



UNIVERSITY OF HELSINKI



<https://helda.helsinki.fi>

Helda

---

## A Bipartite Graph is All We Need for Enhancing Emotional Reasoning with Commonsense Knowledge

Yang, Kailai

2023-10-21

---

Yang, K, Zhang, T, Ji, S & Ananiadou, S 2023, A Bipartite Graph is All We Need for Enhancing Emotional Reasoning with Commonsense Knowledge. in Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. ACM Digital library, New York, pp. 2917-2927, ACM International Conference on Information and Knowledge Management, Birmingham, United Kingdom, 21/10/2023. <https://doi.org/10.1145/3583780.3614758>

---

<http://hdl.handle.net/10138/567746>

10.1145/3583780.3614758

---

cc\_by

publishedVersion

---

*Downloaded from Helda, University of Helsinki institutional repository.*

*This is an electronic reprint of the original article.*

*This reprint may differ from the original in pagination and typographic detail.*

*Please cite the original version.*



# A Bipartite Graph is All We Need for Enhancing Emotional Reasoning with Commonsense Knowledge

Kailai Yang  
The University of Manchester  
Manchester, United Kingdom  
kailai.yang@postgrad.manchester.ac.uk

Tianlin Zhang  
The University of Manchester  
Manchester, United Kingdom  
tianlin.zhang@postgrad.manchester.ac.uk

Shaoxiong Ji  
University of Helsinki  
Helsinki, Finland  
shaoxiong.ji@helsinki.fi

Sophia Ananiadou  
The University of Manchester  
Manchester, United Kingdom  
sophia.ananiadou@manchester.ac.uk

## ABSTRACT

The context-aware emotional reasoning ability of AI systems, especially in conversations, is of vital importance in applications such as online opinion mining from social media and empathetic dialogue systems. Due to the implicit nature of conveying emotions in many scenarios, commonsense knowledge is widely utilized to enrich utterance semantics and enhance conversation modeling. However, most previous knowledge infusion methods perform empirical knowledge filtering and design highly customized architectures for knowledge interaction with the utterances, which can discard useful knowledge aspects and limit their generalizability to different knowledge sources. Based on these observations, we propose a **Bipartite Heterogeneous Graph (BHG)** method for enhancing emotional reasoning with commonsense knowledge. In BHG, the extracted context-aware utterance representations and knowledge representations are modeled as heterogeneous nodes. Two more knowledge aggregation node types are proposed to perform automatic knowledge filtering and interaction. BHG-based knowledge infusion can be directly generalized to multi-type and multi-grained knowledge sources. In addition, we propose a **Multi-dimensional Heterogeneous Graph Transformer (MHGT)** to perform graph reasoning, which can retain unchanged feature spaces and unequal dimensions for heterogeneous node types during inference to prevent unnecessary loss of information. Experiments show that BHG-based methods significantly outperform state-of-the-art knowledge infusion methods and show generalized knowledge infusion ability with higher efficiency. Further analysis proves that previous empirical knowledge filtering methods do not guarantee to provide the most useful knowledge information. Our code is available at: <https://github.com/SteveKGYang/BHG>.

## KEYWORDS

emotion recognition in conversations, casual emotion entailment, knowledge infusion, bipartite heterogeneous graph

## ACM Reference Format:

Kailai Yang, Tianlin Zhang, Shaoxiong Ji, and Sophia Ananiadou. 2023. A Bipartite Graph is All We Need for Enhancing Emotional Reasoning with Commonsense Knowledge. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583780.3614758>

## 1 INTRODUCTION

Understanding human emotions is at the core of affective computing. In natural language processing, recent years have witnessed growing research interests in machine's context-aware emotional reasoning ability, especially in conversations, due to its vital importance in scenarios such as empathetic dialogue systems [17] and online opinion mining from social media [4]. This goal is mostly specified as recognizing the emotion [23] or the emotion cause [21] of certain utterances within a conversation.

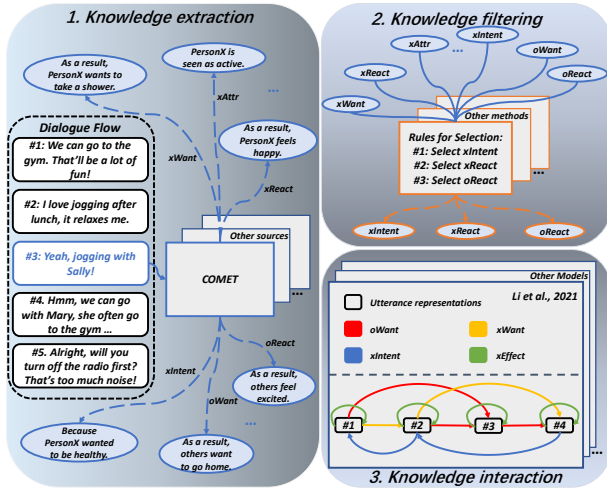
There are two key challenges for enhancing conversational emotional reasoning. Firstly, the emotion of the target speaker is influenced by both his own mental state and other participants' behaviors. Current methods mainly build conversation models [27, 28, 38] based on Pre-trained Language Models (PLMs) [15, 41] to tackle these dependencies. Secondly, emotions are often conveyed implicitly with metaphor, sarcasm, and underlying common sense. To mine the related information, a mainstream solution infuses commonsense knowledge to provide emotional clues and help model the inter-utterance relations, which mostly follows a three-step paradigm. We provide an illustration in Figure 1 and summarize this paradigm as follows: Firstly, **knowledge extraction** obtains commonsense knowledge items with the target conversation as queries. This process is closely dependent on the granularity of the queries and the characteristics of the knowledge source. As shown in the example, with the utterance-level query "Yeah, jogging with Sally!" and the generative knowledge source COMET [2], we can extract sentence-level social commonsense knowledge such as "PersonX is seen as active" and "As a result, others feel excited". Secondly, **knowledge filtering** selects the most relevant knowledge aspects from the extracted knowledge items based on task-specific priors. The most widely adopted methods include rule-based [6, 44, 47] and distantly supervised [13, 37] filtering. Finally, **knowledge interaction** introduces the filtered knowledge to the conversation model via customized architectures to enhance its emotional reasoning



This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0124-5/23/10.  
<https://doi.org/10.1145/3583780.3614758>

ability. For example, a typical architecture [12, 13] assigns knowledge aspects to the edges of the conversation graphs and utilizes the corresponding knowledge features as edge representations to model the inter-utterance dependencies.



**Figure 1: Illustration of the three-step paradigm of knowledge infusion. An example based on the knowledge source COMET is provided for each step.**

In this paper, we dive into the knowledge filtering and knowledge interaction steps and raise the following research questions (RQ):

- **RQ1:** Most knowledge filtering methods select knowledge aspects empirically without further evaluations. What knowledge aspects are most effective in enhancing conversational emotional reasoning?
- **RQ2:** Current knowledge interaction methods are highly coupled with both the knowledge sources and conversation models. Are these customized architectures necessary?

In addressing these research questions, we propose a simple yet effective **Bipartite Heterogeneous Graph (BHG)** method for enhancing emotional reasoning with commonsense knowledge. Firstly, we extract context-aware utterance representations via a PLM-based conversation model and relevant commonsense knowledge from three multi-type and multi-grained knowledge sources. Considering the complementary nature [40] between utterance semantics and related common sense, we model the extracted utterance and knowledge representations as heterogeneous nodes. In addition, we introduce a forward and a backward knowledge aggregation node type to perform automatic knowledge filtering and knowledge interaction, because for each target utterance, effective knowledge aspects for modeling inter-utterance relations are usually different in the past context and future context along the dialogue flow, as shown in previous works [12, 13, 44]. Then a bipartite graph is built on these four heterogeneous node types, where messages from the utterance and knowledge are passed to the knowledge aggregation nodes for semantic-aware knowledge filtering, and the aggregated knowledge messages interact with the utterance nodes to enrich their semantics and model inter-utterance dependencies. The BHG is decoupled from the conversation models and knowledge

sources and enables a unified model architecture for multi-type and multi-grained knowledge infusion. Its simple bipartite structure also facilitates the graph reasoning process.

In a BHG, heterogeneous nodes usually possess different feature spaces with unequal dimensions as conversation models/knowledge sources change. During the graph reasoning process, existing heterogeneous graph neural networks [8, 26, 32] mostly project all types of nodes into a unified feature space to facilitate the interaction between neighbors. However, the projections disrupt the original feature spaces and can lead to unnecessary loss of information for high-dimensional node types. Based on the heterogeneous graph Transformer [8], we propose a **Multi-dimensional Heterogeneous Graph Transformer (MHGT)** for graph reasoning, which utilizes a multi-dimensional edge-dependent matrix to enable direct attention calculation and message passing between heterogeneous nodes with different feature dimensions. MHGT allows all node types to retain the original feature spaces and potentially useful information for the knowledge filtering and knowledge interaction processes during inference.

We evaluate the effectiveness of our proposed BHG and MHGT methods on five datasets across two conversational emotional reasoning tasks: **Emotion Recognition in Conversations (ERC)** and **Casual Emotion Entailment (CEE)**. The experimental results show that the BHG-based methods outperform previous state-of-the-art knowledge infusion models, and show generalized knowledge infusion ability on multi-type and multi-grained knowledge sources with higher efficiency than previous customized methods. We also analyze the effectiveness of different knowledge aspects, and the results show that previous empirical knowledge filtering methods can introduce less useful knowledge aspects and discard knowledge aspects that benefit the emotional reasoning process.

In summary, this paper makes the following contributions: (1) we propose a bipartite heterogeneous graph-based method to enhance emotional reasoning with commonsense knowledge, which enables a unified framework for multi-type and multi-grained knowledge infusion; (2) we propose a multi-dimensional heterogeneous graph Transformer for graph reasoning, which allows unchanged feature spaces and unequal dimensions for heterogeneous node representations during inference; (3) the BHG-based methods significantly outperform state-of-the-art knowledge infusion methods and show generalized knowledge infusion ability with higher efficiency.

## 2 METHODOLOGY

We introduce the utterance and knowledge feature extraction process in Sec. 2.1. Then the BHG construction process and the multi-dimensional HGT methods are introduced in Sec. 2.2 and 2.3. Finally, the examined tasks and prediction process are described in Sec. 2.4.

### 2.1 Feature Extraction

In this section, we extract the utterance features via a PLM-based conversation model. We also extract multi-type (generative and extractive knowledge) and multi-grained (utterance-level and phrase-level) commonsense knowledge to enhance the emotional reasoning process. These processes are illustrated in Fig. 2 (a).

**2.1.1 Conversation Model.** We decouple conversation modeling from knowledge infusion by utilizing existing conversation models

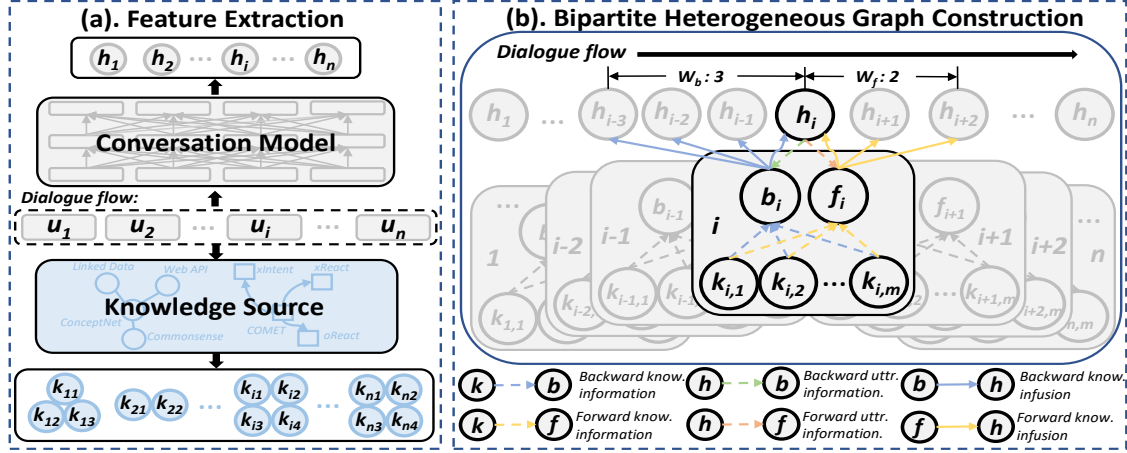


Figure 2: An overview of the utterance/knowledge feature extraction and BHG construction processes. In (b), the graph construction process for the  $i$ -th utterance is presented.  $W_f = 2$  and  $W_b = 3$  are examples of the forward and backward knowledge infusion window sizes. *know.* and *uttr.* denote "knowledge" and "utterance".

to obtain context-aware utterance representations. We utilize a PLM-based conversation model. For a conversation  $\mathcal{D} = [u_1, u_2, \dots, u_i, \dots, u_N]$  with  $N$  utterances,  $u_i$  is the target utterance pre-pended with its speaker. The model concatenates both past and future contexts as the input  $\hat{u}_i = \parallel_{j \in [1, N]} u_j$ , where  $\parallel$  is the concatenation operation. Considering situations where only the dialogue history is available, we also test our methods on only past contexts:  $\hat{u}_i = \parallel_{j \in [1, i]} u_j$ . We use a RoBERTa-Large [15] encoder to model the dialogue:

$$\hat{h}_i = \text{RoBERTa}(\hat{u}_i) \quad (1)$$

where  $\hat{h}_i \in \mathbb{R}^{N_i \times D_h}$  denotes the token-level representations,  $N_i$  is the token number of  $\hat{u}_i$ , and  $D_h$  denotes the dimension of hidden states. We further obtain the utterance-level context-aware representation  $h_j$  for each utterance  $u_j$  within the dialogue flow via mean-pooling:

$$h_j = \text{MeanPooling}(\hat{h}_i[l_j : r_j]) \quad (2)$$

where  $l_j$  and  $r_j$  denote the start and end positions of  $u_j$  in  $\hat{u}_i$ .

**2.1.2 Knowledge Extraction.** In previous works, two commonsense knowledge graphs have been proven most effective in enhancing the emotional reasoning process: the social commonsense knowledge graph ATOMIC [25] and the taxonomic/lexical knowledge graph ConceptNet [29]. Therefore, we separately obtain relevant knowledge from three knowledge sources expanded from ATOMIC and ConceptNet.

With the development of automatic knowledge graph construction, generative knowledge sources receive increasing interest due to their flexibility and convenience in knowledge extraction. We extract the ATOMIC knowledge via two COMET models: the first one is COMET<sub>2019</sub> [2], which pre-trains a GPT [24] model on ATOMIC. During knowledge extraction, each utterance  $u_i$  is constructed into the query:  $(u_i \parallel a_j \parallel [GEN])$ .  $a_j$  denotes a knowledge aspects term representing an if-then relation type for the speaker's actions/mental states. COMET<sub>2019</sub> provides nine knowledge aspects from ATOMIC, where their interpretations are listed in Table 1. The

constructed query is then input to the COMET model, and the final-layer hidden representation of the decoder:  $k_{ij} \in \mathbb{R}^{D_k^{2019}}$  is used as the utterance-level knowledge representation corresponding to  $a_j$ , where  $D_k^{2019}$  denotes the dimension of  $k_{ij}$ .

Table 1: Interpretations of the selected knowledge aspects/relations. "X" denotes the target utterance speaker.

| COMET              |  |
|--------------------|--|
| Aspect             | Interpretation                                       |
| <i>xIntent</i>     | Why does X cause the event?                          |
| <i>xAttr</i>       | How would X be described?                            |
| <i>xNeed</i>       | What does X need to do before the event?             |
| <i>xWant</i>       | What would X likely want to do after the event?      |
| <i>xEffect</i>     | What effects does the event have on X?               |
| <i>oReact</i>      | How does X feel after the event?                     |
| <i>oWant</i>       | What would others likely want to do after the event? |
| <i>oEffect</i>     | What effects does the event have on others?          |
| <i>oReact</i>      | How do others feel after the event?                  |
| ConceptNet         |  |
| Relation           | Interpretation                                       |
| <i>IsA</i>         | is a   |
| <i>HasProperty</i> | has the property of                                  |
| <i>DefinedAs</i>   | is defined as  |
| <i>AtLocation</i>  | is located at  |
| <i>RelatedTo</i>   | is related to  |

The second knowledge source COMET<sub>2020</sub> [9] further extends the generative pre-training to an expanded ATOMIC knowledge graph and part of ConceptNet based on a larger BART [11] model. COMET<sub>2020</sub> follows a similar knowledge extraction process as COMET<sub>2019</sub> and each knowledge representation has dimension  $D_k^{2020}$ .

On the other hand, we obtain phrase-level ConceptNet knowledge in an extractive manner. Specifically, we tokenize each utterance  $u_i$  and concatenate the tokens into n-gram phrases. For each phrase, we extract all its immediate neighbors in the English sub-graph of ConceptNet. Each neighbor assertion is a  $\langle source, relation, target, weight \rangle$  quadruple where *source* and *target* denote the query phrase and the neighbor concept, *relation* denotes the corresponding relation type, and *weight* denotes a confidence score assigned to the assertion. For example, we can extract the assertion:  $\langle think, HasPrerequisite, use\ brains, 2.375 \rangle$  for the input phrase *think*. There are 28 different relation types in the extracted quadruples. We further remove all assertions with confidence scores less than 2.0 for denoising. To facilitate the extraction of the knowledge representations, we manually design an interpretation for each relation type and convert all assertions into natural language by concatenating the concepts to the relation interpretation. For example,  $\langle think, HasPrerequisite, use\ brains \rangle$  is converted to "think has the prerequisite of use brains". We present the interpretations of the five most common relations in Table 1. As most of these knowledge items have simple semantics, we use a RoBERTa-Base encoder to extract the features for the converted natural language knowledge and each knowledge representation has dimension  $D_k^{Concept}$ .

Overall, for each utterance  $u_i$  and a knowledge source, we obtain a set of knowledge representations:  $\{k_{i1}, k_{i2}, \dots, k_{im_i}\}$ , where  $k_{ij} \in \mathbb{R}^{D_k}$ , the representation dimension  $D_k \in \{D_k^{2019}, D_k^{2020}, D_k^{Concept}\}$  depends on the knowledge source, and  $m_i$  denote the number of extracted knowledge items for  $u_i$ .

## 2.2 Bipartite Heterogeneous Graph

We propose a directed bipartite heterogeneous graph (BHG) to complement utterance representations with extracted commonsense knowledge representations. Formally, the BHG is denoted as  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{N}, \mathcal{R})$ , where each node  $v \in \mathcal{V}$  and each edge  $e \in \mathcal{E}$ .  $\mathcal{N}, \mathcal{R}$  denote the set of node and relation types, and each node or edge is projected to its type via a mapping function:  $\gamma(v) : \mathcal{V} \rightarrow \mathcal{N}$ ,  $\tau(e) : \mathcal{E} \rightarrow \mathcal{R}$ , respectively. Specifically, the node and edge types are defined as follows:

**2.2.1 Node Types.** Firstly, each utterance within a conversation is modeled as a node in the BHG. Its representation obtained from the conversation model is used as the node feature. The utterance node type is denoted as  $h$ . The extracted knowledge representations are modeled as another type of node:  $k$ . In addition, we introduce extra knowledge aggregation node types to perform automatic knowledge filtering and knowledge interaction. As previous works have shown that effective knowledge aspects can be different for performing knowledge interactions with the past context and future context [12, 13, 44], we introduce a forward and a backward aggregation node type, which is separately denoted as  $f \in \mathbb{R}^{D_f}$  and  $b \in \mathbb{R}^{D_b}$ , where  $D_f$  and  $D_b$  are their representation dimensions. Overall, we define a set with four types of nodes:  $\mathcal{N} = \{h, k, f, b\}$ .

**2.2.2 Relation Types.** For each edge  $e = \langle v_1, v_2 \rangle$  from the source node  $v_1$  to the target node  $v_2$ , its relation type is defined as  $\tau(e) : \langle \gamma(v_1), \gamma(v_2) \rangle$ . Specifically, we introduce six types of relations to perform knowledge filtering and knowledge interaction. Firstly, the forward/backward knowledge information relations:  $\tau_{kf} : \langle k, f \rangle$  and

$\tau_{kb} : \langle k, b \rangle$  are proposed to separately provide the extracted knowledge information to both forward and backward knowledge aggregation nodes. Secondly, the knowledge filtering process naturally requires considering the corresponding utterance information. Therefore, we utilize the utterance information relations:  $\tau_{uf} : \langle h, f \rangle$  and  $\tau_{ub} : \langle h, b \rangle$  to introduce utterance semantics to the knowledge aggregation nodes. In addition, commonsense knowledge has been proven not only useful for enriching the semantics of each utterance but in modeling inter-utterance dependencies [35, 44, 45]. Therefore, we further incorporate the filtered commonsense knowledge to the utterance nodes by designing a forward infusion relation:  $\tau_f : \langle f, h \rangle$  and a backward infusion relation:  $\tau_b : \langle b, h \rangle$ . Overall, the relation set contains six types of relations:  $\mathcal{R} = \{\tau_{kf}, \tau_{kb}, \tau_{uf}, \tau_{ub}, \tau_f, \tau_b\}$ .

**2.2.3 BHG Construction.** For each utterance node  $h_i$ , we create a forward aggregation node  $f_i$  and a backward aggregation node  $b_i$ . Then we build a BHG for each conversation by considering the following criteria for each relation type:

- $\forall i \neq j$  and  $m, \langle k_{im}, f_j \rangle \notin \mathcal{E}$  and  $\langle k_{im}, b_j \rangle \notin \mathcal{E}$ . Each knowledge aggregation node only receives and filters the extracted knowledge from its corresponding utterance.
- $\forall i \neq j, \langle h_i, f_j \rangle \notin \mathcal{E}$  and  $\langle h_i, b_j \rangle \notin \mathcal{E}$ . Each knowledge aggregation node only considers the information from its corresponding utterance when performing knowledge filtering.
- $\forall i \leq j \leq i + W_f, \langle f_i, h_j \rangle \in \mathcal{E}$ , where  $W_f$  is the pre-defined forward knowledge infusion window size. Each forward aggregated knowledge is only used to enrich the semantics of its own utterance and model the inter-utterance relations with future utterances.
- $\forall i - W_b \leq j \leq i, \langle b_i, h_j \rangle \in \mathcal{E}$ , where  $W_b$  is the pre-defined backward knowledge infusion window size. Each backward aggregated knowledge is only used to enrich the semantics of its own utterance and model the inter-utterance relations with previous utterances.

Based on these criteria, the BHG building process is described in Algorithm 1, and an example is illustrated in Figure 2(b).

## 2.3 Multi-dimensional HGT

In heterogeneous graphs, different types of nodes/relations do not share feature spaces, which requires node and relation-dependent architectures for GNN-based graph modeling. The Heterogeneous Graph Transformer (HGT) [8] separately initializes and optimizes a set of Transformer [31] parameters for each homogeneous sub-graph. In addition, HGT unifies the dimensions of all node types before aggregation via linear transformation due to the complicated structures of most large-scale heterogeneous graphs. In our case, the utterance and knowledge representations often have unequal dimensions:  $D_h \neq D_k$  with the substitution of conversation models and knowledge sources. Considering the simple structure of BHG, we believe unifying dimensions can lead to unnecessary loss of information for high-dimensional node types. Therefore, we propose a Multi-dimensional HGT (MHGT) to model the BHG, which allows the dimensions of heterogeneous nodes to remain unchanged during the attention calculation and message-passing processes.

Based on the structure of vanilla HGT, MHGT improves the heterogeneous dot-product attention process to enable interactions

**Algorithm 1** Construction of the BHG

---

**Input:** dialogue representations:  $\mathcal{H} = [h_1, h_2, \dots, h_n]$ , knowledge representations:  $\mathcal{K} = [k_{11}, \dots, k_{1m_1}, \dots, k_{n1}, \dots, k_{nm_n}]$ , forward and backward knowledge infusion window sizes:  $W_f, W_b$

**Output:** the BHG  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{N}, \mathcal{R})$

- 1:  $\mathcal{V} \leftarrow \mathcal{H} \cup \mathcal{K}, \mathcal{E} \leftarrow \emptyset$  ▷ Initialize the BHG.
- 2:  $\mathcal{N} \leftarrow \{h, k, f, b\}$
- 3:  $\mathcal{R} \leftarrow \{\tau_{kf}, \tau_{kb}, \tau_{uf}, \tau_{ub}, \tau_f, \tau_b\}$
- 4: **for**  $i \in \{1, \dots, n\}$  **do**
- 5:    $\mathcal{V} \leftarrow \mathcal{V} \cup \{f_i, b_i\}$  ▷ Add the aggregation nodes.
- 6:   **for**  $j \in \{1, \dots, m_i\}$  **do**
- 7:      $\mathcal{E} \leftarrow \mathcal{E} \cup \{(k_{ij}, f_i)\}$  ▷ Add knowledge info. relations.
- 8:      $\mathcal{E} \leftarrow \mathcal{E} \cup \{(k_{ij}, b_i)\}$
- 9:   **end for**
- 10:    $\mathcal{E} \leftarrow \mathcal{E} \cup \{(h_i, f_i)\}$  ▷ Add utterance info. relations.
- 11:    $\mathcal{E} \leftarrow \mathcal{E} \cup \{(h_i, b_i)\}$
- 12:    $c_f = 0$
- 13:   **while**  $c_f \leq W_f$  and  $i + c_f \leq n$  **do**
- 14:      $\mathcal{E} \leftarrow \mathcal{E} \cup \{(f_i, h_{i+c_f})\}$  ▷ Forward knowledge infusion.
- 15:      $c_f \leftarrow c_f + 1$
- 16:   **end while**
- 17:    $c_b = 0$
- 18:   **while**  $c_b \leq W_b$  and  $i - c_b \geq 0$  **do**
- 19:      $\mathcal{E} \leftarrow \mathcal{E} \cup \{(b_i, h_{i-c_b})\}$  ▷ Backward knowledge infusion.
- 20:      $c_b \leftarrow c_b + 1$
- 21:   **end while**
- 22: **end for**

---

between nodes with different dimensions. For a target node representation  $v_t^l \in \mathbb{R}^{D_{\gamma(v_t)}}$  at the  $l$ -th MHGT layer, we calculate the multi-head attention weights for its neighbor  $v_n^l \in \mathbb{R}^{D_{\gamma(v_n)}}$  as follows:

$$Att\_head^i(v_n^l, v_t^l) = \frac{K_{\gamma(v_n)}^i(v_n^l) W_{\tau((v_n, v_t))}^{Att} Q_{\gamma(v_t)}^i(v_t^l)}{\sqrt{D_{\gamma(v_t)}}} \quad (3)$$

$$Att(v_n^l, v_t^l) = Softmax_{(v_n, v_t) \in \mathcal{E}} \left( \parallel Att\_head^i(v_n^l, v_t^l) \right) \quad (4)$$

where  $K_{\gamma(v_n)}^i(v_n) \in \mathbb{R}^{\frac{D_{\gamma(v_n)}}{h}}$  and  $Q_{\gamma(v_t)}^i(v_t) \in \mathbb{R}^{\frac{D_{\gamma(v_t)}}{h}}$  denotes the  $i$ -th head of a equal-dimensional linear projection from  $v_n^l$  as the key and  $v_t^l$  as the query,  $\parallel$  denotes concatenation, and  $Softmax$  denotes the softmax operation. Different from the vanilla HGT, MHGT proposes a multi-dimensional edge-dependent matrix  $W_{\tau((v_n, v_t))}^{Att} \in \mathbb{R}^{\frac{D_{\gamma(v_n)}}{h} \times \frac{D_{\gamma(v_t)}}{h}}$  to project nodes of  $\gamma(v_n)$  to the representation space of  $\gamma(v_t)$ , which enables attention calculations between heterogeneous node features and allows situations where  $D_{\gamma(v_n)} \neq D_{\gamma(v_t)}$ . We perform the message-passing process for  $v_n^l$  with the target node  $v_t^l$  as follows:

$$Mes\_head^i(v_n^l, v_t^l) = V_{\gamma(v_n)}^i(v_n^l) W_{\tau((v_n, v_t))}^{Mes} \quad (5)$$

$$Mes(v_n^l, v_t^l) = \parallel_{i \in [1, h]} Mes\_head^i(v_n^l, v_t^l) \quad (6)$$

where  $V_{\gamma(v_n)}(v_n)$  is another equal-dimensional linear projection from  $v_n$  as the value. Similarly, MHGT uses a multi-dimensional

edge-dependent matrix  $W_{\tau((v_n, v_t))}^{Mes} \in \mathbb{R}^{\frac{D_{\gamma(v_n)}}{h} \times \frac{D_{\gamma(v_t)}}{h}}$  to align the neighbor node  $v_n$  to the feature space of the target node  $v_t$ . Same as in vanilla dot-product attention, the message is aggregated using the calculated attention weights:

$$\hat{v}_t^{l+1} = \bigoplus_{(v_n, v_t) \in \mathcal{E}} \left( Att(v_n^l, v_t^l) \cdot Mes(v_n^l, v_t^l) \right) \quad (7)$$

where  $\bigoplus$  denotes the element-wise sum operation. Finally,  $\hat{v}_t^{l+1}$  is used to update  $v_t^l$  in the following manner:

$$v_t^{l+1} = T_{\gamma(v_t)}(\sigma(\hat{v}_t^{l+1})) + v_t^l \quad (8)$$

where  $\sigma$  denotes the *Gelu* activation function and  $T_{\gamma(v_t)}$  denotes an equal-dimensional linear projection from  $\gamma(v_t)$ . We stack  $L$  MHGT layers to allow interactions between non-adjacent nodes. For utterance  $u_i$ , we obtain the  $L$ -th layer output from MHGT:  $h_i^L$  as its final knowledge-enhanced representation.

## 2.4 Prediction and Training

We examine the BHG-based knowledge infusion method on two complex emotional reasoning tasks: Emotion Recognition in Conversations (ERC) and Casual Emotion Entailment (CEE).

**2.4.1 Emotion Recognition in Conversations.** ERC aims to identify each utterance  $u_i$ 's emotion within a dialogue  $\mathcal{D}$  from a pre-defined emotion category set  $E$ , which is modeled as a text classification task on each utterance [23]. Specifically, we utilize a feed-forward neural network to project the knowledge-enhanced utterance representations to the classification space:

$$\hat{y}_i = Softmax(h_i^L W_{ERC} + b_{ERC}) \quad (9)$$

where  $W_{ERC} \in \mathbb{R}^{D_h \times |E|}$  and  $b_{ERC} \in \mathbb{R}^{|E|}$  are learnable parameters. We optimize the standard cross-entropy loss to train the ERC model:

$$\mathcal{L}_{ERC} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{|E|} y_{ij} \log(\hat{y}_{ij}) \quad (10)$$

where  $\hat{y}_{ij}$  and  $y_{ij}$  denote the  $j$ -th element of  $\hat{y}_i$  and the one-hot emotion label  $y_i$  of utterance  $u_i$ , and  $N$  denotes the batch size.

**2.4.2 Casual Emotion Entailment.** CEE aims to identify the causes behind the non-neutral emotion of target utterances and locate their positions from the conversational history. Given a dialogue  $\mathcal{D}$  and the emotion label  $y$  of each utterance, CEE is modeled as a binary classification task to predict whether the candidate utterance  $u_j$  contains the emotion cause for target utterance  $u_i$ , where  $1 \leq j \leq i$ . During inference, we concatenate the knowledge-enhanced utterance representations of  $u_j$  and  $u_i$  to calculate the logits:

$$\hat{z}_{ji} = \frac{1}{1 + e^{-([h_j^L; h_i^L] W_{CEE} + b_{CEE})}} \quad (11)$$

where  $W_{CEE} \in \mathbb{R}^{2D_h \times 1}$  and  $b_{CEE} \in \mathbb{R}^1$  are learnable parameters, respectively. A BCE loss is used to incorporate the supervision signals:

$$\mathcal{L}_{BCE} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^i z_{ji} \log(\hat{z}_{ji}) \quad (12)$$

where  $z_{ji} \in \{0, 1\}$  is the binary label. We further introduce the emotion labels to enhance the learning of the CEE model using the

same prediction and training paradigm as in Eqn. 9 and Eqn. 10, and jointly optimize the two tasks in a multi-task learning manner:

$$\mathcal{L}_{CEE} = \mathcal{L}_{BCE} + \alpha \mathcal{L}_{ERC} \quad (13)$$

where  $\alpha$  denotes a hyper-parameter controlling the weight of the ERC loss.

## 3 EXPERIMENTS

### 3.1 Datasets

We test our method on four Emotion Recognition in Conversations (ERC) datasets and one Causal Emotion Entailment (CEE) dataset. For all datasets, we only utilize the text modality in our experiments.

**IEMOCAP** [3]: A two-party multi-modal ERC dataset derived from the scenarios in the scripts of the two actors. The pre-defined emotion category set  $E$  consists of: *neutral, sad, anger, happy, frustrated, excited*.

**MELD** [20]: A multi-party multi-modal ERC dataset collected from the scripts of American TV show *Friends*. The pre-defined emotions are *neutral, sad, anger, disgust, fear, happy, surprise*.

**DailyDialog** [14]: A ERC dataset compiled from human-written daily conversations with only two parties involved and no speaker information. The pre-defined emotion labels are *neutral, happy, surprise, sad, anger, disgust, fear*.

**EmoryNLP** [42]: Another ERC dataset collected from TV show *Friends*. It is annotated with the following emotion categories: *neutral, sad, mad, scared, powerful, peaceful, joyful*.

**RECCON** [22]: A CEE dataset collected from the scripts of DailyDialog with both utterance-level emotion labels and binary emotion cause labels, with the same emotion category set as DailyDialog.

### 3.2 Baseline Models

We compare our method with strong baselines models and categorize them into four groups according to their characteristics:

*PLM-based methods.* For the ERC task, we select the following two methods: **RoBERTa-Large** [15] used the PLM RoBERTa-Large to directly model the conversation. The utterance representations are used to fine-tune the weights. **DialogXL** [27] improved the XLNet [41] with the enhanced memory and dialog-aware self-attention mechanism to capture long historical context and dependencies between multiple parties. For the CEE task, we select: **RoBERTa-Base/Large** [22] concatenated the conversation as input to the PLM RoBERTa. Then CEE was modeled as a binary classification problem for each utterance pair.

*Graph-based Methods.* For ERC, **RGAT** [10] improved the relation modeling of conversations and added relational position encodings as sequential information. **DAG-ERC** [28] built a directed acyclic graph on the dialogue and used a graph neural network to aggregate the information. For CEE, we select two methods: **ECPE-2D** [5] represented the emotion-cause pairs as 2D representations and utilized the Transformer to model them. **RankCP** [33] ranked the clause pairs and performed end-to-end extraction with inter-clause modeling.

*Knowledge-based Methods.* For ERC, five methods are selected: **KI-Net** [35] infused both commonsense and sentiment lexicon knowledge and proposed a self-matching module to enhance the

knowledge interaction. **COSMIC** [6] used the RNN to model the dialogue history and extracted utterance-level commonsense knowledge to model the speakers’ mental states. **TODKAT** [47] modeled topic information via PLMs and explicitly infused event-centered knowledge. **SKAIG** [12] extracted psychological commonsense knowledge and infused the knowledge to enhance the edge representations of the knowledge graph. **CauAIN** [45] used the emotion cause knowledge to guide the traceback process of context modeling. For CEE, **KAG** [36] utilized the entity-related commonsense knowledge to model the semantic dependencies between the candidates and emotions. **AKM** [30] combined knowledge via an adapted knowledge model in a multi-task learning manner. **KEC** [13] utilized the directed acyclic graph incorporating social commonsense knowledge to improve the causal reasoning ability. **KBCIN** [44] proposed the knowledge-bridged causal interaction network to capture context dependencies of conversations and make emotional cause reasoning.

*Zero-shot Method with ChatGPT.* We also include the zero-shot evaluation results of the latest large language model ChatGPT<sup>1</sup> on all datasets, provided by Yang et al. [37].

### 3.3 Implementation and Evaluation Settings

We conduct all experiments using a single Nvidia Tesla A100 GPU with 80GB of memory. We initialize the pre-trained weights of RoBERTa and use the tokenization tools provided by Huggingface [34]. We leverage AdamW optimizer [16] to train the model. The batch size of experiments on all datasets is set to 16 except in DailyDialog, which is 24. We use a linear warm-up learning rate scheduling of warm-up ratio 20% and a peak learning rate 1e-5. We set a dropout rate of 0.1 and an L2-regularisation rate of 0.01 to avoid over-fitting. For hyper-parameters,  $D_h = D_k^{2020} = 1024$ ,  $D_k^{2019} = D_k^{Concept} = 768$ ,  $L = 3$ , and  $\alpha$  is set to 0.8. For each knowledge source with dimension  $D_k$ , we set  $D_f = D_b = D_k$  and randomly initialize the forward/backward knowledge aggregation node representations. During knowledge infusion, we set  $W_f = W_b = 5$  when both past and future contexts are provided, and  $W_f = 0, W_b = 10$  when only the past context is available.

For ERC, we select the Weighted F1 score as the evaluation metric for IEMOCAP, MELD, and EmoryNLP. Since “neutral” occupies most of DailyDialog, we utilize the Micro F1 score excluding the “neutral” utterances to reflect the performances in non-neutral emotions, as in previous works [12, 28, 47]. We also calculate Macro F1 scores on all classes for DailyDialog to evaluate the overall performances. For CEE, we report the F1 scores on positive and negative utterances and Macro F1 scores as the overall evaluation. All reported results are averages of five random runs.

## 4 RESULTS AND ANALYSIS

### 4.1 Main Results

*4.1.1 Overall Performance.* The performance of our BHG-based methods and all baseline models on the five datasets are presented in Table 2. Firstly, the zero-shot prompting results show that ChatGPT still bears a huge gap with advanced conversation models

<sup>1</sup><https://openai.com/blog/chatgpt>

**Table 2: Test results of our BHG-based methods and baseline models on the five datasets. "Know. Source" lists the commonsense knowledge sources used by each knowledge-based method. In "w/o future context", only the dialogue history is introduced as context for each target utterance. The "w/o emotions" setting removes the utterance emotion information for CEE. Best values are highlighted in bold.**

| ERC                            |                             | IEMOCAP      | MELD         | DailyDialog  |              | EmoryNLP     | CEE                        |                             | RECCON       |              |              |
|--------------------------------|-----------------------------|--------------|--------------|--------------|--------------|--------------|----------------------------|-----------------------------|--------------|--------------|--------------|
| Model                          | Know. Source                | Weighted F1  | Weighted F1  | Micro F1     | Macro F1     | Weighted F1  | Model                      | Know. Source                | Neg. F1      | Pos. F1      | Macro F1     |
| <b>Zero-shot Methods</b>       |                             |              |              |              |              |              |                            |                             |              |              |              |
| ChatGPT <sub>ZS</sub> [37]     | –                           | 53.35        | 61.18        | 43.27        | 38.19        | 32.64        | ChatGPT <sub>ZS</sub> [37] | –                           | 67.18        | 51.35        | 59.26        |
| <b>PLM-based Methods</b>       |                             |              |              |              |              |              |                            |                             |              |              |              |
| RoBERTa-Large [15]             | –                           | 55.67        | 62.75        | 55.16        | 48.2         | 37.0         | RoBERTa-Base [22]          | –                           | 88.74        | 64.28        | 76.51        |
| DialogXL [27]                  | –                           | 65.94        | 62.41        | 54.93        | –            | 34.73        | RoBERTa-Large [22]         | –                           | 87.89        | 66.23        | 77.06        |
| <b>Graph-based Methods</b>     |                             |              |              |              |              |              |                            |                             |              |              |              |
| RGAT [10]                      | –                           | 66.36        | 62.8         | 59.02        | –            | 34.42        | ECPE-2D [5]                | –                           | 94.96        | 55.50        | 75.23        |
| DAG-ERC [28]                   | –                           | 68.03        | 63.65        | 59.33        | –            | 39.02        | RankCP [33]                | –                           | <b>97.30</b> | 33.00        | 65.15        |
| <b>Knowledge-based Methods</b> |                             |              |              |              |              |              |                            |                             |              |              |              |
| CauAIN [45]                    | <i>COMET<sub>2019</sub></i> | 67.61        | 65.46        | 58.21        | 53.85        | –            | KAG [36]                   | <i>ConceptNet</i>           | 86.35        | 58.18        | 72.26        |
| KI-Net [35]                    | <i>ConceptNet</i>           | 66.98        | 63.24        | 57.30        | 50.8         | –            | AKM [30]                   | <i>COMET<sub>2019</sub></i> | 88.18        | 64.53        | 76.36        |
| COSMIC [6]                     | <i>COMET<sub>2019</sub></i> | 65.28        | 65.21        | 58.48        | 51.05        | 38.11        | KEC [13]                   | <i>COMET<sub>2020</sub></i> | 88.85        | 66.55        | 77.70        |
| TODKAT [47]                    | <i>COMET<sub>2019</sub></i> | 61.33        | 65.47        | 58.47        | 52.56        | 38.69        | KBCIN [44]                 | –                           | –            | 67.51        | 78.43        |
| SKAIG [12]                     | <i>COMET<sub>2019</sub></i> | 66.96        | 65.18        | 59.75        | 51.95        | 38.88        | w/o emotions               | <i>COMET<sub>2020</sub></i> | –            | 64.05        | 76.73        |
| <b>BHG-based Methods</b>       |                             |              |              |              |              |              |                            |                             |              |              |              |
| MHGT                           | <i>COMET<sub>2020</sub></i> | <b>71.20</b> | 65.54        | <b>62.37</b> | <b>54.11</b> | <b>39.06</b> | MHGT                       | <i>COMET<sub>2020</sub></i> | 90.34        | <b>69.13</b> | <b>79.73</b> |
| w/o future context             |                             | 70.93        | <b>66.3</b>  | 60.41        | 53.29        | 38.85        | w/o emotions               |                             | 89.89        | 68.58        | 79.23        |
| MHGT                           |                             | 71.07        | 65.57        | 61.5         | 53.0         | 38.62        | MHGT                       |                             | 90.26        | 68.85        | 79.55        |
| w/o future context             | <i>COMET<sub>2019</sub></i> | 70.64        | 65.77        | 60.68        | 52.51        | 38.33        | w/o emotions               | <i>COMET<sub>2019</sub></i> | 89.95        | 68.49        | 79.22        |
| HGT                            |                             | <b>67.2</b>  | <b>62.86</b> | 60.0         | 51.27        | 36.5         | HGT                        |                             | 89.79        | 62.94        | 76.37        |
| MHGT                           |                             | 70.55        | 65.99        | 61.41        | 52.7         | 38.07        | MHGT                       |                             | 89.53        | 68.6         | 79.07        |
| w/o future context             | <i>ConceptNet</i>           | 70.68        | 65.59        | 60.2         | 51.74        | 37.38        | w/o emotions               | <i>ConceptNet</i>           | 90.03        | 67.72        | 78.88        |
| HGT                            |                             | 67.1         | <b>64.29</b> | 59.28        | 51.09        | 36.06        | HGT                        |                             | 90.12        | 62.47        | 76.29        |

and knowledge-based methods in performing emotional reasoning, possibly because these tasks are very subjective even to humans, showing the necessity of exploring few-shot prompting and knowledge infusion to further calibrate ChatGPT’s understanding of these subjective emotion concepts [37]. These results also motivate continual research on supervised task-specific methods. Secondly, the knowledge-based methods significantly improve model performance on most datasets compared to the PLM-based and graph-based conversation models. These results empirically prove the effectiveness of commonsense knowledge infusion to emotional reasoning tasks. Thirdly, the BHG-based methods outperform all baseline models on all ERC and CEE datasets, including the new state-of-the-art performance of 71.2% on IEMOCAP, 62.37% (Micro-F1) on DailyDialog, and 79.73% (Macro-F1) on RECCON. These outstanding performances quantify the advantages of BHG-based knowledge infusion over previous methods.

**4.1.2 Knowledge-based Methods Comparisons.** In the comparison of knowledge-based methods, the unified BHG architecture shows generalized knowledge infusion ability by outperforming all previous highly customized knowledge infusion methods on all three tested knowledge sources. For the phrase-level extractive knowledge source ConceptNet, BHG outperforms KI-Net by at least 2% on IEMOCAP, MELD, and DailyDialog. It also has an impressive 6.81% improvement on RECCON compared to KAG. These results show that the BHG structure can effectively interact with multi-grained heterogeneous nodes to infuse knowledge. For the utterance-level generative knowledge source *COMET<sub>2019</sub>*, BHG possesses a similar advantage over other knowledge-based methods on most datasets, especially on IEMOCAP, DailyDialog, and RECCON. A possible reason is that some useful knowledge is mistakenly discarded by

these previous methods as they all perform empirical knowledge filtering. On the other hand, we provide all knowledge aspects to construct the BHG and perform knowledge filtering automatically, which enables the aggregation nodes to retain all useful knowledge. The knowledge source *COMET<sub>2020</sub>* is proven most effective in enhancing emotional reasoning as it outperforms all other knowledge sources on all five datasets. Compared to *COMET<sub>2019</sub>*, *COMET<sub>2020</sub>* trains a larger language model on a bigger social commonsense knowledge graph. Therefore, *COMET<sub>2020</sub>* is expected to generate more reliable knowledge for each utterance. These results show the importance of high-quality commonsense knowledge sources.

**4.1.3 BHG Variants.** We further compare the BHG-based methods in experimental settings such as "w/o future context" and "w/o emotions" to test model performance in scenarios where future contexts and emotion labels are unavailable. For ERC, the BHG’s performance drops with only past context to a limited extent on most datasets. These results show that future dialogue can provide useful clues for the emotional reasoning of the current utterance in most cases. For CEE, the BHG performances on all knowledge resources drop less than 0.5% without the emotional supervision signals, which is much less than the 1.7% decrease on the previous state-of-the-art model KBCIN. We believe the appropriately infused commonsense knowledge from BHG can make up for the missing emotional information introduced by the labels. Finally, we compare the performance of MHGT and the vanilla HGT in modeling the BHG. During experiments for HGT, we unify the heterogeneous node dimensions by linearly projecting the high-dimension representations to low-dimensional spaces. According to the results, HGT significantly underperforms MHGT on all datasets. For example, we observe an over 3% drop in Weighted-F1 for both

*COMET*<sub>2019</sub> and *ConceptNet* on IEMOCAP. MHGT retains the original semantic space for utterance representations, which preserves useful information to perform semantic-aware knowledge filtering and emotional reasoning, while HGT projects the representations to low-dimensional spaces and causes unnecessary loss of information.

**Table 3: Ablation studies for the BHG and other knowledge-based baseline models. "know." denotes the whole knowledge infusion architectures. "backward/forward aggr." denotes the backward/forward knowledge aggregation designs in BHG. We highlight top-2 performance drops in bold.**

| Model              | IEMOCAP                | MELD                   | DailyDialog            | EmoryNLP               |
|--------------------|------------------------|------------------------|------------------------|------------------------|
| CauAIn             | 67.61                  | 65.46                  | 58.21                  | –                      |
| w/o know.          | 63.77 ( <b>↓3.84</b> ) | 65.2 ( <b>↓0.26</b> )  | 57.2 ( <b>↓1.01</b> )  | –                      |
| COSMIC             | 65.28                  | 65.21                  | 58.48                  | 38.11                  |
| w/o know.          | 63.05 ( <b>↓2.23</b> ) | 64.28 ( <b>↓0.93</b> ) | 56.16 ( <b>↓2.32</b> ) | 37.10 ( <b>↓1.01</b> ) |
| TODKAT             | 61.33                  | 65.47                  | 54.62                  | 38.69                  |
| w/o know.          | 58.96 ( <b>↓2.37</b> ) | 63.97 ( <b>↓1.5</b> )  | 53.44 ( <b>↓1.18</b> ) | 37.1 ( <b>↓1.59</b> )  |
| SKAIG              | 66.96                  | 65.18                  | 59.75                  | 38.88                  |
| w/o know.          | 64.28 ( <b>↓2.68</b> ) | –                      | 58.86 ( <b>↓0.89</b> ) | –                      |
| BHG                | 71.07                  | 65.57                  | 61.5                   | 38.62                  |
| w/o backward aggr. | 68.54 ( <b>↓2.53</b> ) | 64.81 ( <b>↓0.76</b> ) | 60.18 ( <b>↓1.32</b> ) | 38.07 ( <b>↓0.55</b> ) |
| w/o forward aggr.  | 68.07 ( <b>↓3.0</b> )  | 64.41 ( <b>↓1.16</b> ) | 59.3 ( <b>↓2.3</b> )   | 36.95 ( <b>↓1.67</b> ) |
| w/o know.          | 67.93 ( <b>↓3.14</b> ) | 63.7 ( <b>↓1.87</b> )  | 58.76 ( <b>↓2.74</b> ) | 36.26 ( <b>↓2.36</b> ) |

## 4.2 Ablation Studies for Knowledge Infusion

To further investigate the BHG’s efficiency in knowledge infusion and the contributions of its components, we perform ablation studies on the BHG and other knowledge-based baseline methods, and the results are presented in Table 3. To ensure fair comparisons, we select the BHG and four knowledge infusion methods that are all tested with the *COMET*<sub>2019</sub> knowledge source on the four ERC datasets.

**4.2.1 Knowledge Infusion Efficiency.** Firstly, we compare the performance drops between the models when the whole knowledge infusion architecture is removed, where more significant drops reflect the higher efficiency of the method in leveraging the same knowledge source. According to the results, "BHG w/o know." achieves top-2 performance drops on all datasets, and the drops exceed 2% on three out of four datasets, while other knowledge infusion methods normally make substantial contributions to only one or two datasets. These results show that our unified BHG method can utilize the same knowledge source to enhance emotional reasoning on differently distributed data more efficiently than previous customized knowledge infusion methods.

**4.2.2 BHG Modules.** We further investigate the backward/forward knowledge infusion architectures in BHG by removing these modules and comparing the performance. According to the results, "BHG w/o backward aggr." and "BHG w/o forward aggr." both perform worse than BHG on all datasets. These results not only strengthen that infusing commonsense knowledge can enhance the modeling of inter-utterance relations, but further prove that a unified knowledge infusion module can be decoupled from the utterance modeling process and used for any conversation models in a


plug-in manner. They also prove the necessity of splitting the knowledge filtering and interaction processes for previous and future contexts in the BHG architecture. In addition, "BHG w/o forward aggr." suffers from higher performance drops on all datasets than "BHG w/o backward aggr.", which shows that social commonsense knowledge is more useful in modeling inter-utterance relations for future contexts. These observations correspond with the widely recognized prior that the current utterance has more significant influences on future utterances along the dialogue flow [28].

## 4.3 Knowledge Filtering Analysis

To provide an intuitive view of the knowledge filtering process, we record the dot-product attention weights between the knowledge nodes and knowledge aggregation nodes in the last MHGT layer for *COMET*<sub>2019</sub> on IEMOCAP test set and visualize their distributions in box plots. The results are presented in Figure 3. Following the BHG structure, we split the attention weights from forward and backward knowledge filtering, and higher weight distributions reflect more contributions from the corresponding knowledge aspect.

**4.3.1 Overall Analysis.** For forward knowledge filtering, the most attended knowledge aspect *xWant* has a median number over 0.3. *xIntent* and *xEffect* are also frequently utilized, with their median numbers over 0.2. These results show that forward knowledge interaction can benefit more from knowledge aspects reflecting the potential effect of the target utterance on the target speaker. On the other hand, backward knowledge filtering pays attention to the knowledge aspects reflecting the effect on both the target speaker and other participants. For example, *xEffect* and *oWant* are highly attended with median numbers over 0.3, and *xNeed* and *oEffect* also have median numbers of over 0.2. In comparisons of the box plots for forward and backward knowledge filtering, the BHG has very different patterns for selecting knowledge aspects. For example, *xWant* knowledge is considered useful in forward knowledge filtering but receives much less attention in backward knowledge filtering. These observations further show the necessity of designing forward and backward knowledge aggregation nodes in BHG. In addition, some knowledge aspects such as *xAttr* and *oReact* have low attention distributions in both forward and backward knowledge filtering, possibly because these knowledge aspects can provide less relevant information for emotional reasoning.

**4.3.2 Comparison with Previous Methods.** Previous knowledge-based methods mostly perform knowledge filtering manually in the utilization of *COMET*<sub>2019</sub>, where frequently utilized knowledge aspects include *xWant*, *xEffect*, *xIntent*, *oReact*, *oWant*, and *oEffect* [6, 12, 45, 47]. Most of these knowledge aspects are also proven effective in the automatic knowledge filtering of BHG, while some aspects such as *xReact* and *oReact* are less attended in both forward and backward knowledge filtering. These results indicate that some widely used knowledge aspects may not provide much useful information as expected. On the other hand, some previously ignored knowledge aspects are assigned higher attention scores in automatic knowledge filtering. For example, *xNeed*, a less-used knowledge aspect in previous works, receives considerable attention in backward knowledge filtering. We expect the above analysis to guide the adjustment of future strategies for empirical methods.



**Figure 3: Box plots of the knowledge filtering attention weights for nine knowledge aspects on the IEMOCAP test set. We use *COMET*<sub>2019</sub> as the knowledge source and MHGT as the BHG encoder. Orange lines denote the median numbers.**

## 5 RELATED WORK

We discuss relevant background on conversational emotion reasoning, including conversation models and knowledge-based methods.

### 5.1 Conversation Models

Emotion reasoning in a conversation naturally requires modeling interactions between dialogue participants, known as intra- and inter-speaker dependencies [27]. Early works regarded dialogues as temporal flows and utilized RNNs to model the dialogue history and emotional dynamics of each dialogue participant [6, 18]. In other works, Transformer variants were also leveraged to model long-range dependencies in conversations [43, 46]. Recent works mostly relied on PLMs such as RoBERTa [15] and XLNet [41] to obtain utterance-level or conversation-level features [1, 27]. Some other works [7, 10, 28] modeled utterances as nodes and carefully designed graph structures on the dialogue to enable efficient message passing among utterances during the graph aggregation process. Other self-supervised architectures such as the Variational Autoencoder [19, 39] were also used for modeling conversations or discourse information.

### 5.2 Commonsense Knowledge-based Methods

Due to the implicitness of emotional expression in many scenarios, commonsense knowledge has been widely utilized to enhance emotional reasoning. One line of works [35, 43, 46] extracted phrase-level concepts from the large-scale knowledge graph ConceptNet [29] to concatenate them with the token-level utterance representations, and Transformer-based models were used to perform utterance-knowledge interactions. More recent works leveraged the utterance-level knowledge from the generative knowledge source COMET [2, 9]. Some methods directly infused the knowledge into utterance representations [6, 47], followed by neural network-based knowledge infusion modules. Other works utilized

the knowledge as edge representations in the dialogue graphs to model inter-utterance relations between utterance nodes [12, 13, 44] or trace emotion casual clues [45], which obtained outstanding performance in both emotion detection [12] and emotion casual detection tasks [13, 44]. In addition, the pre-trained knowledge adapters were also used for knowledge infusion into PLM-based conversation models [38].

## 6 CONCLUSION AND FUTURE WORK

This paper proposes a bipartite heterogeneous graph for enhancing emotional reasoning with commonsense knowledge. We model the utterance representations and knowledge representations as heterogeneous nodes and design a BHG for commonsense knowledge infusion. In addition, we propose a multi-dimensional heterogeneous graph Transformer to perform graph reasoning to retain unchanged feature spaces for heterogeneous node types. Experiments show that BHG-based methods outperform state-of-the-art knowledge infusion methods on five datasets across two conversational emotional reasoning tasks. The BHG also shows generalized knowledge infusion ability with higher efficiency. Further analysis proves that previous empirical knowledge filtering methods do not guarantee to provide the most useful knowledge information.

In future work, we will test our BHG-based method on more conversation models, such as graph-based models, to further examine their generalizability. We will also explore simultaneous knowledge infusion from multiple knowledge sources in a unified BHG framework, which enables the model to reason on several knowledge types to enhance its performance.

## ACKNOWLEDGMENTS

This work is supported in part by the New Energy and Industrial Technology Development Organization (NEDO) project under grant JPNP20006 and the University of Manchester President’s Doctoral Scholar award.

## REFERENCES

- [1] Yanan Bao, Qianwen Ma, Lingwei Wei, Wei Zhou, and Songlin Hu. 2022. Speaker-Guided Encoder-Decoder Framework for Emotion Recognition in Conversation. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, Luc De Raedt (Ed.). ijcai.org, 4051–4057. <https://doi.org/10.24963/ijcai.2022/562>
- [2] Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. COMET: Commonsense Transformers for Automatic Knowledge Graph Construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, 4762–4779. <https://doi.org/10.18653/v1/P19-1470>
- [3] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. 2008. IEMOCAP: interactive emotional dyadic motion capture database. *Lang. Resour. Evaluation* 42, 4 (2008), 335–359. <https://doi.org/10.1007/s10079-008-9076-6>
- [4] Ankush Chatterjee, Umang Gupta, Manoj Kumar Chinnakotla, Radhakrishnan Srikanth, Michel Galley, and Puneet Agrawal. 2019. Understanding Emotions in Text Using Deep Learning and Big Data. *Comput. Hum. Behav.* 93 (2019), 309–317. <https://doi.org/10.1016/j.chb.2018.12.029>
- [5] Zixiang Ding, Rui Xia, and Jianfei Yu. 2020. ECPE-2D: Emotion-Cause Pair Extraction based on Joint Two-Dimensional Representation, Interaction and Prediction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 3161–3170. <https://doi.org/10.18653/v1/2020.acl-main.288>
- [6] Deepanway Ghosal, Navonil Majumder, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. COSMIC: COMmonSense knowledge for eMotion Identification in Conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 2470–2481. <https://doi.org/10.18653/v1/2020.findings-emnlp.224>
- [7] 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)*. Association for Computational Linguistics, Hong Kong, China, 154–164. <https://doi.org/10.18653/v1/D19-1015>
- [8] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, Yennun Huang, Irwin King, Tie-Yan Liu, and Maarten van Steen (Eds.). ACM / IW3C2, 2704–2710. <https://doi.org/10.1145/3366423.3380027>
- [9] Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. (Comet-) Atomic 2020: On Symbolic and Neural Commonsense Knowledge Graphs. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*. AAAI Press, 6384–6392. <https://ojs.aaai.org/index.php/AAAI/article/view/16792>
- [10] Taichi Ishiwatari, Yuki Yasuda, Taro Miyazaki, and Jun Goto. 2020. 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)*. Association for Computational Linguistics, Online, 7360–7370. <https://doi.org/10.18653/v1/2020.emnlp-main.597>
- [11] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 7871–7880. <https://doi.org/10.18653/v1/2020.acl-main.703>
- [12] Jiangnan Li, Zheng Lin, Peng Fu, and Weiping Wang. 2021. Past, Present, and Future: Conversational Emotion Recognition through Structural Modeling of Psychological Knowledge. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, Punta Cana, Dominican Republic, 1204–1214. <https://doi.org/10.18653/v1/2021.findings-emnlp.104>
- [13] Jiangnan Li, Fandong Meng, Zheng Lin, Rui Liu, Peng Fu, Yanan Cao, Weiping Wang, and Jie Zhou. 2022. Neutral Utterances are Also Causes: Enhancing Conversational Causal Emotion Entailment with Social Commonsense Knowledge. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, Luc De Raedt (Ed.). ijcai.org, 4209–4215. <https://doi.org/10.24963/ijcai.2022/584>
- [14] Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Asian Federation of Natural Language Processing, Taipei, Taiwan, 986–995. <https://aclanthology.org/I17-1099>
- [15] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [16] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *ICLR*. OpenReview.net.
- [17] Yukun Ma, Khanh Linh Nguyen, Frank Z. Xing, and Erik Cambria. 2020. A survey on empathetic dialogue systems. *Inf. Fusion* 64 (2020), 50–70. <https://doi.org/10.1016/j.inffus.2020.06.011>
- [18] Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander F. Gelbukh, and Erik Cambria. 2019. DialogueRNN: An Attentive RNN for Emotion Detection in Conversations. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. AAAI Press, 6818–6825. <https://doi.org/10.1609/aaai.v33i01.33016818>
- [19] Donovan Ong, Jian Su, Bin Chen, Anh Tuan Luu, Ashok Narendranath, Yue Li, Shuqi Sun, Yingzhan Lin, and Haifeng Wang. 2022. Is discourse role important for emotion recognition in conversation?. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 11121–11129.
- [20] 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*. Association for Computational Linguistics, Florence, Italy, 527–536. <https://doi.org/10.18653/v1/P19-1050>
- [21] Soujanya Poria, Navonil Majumder, Devamanyu Hazarika, Deepanway Ghosal, Rishabh Bhardwaj, Samson Yu Bai Jian, Pengfei Hong, Romila Ghosh, Abhinaba Roy, Niyati Chhaya, Alexander F. Gelbukh, and Rada Mihalcea. 2021. Recognizing Emotion Cause in Conversations. *Cogn. Comput.* 13, 5 (2021), 1317–1332. <https://doi.org/10.1007/s12559-021-09925-7>
- [22] Soujanya Poria, Navonil Majumder, Devamanyu Hazarika, Deepanway Ghosal, Rishabh Bhardwaj, Samson Yu Bai Jian, Pengfei Hong, Romila Ghosh, Abhinaba Roy, Niyati Chhaya, Alexander F. Gelbukh, and Rada Mihalcea. 2021. Recognizing Emotion Cause in Conversations. *Cogn. Comput.* 13, 5 (2021), 1317–1332. <https://doi.org/10.1007/s12559-021-09925-7>
- [23] Soujanya Poria, Navonil Majumder, Rada Mihalcea, and Eduard H. Hovy. 2019. Emotion Recognition in Conversation: Research Challenges, Datasets, and Recent Advances. *IEEE Access* 7 (2019), 100943–100953. <https://doi.org/10.1109/ACCESS.2019.2929050>
- [24] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [25] Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 3027–3035.
- [26] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings 15*. Springer, 593–607.
- [27] Weizhou Shen, Junqing Chen, Xiaojun Quan, and Zhixian Xie. 2021. DialogXL: All-in-One XLNet for Multi-Party Conversation Emotion Recognition. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*. AAAI Press, 13789–13797. <https://ojs.aaai.org/index.php/AAAI/article/view/17625>
- [28] 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)*. Association for Computational Linguistics, Online, 1551–1560. <https://doi.org/10.18653/v1/2021.acl-long.123>
- [29] Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, Satinder Singh and Shaul Markovitch (Eds.). AAAI Press, 4444–4451. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14972>
- [30] Elsbeth Turcan, Shuai Wang, Rishita Anubhai, Kasturi Bhattacharjee, Yaser Al-Onaizan, and Smaranda Muresan. 2021. Multi-Task Learning and Adapted Knowledge Models for Emotion-Cause Extraction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Association for Computational Linguistics, Online, 3975–3989. <https://doi.org/10.18653/v1/2021.findings-acl.348>
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [32] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. 2019. Heterogeneous graph attention network. In *The world wide web conference. 2022–2032*.
- [33] Penghui Wei, Jiahao Zhao, and Wenji Mao. 2020. Effective Inter-Clause Modeling for End-to-End Emotion-Cause Pair Extraction. In *Proceedings of the 58th Annual*

- Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 3171–3181. <https://doi.org/10.18653/v1/2020.acl-main.289>
- [34] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771* (2019).
- [35] Yunhe Xie, Kailai Yang, Chengjie Sun, Bingquan Liu, and Zhenzhou Ji. 2021. Knowledge-Interactive Network with Sentiment Polarity Intensity-Aware Multi-Task Learning for Emotion Recognition in Conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, Punta Cana, Dominican Republic, 2879–2889. <https://doi.org/10.18653/v1/2021.findings-emnlp.245>
- [36] Hanqi Yan, Lin Gui, Gabriele Pergola, and Yulan He. 2021. Position Bias Mitigation: A Knowledge-Aware Graph Model for Emotion Cause Extraction. 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)*. Association for Computational Linguistics, Online, 3364–3375. <https://doi.org/10.18653/v1/2021.acl-long.261>
- [37] Kailai Yang, Shaoxiong Ji, Tianlin Zhang, Qianqian Xie, and Sophia Ananiadou. 2023. On the Evaluations of ChatGPT and Emotion-enhanced Prompting for Mental Health Analysis. *arXiv preprint arXiv:2304.03347* (2023).
- [38] Kailai Yang, Tianlin Zhang, Hassan Alhuzali, and Sophia Ananiadou. 2023. Cluster-Level Contrastive Learning for Emotion Recognition in Conversations. *IEEE Transactions on Affective Computing* (2023), 1–12. <https://doi.org/10.1109/TAFFC.2023.3243463>
- [39] Kailai Yang, Tianlin Zhang, and Sophia Ananiadou. 2023. Disentangled Variational Autoencoder for Emotion Recognition in Conversations. *IEEE Transactions on Affective Computing* (2023), 1–12. <https://doi.org/10.1109/TAFFC.2023.3280038>
- [40] Liang Yang, Fan Wu, Zichen Zheng, Bingxin Niu, Junhua Gu, Chuan Wang, Xiaochun Cao, and Yuanfang Guo. 2021. Heterogeneous Graph Information Bottleneck. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, Zhi-Hua Zhou (Ed.), ijcai.org, 1638–1645. <https://doi.org/10.24963/ijcai.2021/226>
- [41] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems* 32 (2019).
- [42] Sayyed M Zahiri and Jinho D Choi. 2017. Emotion detection on tv show transcripts with sequence-based convolutional neural networks. *arXiv preprint arXiv:1708.04299* (2017).
- [43] Duzhen Zhang, Xiuyi Chen, Shuang Xu, and Bo Xu. 2020. Knowledge Aware Emotion Recognition in Textual Conversations via Multi-Task Incremental Transformer. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 4429–4440. <https://doi.org/10.18653/v1/2020.coling-main.392>
- [44] Weixiang Zhao, Yanyan Zhao, Zhuojun Li, and Bing Qin. 2022. Knowledge-Bridged Causal Interaction Network for Causal Emotion Entailment. *arXiv preprint arXiv:2212.02995* (2022).
- [45] Weixiang Zhao, Yanyan Zhao, and Xin Lu. 2022. CauAIN: Causal Aware Interaction Network for Emotion Recognition in Conversations. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, Luc De Raedt (Ed.), ijcai.org, 4524–4530. <https://doi.org/10.24963/ijcai.2022/628>
- [46] Peixiang Zhong, Di Wang, and Chunyan Miao. 2019. Knowledge-Enriched Transformer for Emotion Detection in Textual Conversations. 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)*. Association for Computational Linguistics, Hong Kong, China, 165–176. <https://doi.org/10.18653/v1/D19-1016>
- [47] Lixing Zhu, Gabriele Pergola, Lin Gui, Deyu Zhou, and Yulan He. 2021. Topic-Driven and Knowledge-Aware Transformer for Dialogue Emotion Detection. 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)*. Association for Computational Linguistics, Online, 1571–1582. <https://doi.org/10.18653/v1/2021.acl-long.125>