

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

InferEM: Inferring the Speaker's Intention for Empathetic Dialogue Generation

#### **Permalink**

<https://escholarship.org/uc/item/0dk8x579>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

#### **Authors**

Lv, Guoqing

Li, Jiang

Wang, XiaoPing

et al.

#### **Publication Date**

2023

Peer reviewed

# InferEM: Inferring the Speaker’s Intention for Empathetic Dialogue Generation

Guoqing Lv<sup>1</sup>, Jiang Li<sup>1,2</sup>, Xiaoping Wang<sup>2</sup>, Zhigang Zeng

The School of Artificial Intelligence and Automation, and the Key

Laboratory of Image Processing and Intelligent Control of Education Ministry of China, and the Hubei Key

Laboratory of Brain-inspired Intelligent Systems, Huazhong University of Science and Technology, Wuhan 430074, China

{guoqinglv, lijfrank, wangxiaoping, zgzen}@hust.edu.cn

## Abstract

Current approaches to empathetic response generation typically encode the entire dialogue history directly and put the output into a decoder to generate friendly feedback. These methods focus on modelling contextual information but neglect capturing the direct intention of the speaker. We argue that the last utterance in the dialogue empirically conveys the intention of the speaker. Consequently, we propose a novel model named InferEM for empathetic response generation. We separately encode the last utterance and fuse it with the entire dialogue through the multi-head attention based intention fusion module to capture the speaker’s intention. Besides, we utilize previous utterances to predict the last utterance, which simulates human’s psychology to guess what the interlocutor may speak in advance. To balance the optimizing rates of the utterance prediction and response generation, a multi-task learning strategy is designed for InferEM. Experimental results demonstrate the plausibility and validity of InferEM in improving empathetic expression.

**Keywords:** empathetic response generation, multi-head attention, multi-task learning

## Introduction

Empathetic emotion, a reaction of one individual to the observed experiences of the interlocutor (Davis, 1983), is a fundamental trait of human beings. Empathy is a contentious and complex emotion that emerges when we respond to the needs or suffering from another. Empathy has been applied in a wide range of domains, such as volunteering (Batson et al., 1997), psychotherapy (Bohart & Greenberg, 1997), charitable giving (Pavey, Greitemeyer, & Sparks, 2012) and longevity (Poulin, Brown, Dillard, & Smith, 2013). Equipping AI systems with the ability to express empathy can improve the human-machine interaction in many fields such as mental health support (Sharma, Lin, Miner, Atkins, & Althoff, 2023) and social chatbot (Zhou, Gao, Li, & Shum, 2020).

Empathetic dialogue systems are designed to comprehend the user’s emotions and generate humanized responses, which facilitate enabling a better user experience (Liu et al., 2021). There have been a great deal of fruitful studies (Hu, Liu, Zhao, & Jin, 2021; J. Li, Wang, Lv, & Zeng, 2023) on emotion recognition. Recently, some attempts have been made to improve the quality of the generated empathetic response. MoEL (Lin, Madotto, Shin, Xu, & Fung, 2019) computes

the emotion distribution of the dialogue history and softly fuses the outputs of decoders, which is more interpretable than previous models and can select the correct decoder to effectively generate empathetic responses. MIME (Majumder et al., 2020) splits 32 emotions into positivity and negativity groups, and mimics the speaker’s emotion by appropriately mixing emotions from distinct groups. Gao et al. (2021) design an emotion reasoner to judge context emotion labels and emotion cause-oriented labels, which are employed through gated attention networks to enhance the generation of empathetic responses. EmpHi (Chen, Li, & Yang, 2022) generates human-like responses with empathetic intents by means of computing the empathetic intent distribution of responses. KEMP (Q. Li, Li, Ren, Ren, & Chen, 2022) and CEM (Sabour, Zheng, & Huang, 2022) leverage common-sense knowledge to enrich the dialogue history and generate more empathetic and informative responses. These approaches typically concatenate on encoding the representations of the entire dialogue history and recognizing the user’s emotion but don’t attach enough importance to capturing the direct intention of the speaker to generate a more targeted response.

Zhang, Li, Zhu, Zhao, and Liu (2018) point out that the last utterance in a dialogue plays a crucial function, which empirically conveys the intention of interlocutor while previous utterances describe the dialogue in different aspects. Garrod and Pickering (2009) assume that people can predict the interlocutor’s response to enhance mutual understanding and promote successful communication. Inspired by their opinions, we propose a novel model to Infer the speaker’s intention for EMpathetic dialogue generation (InferEM). The proposed model encodes separately the last utterance in the dialogue to capture the speaker’s intention, then the multi-head attention (Vaswani et al., 2017) is utilized to fuse the entire dialogue and the last utterance. Besides, InferEM leverages previous utterances to predict a virtual last utterance of the speaker, which is combined with the real one to improve the diversity of the speaker’s intention. We design a multi-task learning strategy for InferEM to better optimize the parameters of the utterance prediction module and the response generation module. Figure 1 is an example to illustrate how InferEM generates an empathetic response to the user based on the given dialogue history. The orange lines mean that the predicted Utterance<sub>3</sub> is generated based on Utterance<sub>1</sub> and

<sup>1</sup>Guoqing Lv and Jiang Li contributed equally to this work.

<sup>2</sup>Corresponding authors: Jiang Li and Xiaoping Wang.

Utterance<sub>2</sub>. The blue lines mean that the Response is generated based on all utterances in dialogue history and the predicted Utterance<sub>3</sub>. The red words imply the current intention of the user, so we separately encode the utterances containing these elements and offer them special attention.

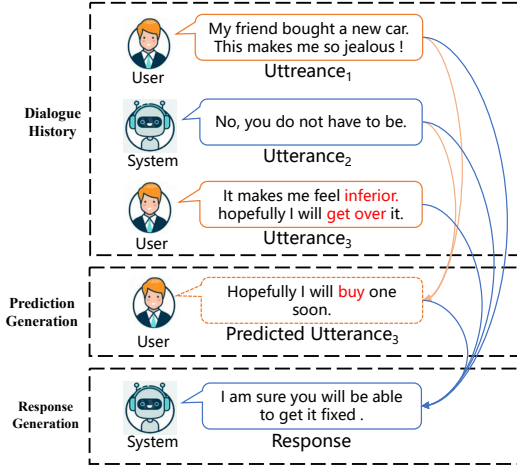


Figure 1: An example to illustrate how InferEM generates an empathetic response based on the dialogue history.

Our contributions are as follows: (1) We encode simultaneously the entire dialogue and the last utterance, then the multi-head attention based intention fusion network is adopted to integrate them. This process leads to a better understanding of the speaker’s intention, which allows the model to generate more reliable empathetic responses. (2) Aimed at promoting the diversity of the speaker’s intention and generated response, we add a prediction module that simulates humans guessing the speaker’s intention in advance based on the dialogue history. To the best of our knowledge, InferEM is the first attempt to leverage the predictive function to improve empathetic dialogue generation. (3) We tailor a multi-task learning strategy for InferEM to balance the optimizing rates of the prediction module and the response module. Extensive experiments on a benchmark dataset verify the effectiveness of our proposed method.

## Preliminaries

### Task Definition

The dialogue history can be denoted as  $C = [U_1, U_2, \dots, U_n]$ , where each  $U_i$  is the  $i$ -th utterance and  $U_i$  includes  $m_i$  words, denoted as  $U_i = [w_i^1, w_i^2, \dots, w_i^{m_i}]$ .  $[U_1, U_2, \dots, U_n]$  may be uttered by the speaker or the listener. The target task is to generate an empathetic response based on dialogue history  $C$ . There is a preparatory task in our model that generate a prediction of  $U_n$  according to the  $n - 1$  past utterances  $C' = [U_1, U_2, \dots, U_{n-1}]$ .

### Partial Notation

As external knowledge can promote emotional understanding and expression, InferEM is furnished with the same exter-

nal knowledge as KEMP (Q. Li et al., 2022): (1) we introduce emotion-related concepts (Speer, Chin, & Havasi, 2017) by constructing emotional context graph, and leverage both graph attention (Veličković et al., 2018) and Transformer encoder (Vaswani et al., 2017) to extract contextual semantic features. This process is called Emotion context encoder according to KEMP (see Figure 2a), which we simply denote as **EcEnc**; (2) we leverage emotion intensity values  $\eta$  provided by NRC\_VAD (Zhong, Wang, & Miao, 2019) to calculate the emotion signal representation.

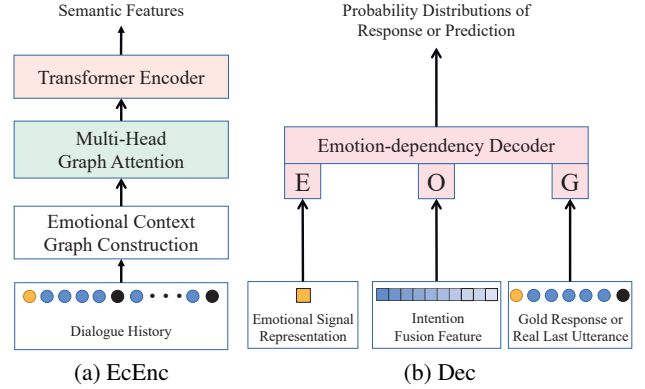


Figure 2: The illustration of partial components.

Referring to KEMP, the Emotion-dependency decoder (we simply denote as **Dec** ( $E, O, G$ )) is employed to generate the probability distributions of prediction or response, where  $E$  indicates the emotion signal representation,  $O$  is the semantic feature encoded by the Emotional context encoder, and  $G$  indicates the gold prediction or response during training and the decoded words during testing. The Emotion-dependency decoder is shown in Figure 2b.

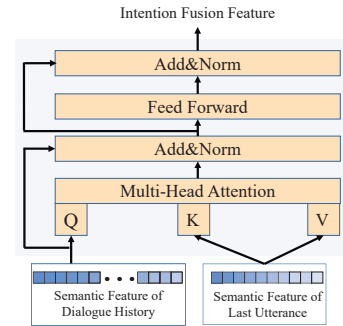


Figure 3: Multi-head attention based intent fusion network.

In our model, we adopt a Multi-head Attention based Intent Fusion Network (MAIFNet) by drawing on the structure of Transformer (Vaswani et al., 2017). MAIFNet mainly consists of a multi-head attention layer, a feedforward layer, two residual connections and two normalization layers, and the structure of MAIFNet is shown in Figure 3. MAIFNet are utilized to fuse the entire dialogue history and the last utter-

ance in the dialogue, which can perceive the intention of the person uttering the last utterance. We denote the multi-head attention as  $\mathbf{Att}(Q, K, V)$ , where  $Q, K$  and  $V$  indicate respectively the representations of query, key and value. The feed-forward layer consists of two fully connected layers and is indicated as  $\mathbf{FF}(X)$ , where  $X$  denotes the input of the feed-forward layer.

Following previous work (Wolf, Sanh, Chaumond, & De-langue, 2019), the embedding of the given dialogue history  $C$  is computed as:

$$E_C = E_w(C) + E_p(C) + E_s(C), \quad (1)$$

where  $E_w(C)$ ,  $E_p(C)$  and  $E_s(C)$  denote the word embedding, positional embedding and state embedding of dialogue history  $C$ .

## Methodology

The proposed InferEM mainly includes virtual last utterance prediction, attention based intention perception and empathetic dialogue generation, as shown in Figure 4. In the stage of last utterance prediction, the semantic features of  $C'$  (the  $n-1$  past utterances) and  $U_{n-1}$  (the  $(n-1)$ -th utterance) are extracted by **EcEnc**; on this basis, we fuse them by MAIFNet and then predict a virtual last utterance  $U_n^p$ . In the stage of intention perception, we first extract the semantic features of  $C$  (the  $n$  past utterances) and  $U_n$  (the  $n$ -th utterance) by **EcEnc**; then the semantic feature of  $U_n$  and that of the virtual one are concatenated as that of the final last utterance  $U_n^{pr}$  to enhance the diversity of the speaker’s intention; finally, the semantic feature of  $U_n^{pr}$  and that of the entire dialogue are fused by MAIFNet to perceive the intention of the speaker uttering the last utterance. And in the stage of empathetic dialogue generation, we feed the above result into **Dec** to generate an empathetic response.

### Virtual Last Utterance Prediction

To extract the semantic feature of the  $n-1$  past utterances  $C'$  and that of the  $(n-1)$ -th utterance  $U_{n-1}$ , we first obtain their embeddings (i.e.,  $E_{C'}$  and  $E_{U_{n-1}}$ ) and furnish them with emotion-related concepts according to previous work (Q. Li et al., 2022); then  $E_{C'}$  and  $E_{U_{n-1}}$  are fed to **EcEnc**:

$$\begin{aligned} S_{C'} &= \mathbf{EcEnc}_1(E_{C'}), \\ S_{U_{n-1}} &= \mathbf{EcEnc}_2(E_{U_{n-1}}), \end{aligned} \quad (2)$$

where  $E_{C'} \in \mathbb{R}^{q_{C'} \times f}$ ,  $q_{C'}$  is the number of words in  $C'$ ,  $q_{C'} = m_1 + \dots + m_{n-1}$ , and  $f$  is the dimension of the word embedding;  $E_{U_{n-1}} \in \mathbb{R}^{m_{n-1} \times f}$ ,  $m_{n-1}$  is the number of words in  $U_{n-1}$ ; **EcEnc**<sub>1</sub> and **EcEnc**<sub>2</sub> are Emotion context encoders with different emotion-related concepts and parameters.

As the virtual last utterance  $U_n^p$  is predicted after the  $(n-1)$ -th utterance  $U_{n-1}$  is spoken, we argue that the generated process of  $U_n^p$  should adequately capture the intention of  $U_{n-1}$ . To this end, we utilize MAIFNet to fuse  $S_{C'}$  and  $S_{U_{n-1}}$ , which is shown in the upper part of Figure 4. The fusion

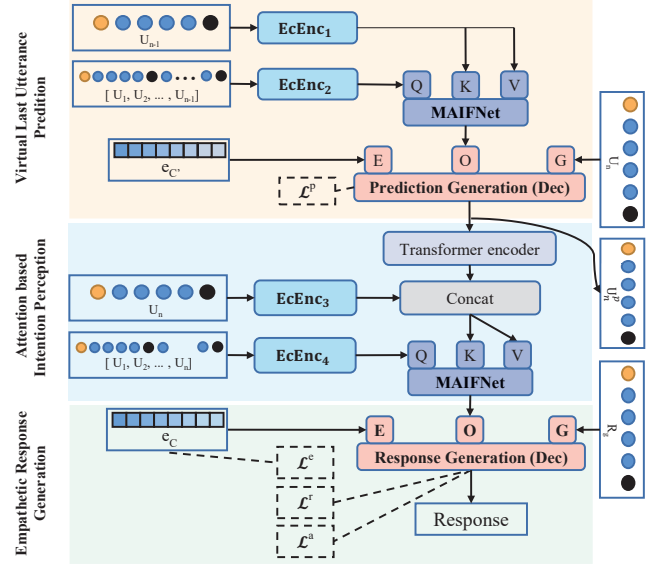


Figure 4: The architecture of InferEM.

process can be formulated as:  $F_{C'}^1 = \mathbf{MAIFNet}(S_{C'}, S_{U_{n-1}})$ . Concretely, **MAIFNet** mainly includes a multi-head attention layer, a feed-forward layer, two residual connections and two normalization layers:

$$\begin{aligned} S_{C'}^1 &= \mathbf{NM}(S_{C'} + \mathbf{Att}(S_{C'}, S_{U_{n-1}}, S_{U_{n-1}})), \\ F_{C'} &= \mathbf{NM}(S_{C'}^1 + \mathbf{FF}(S_{C'}^1)), \end{aligned} \quad (3)$$

where  $F_{C'} \in \mathbb{R}^{q_{C'} \times f_2}$ ,  $f_2$  is the feature dimension of  $F_{C'}$ ; **Att**, **FF** and **NM** denote the multi-head attention layer, feed-forward layer and layer normalization, respectively; we input  $S_{C'}$  as  $Q$  of **Att**, and take  $S_{U_{n-1}}$  as  $K$  and  $V$  of **Att**; the purpose of multi-head attention is to query the intention of the  $(n-1)$ -th utterance  $U_{n-1}$  according to the semantic information of the  $n-1$  past utterances  $C'$ ;  $S_{C'}^1$  contains both the contextual semantic feature of  $C'$  and the intention information of  $U_{n-1}$ . Referring to KEMP, the emotion signal representation  $e_{C'}$  is calculated as follows:

$$e_{C'} = \mathbf{softmax}(\eta(C'))^\top \times S_{C'}, \quad (4)$$

where  $e_{C'}^\top \in \mathbb{R}^{f_1}$ ;  $S_{C'} \in \mathbb{R}^{q_{C'} \times f_1}$  denotes the semantic feature of the  $n-1$  past utterances  $C'$ , and  $\eta(C') \in \mathbb{R}^{q_{C'}}$  is emotion intensity values provided by NRC\_VAD. Finally, we use **Dec** to generate the probability distribution of the virtual last utterance:

$$P_n = \mathbf{Dec}(e_{C'}, F_{C'}, E_w(U_n)), \quad (5)$$

where **Dec** is the Emotion-dependency decoder, and  $E_w(U_n)$  is the word embedding of the real last utterance  $U_n$ ;  $E_w(U_n) \in \mathbb{R}^{m_n \times f}$ ,  $m_n$  is the number of words in the real last utterance  $U_n$ .

During training, the prediction loss is quantified by the negative log-likelihood with respect to the gold prediction  $U_n$ :

$$\mathcal{L}^p = - \sum_{j=1}^{m_n} \log P_n(U_n^p[j] = U_n[j] | (U_n[1], \dots, U_n[j-1])), \quad (6)$$

where  $U_n[j]$  denotes the  $j$ -th ground-truth word in  $U_n$ , and  $U_n^p[j]$  denotes the  $j$ -th predicted word in the virtual last utterance  $U_n^p$ . In order to obtain the virtual last utterance  $U_n^p$ , we calculate the index corresponding to the maximum probability distribution of each element in  $P_n$ :

$$IDX(U_n^p[j]) = \mathbf{argmax}(P_n(U_n^p[j])), \quad (7)$$

where  $IDX(U_n^p[j])$  denotes the predicted index of the  $j$ -th word in  $U_n^p$ , and  $P_n(U_n^p[j])$  is the probability distribution of the  $j$ -th element in  $P_n$ . Then, the corresponding word is looked up in the dictionary based on the index.

### Attention based Intention Perception

In this section, MAIFNet is exploited to integrate the dialogue history and last utterance to capture the speaker’s intention information. We provide the entire dialogue  $C$  and last utterance  $U_n$  with emotion related concepts and extract their semantic features by Emotion context encoder:

$$\begin{aligned} S_C &= \mathbf{EcEnc}_3(E_C), \\ S_{U_n} &= \mathbf{EcEnc}_4(E_{U_n}), \end{aligned} \quad (8)$$

where  $E_C \in \mathbb{R}^{q_C \times f}$ ,  $q_C$  is the number of words in  $C$ , and  $q_C = m_1 + \dots + m_n$ ;  $E_{U_n} \in \mathbb{R}^{m_n \times f}$ .

We believe that the last utterance contains the speaker’s intention, so we take  $S_C$  and  $S_{U_n}$  together as the input to MAIFNet. Before that, we hope to increase the diversity of intention of the last utterance. First, we leverage the Transformer encoder to extract the semantic feature of the virtual last utterance  $U_n^p$ . Here, we denote the semantic feature of  $U_n^p$  as  $S_{U_n^p}$ . Then, we concatenate  $S_{U_n^p}$  with  $S_{U_n}$  to obtain a new semantic feature representation  $S_{U_n^{pr}}$ .

In the most important step, we leverage **MAIFNet** to fuse the semantic information of the dialogue history and that of the last utterance in order to perceive the speaker’s intention. The process can be formulated as follows:

$$\begin{aligned} S_C^1 &= \mathbf{NM}(S_C + \mathbf{Att}(S_C, S_{U_n^{pr}}, S_{U_n^{pr}})), \\ F_C &= \mathbf{NM}(S_C^1 + \mathbf{FF}(S_C^1)), \end{aligned} \quad (9)$$

where  $F_C \in \mathbb{R}^{q_C \times f_3}$ ; we input  $S_C$  as  $Q$  of **Att**, and take  $S_{U_n^p}$  as  $K$  and  $V$  of **Att**;  $S_C^1$  contains both the semantic information of the dialogue history and the intention information of the last utterance.

### Empathetic Response Generation

Like the Virtual Last Utterance Prediction section, we calculate the emotional signal representation  $e_C$  of the entire dialogue as follows:

$$e_C = \mathit{softmax}(\eta(C))^\top \times S_C, \quad (10)$$

where  $e_C^\top \in \mathbb{R}^{f_4}$ ,  $S_C \in \mathbb{R}^{q_C \times f_4}$ , and  $\eta(C) \in \mathbb{R}^{q_C}$ . Finally, we generate the probability distribution of the empathetic response by employing the Emotion-dependency decoder:

$$P_R = \mathbf{Dec}(e_C, F_C, E_w(R_g)), \quad (11)$$

where  $E_w(R_g)$  is the word embedding of the gold response. During training, we use the negative log-likelihood loss to optimize the response generation:

$$\mathcal{L}^r = - \sum_{j=1}^t \log P_R(R[j] = R_g[j] | (R_g[1], \dots, R_g[j-1])), \quad (12)$$

where  $t$  is the number of words in the gold response;  $R_g[j]$  denotes the  $j$ -th word in the gold response, and  $R[j]$  denotes the  $j$ -th word in the generated response.

To improve our model’s emotion perception ability, we project the emotion signal representation  $e_C$  into a vector whose length is same as the number of emotion categories and use the negative log-likelihood loss to optimize emotion perception learning:

$$\begin{aligned} \tilde{e}_C &= W_e \times e_C^\top, \\ P_e(\tilde{e}_C) &= \mathbf{softmax}(\tilde{e}_C), \\ \mathcal{L}^e &= -\log P_e(\tilde{e}_C = e^*), \end{aligned} \quad (13)$$

where  $W_e \in \mathbb{R}^{q \times f_4}$ ,  $q$  is the number of emotion categories, and  $f_4$  is the dimension of the emotion signal representation.  $e^* \in \mathbb{R}^q$  is the emotion label of the entire dialogue. As in KEMP, we adopt the emotional attention loss  $\mathcal{L}^a$  to enforce the decoder to attach more attention to words with higher emotion intensity values:

$$\mathcal{L}^a = \frac{1}{e} \sum_{i=1}^e (\eta(C[i]) - a_i)^2, \quad (14)$$

where  $e$  is the number of words in the entire dialogue history, and  $\eta(C[i])$  is the emotional intensity value of the  $i$ -th word in  $C$ ;  $a_i$  is the average attention score of  $C[i]$ , i.e.,  $a_i = \sum_{n=1}^H a^n(R[j-1], C[i])/H$ . Here,  $H$  is the number of attention heads,  $a^n(R[j-1], C[i])$  is the attention score between the  $(j-1)$ -th generated word  $R[j-1]$  and the  $i$ -th word  $C[i]$  of the dialogue history  $C$ .

### Multi-task Training Objective

We optimize all the parameters by minimizing the weighted summation of the above-mentioned four losses:

$$\mathcal{L} = \mathcal{L}^r + \alpha_1 \mathcal{L}^p + \alpha_2 \mathcal{L}^e + \alpha_3 \mathcal{L}^a, \quad (15)$$

where  $\alpha_1, \alpha_2, \alpha_3$  are trade-off parameters. To better balance the optimizing rate of the parameters of prediction module and response module, we set  $\alpha_1$  in this way:  $\alpha_1 = \alpha_1'$  when  $\mathcal{L}^p > \mathcal{L}^r$ ,  $\alpha_1 = \alpha_1''$  when  $\mathcal{L}^p \leq \mathcal{L}^r$ . Here,  $\alpha_1'$  is a relatively larger value while  $\alpha_1''$  is a relatively smaller value. This strategy can ensure that parameters of the Virtual Last Utterance Prediction section are optimized more quickly and get a more appropriate predicted utterance at the early training stage. On the contrary, at the late training stage, the parameters of the Attention based Intention Perception section and Empathetic Response Generation section are mainly optimized.

## Experiment

### Dataset and Experimental Setup

We evaluate the model on the EmpatheticDialogues (Rashkin, Smith, Li, & Boureau, 2018) dataset, an empathetic dialogue dataset containing 24850 one-to-one open-domain dialogues, where the listener responds emphatically to the last utterance of the speaker. We use 17802/2628/2494 dialogues as training/validation/testing set. The word embedding initialized by using the 300-dimensional pre-trained Glove vectors (Pennington, Socher, & Manning, 2014). The hidden dimension of each component of InferEM is set to 300. The emotion lexical knowledge and commonsense knowledge are obtained from NRC\_VAD and ConceptNet in the same way as KEMP (Q. Li et al., 2022). During testing the maximum of decoding steps is set as 30. The trade-off parameters  $\alpha'_1$ ,  $\alpha''_1$ ,  $\alpha_2$ , and  $\alpha_3$  are set as 1.5, 0.3, 1.2, and 0.12, respectively. We train the model using Adam (Kingma & Ba, 2015) with  $\beta_1 = 0.9$  and  $\beta_2 = 0.98$ . The initial learning rate is set to 0.0001 and varied during training in accordance with Vaswani et al. (2017), and the mini-batch size is set to 16. All experiments are implemented in PyTorch on a single NVIDIA GeForce RTX 3090 with early stopping applied.

### Baselines and Evaluation Metrics

**MoEL** (Lin et al., 2019) is a Transformer-based model which designs different decoders to focus on different type of emotions. **MIME** (Majumder et al., 2020) is a Transformer-based model which separates all emotions into positivity and negativity groups and decodes mimicking as well as non-mimicking response representations. **KEMP** (Q. Li et al., 2022) is a GAT and Transformer based model which uses commonsense knowledge to enrich dialogue history. **CEM** (Sabour et al., 2022) is a Transformer-based model which focus on both affection and cognition aspects of empathy. To evaluate the effectiveness of the utterance prediction and intention perception components in our model, we design ablation studies as follows. **w/o SIP**: the model without the speaker’s intention perception. **w/o LUP**: the model without the last utterance prediction.

**Automatic Evaluations:** With reference to KEMP, we adopt *Emotion Accuracy*, *Perplexity* (Serban, Sordoni, Bengio, Courville, & Pineau, 2015), *Distinct-1*, and *Distinct-2* (J. Li, Galley, Brockett, Gao, & Dolan, 2015) to evaluate the performance of our model. **Human Evaluations:** following CEM, we carry out an aspect-based pairwise response comparison. That is, given 100 randomly sampled dialogue inputs, we ask 3 annotators to compare responses from InferEM and baselines, and choose the better one on the following three aspects, including *Empathy* (i.e., which response expresses more appropriate emotions), *Relevance* (i.e., which response is more relevant with the given dialogue history) and *Fluency* (i.e., which response has more readability and grammatical correctness). When they think two responses have much similar qualities and can’t determine which one is better, they will give a judgment of *Tie*.

## Results and Analysis

**Comparison on Automatic Evaluation** Table 1 shows that all automatic evaluation metrics of InferEM outperform the baseline KEMP by an obvious margin, which verifies the effectiveness of the last utterance prediction and intention perception components of our model to improve emotion perception, empathy expression and response diversity. Our model achieves the State-Of-The-Art performance in terms of Emotion Accuracy and Perplexity. The CEM performs better than our model in terms of Distinct-1 and Distinct-2 because it uses an additional Frequency-Aware Cross-Entropy loss (Jiang, Ren, Monz, & de Rijke, 2019) to penalize frequently occurring tokens. As InferEM obtains external knowledge in the same way as KEMP and is mainly compared with KEMP, we don’t consider this loss. Therefore, we can compare InferEM with KEMP more fairly and verify the effectiveness of our proposed modules more convincingly.

Table 1: Comparison on automatic evaluation.

Models	Accuracy(%)	Perplexity	Distinct-1	Distinct-2
MoEL	32.00	38.04	0.44	2.10
MIME	34.24	37.09	0.47	1.91
KEMP	39.31	36.89	0.55	2.29
CEM	39.11	36.11	<b>0.66</b>	<b>2.99</b>
InferEM	<b>39.98</b>	<b>31.26</b>	0.59	2.60
w/o SIP	39.14	31.43	0.52	2.08
w/o LUP	38.15	33.73	0.52	2.09

Table 1 shows that when the last utterance prediction or the intention perception module is removed, all the automatic evaluation metrics of our model become obviously worse, which means that the model’s abilities to perceive intention and understand dialogue context decrease to some extent. This indicates that our proposed modules are beneficial for emotion perception and empathetic expression.

Table 2: Comparison on human evaluation.

Comparisons	Aspects	Win(%)	Loss(%)	Tie(%)
InferEM vs MoEL	Emp.	42.3	23.7	34.0
	Rel.	45.0	22.3	32.7
	Flu.	24.7	19.3	56.0
InferEM vs MiME	Emp.	41.3	26.7	33.0
	Rel.	42.3	23.7	34.0
	Flu.	22.3	17.7	60.0
InferEM vs KEMP	Emp.	43.3	23.0	33.7
	Rel.	44.3	20.7	35.0
	Flu.	25.0	19.7	62.3
InferEM vs CEM	Emp.	32.7	23.3	44.0
	Rel.	42.7	29.0	28.3
	Flu.	24.0	20.7	55.3

**Comparison on Human Evaluation** Table 2 shows that our model performs best on Empathy and Relevance. This

suggests that our proposed modules can promote successful communication and generate more appropriate responses. As Transformer-based models can already generate fluent responses, our model is not noticeably different with baselines on Fluency.

### Case Studies

Table 3 shows two generated responses of our model and baselines. In the first case, InferEM outputs a suitable prediction of the last utterance  $U_3^p$  with context-related words “try my best”, and generates a proper response with the excited emotion by replying with “get it”. In the second case, InferEM generates the most appropriate response, containing the context-related words “help them up” and emotion-rated word “great”. These cases show that the last utterance prediction and intention perception modules of InferEM can improve mutual understanding and empathetic expression.

Table 3: Generated responses from InferEM and baselines.

<b>Emotion</b>	<b>Excited</b>
<b>History</b>	$U_1$ : I have a huge chance to win a ps4! $U_2$ : Wow, please make it real. $U_3$ : There is only 2 other people in the final round! winning a ps4 will be awesome!
<b>Gold</b>	Wow, i belief u can make it.
MoEL	I am sure you will have a great time!
MIME	That is a great attitude to have!
KEMP	I agree with you. I have a huge fan of a few years ago.
CEM	I am sure you will do great!
$U_3^p$	I will <u>try my best</u> for it!
InferEM	That is great! I am sure you will <u>get it</u> !
<b>Emotion</b>	<b>Proud</b>
<b>History</b>	$U_1$ : I am so happy for my son, went to his graduation from harvard last week. $U_2$ : Wao, thats great, congrats! $U_3$ : Thank you! looks like there is going to be another doctor in the family!
<b>Gold</b>	Thats good to know, keep it up.
MoEL	I am so happy for you!
MIME	I am sure he will be fine.
KEMP	I am sure you will do great!
CEM	That is great! I hope he’s doing well
$U_3^p$	I am so happy for him!
InferEM	That is <u>great</u> ! I am sure you will be able to <u>help them up</u> !

Table 4 is a case where InferEM generates a different response when we remove the last utterance prediction module from it. Considering the impact of the predicted utterance  $U_3^p$ , InferEM generates a sounder response containing the context

related and emotionally reasonable words “cooperate better”. It is obvious that if only the dialogue history is considered, the model generates a relatively monotonous response.

Table 4: Generated responses from InferEM with and without the last utterance prediction module.

<b>Emotion</b>	<b>Impressed</b>
<b>History</b>	$U_1$ : I admire my boss because he does not act like a boss. He acts more like a friend and he is so down to earth. $U_2$ : That is nice to have a friendly boss. $U_3$ : He is not just friendly, he is a good leader and knows how to treat his workers.
<b>Gold</b>	I envy you. That is cool.
$U_3^p$	Sure. It is my luck to work together with him.
InferEM	That is good to hear. I believe that you will <u>cooperate better</u> in the future.
w/o LUP	I am glad you have a good leader.

### Limitation

Compared with recent large language models such as ChatGPT, InferEM may output flat responses such as “I am sorry to hear that” in some cases. InferEM remains to be combined with pre-trained large language models to overcome this deficiency and generate more meaningful responses. It has to be mentioned that even though large language models show powerful performance, they are computationally very expensive and also require some specific components to further boost performance and enhance their interpretability. Therefore, the shallow frameworks like InferEM still deserve to be studied.

Moreover, since the predicted last utterance and the real last utterance have different importance for generating responses, we believe that a proper attention mechanism designed to provide them with different weights will lead to better experimental results.

### Conclusion and Future Work

Based on two key ideas of utterance prediction and intention perception, we propose a novel empathetic dialogue generation model named InferEM. The proposed intention fusion module is conducive to understanding the speaker’s intention to generate more targeted responses. The predictive function of InferEM is beneficial for enhancing mutual understanding and generating more diverse responses. Furthermore, we design a multi-task learning strategy to better optimize the parameters of InferEM. Experimental results of automatic and human evaluations indicate that our proposed approach is effective in empathetic expression. In the future work, we will try utilizing superior information fusion methods to perceive intention.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62236005, 61936004 and U1913602. We would like to thank the anonymous reviewers for their helpful remarks.

## References

- Batson, C. D., Polycarpou, M. P., Harmon-Jones, E., Imhoff, H. J., Mitchener, E. C., Bednar, L. L., ... Highberger, L. (1997). Empathy and attitudes: Can feeling for a member of a stigmatized group improve feelings toward the group? *Journal of Personality and Social Psychology*, 72(1), 105–118. doi: 10.1037/0022-3514.72.1.105
- Bohart, A. C., & Greenberg, L. S. (1997). *Empathy reconsidered: New directions in psychotherapy* (A. C. Bohart & L. S. Greenberg, Eds.). American Psychological Association. doi: 10.1080/10503309912331332721
- Chen, M. Y., Li, S., & Yang, Y. (2022). Emphi: Generating empathetic responses with human-like intents. In *Proceedings of the conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 1063–1074). Seattle, United States: Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.78
- Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, 44(1), 113–126. doi: 10.1037/0022-3514.44.1.113
- Gao, J., Liu, Y., Deng, H., Wang, W., Cao, Y., Du, J., & Xu, R. (2021). Improving empathetic response generation by recognizing emotion cause in conversations. In *Findings of the association for computational linguistics: Emnlp 2021* (pp. 807–819). doi: 10.18653/v1/2021.findings-emnlp.70
- Garrod, S., & Pickering, M. J. (2009). Joint action, interactive alignment, and dialog. *Topics in Cognitive Science*, 1(2), 292–304. doi: 10.1111/j.1756-8765.2009.01020.x
- Hu, J., Liu, Y., Zhao, J., & Jin, Q. (2021). Mmgen: Multimodal fusion via deep graph convolution network for emotion recognition in conversation. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing* (p. 5666–5675). Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.440
- Jiang, S., Ren, P., Monz, C., & de Rijke, M. (2019). Improving neural response diversity with frequency-aware cross-entropy loss. In *The world wide web conference* (pp. 2879–2885). doi: 10.1145/3308558.3313415
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations* (pp. 1–13).
- Li, J., Galley, M., Brockett, C., Gao, J., & Dolan, B. (2015). A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 110–119). doi: 10.18653/v1/n16-1014
- Li, J., Wang, X., Lv, G., & Zeng, Z. (2023). Graphcfc: A directed graph based cross-modal feature complementation approach for multimodal conversational emotion recognition. *IEEE Transactions on Multimedia*, 1-13. doi: 10.1109/TMM.2023.3260635
- Li, Q., Li, P., Ren, Z., Ren, P., & Chen, Z. (2022). Knowledge bridging for empathetic dialogue generation. In *Proceedings of the aaai conference on artificial intelligence*. doi: 10.1609/aaai.v36i10.21347
- Lin, Z., Madotto, A., Shin, J., Xu, P., & Fung, P. (2019). Moel: Mixture of empathetic listeners. In *Proceedings of the conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (emnlp-ijcnlp)* (pp. 121–132). doi: 10.18653/v1/d19-1012
- Liu, S., Zheng, C., Demasi, O., Sabour, S., Li, Y., Yu, Z., ... Huang, M. (2021). Towards emotional support dialog systems. 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)*. doi: 10.18653/v1/2021.acl-long.269
- Majumder, N., Hong, P., Peng, S., Lu, J., Ghosal, D., Gelbukh, A., ... Poria, S. (2020). Mime: Mimicking emotions for empathetic response generation. In *Proceedings of the 2020 conference on empirical methods in natural language processing (emnlp)*. doi: 10.18653/v1/2020.emnlp-main.721
- Pavey, L., Greitemeyer, T., & Sparks, P. (2012). “i help because i want to, not because you tell me to”: Empathy increases autonomously motivated helping. *Personality and Social Psychology Bulletin*, 38(5), 681–689. doi: 10.1177/0146167211435940
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (emnlp)* (pp. 1532–1543). doi: 10.3115/v1/d14-1162
- Poulin, M. J., Brown, S. L., Dillard, A. J., & Smith, D. M. (2013). Giving to others and the association between stress and mortality. *American Journal of Public Health*, 103(9), 1649–1655. doi: 10.2105/AJPH.2012.300876
- Rashkin, H., Smith, E. M., Li, M., & Boureau, Y.-L. (2018). Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 5370–5381). doi: 10.18653/v1/p19-1534
- Sabour, S., Zheng, C., & Huang, M. (2022). Cem: Commonsense-aware empathetic response generation. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 36, pp. 11229–11237). doi: 10.1609/aaai.v36i10.21373



- Serban, I. V., Sordoni, A., Bengio, Y., Courville, A., & Pineau, J. (2015). Hierarchical neural network generative models for movie dialogues. , 7(8), 434–441.
- Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C., & Althoff, T. (2023). Human-ai collaboration enables more empathic conversations in text-based peer-to-peer mental health support. *Nature Machine Intelligence*, 1–12. doi: 10.1038/s42256-022-00593-2
- Speer, R., Chin, J., & Havasi, C. (2017). Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the aaai conference on artificial intelligence*. doi: 10.1609/aaai.v31i1.11164
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (Vol. 30).
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. In *International conference on learning representations* (pp. 1–12).
- Wolf, T., Sanh, V., Chaumond, J., & Delangue, C. (2019). Transfertransfo: A transfer learning approach for neural network based conversational agents. In *Proceedings of the conference on neural information processing systems*.
- Zhang, Z., Li, J., Zhu, P., Zhao, H., & Liu, G. (2018). Modeling multi-turn conversation with deep utterance aggregation. In *Proceedings of the 27th international conference on computational linguistics* (pp. 3740–3752).
- Zhong, P., Wang, D., & Miao, C. (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)* (pp. 165–176). doi: 10.18653/v1/d19-1016
- Zhou, L., Gao, J., Li, D., & Shum, H.-Y. (2020). The design and implementation of xiaoice, an empathetic social chatbot. *Computational Linguistics*, 46(1), 53–93. doi: 10.1162/coli\_a.00368