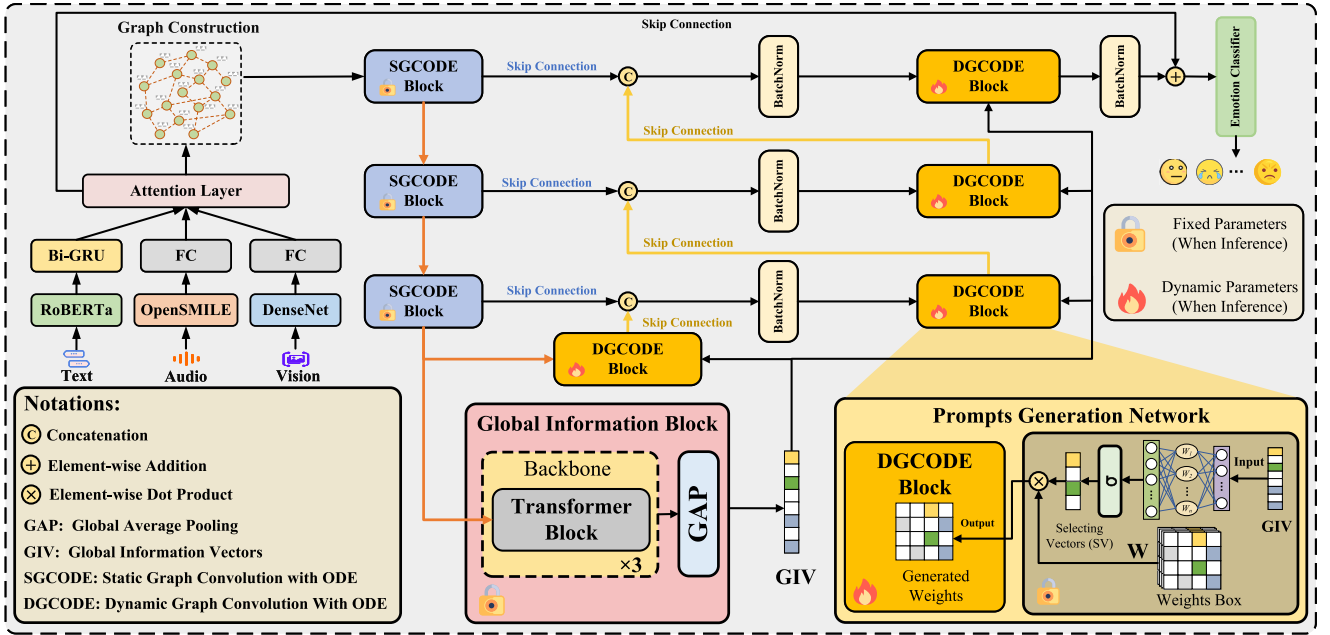
Multimodal emotion recognition in conversation (MERC) requires processing and analyzing speech, facial expressions, and text data to discern and understand human emotional states [17],[18]. In this paper, we roughly classify MERC methods into three categories: recurrent neural networks (RNNs), Transformers, and GCN.

RNN-based MERC methods mainly extract contextual semantic information from multimodal features through recurrent memory units. For example, DialogueRNN [19] considers emotions in conversation by constructing GRUs. The CNN-RNN-based methods [20] combine multiple CNN features with RNN subnetworks through pre-training and feature fusion to achieve dimensional emotion recognition.

Transformer-based MERC methods effectively extract and integrate emotional features from different modalities through self-attention mechanisms and multimodal fusion strategies. For example, the Transformer-based self-distillation model (SDT) [1] aims to apply self-distillation technology and hierarchical gated fusion strategy to adaptively learn weights related to different modalities to effectively capture intra-modal and inter-modal interactions in dialogue data. The main modality transformer (MMTr) [21] learns information interaction between modalities through the multi-head attention mechanism to enhance the representation of weak modalities. HTNet [22] improves the accuracy of micro-expression recognition by fusing CNN and Transformer to achieve joint spatial-temporal modeling of micro-expressions.

GCN-based MERC methods mainly model the interaction relationships among multiple modalities in the dialogue through graph structure and captures and analyzes the complex dependencies between these modalities through the graph convolutional network. For example, DialogueGCN [23] captures the contextual information of the dialogue by modeling the mutual dependencies between each discourse in the dialogue. C-GCN [24] constructs association graphs and uses intra-class and inter-class correlations between videos for feature learning and information fusion. DER-GCN [15] captures the dependencies between speakers and events in a conversation by constructing multiple relationship graphs and uses self-supervised learning and multi-information transformers to optimize feature fusion and loss functions. By introducing context-aware graphs and knowledge graphs into graph convolutional networks [25], multimodal feature representations are enhanced, enabling more accurate emotion recognition.

Although these methods have achieved good performance on MERC, they often rely on fixed parameters to fuse multimodal features across different emotion types during the inference stage. Specifically, due to the limitation of fixed parameters, it is difficult for the model to adaptively adjust according to the specific characteristics of each emotion type, resulting in the inability to fully explore and express the unique characteristics of each emotion. This limitation prevents models from accurately capturing the subtle differences and nuances among emotions when dealing with complex emotional situations, thus affecting the accuracy of emotion recognition and the generalization ability of the model.



**Fig. 2.** The overall architecture of the proposed DF-GCN model. The framework consists of three main components: (i) Graph Construction, where multimodal utterances from text (RoBERTa), audio (OpenSMILE), and vision (DenseNet) are encoded and integrated through an attention layer and Bi-GRU to form discourse-level interaction graphs; (ii) Global Information Block, which employs stacked Transformer layers and global average pooling to generate GIV that guide context-aware fusion; and (iii) Prompts Generation Network, where DGCODE blocks dynamically generate weight prompts for adaptive multimodal feature integration. SGCODE captures structural dependencies, while DGCODE models temporal and contextual dynamics. Skip connections and batch normalization enhance training stability, and the final fused representations are used for emotion classification.

## 2.2. Prompt learning

In natural language processing (NLP), prompt-based techniques are effective ways to provide contextual information to models for fine-tuning on target tasks [26]. This approach provides task-specific information to models through prompts, enabling them to adjust their output behavior according to the context [27]. However, unlike traditional manually designed instruction sets, learnable prompts can automatically adapt to task requirements through training and adjust model parameters more efficiently [28]. Compared with manually written static instructions, learnable prompts not only have advantages in adaptability but also can more fully exploit the potential of the model performance and computational cost. As prompt learning techniques can simulate task-specific contextual information, they have not only made significant progress in NLP, but have also been successfully applied to other fields (e.g., fine-tuning to visual tasks [29–31], and incremental learning [32,33]. In these applications, prompts act as a bridge between pre-trained models and downstream tasks, enabling models to respond more flexibly to changing task requirements. In addition, prompt learning techniques have also shown great potential in multi-task learning scenarios [34]. The correct choice of prompts for different tasks is crucial because the contextual information required for each task may be different, which has a direct impact on the performance of the model. For example, the LLM (large language model)-based MERC method, such as leveraging visual and textual prompts to guide emotion understanding within a large visual language model [35], improves multimodal emotion recognition performance. Visual Prompting enhances cross-modal emotion recognition capabilities within large language models [36] and improves recognition performance in complex scenarios.

## 2.3. Continuous graph neural networks

Neural Ordinary Differential Equations (ODEs) represent a recent advance in modeling continuous-time dynamic systems [37]. Instead of

relying on the discrete stacking of hidden layers as in conventional neural networks, neural ODEs parameterize the derivatives of hidden states with a neural function, enabling continuous inference along the temporal dimension. This paradigm provides a more precise characterization of temporal evolution and is particularly well-suited for tasks that involve complex time-varying behaviors. Building upon this idea, Continuous Graph Neural Networks (CGNNs) [38] extend the ODE framework to graph-structured data. CGNNs introduce a continuous message-passing mechanism, where the evolution of node representations is governed by an ordinary differential equation rather than a fixed number of propagation layers. This allows information to flow across nodes in a continuous manner, alleviating the constraints imposed by discrete architectures. To further address the problem of over-smoothing commonly observed in graph neural networks, CGNN incorporates a restart distribution mechanism, which periodically reset node representations to initial states during propagation. This design preserves informative distinctions among nodes while maintaining the benefits of continuous-time modeling.

## 3. Preliminaries

### 3.1. Graph neural networks

Consider a graph  $G = (V, E)$ , where  $V$  denotes the set of nodes and  $E$  represents the set of edges. The goal is to derive meaningful representations for each node within the graph. Nodes in  $V$  are represented by a set of features stored in the node feature matrix  $X \in \mathbb{R}^{|V| \times d}$ , where  $d$  represents the dimension of the feature vector. Each row of  $X$  corresponds to the feature representation of node  $v \in V$ . To model the connectivity between nodes in the graph, we utilize the binary adjacency matrix  $A \in \mathbb{R}^{|V| \times |V|}$ , where each element  $a_{ij}$  represents whether there is an edge between node  $i$  and node  $j$ . Specifically,  $a_{ij} = 1$  indicates that there is an edge between nodes  $i$  and  $j$ , while  $a_{ij} = 0$  indicates that there is no edge between the two nodes. The adjacency matrix  $A$  represents

the dependency relationship between nodes. The main goal of the graph learning task is to learn a node representation matrix  $H$  of size  $|V| \times k$ , where  $k$  is the dimension of the learned node embedding.

### 3.2. Neural ordinary differential equation

ODEs offer an innovative approach to continuous-time dynamic modeling, where the forward propagation process in neural networks is framed as a solution to differential equations. Unlike traditional neural network architectures, neural differential equations can simulate more complex and continuous-time dynamic behaviors by directly describing the evolution of input data as a dynamic system of differential equations. Specifically, consider an input data  $x(t)$ , whose differential equation evolving over time  $t$  is defined as follows:

$$\frac{d}{dt}x(t) = f(x(t), t, \theta) \quad (1)$$

where  $x(t)$  is the state at time  $t$ , and  $f(x(t), t, \theta)$  is a function that describes the evolution of the system, usually approximated by a neural network.

## 4. Methodology

The proposed DF-GCN contains five modules: multimodal feature encoding, static graph convolution, global information vector generation, dynamic graph convolution, and emotion classification. The detailed architecture of the proposed DF-GCN is shown in Fig. 2.

### 4.1. Task definition

In MERC, the number of speakers  $M$  ( $M \geq 2$ ) is not fixed. We assume that  $M$  speakers participate in the conversation, represented as  $S = \{S_1, S_2, \dots, S_M\}$ . The utterances of the speakers in the conversation are ordered in time, represented as  $T = \{t_1, t_2, \dots, t_N\}$ , where  $N$  represents the total number of utterances. The task of MERC is to identify the emotion label (e.g., happy, sad, frustrated, neutral, angry, and depressed) of each utterance in  $T$ .

### 4.2. Multimodal feature encoding

Following the previous work [1], we first use the RoBERTa [39] Large model, the openSMILE [40] toolkit, and the DenseNet [41] model to encode the text, audio, and video data in each utterance respectively, and obtain the initial embedding representation as  $X_t$ ,  $X_a$ , and  $X_v$ .

Then, we further use Bi-GRU to capture the contextual information of text features and use fully connected (FC) networks to transform audio and video features, respectively, to facilitate the understanding of spatiotemporal information. The formula is as follows:

$$\mathbf{h}_t = \overline{\text{Bi-GRU}(X_t)}, \mathbf{h}_a = \overline{\text{FC}(X_a)}, \mathbf{h}_v = \overline{\text{FC}(X_v)}. \quad (2)$$

To obtain preliminary fused multimodal features, we define a learned attention weight  $\alpha_t, \alpha_a, \alpha_v$  for each modality to represent their importance. For each modality, a linear transformation is used to compute its attention score. This is done by computing the dot product of  $\mathbf{h}_m$  as follows:

$$\begin{aligned} \alpha_m &= \mathbf{q}^T \mathbf{h}_m, \\ \hat{\alpha}_m &= \frac{\exp(\alpha_m)}{\exp(\alpha_t) + \exp(\alpha_a) + \exp(\alpha_v)} \end{aligned} \quad (3)$$

where  $m \in \{t, a, v\}$  represents any modality from text, audio, or vision,  $\mathbf{q}$  is a learnable query vector used to compute the attention score for each modality, with the same dimension as the modality feature  $\mathbf{h}_m$ ,  $\hat{\alpha}_m$  represents the attention score after normalization of the  $m$ -th modality. The scores are converted into attention weights using the softmax function to ensure that the sum of the weights is 1. Then, the features from different modalities are weighted and summed to obtain the fused multimodal feature representation  $\mathbf{h}_f$ .

To align heterogeneous multimodal features within a unified semantic space, we adopt an attention-based soft alignment strategy. Instead of requiring strict temporal or frame-level correspondence, the learnable query vector dynamically evaluates the contribution of each modality to the emotional semantics of the current utterance. Modalities that provide more emotionally salient cues receive higher attention weights and therefore dominate the fused representation. This semantic-level alignment allows the model to robustly integrate multimodal information while mitigating noise or inconsistencies across modalities. The fused feature representation is then obtained through a weighted sum as follows:

$$\mathbf{h}_f = \hat{\alpha}_t \mathbf{h}_t + \hat{\alpha}_a \mathbf{h}_a + \hat{\alpha}_v \mathbf{h}_v. \quad (4)$$

### 4.3. Static graph convolution

Traditional Graph Convolutional Networks (GCNs) aggregate information in a discrete, layer-wise fashion, updating node representations over a fixed number of propagation steps. Such a formulation implicitly assumes that information propagation occurs at discrete time points. However, in conversational emotion recognition, emotional states evolve continuously over time rather than through abrupt layer-wise transitions. Modeling emotional dependency propagation using discrete graph layers may therefore limit the ability to capture long-term and smooth temporal dynamics.

To address this issue, we aim to reinterpret the discrete information propagation process of GCNs as a continuous-time dynamical system. By doing so, emotional state evolution can be explicitly modeled as a temporal process governed by an ordinary differential equation (ODE), enabling more flexible and stable information propagation over extended time horizons.

Let  $E$  denote the initial node embeddings produced by the encoder,  $A$  the normalized adjacency matrix, and  $W$  the learnable weight matrix. After multiple propagation steps, the accumulated node representation can be written as:

$$H_n = \sum_{n=1}^N \sum_{k=0}^n A^k E W^k. \quad (5)$$

To bridge discrete propagation and continuous dynamics, we observe that the summation over propagation depth can be interpreted as a Riemann sum over a temporal interval. Letting the propagation index approach a continuous variable  $t$ , Eq. (5) can be rewritten in integral form as:

$$H(t) = \frac{1}{N} \sum_{n=1}^N \int_0^{t+1} A^s E W^s ds. \quad (6)$$

This transformation ensures that the continuous-time formulation remains theoretically consistent with standard GCN behavior.

Taking the derivative of  $H(t)$  with respect to time yields:

$$\frac{dH(t)}{dt} = \frac{1}{N} \sum_{n=1}^N A^{t+1} E W^{t+1}. \quad (7)$$

While this expression captures the instantaneous rate of information propagation, shallow aggregation may lead to information loss. To obtain a more stable formulation, we further analyze the second-order derivative:

$$\frac{d^2H(t)}{dt^2} = \frac{1}{N} \sum_{n=1}^N (A^{t+1} E W^{t+1} + A^{t+1} E W^{t+1} \ln W). \quad (8)$$

Rearranging and integrating both sides with respect to time leads to the following first-order ODE:

$$\frac{dH(t)}{dt} = \ln A H(t) + H(t) \ln W + c. \quad (9)$$

where  $c$  is a constant determined by the initial condition.

Let the eigen-decompositions of the adjacency matrix and weight matrix be given by:

$$A = P\Lambda P^{-1}, \quad W = Q\Phi Q^{-1}. \quad (10)$$

The initial condition  $H(0)$  is defined as:

$$(P^{-1}H(0)Q)_{ij} = \frac{\bar{A}_{ii}\bar{E}_{ij}\bar{\Phi}_{jj} - \bar{E}_{ij}}{\ln \bar{A}_{ii} + \ln \bar{\Phi}_{jj}}. \quad (11)$$

When  $t = 0$ , we obtain:

$$\left. \frac{dH(t)}{dt} \right|_{t=0} = AEW. \quad (12)$$

Combining Eqs. (9)–(12), the constant term can be derived as:

$$c = P\bar{E}Q^{-1} = E. \quad (13)$$

Substituting Eq. (13) into Eq. (9), the final Graph ODE is obtained:

$$\frac{dH(t)}{dt} = \frac{1}{N} \sum_{n=1}^N (\ln \hat{A} H(t) + H(t) \ln W + E). \quad (14)$$

The Graph ODE in Eq. (14) describes the continuous evolution of node representations over time. In practice, this ODE is solved numerically using standard ODE solvers (e.g., the Runge–Kutta method):

$$H(t) = \text{ODESolver} \left( \frac{dH(t)}{dt}, H_0, t \right). \quad (15)$$

This continuous-time formulation enables smooth and long-range emotional dependency propagation and serves as the theoretical foundation for the subsequent dynamic graph convolution module.

The Graph ODE derived above provides a principled continuous-time formulation of graph information propagation that is theoretically consistent with discrete GCNs. To instantiate this formulation within our model, we parameterize the vector field of the ODE using a graph convolution operator, resulting in the proposed SGCODE. Specifically, we adopt the normalized adjacency matrix to encode structural dependencies among utterances and employ a learnable weight matrix that remains fixed during inference to ensure stable and consistent message propagation. Under this setting, the continuous evolution of node representations is governed by a static ODE system, where the graph topology and convolutional parameters jointly determine the propagation dynamics. By grounding SGCODE in the derived Temporal Graph ODE, we ensure that static graph convolution is not merely an architectural choice but a direct realization of the continuous-time dynamics implied by discrete GCN propagation. This static ODE-based formulation serves as the foundational encoding stage of our framework, upon which more flexible and context-adaptive dynamic graph convolution mechanisms are subsequently built.

**Emotional Interaction Graph Construction.** For node construction, we take the multi-modal data as the input. We represent each utterance  $u_i$  as a node  $v_i$ . For edge construction, we consider each node  $v_i$  only interacts with the nodes within the context window  $\{v_j\}_{j \in [\max(i-w, 1), \min(i+w, L)]}$ , where  $w$  is set to 10,  $L$  is the number of utterances. We use  $e_{ij}$  to represent the edge from  $v_i$  to  $v_j$ . For edge type, we assign each edge  $e_{ij}$  with a speaker type identifier  $\alpha_{ij}(p_{s(u_i)} \rightarrow p_{s(u_j)})$ , where  $p_{s(u_i)}$  and  $p_{s(u_j)}$  represent the speaker identity of  $u_i$  and  $u_j$ , respectively.

The construction of the graph mainly depends on the adjacency matrix  $A$ . The element  $A_{ij}$  of the adjacency matrix represents the intensity of emotional dependence between multimodal utterances  $i$  and  $j$ . Assuming that the features of utterances  $i$  and  $j$  are  $\mathbf{h}_i$  and  $\mathbf{h}_j$ , the adjacency matrix  $A$  is constructed by calculating the cosine similarity between them. The formula is as follows:

$$\text{sim}(\mathbf{h}_i, \mathbf{h}_j) = \frac{\mathbf{h}_i^\top \mathbf{h}_j}{\|\mathbf{h}_i\| \|\mathbf{h}_j\|}, \quad A_{ij} = \mathbb{I}(\text{sim}(\mathbf{h}_i, \mathbf{h}_j) > \theta), \quad (16)$$

where  $\mathbb{I}(\cdot)$  is the indicator function and  $\theta$  is the similarity threshold.

**Graph ODE Modeling.** Graph ordinary differential equation (ODE) is used to describe the changes in node features over time. The features

of each node at time  $t$  are updated through the graph ODE. In general, the form of graph ODE is as follows:

$$\frac{d\mathbf{h}_i(t+1)}{dt} = f(\mathbf{h}_i(t), \mathbf{A}, t), \quad (17)$$

where  $\mathbf{h}_i(t)$  is the feature vector of node  $i$  at time  $t$ .  $\mathbf{A}$  is the adjacency matrix of the graph.  $f(\cdot)$  is a GCN used to learn the relationship between nodes. Time  $t$  represents the node features that change over time in a dynamic system.

**Convolution with Graph ODE.** Under the framework of graph ODE, the features of each node  $\mathbf{h}_i(t)$  are propagated through graph convolution according to the features of neighboring nodes and updated over time. The graph ODE models the changes in the features of each node in the form of time derivatives, while graph convolution provides information transfer between nodes. The formula for SGCODE is as follows:

$$\frac{d\mathbf{h}_i(t+1)}{dt} = \sum_{j \in \mathcal{N}(i)} \ln A_{ij} \mathbf{h}_j(t) + \mathbf{h}_i(t) \ln W_s, \quad (18)$$

where  $\mathcal{N}(i)$  represents the first-order neighbors of node  $i$ , and  $W_s$  is a trainable parameter matrix that will remain fixed during the inference phase,  $\sum_{j \in \mathcal{N}(i)}$  represents the summation of all nodes  $j$  belonging to  $\mathcal{N}(i)$ .

#### 4.4. Global information vector

Drawing inspiration from the Transformer’s ability to model global dependencies, we generate a GIV using global attention mechanisms and global average pooling (GAP), which guides the subsequent prompt generation and dynamic fusion process.

**Global Context Modeling.** SGCODE obtains the representation of each utterance at time  $t$  by simulating the evolution of the utterance representation of the graph over time  $\mathbf{h}_i(t)$ . These utterance representations can be used as input for global attention to fuse important emotional information in the context of global conversation. Next, we calculate the attention weights and fuse the conversation context information:

$$\mathbf{A}_{score} = \text{softmax} \left( \frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} \right), \quad (19)$$

$$\mathbf{Z} = \mathbf{A}_{score} \mathbf{V},$$

where  $\mathbf{A}_{score}$  is the attention score matrix, which represents the relationship between each utterance and other utterances;  $d_k$  is the dimension of the query vector, which is used for scaling calculations;  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  denote the query, key, and value matrices obtained by linearly projecting the input features; and  $\mathbf{Z}$  is the fused feature matrix of all utterances.

**Global Average Pooling.** The global information vector represents the global conversation information by averaging the features of all input utterances. The generated vector  $\mathbf{g}$  can represent the global features of the entire emotional interaction graph. The GAP operation synthesizes the representation of all utterances into a fixed-length vector. The specific formula is as follows:

$$\mathbf{g} = \frac{1}{N} \sum_{i=1}^N \mathbf{z}_i, \quad (20)$$

where  $\mathbf{g}$  is the global information vector, representing the global features of the entire graph, and  $\mathbf{z}_i \in \mathbf{Z}$  is the representation of each utterance output by the Transformer,  $N$  represents the total number of utterances.

#### 4.5. Dynamic graph convolution

In DGCODE, the core idea is to dynamically generate model parameters conditioned on the GIV. We clarify that the GIV, generated from the internal features of the input conversation, functions as an internal prompt or a conditioning signal. Unlike traditional prompt learning, which often employs external, human-designed task instructions, our internal prompt (GIV) encapsulates the global conversational context. This prompt then guides the Prompt Generation Network (PGN) to

produce dynamic weights, effectively instructing the model on how to adaptively fuse multimodal features for the current emotional and contextual scenario. This conditioned dynamic fusion is applied to the graph ODE. Consequently, the ODE system in DGCODE is not a fixed dynamic system but one whose evolution is explicitly instructed by the unique context of the input, allowing it to equip different network parameters for different emotion categories during inference.

**Dynamic Weight Generation.** The PGN generates parameters related to the dynamic evolution of graph ODE based on the GIV, including node update rules and edge influences. Specifically, we first input the global information vector  $\mathbf{g}$  into the MLP with a fully connected layer and use the activation function to perform nonlinear changes to obtain the selection vector  $\mathbf{s}$ :

$$\mathbf{s} = \sigma(w_l \mathbf{g} + b_l), \quad (21)$$

where  $w_l, b_l$  are the trainable parameters. Then, we introduce the learnable weights box  $W_b$  and perform dot product with the selection vector  $\mathbf{s}$  to generate a dynamic weight matrix, denoted as  $W_d$ :

$$W_d = \mathbf{s} \otimes W_b. \quad (22)$$

**Dynamic Convolution with Graph ODE.** Unlike SGCODE, which employs learnable weights during training and fixed weights during inference, DGCODE utilizes dynamically generated weights for both training and inference phases, enabling the fusion and updating of utterance features. In addition, to avoid the over-smoothing problem of GCN, we concatenate the outputs of each layer of SGCODE and DGCODE and batch normalize them to construct residual connections to improve the model's generalization ability. The formula for DGCODE is as follows:

$$\frac{d\mathbf{u}_i(t+1)}{dt} = \sum_{j \in \mathcal{N}(i)} \ln \mathbf{A}_{ij} \mathbf{u}_j(t) + \mathbf{u}_i(t) \ln W_d, \quad (23)$$

$$\mathbf{u}_i(t) = \text{BN}(\text{Concat}(\mathbf{h}_i(t-1), \mathbf{u}_i(t-1)); \gamma, \beta),$$

where  $W_d$  is the generated dynamic parameter,  $\mathbf{h}_i$  is the output feature of SGCODE,  $\mathbf{u}_i$  is the output feature of DGCODE,  $\text{BN}(\cdot)$  is the batch normalization operation,  $\gamma$  is the learnable scaling parameter, and  $\beta$  is the learnable offset parameter.

Unlike static graph convolution, where a fixed weight matrix is shared across all samples,  $W_d$  dynamically modulates the node update process at each infinitesimal time step. Specifically,  $W_d$  controls the relative importance of different feature dimensions during temporal propagation, thereby shaping the continuous-time evolution of node representations. As a result, the same graph structure can exhibit different propagation behaviors under different conversational contexts. The influence of  $W_d$  on the final classification is accumulated throughout the ODE integration process. Since the node representations produced by DGCODE are directly fed into the emotion classifier,  $W_d$  affects which semantic or modality-related features are amplified or suppressed before prediction, shaping the decision process rather than acting as a post-hoc transformation. Importantly,  $W_d$  is not an unconstrained parameter. It is generated by the prompt generation network conditioned on the GIV, which summarizes the global emotional context of the conversation. This design allows  $W_d$  to be interpreted as an emotion-aware convolutional kernel that adapts the ODE dynamics to different emotional trends. Moreover, while  $W_d$  modulates node self-dynamics, the graph topology and neighborhood aggregation remain unchanged, preserving interpretability and ensuring that  $W_d$  modifies how information is integrated over time rather than which information is propagated.

#### 4.6. Emotion classifier

The feature  $\mathbf{u}_i$  obtained by DGCODE is batch normalized and input into the linear layer through residual connection with the feature  $\mathbf{h}_i$  output by the encoder, and then through the softmax layer to obtain the probability distribution  $\mathbf{p}_i$  of the emotional label. The formula is defined

as follows:

$$\begin{aligned} \mathbf{v}_i &= \text{Concat}(\mathbf{h}_i, \text{BN}(\mathbf{u}_i; \gamma, \beta)), \\ \mathbf{r}_i &= \text{ReLU}(\mathbf{v}_i W_r + b_r), \\ \mathbf{p}_i &= \text{softmax}(\mathbf{r}_i W_c + b_c), \end{aligned} \quad (24)$$

where  $W_r, b_r, W_c, b_c, \gamma$  and  $\beta$  are learnable parameters.

Finally, we use the argmax function to obtain the emotional label with the highest probability of utterance  $i$  and use the classification cross-entropy loss for model training.

$$\begin{aligned} \hat{y}_i &= \text{argmax}(\mathbf{p}_i), \\ \mathcal{L}_{ce} &= -\frac{1}{n} \sum_{i=1}^n y_i \log(\hat{y}_i), \end{aligned} \quad (25)$$

where  $y_i$  and  $\hat{y}_i$  are the predicted and true emotion labels of utterance  $i$ , respectively.

## 5. Experiments

### 5.1. Database used

IEMOCAP [42] and MELD [43] are the two most commonly used databases in the MERC. They are widely used to evaluate the effectiveness of models.

The IEMOCAP dataset is collected by the SAIL laboratory of the University of Southern California. It is a multimodal conversation dataset that contains about 12 hours of multimodal data such as video, audio, facial motion capture, and transcribed text. It is recorded through the improvisation or scripted dialogues of actors. The data is annotated with emotion categories and dimension labels such as anger, happiness, and sadness, providing rich materials for multimodal emotion analysis.

The MELD dataset is derived from the EmotionLines dataset. It adds audio and visual modalities based on the plain text dialogues of "Friends". It contains 1443 dialogues and 13,708 sentences. Each sentence is annotated with one of the seven emotion labels, such as anger, disgust, sadness, positive, negative, and neutral emotion labels. It is used for emotion recognition in multi-party conversations. Both datasets have important value in multimodal emotion analysis.

### 5.2. Evaluation metrics

In this section, we mainly explain the evaluation metrics to verify the effectiveness of the model in this paper. This paper mainly uses the following four evaluation metrics: 1) accuracy (ACC); 2) weighted accuracy (WA); 3) F1; 4) weighted F1 (WF1). To consider the problem of data set category imbalance, more comprehensively evaluate the model performance, and meet the actual application needs, we use WA and WF1 as the main evaluation metrics.

### 5.3. Baselines

To verify the superior performance of our proposed DFGCN method, we compare it with other comparative methods as follows:

**QMNN** [44] uniformly handles multimodal fusion and dialog context modeling through quantum analogy, and performs end-to-end training using complex-valued neural networks.

**A-DMN** [45] reconstructs the emotion recognition in conversation task into a question-answering-like reasoning problem, and models the speaker's self-dependency (emotional inertia) and mutual dependency (emotional contagion) through a collaborative architecture.

**CoMPM** [46] transforms the Dynamic Memory Network (DMN) into a four-module architecture. It efficiently processes multimodal inputs using hierarchical convolutional fusion, models separately the global BiLSTM states and individual LSTM states, and directionally extracts emotional clues of target utterances through multi-round attention memory iteration.

**MGLRA** [47] iteratively refines multimodal features through a memory-enhanced recurrent alignment module, fuses semantics using a masked graph convolutional network, and finally classifies emotions through MLP.

**MMGCN** [48] uniformly models the dependencies among text, audio, and visual modalities in conversations through a fully connected heterogeneous graph, enhances context modeling by combining speaker embeddings, and finally performs emotion fusion and classification via deep spectral domain GCN.

**AdaIGN** [11] models the temporal dependencies of multimodality in conversations through a directed heterogeneous graph, dynamically filters conflicting modal information by combining node/edge-level selection strategies, and outputs emotion classification by fusing dual-graph features through a graph-level selection strategy.

**DER-GCN** [15] proposes a multi-relational emotional interaction graph network, which fuses dialogue relations and event relations. It jointly reconstructs node features and edge structures through a self-supervised masked graph autoencoder, dynamically aggregates cross-relational semantics using a multi-informative Transformer, and alleviates the long-tail data problem with the assistance of contrastive learning loss optimization.

**RGAT** [49] proposes relational positional encoding, which integrates sequence positional information into the relational graph attention network for the first time, solving the problem that traditional graph models ignore temporal dependencies.

**CBERL** [50] generates minority samples that conform to the original distribution through a dual-path generative adversarial network, constructs a deep joint variational autoencoder, constrains the latent space via KL divergence, and fuses complementary semantics of multimodal modalities.

**M3Net** [16] constructs a hypergraph to model high-order multimodal relationships and designs low-pass/high-pass filters to separate emotional common/differential signals, aiming to address the issues of insufficient modeling of complex multivariate relationships and neglect of high-frequency emotional signals in multimodal conversation emotion recognition.

**CT-Net** [51] solves the three major bottlenecks in traditional methods (i.e., multimodal misalignment, lack of long-range dependencies, and neglect of speaker influence) through cross-modal temporal alignment, context attention mechanism, and speaker-aware fusion.

**EmoBERTa** [39] replaces the traditional complex multi-module architecture by adding speaker labels and separating historical/current/future utterances, that is, constructing input sequences that can perceive speaker identities and dialogue contexts.

**FrameERC** [52] leverages graph framelet transforms and a dual-reminder fusion mechanism to capture both fine-grained emotional cues and balanced cross-modal interactions.

**HiMul-LGG** [53] proposes a hierarchical decision fusion-based local-global graph neural network for multimodal emotion recognition in conversation. It aligns multimodal features through a hierarchical fusion strategy, models speaker dependencies via a local-global graph structure, and enhances cross-modal interaction using a multi-head attention mechanism.

#### 5.4. Experimental setup

All experiments were conducted using Python 3.8 and PyTorch 1.8, and all computations were performed on a single NVIDIA RTX 4090 GPU with 24 GB memory. The proposed model was trained using the AdamW optimizer with an initial learning rate of  $5 \times 10^{-4}$ , a weight decay of  $1 \times 10^{-4}$ , and the cross-entropy loss function. The batch size was set to 32, and the maximum number of training epochs was set to 50, with early stopping applied based on validation performance. For the model architecture, the hidden dimension of node representations was set to 256, and three stacked graph convolution blocks were used for both SGCODE and DGCODE. The context window size for graph con-

struction was fixed to  $w = 10$ . For the ODE-based modules, the initial state was defined by the encoder output, and the ODE system was solved using a fourth-order Runge–Kutta method with a fixed step size of 0.1. The prompt generation network consisted of a two-layer multilayer perceptron with ReLU activation, and the dimension of the GIV was set to 256. For visualization, t-SNE was applied to the learned utterance-level representations extracted from the test set only. We used the scikit-learn implementation of t-SNE with two output dimensions, random initialization, a perplexity value of 30, and 1000 optimization iterations, while other parameters were set to their default values. To determine the optimal hyperparameters, leave-one-out cross-validation was employed on the validation set. All reported results represent the average performance over 10 independent runs with different random weight initializations. To assess statistical significance, paired *t*-tests were conducted across the 10 runs, and all improvements were found to be statistically significant ( $p < .05$ ), indicating that the observed performance gains are unlikely to be caused by random variation.

#### 5.5. Overall result

The experimental results, as shown in [Tables 1 and 2](#), clearly demonstrate that the proposed DFGCN method delivers significant performance improvements across various emotion recognition tasks. This consistent improvement can be primarily attributed to the design of the ODE-injected graph convolutional network, which explicitly models the dynamic evolution of emotional dependencies within discourse interaction networks. By leveraging the global discourse-level information vector as guidance, our model is capable of dynamically fusing multimodal features in a context-aware manner, thereby capturing subtle variations in emotional transitions that traditional static fusion strategies often fail to recognize.

To further illustrate the advantages of our framework, we report the weighted-F1 (W-F1) scores for individual emotion categories. On the IEMOCAP dataset, DFGCN consistently outperforms strong baseline models in critical categories such as happy, neutral, and depressed. These categories are known to be challenging due to their overlapping linguistic and acoustic patterns, and the observed improvements suggest that our model is particularly effective at disentangling such fine-grained distinctions. Similarly, on the MELD dataset, DFGCN achieves the highest W-F1 score for the sad category, while also delivering competitive results in surprise and joy, which are emotions typically sensitive to both verbal and non-verbal cues. These results indicate that the model not only excels in recognizing primary emotions but also generalizes well to subtle and context-dependent emotional states.

While DF-GCN achieves state-of-the-art overall performance on the MELD dataset, particularly in terms of weighted average F1, a closer examination of the per-class results in [Table 2](#) shows that recognition accuracy for certain minority emotion categories, especially Sadness and Joy, remains relatively limited. This limitation does not arise from deficiencies in our dynamic fusion mechanism. Instead, it is primarily driven by two interconnected challenges inherent to the MELD dataset. The first challenge is the severe class imbalance. The Neutral category constitutes more than half of all utterances, whereas Sadness and Joy contribute only about 3 percent and 4.7 percent of the test set. This extreme imbalance restricts the model's exposure to representative samples of minority emotions and significantly constrains its ability to learn stable and discriminative features, even when the fusion mechanism adapts to contextual cues. The second challenge relates to the highly heterogeneous and context-sensitive nature of multimodal emotion expression in natural dialogue. Sadness may manifest through muted prosodic patterns, subtle or minimal facial expressions, or may be concealed by irony or social politeness. Joy can appear as intense laughter or quiet and restrained expressions of positive affect. These diverse realizations often produce inconsistent or weak alignment among textual, acoustic, and visual modalities, thereby increasing intra-class variability and making minority emotions particularly prone to misclassification as the

**Table 1**

The performance of different methods on IEMOCAP Dataset. The values reported in each cell represent the Acc/F1. The best result in each column is in bold.

Method	Happy		Sad		Neutral		Angry		Excited		Frustrated		Average(w)	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
A-DMN [45]	43.1	50.6	69.4	76.8	63.0	62.9	63.5	56.5	<b>88.3</b>	77.9	53.3	55.7	64.6	64.3
RGAT [49]	42.3	51.6	70.1	77.3	61.6	65.4	64.2	63.0	69.3	68.0	57.8	61.2	66.5	65.2
QMNN [44]	41.4	39.7	72.9	68.3	54.1	55.3	65.4	62.6	66.0	66.7	55.6	62.2	60.8	59.9
CoMPM [46]	55.4	60.7	78.3	82.2	64.5	63.0	66.2	59.9	77.4	78.2	57.3	59.5	64.2	67.3
EmoBERTa [39]	56.9	56.4	79.1	<b>83.0</b>	64.0	61.5	70.6	69.6	86.0	78.0	63.8	68.5	67.3	67.3
CTNet [51]	47.9	51.3	78.0	79.9	69.0	65.8	65.7	67.2	85.3	<b>78.7</b>	52.2	58.8	68.0	67.5
MMGCN [48]	48.9	47.1	<b>82.3</b>	81.9	64.1	66.4	64.2	63.5	73.9	76.2	60.0	59.1	66.8	66.8
M3Net [16]	59.1	60.9	74.5	78.8	65.4	70.1	69.2	68.1	73.4	77.1	72.8	67.0	72.5	71.1
MGLRA [47]	<b>62.9</b>	<b>63.5</b>	81.1	81.5	70.9	71.5	60.2	61.1	74.4	76.3	69.2	67.8	71.3	70.1
CBERL [50]	58.8	67.3	63.3	72.8	56.4	60.8	75.3	73.5	70.3	70.8	<b>78.2</b>	71.2	69.4	69.3
AdaIGN [11]	51.1	53.0	80.5	81.5	72.3	71.3	63.3	65.9	72.8	76.3	63.5	67.8	72.1	70.7
DER-GCN [15]	60.7	58.8	75.9	79.8	66.5	61.5	71.3	72.1	71.1	73.3	66.1	67.8	69.7	69.4
FrameERC [52]	60.3	56.2	80.1	80.9	66.4	67.7	<b>75.5</b>	<b>78.6</b>	70.0	69.7	71.3	70.8	70.0	70.7
HiMul-LGG [53]	53.2	54.0	79.4	79.9	71.2	71.7	69.3	67.6	71.1	72.0	69.1	68.5	70.1	70.2
DF-GCN	<b>62.9</b>	62.0	72.6	81.3	<b>72.9</b>	<b>71.9</b>	69.1	70.7	86.6	75.5	70.1	<b>68.6</b>	<b>73.4</b>	<b>72.2</b>

**Table 2**

The performance of different methods on MELD dataset. The values reported in each cell represent the Acc/F1. The best result in each column is in bold.

Method	Neutral		Surprise		Fear		Sadness		Joy		Disgust		Anger		Average(w)	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
A-DMN [45]	78.8	78.9	55.4	55.3	8.7	6.8	24.7	23.9	24.6	22.7	3.5	3.2	41.3	40.9	55.5	55.4
RGAT [49]	77.6	78.1	43.2	41.5	5.6	2.4	29.8	30.7	59.1	58.6	3.4	2.2	45.2	44.6	60.2	59.5
QMNN [44]	71.2	77.0	45.8	49.8	–	–	24.3	16.5	53.5	52.1	–	–	42.9	43.2	60.8	58.0
CoMPM [46]	79.1	82.0	51.5	49.2	5.2	2.9	33.6	32.3	62.3	61.5	8.4	2.8	42.7	45.8	64.3	63.0
EmoBERTa [39]	78.9	<b>82.5</b>	50.2	50.2	1.8	1.9	33.3	31.2	<b>72.1</b>	61.7	9.1	2.5	43.3	46.4	64.1	63.3
CTNet [51]	75.6	77.4	51.3	50.3	5.1	10.0	30.9	32.5	54.3	56.0	11.6	11.2	42.5	44.6	61.9	60.2
MMGCN [48]	75.8	77.0	52.5	49.6	4.4	3.6	18.8	20.4	54.9	53.8	10.2	2.8	44.1	45.2	58.3	58.4
M3Net [16]	73.6	79.1	57.6	59.5	23.1	13.3	45.6	42.9	63.2	65.1	24.3	21.7	50.4	53.5	63.9	65.8
MGLRA [47]	78.2	80.8	59.8	59.5	0.0	0.0	30.8	27.8	68.5	66.5	0.0	0.0	43.8	48.4	65.4	64.9
CBERL [50]	79.8	82.0	56.4	57.9	26.3	<b>22.2</b>	44.2	41.4	67.1	65.7	10.2	14.7	52.8	55.3	64.4	65.5
AdaIGN [11]	76.1	77.8	<b>59.9</b>	60.5	22.4	15.2	44.8	43.7	66.2	64.5	33.2	<b>29.3</b>	54.3	56.2	65.3	65.8
DER-GCN [15]	76.8	80.6	50.5	51.0	14.8	10.4	56.7	41.5	69.3	64.3	17.2	10.3	52.5	<b>57.4</b>	65.2	65.5
FrameERC [52]	80.7	80.0	56.4	58.1	12.3	10.5	41.2	42.5	60.8	63.4	18.4	20.1	55.1	56.3	65.5	66.1
HiMul-LGG [53]	77.2	75.5	55.1	53.2	10.0	9.1	<b>64.4</b>	<b>65.9</b>	69.3	<b>68.1</b>	8.7	9.3	52.2	52.9	66.2	65.2
DF-GCN	<b>81.5</b>	80.0	56.8	<b>62.7</b>	<b>27.6</b>	20.3	40.8	43.5	64.7	64.9	<b>34.1</b>	26.8	<b>56.1</b>	55.4	<b>67.4</b>	<b>67.6</b>

dominant Neutral class. Although the DGCODE module uses global contextual prompts to adaptively reweight modality contributions and our confusion matrices confirm reduced misclassification of Sadness as Neutral compared with static-fusion baselines, the benefit remains bounded by the scarcity and ambiguity of minority samples. These findings highlight a broader issue common to current multimodal emotion recognition in conversation. Without incorporating dedicated techniques for handling long-tailed distributions, including focal loss, balanced resampling, or decoupled representation learning, even advanced dynamic fusion architectures face inherent challenges in achieving uniformly strong performance across all emotion categories. We recognize this as an important direction for future research.

Furthermore, it is noteworthy that the performance gains achieved by the model come with only a marginal increase in computational cost when compared to existing graph-based methods. Although the proposed framework introduces ODE-based propagation and dynamic parameter generation, the overall model complexity and inference time remain comparable to strong graph-based baselines. The integration of ODE dynamics enables efficient representation learning while keeping the computational overhead manageable. This favorable balance between accuracy and efficiency highlights the practicality of DF-GCN for real-world multimodal emotion recognition systems, where both predictive performance and computational cost are critical considerations.





Fig. 3. Verifying the effectiveness of multimodal features.

### 5.6. The effectiveness of multimodal features

To assess the impact of different combinations of modality features on emotion recognition performance, we performed a detailed ablation study, analyzing the contributions of text, audio, and video features both individually and in various combinations. Specifically, we designed experiments in which the model was provided with different subsets of modalities, and we carefully examined the resulting changes in recognition accuracy. The outcomes of these experiments are summarized in Fig. 3.



**Fig. 4.** Sensitivity analysis of DF-GCN with respect to key graph construction hyperparameters. (a) Effect of the similarity threshold  $\theta$  on WF1. (b) Effect of the context window size  $w$  on WF1.

In the first set of experiments, where each modality was tested in isolation, we observed that the text modality consistently delivered the highest recognition accuracy compared with the audio and video modalities. This finding indicates that textual features contain the most discriminative cues for emotion classification, likely because they directly encode semantic content and contextual information that are critical for distinguishing between nuanced emotional states. By contrast, audio and visual features, although useful, proved less effective when used independently, as they tend to capture paralinguistic or nonverbal signals that may be more ambiguous or context-dependent.

Subsequently, we evaluated the performance of the model when combining two modalities. The results revealed that both the text-audio combination and the text-video combination yielded significant improvements over the single-modality settings. These enhancements suggest that the integration of multiple modalities enables the model to capture complementary information. While the text modality provides a strong semantic foundation, the audio modality contributes prosodic and acoustic cues, and the video modality adds visual expressions such as facial movements and gestures. The fusion of these heterogeneous features therefore strengthens the model’s ability to accurately infer emotions, particularly in cases where a single modality may provide incomplete or ambiguous signals.

Finally, when all three modalities (e.g., text, audio, and video) were used jointly, the model achieved its best overall performance. This outcome highlights the fundamental advantage of multimodal learning in emotion recognition tasks. By simultaneously exploiting information from linguistic, acoustic, and visual channels, the model is able to construct a more comprehensive and context-aware representation of emotional states, leading to predictions that are both more accurate and more robust.

### 5.7. Sensitivity analysis of graph construction hyperparameters

Fig. 6 presents the sensitivity analysis of DF-GCN with respect to two key graph construction hyperparameters: the similarity threshold  $\theta$  and the context window size  $w$ . Results from both the IEMOCAP and MELD datasets indicate that the proposed model maintains stable performance over a wide range of parameter values, suggesting that DF-GCN is not overly sensitive to the selection of these hyperparameters. As shown in Fig. 1(a), varying the similarity threshold  $\theta$  results in only moderate performance fluctuations. When  $\theta$  is too small, the constructed graph may become overly dense and introduce noisy connections, whereas excessively large values of  $\theta$  tend to produce sparse graphs that limit effective information propagation. The model achieves consistently strong performance when  $\theta$  is set to an intermediate range, indicating a balanced trade-off between graph connectivity and noise suppression. Fig. 1(b) illustrates the effect of the context window size  $w$ . A small window size restricts the available conversational context, while an overly large window may introduce redundant or less relevant interactions. The results show that DF-GCN maintains robust performance across different window sizes, with moderate values of  $w$  providing an effective balance between contextual coverage and modeling efficiency.

**Table 3**

Comparison of model parameters and inference time on IEMOCAP and MELD datasets.

Model	# Params.	IEMOCAP	MELD
		Inference Time	Inference Time
MMGCN	5.63M	53.7s	75.3s
RGAT	15.28M	68.5s	146.3s
DER-GCN	78.59M	125.5s	189.7s
M3Net	7.88M	57.0s	87.6s
DF-GCN (Ours)	9.19M	60.2s	94.1s

### 5.8. Error analysis

To better understand the model’s limitations, we analyzed the confusion matrices of the IEMOCAP and MELD test sets. As shown in Fig. 5 (a) and (b), DFGCN demonstrates stronger diagonal concentration compared with DER-GCN, yet it still exhibits several notable misclassification patterns. On the IEMOCAP dataset, DFGCN often predicts excited as happy and confuses angry with frustrated or neutral. While similar errors are also present in DER-GCN (Fig. 5 (c)), the overall distribution of DFGCN shows fewer cross-category confusions, particularly for the sad and frustrated classes, indicating its improved ability to capture fine-grained distinctions.

On the MELD dataset, DFGCN also reduces confusion compared to DER-GCN, especially in the recognition of joy and sadness. Nevertheless, challenges remain due to class imbalance: the dominance of the neutral category biases predictions, leading to frequent misclassification of minority emotions, such as fear into neutral. Moreover, categories like fear and disgust, which are underrepresented in MELD, remain difficult to classify correctly for both models. Although DFGCN leverages dynamic weight allocation to mitigate imbalance, its recognition of these minority classes is still limited, albeit consistently superior to DER-GCN.

### 5.9. Computational efficiency analysis

Table 3 presents a comparison of model size and inference efficiency among different graph-based approaches on the IEMOCAP and MELD datasets. Overall, the proposed DF-GCN achieves a favorable balance between recognition performance and computational efficiency. From the perspective of model complexity, DF-GCN maintains a moderate parameter scale compared with existing methods. While it introduces additional components such as ODE-based graph propagation and dynamic parameter generation, the overall parameter count remains comparable to mainstream graph-based models and is substantially smaller than that of more complex relational graph architectures. This indicates that the proposed design avoids excessive parameter growth while enhancing modeling capacity. In terms of inference efficiency, DF-GCN demonstrates competitive runtime performance across both datasets. Its inference speed is on par with commonly used graph convolutional models and notably faster than heavier graph-based baselines that rely on complex relational modeling. This suggests that the additional computations introduced by dynamic parameter generation and ODE integration are effectively controlled and do not lead to a prohibitive increase in inference time. Overall, these results show that the performance gains achieved by DF-GCN come with only marginal additional computational cost. The lightweight prompt generation network and the fixed-depth ODE integration ensure that the proposed model remains efficient in practice.

### 5.10. Stability and robustness analysis

As illustrated in Table 4, our proposed DF-GCN not only achieves the state-of-the-art (SOTA) performance in terms of average Accuracy and F1-score across both IEMOCAP and MELD datasets but also exhibits superior stability compared to all baseline methods. Specifically,




Fig. 5. Confusion matrices of emotion classification results on the IEMOCAP and MELD test sets. Subfigures (a) and (b) present the performance of the proposed DF-GCN model on IEMOCAP and MELD, respectively, while subfigures (c) and (d) show the results of the baseline DER-GCN under the same settings. Darker diagonal blocks indicate higher accuracy in correctly identifying specific emotion categories, whereas off-diagonal values represent misclassifications.

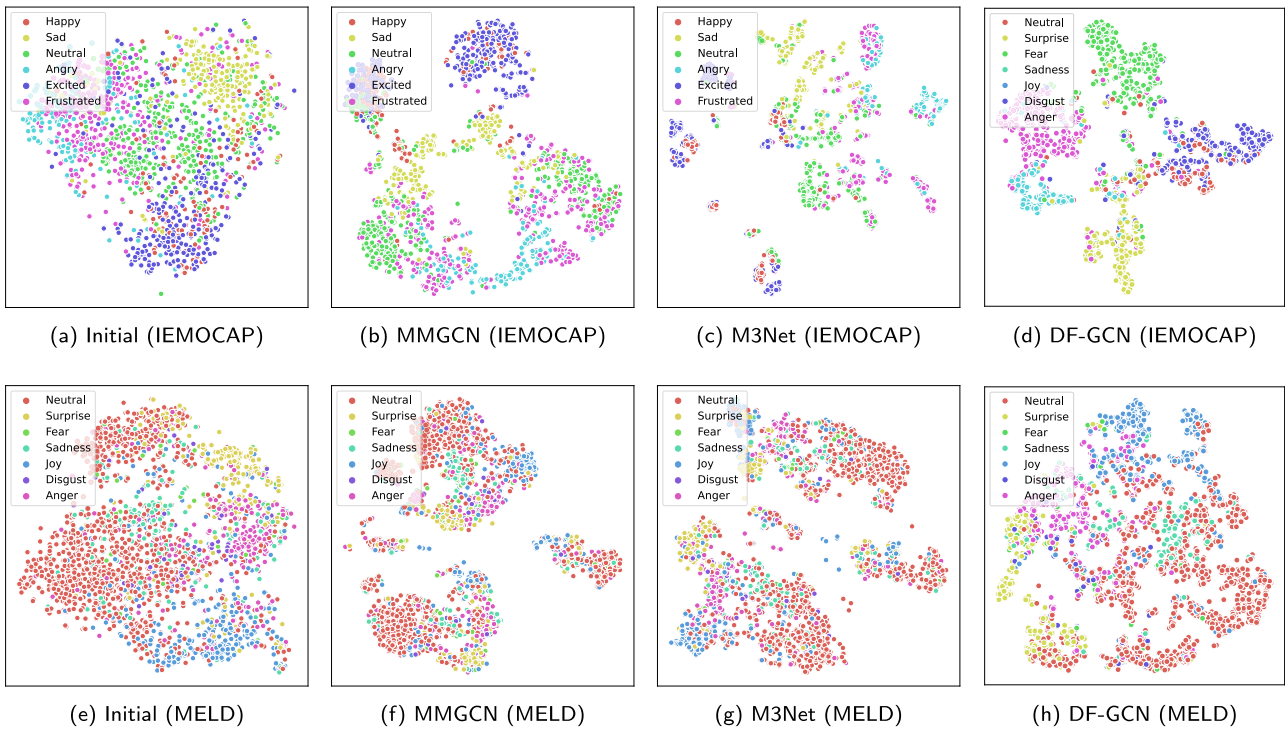


Fig. 6. Visualization of the learned embeddings.

on the IEMOCAP dataset, DF-GCN yields the lowest standard deviation for both metrics, significantly outperforming other competitive models like M3Net and AdaIGN. A similar trend is observed on the MELD dataset, where DF-GCN maintains an remarkably low variance, while the variances of most baselines range from 0.45 to 1.58. The consistently smaller standard deviation after 10 independent runs indicates that DF-GCN is less sensitive to random initialization and noise within the multimodal data. These results strongly demonstrate the robustness of our architecture, confirming that the performance gains of DF-GCN are reliable and reproducible rather than the result of advantageous random seeding.

5.11. Ablation experiment

To analyze the contribution of each component in DF-GCN, we conducted a series of ablation experiments on the IEMOCAP and MELD datasets. As shown in Table 5, DF-GCN consistently outperforms all ablated variants in terms of both Accuracy (Acc) and Weighted F1 score (W-F1), demonstrating the effectiveness of the complete model architecture. These results indicate that removing any individual com-

Table 4

Average performance after 10 runs on the IEMOCAP and MELD datasets. The values reported in each cell represent the average Acc/F1 ± std. The best result in each column is in bold.

Methods	IEMOCAP		MELD	
	Acc	F1	Acc	F1
QMNN	60.8 ± 0.87	59.9 ± 1.23	60.8 ± 0.45	58.0 ± 1.58
CoMPM	64.2 ± 0.19	67.3 ± 1.02	64.3 ± 0.66	63.0 ± 1.41
EmoBERTa	67.3 ± 0.33	67.3 ± 1.15	64.1 ± 0.92	63.3 ± 0.54
MMGCN	66.8 ± 1.37	66.8 ± 0.28	58.3 ± 1.49	58.4 ± 0.71
M3Net	72.5 ± 1.08	71.1 ± 1.55	63.9 ± 0.83	65.8 ± 1.29
MGLRA	71.3 ± 0.41	70.1 ± 0.96	65.4 ± 1.12	64.9 ± 0.62
CBERL	69.4 ± 1.44	69.3 ± 1.33	64.4 ± 0.78	65.5 ± 1.05
AdaIGN	72.1 ± 0.51	70.7 ± 1.52	65.3 ± 1.19	65.8 ± 0.89
DER-GCN	69.7 ± 1.46	69.4 ± 1.31	65.2 ± 0.68	65.5 ± 1.01
DF-GCN	<b>73.4 ± 0.21</b>	<b>72.2 ± 0.17</b>	<b>67.4 ± 0.13</b>	<b>67.6 ± 0.25</b>

ponent leads to a noticeable performance drop, confirming that each module contributes meaningfully to the overall performance. In particular, removing the GIV results in a clear degradation in performance,

**Table 5**  
Ablation results of DF-GCN on the IEMOCAP and MELD benchmark dataset.

GIV	PGN	DGCODE	IEMOCAP		MELD	
			Acc.	F1	Acc.	F1
+	+	+	<b>73.4</b>	<b>72.2</b>	<b>67.4</b>	<b>67.6</b>
-	+	+	71.3	71.6	67.2	65.8
+	-	+	70.6	70.8	66.9	64.8
+	+	-	69.4	68.7	65.4	64.4
-	-	+	70.6	70.8	66.8	65.0
-	+	-	68.9	68.4	64.1	63.5
+	-	-	68.5	68.0	63.3	62.4
-	-	-	66.7	67.0	62.3	61.6

highlighting its critical role in capturing and integrating global contextual information across the conversation. GIV enables the model to maintain a holistic understanding of emotional context, which is essential for robust multimodal emotion recognition. Moreover, eliminating the PGN or DGCODE also leads to reduced performance, underscoring the importance of dynamically adapting model parameters to emotion-specific contexts. The joint operation of PGN and DGCODE allows the model to generate adaptive weights conditioned on emotional dynamics, thereby enhancing classification accuracy and robustness.

### 5.12. Visualization

To more intuitively demonstrate the classification effect of the DF-GCN method on the MERC task, we used the T-SNE technology to reduce the dimension of the generated sentence vectors and visualize them. As shown in Fig. 6, on the IEMOCAP dataset, the DF-GCN model shows excellent classification performance. We can see that different emotion categories form obvious cluster distributions, and the distances between different clusters are large, indicating that DF-GCN can well capture the differences between different emotions. In contrast, although the MMGCN model can also distinguish samples of different emotion categories to some extent, its classification performance is obviously inferior to that of DF-GCN. From the visualization results, it can be seen that the distribution of samples generated by MMGCN is relatively chaotic, the boundaries between different emotion categories are not clear enough, and samples of some categories even overlap. For example, samples of happy and neutral emotions are distributed close to each other in the figure, which may cause confusion during classification. Furthermore, we also compared the classification effect of the M3Net model on the IEMOCAP dataset. Similar to DF-GCN, M3Net also shows good classification performance and can clearly separate samples of different emotion categories. It can be seen that the samples generated by M3Net are more concentrated, the boundaries between categories are clear, and the clustering effect is comparable to that of DF-GCN. To further verify the generalization ability of DF-GCN, we conducted similar experiments on the MELD dataset and observed the same phenomenon as that on the IEMOCAP dataset. DF-GCN also showed excellent classification performance, and samples of different emotion categories were effectively separated in the two-dimensional space.

## 6. Conclusion

This paper presents DF-GCN, a robust framework for recognizing multimodal emotional features in conversations. Specifically, DF-GCN injects ordinary differential equations into GCN to explore the dynamics of emotion dependencies in utterance interaction networks and leverages the prompts generated by the global information vector of the utterance to guide the dynamic fusion of multimodal features. This allows our model to dynamically change parameters when processing each utterance feature, so that different network parameters can be equipped for different emotion categories in the inference stage, thereby achiev-

ing more flexible emotion classification and enhancing the generalization ability of the model. Extensive experiments on two public multimodal conversation datasets demonstrate that the proposed DF-GCN model achieves superior performance and explicitly reveals the uncertainty of emotion variations through dynamics fusion. To the best of our knowledge, we make the first attempt to assign different weights to different emotion categories in the inference stage to achieve more effective information fusion.

### CRedit authorship contribution statement

**Tao Meng:** Conceptualization, Methodology, Investigation, Data curation, Writing – Original Draft; **Weilun Tang:** Supervision, Investigation & Review; **Yuntao Shou:** Supervision, Investigation, Writing – Review & Editing; **Yilong Tan:** Writing – Review & Editing; **Jun Zhou:** Supervision, Investigation, Writing – Review & Editing; **Wei Ai:** Supervision, Investigation, Writing – Review & Editing; **Keqin Li:** Supervision, Investigation, Writing – Review & Editing.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Data availability

Data will be made available on request.

### Acknowledgement

The authors deepest gratitude goes to the anonymous reviewers and AE for their careful work and thoughtful suggestions that have helped improve this paper substantially. This work is supported by [National Natural Science Foundation of China](#) (Grant No. 69189338), the [Hunan Provincial Natural Science Foundation General Project](#) (Grant No. 2025JJ50380), and the [Hunan Provincial Natural Science Foundation Youth Project](#) (Grant No. 2025JJ60420).

### References

- [1] H. Ma, J. Wang, H. Lin, B. Zhang, Y. Zhang, B. Xu, A transformer-based model with self-distillation for multimodal emotion recognition in conversations, *IEEE Trans. Multimedia* 26 (2023) 776–788.
- [2] W. Lu, Z. Hu, J. Lin, L. Wang, LECM: a model leveraging emotion cause to improve real-time emotion recognition in conversations, *Knowl. Based Syst.* 309 (2025) 112900.
- [3] S. Wei, X. Huang, C. Zhang, Emotional inverse reasoning trees and dominant modality focus for emotion recognition in conversations, *Knowl. Based Syst.* 326 (2025) 114035.
- [4] M.L.B. Estrada, R.Z. Cabada, R.O. Bustillos, M. Graff, Opinion mining and emotion recognition applied to learning environments, *Expert Syst. Appl* 150 (2020) 113265.
- [5] F. Fu, W. Ai, F. Yang, Y. Shou, T. Meng, K. Li, SDR-GNN: spectral domain reconstruction graph neural network for incomplete multimodal learning in conversational emotion recognition, *Knowl. Based Syst.* 309 (2025) 112825.
- [6] F.A. Pujol, H. Mora, A. Martínez, Emotion recognition to improve e-healthcare systems in smart cities, in: *Research & Innovation Forum 2019: Technology, Innovation, Education, and Their Social Impact 1*, Springer, 2019, pp. 245–254.
- [7] L. Zhou, J. Gao, D. Li, H.-Y. Shum, The design and implementation of xiaoice, an empathetic social chatbot, *Computat. Ling.* 46 (1) (2020) 53–93.
- [8] S.K. Khare, V. Blanes-Vidal, E.S. Nadimi, U.R. Acharya, Emotion recognition and artificial intelligence: a systematic review (2014–2023) and research recommendations, *Inf. Fusion* 102 (2024) 102019.
- [9] V. Chudasama, P. Kar, A. Gudmalwar, N. Shah, P. Wasnik, N. Onoe, M2FFNet: multimodal fusion network for emotion recognition in conversation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4652–4661.
- [10] S. Liu, P. Gao, Y. Li, W. Fu, W. Ding, Multi-modal fusion network with complementarity and importance for emotion recognition, *Inf. Sci.* 619 (2023) 679–694.
- [11] G. Tu, T. Xie, B. Liang, H. Wang, R. Xu, Adaptive graph learning for multimodal conversational emotion detection, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 38, 2024, pp. 19089–19097.

- [12] Z. Yang, X. Li, Y. Cheng, T. Zhang, X. Wang, Emotion recognition in conversation based on a dynamic complementary graph convolutional network, *IEEE Trans. Affect. Comput. Syst.* 15 (2024) 1567–1579.
- [13] Z. Yi, Z. Zhao, Z. Shen, T. Zhang, Multimodal fusion via hypergraph autoencoder and contrastive learning for emotion recognition in conversation, in: *Proceedings of the 32nd ACM International Conference on Multimedia*, 2024, pp. 4341–4348.
- [14] J.-B. Delbrouck, N. Tits, M. Brousmiche, S. Dupont, A transformer-based joint-encoding for emotion recognition and sentiment analysis, *ACL*, 2020, pp. 1–7.
- [15] W. Ai, Y. Shou, T. Meng, K. Li, DER-GCN: dialog and event relation-aware graph convolutional neural network for multimodal dialog emotion recognition, *IEEE Trans. Neural Netw. Learn. Syst.* 36 (2024) 4908–4921.
- [16] F. Chen, J. Shao, S. Zhu, H.T. Shen, Multivariate, multi-frequency and multimodal: rethinking graph neural networks for emotion RRecognition in conversation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10761–10770.
- [17] A. Saxena, A. Khanna, D. Gupta, Emotion recognition and detection methods: a comprehensive survey, *J. Artif. Intell. Syst.* 2 (1) (2020) 53–79.
- [18] J. Zhang, Z. Yin, P. Chen, S. Nichele, Emotion recognition using multi-modal data and machine learning techniques: a tutorial and review, *Inf. Fusion* 59 (2020) 103–126.
- [19] N. Majumder, S. Poria, D. Hazarika, R. Mihalcea, A. Gelbukh, E. Cambria, DialogueRNN: an attentive RNN for emotion detection in conversations, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 2019, pp. 6818–6825.
- [20] D. Kollias, S. Zafeiriou, Exploiting multi-cnn features in cnn-rnn based dimensional emotion recognition on the omg in-the-wild dataset, *IEEE Trans. Affect. Comput.* 12 (3) (2020) 595–606.
- [21] S. Zou, X. Huang, X. Shen, H. Liu, Improving multimodal fusion with main modal transformer for emotion recognition in conversation, *Knowl Based. Syst.* 258 (2022) 109978.
- [22] Z. Wang, Q. Zhang, P. Zhang, W. Niu, K. Zhang, R. Sankaranarayanan, S. Caldwell, T. Gedeon, Visual and textual prompts in vlms for enhancing emotion recognition, *IEEE Trans. Circuits Syst. Video Technol.* 35 (2025) 12355–12368.
- [23] D. Ghosal, N. Majumder, S. Poria, N. Chhaya, A. Gelbukh, DialogueGcn: a graph convolutional neural network for emotion recognition in conversation, (2019) arXiv:1908.11540.
- [24] W. Nie, M. Ren, J. Nie, S. Zhao, C-GCN: Correlation based graph convolutional network for audio-video emotion recognition, *IEEE Trans. Multimedia* 23 (2020) 3793–3804.
- [25] Y. Fu, S. Okada, L. Wang, L. Guo, Y. Song, J. Liu, J. Dang, Context-and knowledge-aware graph convolutional network for multimodal emotion recognition, *IEEE Multimedia* 29 (3) (2022) 91–100.
- [26] Y. Niu, Y. Pu, Z. Yang, X. Li, T. Zhou, J. Ren, S. Hu, H. Li, Y. Liu, Lightzero: a unified benchmark for monte carlo tree search in general sequential decision scenarios, *Adv. Neural Inf. Process. Syst.* 36 (2024) 37594–37635.
- [27] W. Guo, S. Li, J. Yang, Scattering prompt tuning: a fine-tuned foundation model for SAR object recognition, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 3056–3065.
- [28] K. Zhou, J. Yang, C.C. Loy, Z. Liu, Learning to prompt for vision-language models, *Int. J. Comput. Vis.* 130 (9) (2022) 2337–2348.
- [29] X.L. Li, P. Liang, Prefix-tuning: optimizing continuous prompts for generation, 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)*, 2021, pp. 4582–4597.
- [30] M. Jia, L. Tang, B.-C. Chen, C. Cardie, S. Belongie, B. Hariharan, S.-N. Lim, Visual prompt tuning, in: *European Conference on Computer Vision*, Springer, 2022, pp. 709–727.
- [31] M.U. Khattak, H. Rasheed, M. Maaz, S. Khan, F.S. Khan, Maple: multi-modal prompt learning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 19113–19122.
- [32] J.S. Smith, L. Karlinsky, V. Gutta, P. Cascante-Bonilla, D. Kim, A. Arbelle, R. Panda, R. Feris, Z. Kira, Coda-prompt: continual decomposed attention-based prompting for rehearsal-free continual learning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 11909–11919.
- [33] Z. Wang, Z. Zhang, S. Ebrahimi, R. Sun, H. Zhang, C.-Y. Lee, X. Ren, G. Su, V. Perot, J. Dy, et al., Dualprompt: complementary prompting for rehearsal-free continual learning, in: *European Conference on Computer Vision*, Springer, 2022, pp. 631–648.
- [34] Y. He, S. Zheng, Y. Tay, J. Gupta, Y. Du, V. Aribandi, Z. Zhao, Y. Li, Z. Chen, D. Metzler, et al., Hyperprompt: prompt-based task-conditioning of transformers, in: *International Conference on Machine Learning*, PMLR, 2022, pp. 8678–8690.
- [35] Z. Wang, K. Zhang, W. Luo, R. Sankaranarayanan, Htnet for micro-expression recognition, *Neurocomputing* 602 (2024) 128196.
- [36] Q. Zhang, Z. Wang, D. Zhang, W. Niu, S. Caldwell, T. Gedeon, Y. Liu, Z. Qin, Visual prompting in llms for enhancing emotion recognition, (2024) arXiv:2410.02244.
- [37] R.T.Q. Chen, Y. Rubanova, J. Bettencourt, D.K. Duvenaud, Neural ordinary differential equations, *Adv. Neural Inf. Process. Syst.* 31 (2018) 1–13.
- [38] L.-P. Xhonneux, M. Qu, J. Tang, Continuous graph neural networks, in: *International Conference on Machine Learning*, PMLR, 2020, pp. 10432–10441.
- [39] T. Kim, P. Vossen, Emoberta: speaker-aware emotion recognition in conversation with roberta, (2021) arXiv:2108.12009.
- [40] F. Eyben, F. Weninger, F. Gross, B. Schuller, Recent developments in opensmile, the munich open-source multimedia feature extractor, in: *Proceedings of the 21st ACM International Conference on Multimedia*, 2013, pp. 835–838.
- [41] G. Huang, Z. Liu, L. Van Der Maaten, K.Q. Weinberger, Densely connected convolutional networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4700–4708.
- [42] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J.N. Chang, S. Lee, S.S. Narayanan, IEMOCAP: interactive emotional dyadic motion capture database, *Lang. Resour. Eval.* 42 (2008) 335–359.
- [43] P.S. Kamath, W.R. Kim, The model for end-stage liver disease (MELD), *Hepatology* 45 (3) (2007) 797–805.
- [44] Q. Li, D. Gkoumas, A. Sordani, J.-Y. Nie, M. Melucci, Quantum-inspired neural network for conversational emotion recognition, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 35, 2021, pp. 13270–13278.
- [45] S. Xing, S. Mai, H. Hu, Adapted dynamic memory network for emotion recognition in conversation, *IEEE Trans. Affect. Comput.* 13 (3) (2020) 1426–1439.
- [46] J. Lee, W. Lee, CoMPM: context modeling with speaker’s pre-trained memory tracking for emotion recognition in conversation, (2021) arXiv:2108.11626.
- [47] T. Meng, F. Zhang, Y. Shou, H. Shao, W. Ai, K. Li, Masked graph learning with recurrent alignment for multimodal emotion recognition in conversation, *IEEE/ACM Trans. Audio Speech Lang. Process.* 32 (2024) 4298–4312.
- [48] J. Hu, Y. Liu, J. Zhao, Q. Jin, MMGCN: multimodal fusion via deep graph convolution network for emotion recognition in conversation, (2021) arXiv:2107.06779.
- [49] T. Ishiwatari, Y. Yasuda, T. Miyazaki, J. Goto, Relation-aware graph attention networks with relational position encodings for emotion recognition in conversations, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 7360–7370.
- [50] T. Meng, Y. Shou, W. Ai, N. Yin, K. Li, Deep imbalanced learning for multimodal emotion recognition in conversations, *IEEE Trans. Artif. Intell.* 5 (2024) 6472–6487.
- [51] Z. Lian, B. Liu, J. Tao, CTNet: conversational transformer network for emotion recognition, *IEEE/ACM Trans. Audio Speech Lang. Process.* 29 (2021) 985–1000.
- [52] M. Li, J. Shi, L. Bai, C. Huang, Y. Jiang, K. Lu, S. Wang, E.R. Hancock, FrameERC: framelet transform based multimodal graph neural networks for emotion recognition in conversation, *Pattern Recognit.* 161 (2025) 111340.
- [53] C. Fu, F. Qian, K. Su, Y. Su, Z. Wang, J. Shi, Z. Liu, C. Liu, C.T. Ishi, Himul-LGG: a hierarchical decision fusion-based local–global graph neural network for multimodal emotion recognition in conversation, *Neural Netw.* 181 (2025) 106764.