

EmpDG: Multi-resolution Interactive Empathetic Dialogue Generation

Qintong Li¹, Hongshen Chen², Zhaochun Ren^{1,*}, Pengjie Ren¹, Zhaopeng Tu³, Zhumin Chen¹

¹School of Computer Science and Technology, Shandong University, Qingdao, China

²Data Science Lab, JD.COM, Beijing, China

³Tencent AI Lab, Shenzhen, China

qintongli@mail.sdu.edu.cn, ac@chenhongshen.com,
{zhaochun.ren, chenzhumin}@sdu.edu.cn,
jay.ren@outlook.com, zptu@tencent.com

Abstract

A humanized dialogue system is expected to generate empathetic replies, which should be sensitive to the users' expressed emotion. The task of empathetic dialogue generation is proposed to address this problem. The essential challenges lie in accurately capturing the nuances of human emotion and considering the potential of user feedback, which are overlooked by the majority of existing work. In response to this problem, we propose a multi-resolution adversarial model – EmpDG, to generate more empathetic responses. EmpDG exploits both the coarse-grained dialogue-level and fine-grained token-level emotions, the latter of which helps to better capture the nuances of user emotion. In addition, we introduce an interactive adversarial learning framework which exploits the user feedback, to identify whether the generated responses evoke emotion perceptivity in dialogues. Experimental results show that the proposed approach significantly outperforms the state-of-the-art baselines in both content quality and emotion perceptivity.

1 Introduction

Studies on social psychology suggest that “empathy” is a crucial step towards a more humanized human-machine conversation, which improves the emotional perceptivity in emotion-bonding social activities (Zech and Rimé, 2005). To design an intelligent automatic dialogue system, it is important to make the chatbot become empathetic within dialogue interactions (Prendinger and Ishizuka, 2005). Therefore, in this paper, we focus on the task of *empathetic dialogue generation* (Rashkin et al., 2019), which automatically tracks and understands the user emotion information in a multi-turn dialogue scenario.

Despite the achieved successes (Rashkin et al., 2019; Lin et al., 2019), obstacles to establishing an empathetic conversational system are still far beyond current progress: (1) It is still difficult to accurately capture the nuances of human emotion in dialogue generation (Ghosal et al., 2019). (2) Merely relying on the dialogue history but overlooking the potential of user feedback for the generated responses further aggravates the aforementioned deficiencies, which causes undesirable responses (Zhang et al., 2018a). In Figure 1, we give an example of the benchmark dataset EMPATHETICDIALOGUES (Rashkin et al., 2019). Note that the emotional words in the sequence of utterances (in Figure 1, they are “new”, “job” in utterance 1, “amazing”, “excited” in utterance 2, and “excited” in utterance 3) have fine-grained emotional connections. Without considering fine-grained emotional words, the responses generated by existing methods are trivial and uninformed, even though they expressed the appropriate emotions. Therefore, explicitly modelling the fine-grained emotional factor is necessary.

In this paper, we propose a multi-resolution adversarial empathetic dialogue generation model, named EmpDG, to address above challenges through generating more appropriate and empathetic responses. To capture the nuances of user emotions sufficiently, EmpDG generates responses by jointly taking coarse-grained dialogue-level emotions and fine-grained token-level emotions into account. The multi-resolution emotions are the prerequisites of the response generator to perceive the user's nuanced emotion states. Furthermore, we propose an interactive adversarial learning framework to augment the user's feedback

*Corresponding author.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

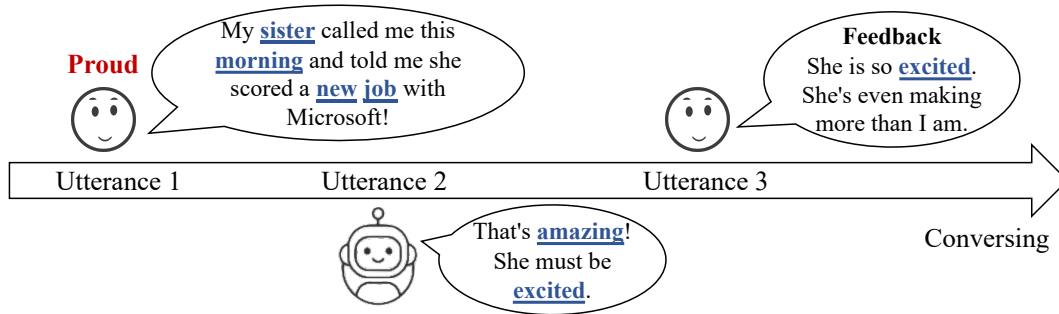


Figure 1: An empathetic dialogue example from dataset EMPATHETICDIALOGUES. The emotional-related words are highlighted in **blue**. “Proud” is the coarse-grained emotional label. These emotional words are labelled by an external emotion lexicon (Mohammad and Turney, 2013).

thoughtfully, where two interactive discriminators are designed to identify whether the generated responses evoke the emotion perceptivity regarding both the dialogue history and the user’s emotions. Conducted on a benchmark dataset EMPATHETICDIALOGUES, extensive experimental results have verified the effectiveness of the EmpDG in terms of both content quality and empathy quality. We also find that the EmpDG outperforms state-of-the-art baselines significantly.

In general, our main contributions are summarized as follows:

- We propose a multi-resolution adversarial neural network, which considers multi-granularity emotion factors in the dialogue context.
- To induce the response generation from users’ feedback, we propose an interactive adversarial learning network with two interactive discriminators.
- Experiments show that EmpDG significantly outperforms state-of-the-art baselines in terms of both content quality and empathy quality in the empathetic dialogue generation.

2 Related Work

Our research is in line with empathetic conversation generation through human-computer interactions, which avoids the additional step of assigning emotion to responses during conversations (Skowron et al., 2013). Several work (Lubis et al., 2018; Rashkin et al., 2018; Zhong et al., 2019; Wei et al., 2019; Shin et al., 2019; Chatterjee et al., 2019; Rashkin et al., 2019; Santhanam and Shaikh, 2019; Lin et al., 2019; Zhong et al., 2020; Lin et al., 2020) have attempted to make dialogue models more empathetic. Rashkin et al. (2019) combine existing models in different ways to produce empathetic responses. Lin et al. (2019) softly combine the possible emotional responses from several separate experts to generate the final empathetic response. Lin et al. (2020) fine-tune a large-scale pre-trained language model with multiple objectives. However, all existing approaches only consider monogranular emotions in the dialogue context. In comparison, EmpDG jointly considers highly correlated coarse-grained emotional labels and fine-grained emotional terms. Moreover, we explicitly consider the effect of user feedback via a novel interactive adversarial mechanism, so that EmpDG is capable to evoke more emotion perceptivity in dialogues.

Besides the advancements in empathetic dialogue models, the emergence of new emotion-labelled dialogue corpora, e.g., DAILYDIALOG (Li et al., 2017b) and EMOTIONLINES (Hsu et al., 2018), have also contributed to this research field. However, only 5% of the utterances in DAILYDIALOG and 16.68% of the utterances in EMOTIONLINES have diverse emotional labels and others are either “none” or “happy” labels. Because of the extremely imbalanced data distribution, they are not suitable to be engaged as the benchmarks of empathetic dialogue generation. Rashkin et al. (2019) consider a richer and evenly distributed set of emotions and release a dataset EMPATHETICDIALOGUES, where a listener responds to a speaker who is under an emotional situation in an empathetic way. Furthermore, several emotion

lexicons (Mohammad and Turney, 2013; Sedoc et al., 2020) have also been shown to be effective in tracking emotions in texts. In our work, we focus on the task of empathetic dialogue generation on the EMPATHETICDIALOGUES dataset and emotion lexicon (Mohammad and Turney, 2013).

The second line of our related work is the emotional dialogue generation, which has received an increasing amount of attention to address emotion factors (Zhou et al., 2018; Huang et al., 2018; Colombo et al., 2019). Prior studies on emotion-related conversational systems mainly focused on rule-based systems, which heavily rely on hand-crafted features (Prendinger and Ishizuka, 2005). Recently, many neural emotional dialogue generation approaches (Ghosh et al., 2017; Zhou and Wang, 2018; Zhou et al., 2018; Huang et al., 2018; Colombo et al., 2019) have been explored to control the emotional expression in the target response. However, as Li et al. (2018) reveal that conventional emotional conversation systems aim to produce more emotion-rich responses according to a given specific user-input emotion, which inevitably leads to an emotional inconsistency problem.

Our research also aligns with recent advances in open-domain dialogue generation models (Vinyals and Le, 2015; Li et al., 2016b; Zhang et al., 2018b; Hancock et al., 2019; Li et al., 2020; Song et al., 2020). These dialogue models usually adopt the sequence-to-sequence (Seq2Seq) (Sutskever et al., 2014) fashion. Adversarial learning achieves considerable success in generating higher-quality responses (Goodfellow et al., 2014; Li et al., 2017a) but often leads to gradient vanishing as the discriminator saturates (Gulrajani et al., 2017). To tackle this problem, Gao et al. (2019) utilize Wasserstein GAN (Arjovsky et al., 2017) to enhance response consistency with external facts. Romanov et al. (2019) propose an adversarial decomposition method for fine-grained text representation. Unlike the previous work, we investigate an adversarial approach to improve the empathy quality of neural dialogue models.

3 Problem Formulation

Before detailing our method, we introduce our key notations and concepts. A multi-turn dialogue context consists of M utterances between two interlocutors. We assume both *semantic context* and *emotional context* exist in such dialogue. The semantic context \mathcal{U} refers to the sequence of utterances, i.e., $\mathcal{U} = [U_1, \dots, U_M]$. Following Lin (2019), we flat \mathcal{U} into a long token sequence and insert a CLS token at the start of the token sentence, i.e., $\mathcal{U} = [\text{CLS}, x_1, \dots, x_m]$. The emotional context \mathcal{E} considers emotions with different granularities, i.e., $\mathcal{E} = [\text{LAB}, w_1, \dots, w_e]$, where w_i is the emotional word in the semantic context \mathcal{U} , LAB is a special emotion token which is used to derive the emotional state of the dialogue context. We extract emotional words using an external emotional vocabulary V_E (Mohammad and Turney, 2013). In the following part, we use x_0 to denote CLS in \mathcal{U} and w_0 to denote LAB in \mathcal{E} .

Given \mathcal{U} and \mathcal{E} , our model aims to generate a n -length response $\mathcal{Y} = \{y_1, \dots, y_n\}$ through maximizing the probability $P(\mathcal{Y}|\mathcal{U}, \mathcal{E}) = \prod_{n=1}^N P(y_n|y_{<n}, \mathcal{U}, \mathcal{E})$.

4 EmpDG

In this section, we detail our proposed multi-resolutional adversarial model, abbreviated as EmpDG. The overview of EmpDG is illustrated in Figure 2. There are two main components in EmpDG: the *empathetic generator* and the *interactive discriminators*.

To summarize, the *empathetic generator* is established based on an encoder-decoder architecture, which are all implemented with Transformer (Vaswani et al., 2017). During encoding procedure, semantic context and the multi-resolution emotional context are encoded; whereas the decoder fuses the semantic context and emotional context to generate responses.

To enhance the empathy of the generator, we design two CNN-based (Kalchbrenner et al., 2014a) discriminators (i.e., the *semantic discriminator* and the *emotional discriminator*). In the training procedure, the two discriminators additionally interact with the user feedback (the next utterance and corresponding emotional words of the next utterance). By minimizing the Wasserstein-1 distance to optimize the discriminators, we use the sum of classification results as a training signal to encourage response generator to evoke more emotion perceptivity.

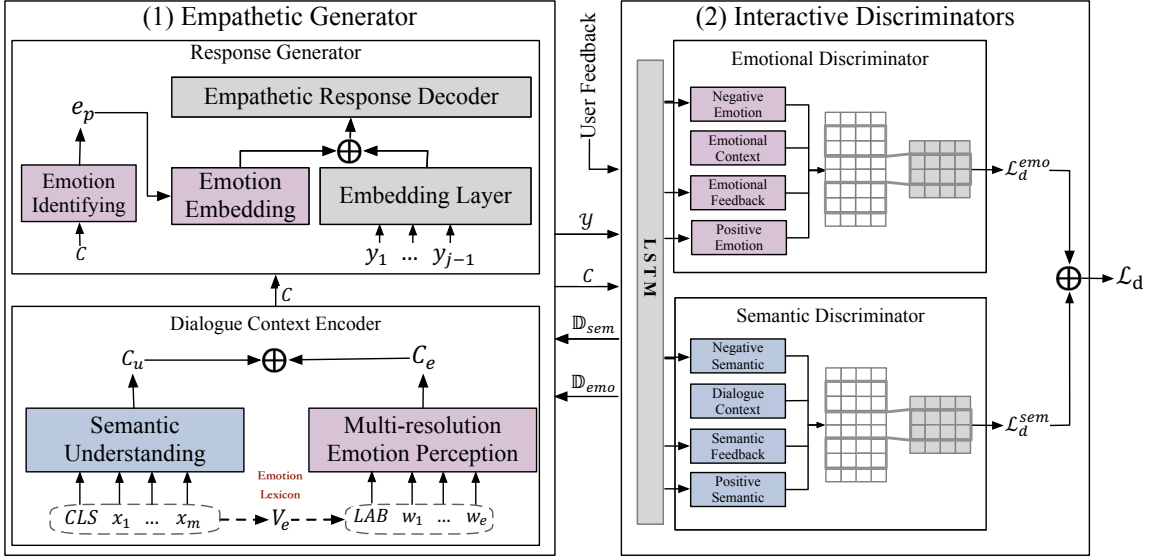


Figure 2: Overview of EmpDG. We divide EmpDG into two parts: (1) *Empathetic generator* generates empathetic responses based on the semantic understanding and multi-resolution emotion perception. (2) *Interactive discriminators* distinguishes whether the generated responses are emotion-appropriate and context-consistent.

4.1 Empathetic Generator

We start by proposing an empathetic generator to generate a response \mathcal{Y} . To better model the multi-granularity emotions, we divide our encoding-decoding process into 3 phases individually: semantic understanding, multi-resolution emotion perception, and empathetic response generation.

Semantic Understanding. We first use a word embedding layer and a positional embedding layer (Vaswani et al., 2017) to convert each token of semantic context \mathcal{U} into vector representations $\mathbf{e}_{x_i}^W \in \mathbb{R}^d$ and $\mathbf{e}_{x_i}^P \in \mathbb{R}^d$, where d is the dimensionality of embeddings. In the multi-turn dialogue settings, distinguishing utterances from speaker and listener is helpful. Therefore, we incorporate the dialogue state embedding $\mathbf{e}_{x_i}^D$ for each token. Our semantic context embedding $[\mathbf{x}_0, \dots, \mathbf{x}_m]$ is the composition of three types of embeddings, where \mathbf{x}_i is computed as follows:

$$\mathbf{x}_i = \mathbf{e}_{x_i}^W + \mathbf{e}_{x_i}^P + \mathbf{e}_{x_i}^D. \quad (1)$$

Then we use the local transformer layer (Vaswani et al., 2017) to encode semantic context information for each tokens \mathbf{x}_i .

$$\mathbf{h}_i = \text{LayerNorm}(\mathbf{x}_i^{l-1} + \text{MHAtt}(\mathbf{x}_i^{l-1})), \quad (2)$$

$$\tilde{\mathbf{x}}_i^l = \text{LayerNorm}(\mathbf{h}_i + \text{FFN}(\mathbf{h}_i)), \quad (3)$$

where LayerNorm is the Layer Normalization trick proposed in (Ba et al., 2016); FFN is a two-layer feed-forward network with ReLU as hidden activation function. The transformer layers are stacked L times. The obtained final context representations are denoted as $\mathbf{C}_u = [\tilde{\mathbf{x}}_0, \dots, \tilde{\mathbf{x}}_m]$.

Multi-resolution Emotion Perception. We use another transformer encoder with a different set of parameters to encode the multi-resolution emotional context \mathcal{E} .

$$\mathbf{w}_i = \mathbf{e}_{w_i}^W + \mathbf{e}_{w_i}^P + \mathbf{e}_{w_i}^E, \quad (4)$$

$$\mathbf{h}_i = \text{LayerNorm}(\mathbf{w}_i^{l-1} + \text{MHAtt}(\mathbf{w}_i^{l-1})), \quad (5)$$

$$\tilde{\mathbf{w}}_i^l = \text{LayerNorm}(\mathbf{h}_i + \text{FFN}(\mathbf{h}_i)), \quad (6)$$

where $\mathbf{e}_{w_i}^E$ is the emotion state embedding to distinguish the emotional context from semantic context. Then the multi-resolution emotional context is represented as $\mathbf{C}_e = [\tilde{\mathbf{w}}_0, \dots, \tilde{\mathbf{w}}_e]$.

To perceive the emotional information in dialogue context, a linear layer with softmax operation projects the concatenation of $\tilde{\mathbf{w}}_0$ and $\tilde{\mathbf{x}}_0$ into an emotion category distribution P_e over the coarsed-grained emotional label e to identify the emotion signal user expressed:

$$\mathbf{e}_p = \mathbf{W}_e[\tilde{\mathbf{w}}_0; \tilde{\mathbf{x}}_0], \quad (7)$$

$$P_e(e|\mathcal{E}) = \text{softmax}(\mathbf{e}_p), \quad (8)$$

where $\mathbf{W}_e \in \mathbb{R}^{2d \times q}$, q is the number of emotion categories, and $[\cdot; \cdot]$ is the concatenation operation. During training, given the ground truth emotional label e^* of dialogue context, we employ negative log-likelihood as the loss function to conduct the parameter learning:

$$\mathcal{L}_{emo} = -\log(P_e(e^*|\mathcal{E})). \quad (9)$$

In addition, the obtained intermediate emotional representation \mathbf{e}_p will be fed into the decoder as a crucial emotional feature to guild the empathetic response generation.

The final dialogue context representations \mathbf{C} is the concatenation of the semantic context vectors \mathbf{C}_u and emotional context vectors \mathbf{C}_e , i.e., $\mathbf{C} = [\mathbf{C}_u; \mathbf{C}_e]$.

Empathetic Response Generation. The predicted emotion signal $\mathbf{e}_p \in \mathbb{R}^{1 \times q}$ is firstly be transformed by a linear transformation into $\mathbf{e}'_p \in \mathbb{R}^{1 \times d}$. Then we concatenate \mathbf{e}'_p with the embeddings of the decoder input tokens $[y_1, \dots, y_{j-1}]$ into representations $\mathbf{E}^Y = [y_0 \dots, y_{j-1}]$ where $y_0 = \mathbf{e}'_p$. We feed \mathbf{E}^Y into the response generator.

The generator is built based on the Transformer layers as well. For each Transformer decoder layer, the decoder inputs \mathbf{E}^Y are first updated into new vector representations \mathbf{Y} . Then a multi-head cross-attention mechanism MH-CAtt (Vaswani et al., 2017) derives a context vector from dialogue context:

$$\mathbf{D} = \mathbf{Y} + \mathbf{W}_m \text{MH-CAtt}(\mathbf{Y}, \mathbf{C}), \quad (10)$$

$$\hat{\mathbf{D}} = \text{LayerNorm}(\mathbf{D}), \quad (11)$$

$$\hat{\mathbf{Y}} = \text{LayerNorm}(\hat{\mathbf{D}} + \text{FFN}(\hat{\mathbf{D}})), \quad (12)$$

where $\hat{\mathbf{Y}} = [\hat{y}_1, \dots, \hat{y}_j]$. The response generator yields the distribution over the vocabulary for the next j -th token:

$$p(y_j|y_{<j}, \mathbf{C}) = \text{softmax}(\mathbf{W}_o \hat{y}_j). \quad (13)$$

As most dialogue generation tasks, we use standard maximum likelihood estimator (MLE) as the optimization objective:

$$\mathcal{L}_{gen} = -\log p(y_j|y_{<j}, \mathbf{C}). \quad (14)$$

Finally, considering all the aforementioned components, we define a joint loss function as follows:

$$\mathcal{L}_g = \gamma_1 \mathcal{L}_{emo} + \gamma_2 \mathcal{L}_{gen}, \quad (15)$$

where γ_1, γ_2 are hyper-parameters that control the weights of the two losses. We set $\gamma_1 = \gamma_2 = 1$. All the parameters are jointly trained in an end-to-end paradigm.

4.2 Interactive Discriminators

To evaluate whether the response is generated in an empathetic and context-consistent way, we introduce two discriminators to provide additional training signals for the empathetic generator. A semantic discriminator measures the semantic distance from the generated response to the gold response. An emotional discriminator specifies whether the generated responses are empathetic enough.

Specially, the next utterance of response could serve as the user's implicit feedback and provide semantic and emotional guidance for target response (Zhang et al., 2018a). Therefore, we regard the next utterance of response as user semantic feedback and its contained emotional words as user emotional feedback. In the training procedure, we utilize user semantic feedback and emotional feedback to optimize the content and empathy ability of the empathetic generator, respectively. Both the semantic discriminator and the emotional one are built on the convolutional neural network (CNN) based classifier, so we detail the semantic discriminator first for convenience.

Semantic Discriminator. First, we apply an LSTM encoder (Hochreiter and Schmidhuber, 1997) to respectively encode the generated response and gold response into hidden representations, i.e., \mathbf{d}_t^N and \mathbf{d}_t^P . We regard \mathbf{d}_t^N as negative vectors and \mathbf{d}_t^P as positive vectors.

Thereafter, a two-dimensional convolutional layer (Kalchbrenner et al., 2014b) convolves the hidden vector \mathbf{d}_t^* with multiple convolutional kernels of different widths, where $*$ \in $\{N, P\}$. Each kernel corresponds to a linguistic feature detector which extracts a specific pattern of multi-grained n-grams (Kalchbrenner et al., 2014b). A convolutional filter \mathbf{W}_s maps hidden states in the receptive field to a single feature. As we slide the filter across the negative or positive sequence, a sequence of new features $\mathbf{F}^* = [\mathbf{f}_1^*, \dots, \mathbf{f}_k^*]$ is obtained:

$$\mathbf{f}_t^* = \text{ReLU}(\mathbf{d}_t^* \otimes \mathbf{W}_s + \mathbf{b}_s), \quad (16)$$

where ReLU is activation function, \otimes denotes the convolution operation, $\mathbf{W}_s \in \mathbb{R}^{d \times k}$ and $\mathbf{b}_s \in \mathbb{R}^k$ are learnable parameters in the convolutional filter. For each convolutional filter, the max-pooling layer takes the maximal value among the convolutional features \mathbf{F}^* and results in a fixed-size vector \mathbf{f}^* . Then we obtain the semantic classification result $D_{sem}(\mathbf{d}_t^*) \in \mathbb{R}$ through interaction among \mathbf{f}^* , semantic feedback representation \mathbf{d}^F , and dialogue context vector $\tilde{\mathbf{x}}_0$ (we use emotional context vector $\tilde{\mathbf{w}}_0$ in emotional discriminator):

$$D_{sem}(\mathbf{d}^*) = \mathbf{W}_d \text{ReLU}(\mathbf{f}^* + \mathbf{d}^F + \tilde{\mathbf{x}}_0) + b_d, \quad (17)$$

where $\mathbf{W}_d \in \mathbb{R}^{1 \times d}$ and $b_d \in \mathbb{R}$ are model parameters. \mathbf{d}^F is the last hidden state derived by the above-mentioned LSTM encoder.

Inspired by previous work (Arjovsky et al., 2017; Gao et al., 2019), we minimize the Wasserstein-1 distance between distributions of positive samples and negative samples. Hence, the discriminator $D \in \mathcal{D}$ becomes a *1-Lipschitz* function, where \mathcal{D} is the set of *1-Lipschitz* functions (Gulrajani et al., 2017). In order to meet the *1-Lipschitz* constraints of discriminator D , we incorporate a gradient penalty of the output of D with respect to its input into the discriminator objective function. The gradient penalty is sampled uniformly along a straight line between points sampled from the negative representations \mathbf{d}_t^N and the positive representations \mathbf{d}_t^P . Then the loss function of semantic discriminator is computed as follows:

$$\mathbf{d}'_t = \alpha \mathbf{d}_t^N + (1 - \alpha) \mathbf{d}_t^P, \quad (18)$$

$$\mathcal{L}_d^{sem} = \frac{1}{n} \sum_t D_{sem}(\mathbf{d}_t^N) - D_{sem}(\mathbf{d}_t^P) + \beta (\|\nabla_{\mathbf{d}'_t} D_{sem}(\mathbf{d}'_t)\|_2 - 1)^2, \quad (19)$$

where $\alpha \sim \mathcal{U}[0, 1]$ is a random value and β denotes a coefficient of gradient penalty term.

Emotional Discriminator. The architecture of the emotional discriminator is the same as that of the semantic discriminator. The main difference is that emotional discriminator conducts on the emotional words of the generated response, gold response, and user feedback (i.e., user emotional feedback). We use \mathcal{L}_d^{emo} to denote the loss function of emotional discriminator.

The total loss \mathcal{L}_d of interactive discriminators is the loss summation of semantic discriminator \mathcal{L}_d^{sem} and emotional discriminator \mathcal{L}_d^{emo} . Meanwhile, we add the summation of $-D_{sem}(\mathbf{d}_t^N)$ and $-D_{emo}(\mathbf{d}_t^N)$ to \mathcal{L}_g to facilitate empathetic generator with more empathy ability.

4.3 Training

At the beginning of the training, we use the maximum likelihood estimation (MLE) to pre-train empathetic generator (Eq. 15). Since pre-trained discriminators is effective to help adjust the empathetic generator (Yu et al., 2017), we pre-train the interactive discriminators as well. After the pre-training, the empathetic generator and interactive discriminators are trained alternatively.

5 Experiment

5.1 Dataset

We evaluate EmpDG on the dataset EMPATHETICDIALOGUES (Rashkin et al., 2019), which is a large-scale multi-turn empathetic dialogue dataset collected on the Amazon Mechanical Turk platform, containing

about 25k one-to-one open-domain conversations. Specifically, Rashkin et al. (2019) pair two crowd-workers: a speaker and a listener. The speaker is asked to talk about the personal emotional situation. The listener infers the underlying emotion through what the speaker says and responds empathetically. The dataset provides 32 evenly distributed emotion labels, which act as the coarse-grained dialogue-level emotions. We use the **NRC Emotion Lexicons** (NRC) (Mohammad and Turney, 2013) to extract the emotional words in dialogue context to conduct fine-grained token-level emotions. In order to supplement the language gap between the training data and NRC, all adjectives not included in NRC are extracted together with NRC emotional words. We treat the dialogue context and fine-grained emotional context as system inputs. The target outputs are a coarse-grained emotion label and the listener’s response. For our model, we reserve the next utterance of the target response as user feedback in the training procedure. Finally, we obtain 20,724 dialogues in the training set, 2,972 in the validation set, and 2,713 in the testing set.

5.2 Evaluation Methods

Automatic Evaluation. Liu et al. (2016) have verified BLEU might be improper to evaluate the conversation generation problem, as it correlates weakly with human judgements of the response quality; METEOR (Banerjee and Lavie, 2005) and ROUGE (Lin, 2004) have the same problem. Therefore, following previous emotion-related studies (Zhou et al., 2018; Rashkin et al., 2019; Song et al., 2019; Wei et al., 2019), we employ three evaluation metrics to automatically evaluate the performance of our EmpDG: Perplexity (Serban et al., 2015) measures the high-level general quality of the generation model; Distinct-1 and Distinct-2 (Li et al., 2016a) measure the proportion of the distinct unigrams / bigrams in all the generated results to indicate diversity. To evaluate the model at the emotional level, we adopt Emotion Accuracy as the agreement between the ground truth emotion labels and the predicted emotion labels by the empathetic generator.

Human Evaluation. To qualitatively examine model performance from both the content and the empathy perspectives, we also conduct widely-adopted human evaluations. We randomly sample 100 dialogues and their corresponding generations from our model as well as the baselines. We recruit three professional annotators from a third-party company to evaluate the responses generated by different models. All models are evaluated in terms of following 3 metrics: Empathy, Relevance and Fluency (Rashkin et al., 2019; Lin et al., 2019). Empathy measures whether the listener’s responses show the understanding of the speaker’s feelings; Relevance evaluates whether the generated responses are on-topic with the dialogue context; Fluency measures the grammatical correctness and readability of the generated responses. Each metric is rated on five-scale, where 1, 3, and 5 indicate unacceptable, moderate, and excellent performance, respectively.

5.3 Baselines

We conducted extensive experiments to compare EmpDG against the following representative baselines:

- **Transformer** (Vaswani et al., 2017): A Transformer-based sequence to sequence model that is trained based on MLE loss.
- **EmoPrepend-1** (Rashkin et al., 2019): An extension of the Transformer model which incorporates an additional supervised emotion classifier. The whole model is jointly trained by optimizing both the classification and generation loss.
- **MoEL** (Lin et al., 2019): Another extension of Transformer model which softly combines the response representations from different transformer decoders. Each decoder is optimized to focus on one type of emotion accordingly.

Additionally, to better analyze the influence of different components in our model, we also conducted ablation tests as follows:

- **w/o G**: The EmpDG model without the multi-resolution emotion factors, where we only consider the coarse-grained dialogue-level emotions and semantic discriminator.

Models	Accuracy	Perplexity ↓	Distinct-1	Distinct-2	Empathy	Relevance	Fluency
Transformer	-	33.91	1.17	4.70	3.15	3.42	3.64
EmoPrepend-1	0.3328	33.35	1.06	4.29	3.29	3.59	3.71
MoEL	0.3200	33.58	1.38	4.66	3.47	3.88	3.68
EmpDG	0.3431	34.18	1.81	6.94	3.58	3.91	3.67

Table 1: Evaluation results between baselines and our models. The first four metrics are automatic metrics, while the next three metrics are human metrics. **Bold face** indicates leading results in terms of the corresponding metric.

Models	Accuracy	Perplexity ↓	Distinct-1	Distinct-2
EmpDG	0.3431	34.18	1.81	6.94
w/o G	0.3281	33.94	1.27	5.21
w/o D	0.3347	32.66	1.00	3.89

Table 2: Ablation test of different components.

- **w/o D**: The EmpDG model without the interactive discriminators.

5.4 Implementation Details

We implement all models using Pytorch (Paszke et al., 2017)¹ and optimize the models using Adam (Kingma and Ba, 2015) with a mini-batch size of 16. We use pre-trained Glove vectors (Pennington et al., 2014) to initialize the word embedding. During the training of empathetic generator, the learning rate is initialised as 0.0001 and we vary the learning rate following Vaswani et al. (2017). Early stopping is applied when training. When inference, we set the maximum decoding step as 30. All common hyperparameters are the same as the work in (Lin et al., 2019). During the interactive adversarial training, D-steps (for two interactive discriminators) is set to 1 and G-steps (for empathetic generator) is set to 5. Hyper-parameter β in interactive discriminators is set to 0.1. Meanwhile, we employ the teacher-forcing technique from Li et al. (2017a) to increase adversarial training efficiency.

5.5 Performance Comparisons

Automatic Evaluation Results. As shown in Table 1, EmpDG achieves the highest scores for most metrics compared with other baselines. The noticeable improvement indicates the effectiveness of our multi-resolution adversarial empathetic dialogue generation model in empathetic expression and response diversity. Although the perplexity score of EmpDG is slightly worse than the EmoPrepend-1 due to the introduction of interactive discriminators, the other scores for EmpDG are obviously better than EmoPrepend-1. EmpPrepend-1 and MoEL have similar performance, as both of them only learn the coarse-grained dialogue-level emotional label to infer emotional situation and generate responses. Without emotion modelling, Transformer only generates fluent responses based on semantic mapping, but fail to express diverse responses.

We also perform an ablation study for better understanding the contributions of the main parts of our model. As shown in Table 2, after we remove the multi-resolution empathetic generator (i.e., w/o G), both the emotion accuracy and distinct metrics performance become obviously worse, demonstrating the effectiveness of multi-resolution emotion factors in emotional understanding and model generation quality. We also investigate the effect of removing interactive discriminators (i.e., w/o D). We notice that the scores of distinct metrics decrease dramatically. This makes sense because user feedback interacts with generated responses to facilitate empathetic generator optimization. Therefore, applying interactive discriminators can bring performance improvement on appropriate emotional expressions.

¹Our code is available at <https://github.com/qtli/EmpDG>.

Models	Win	Loss	Tie
EmpDG vs Transformer	40.2%	27.5%	32.3%
EmpDG vs EmoP	38.7%	29.4%	31.9%
EmpDG vs MoEL	35.7%	31.2%	33.1%

Table 3: Result of human A/B test. Tests are conducted pairwise between EmpDG and baseline models.

Semantic Context MoEL	I am so grateful for my family due to an incident with my friend . Yeah, it is good about being a while.
Semantic Context Emotional Context EmpDG User Feedback	I am so grateful for my family due to an incident with my friend . grateful , due , incident , friend That is good. What kind of <u>incident</u> ? My friend lost her parents and siblings in a fire and their family was very close to mine. So seeing her go through the grief made me realize how thankful I should be.
Semantic Context MoEL	I recently had a scary experience . I heard gunshots on my street . Oh my god, I am not sure what I would be scary.
Semantic Context Emotional Context EmpDG User Feedback	I recently had a scary experience . I heard gunshots on my street . scary , experience , heard , street Oh no, I am <u>sorry</u> to hear that. What <u>happened</u> ? Everything turned out fine. I think somebody in the nearby neighborhoods may have been hurt, though.

Table 4: Two visualization cases of the cross-attention weights in MoEL and EmpDG. Words with the underline in the generated responses of EmpDG contain context-consistent emotions.

Human Evaluation Results. Table 1 illustrates that EmpDG obtains the best performance on both Empathy and Relevance scores. This suggests that the multi-resolution adversarial model helps capture implicit emotions, improve the topic consistency, and evoke more emotion perceptivity. As the generated responses by Transformer are already fluent and grammatical, we observe there is no obvious difference among models in terms of Fluency. Additionally, we carried out pairwise response comparison to directly compare the dialogue quality gains in Table 3. The results confirm that the responses from EmpDG are more preferred by human judges.

5.6 Analysis of Emotion Interactions

To gain an insight into how well the emotion is expressed in the generated responses, we portrait two examples illustrating the cross-attention weights (between the encoder and decoder of the empathetic generator) of the dialogue context in the Table 4. For the first visualization case, by giving additional attention to the multi-resolution emotional information, the response generated by EmpDG shows much more informative and empathetic, such as words “good” and “incident”. For the second visualization case, although MoEL puts the major attention on the emotional words, i.e., “scary”, the content of its generated response is inappropriate, which is not enough to evoke much empathy. EmpDG can generate the more contentful response with better emotional expressions, e.g., “sorry” and “what happened” which are also relevant to the user feedback. We can conclude that considering user feedback in adversarial training can teach the model to generate more empathetic responses. In total, the effectiveness of the multi-resolution mechanism and interactive discriminators are proved to some extent.

5.7 Case Study

Table 5 shows two examples generated by EmpDG and other baseline models. In the first case, EmpDG generates a coherent and informative response with a proper “hopeful” emotion by replying with words “sorry”, “hope”, “gets”, “better”, whereas baselines fail to understand the negative emotions or express the appropriate contents. In the second case, EmpDG generates the most context-consistent response, which

Emo.	Hopeful
S_Cont.	X_1 : My husband lost a job but I am hoping he can find a full time job soon.
E_Cont.	lost, job, hoping, find, full, time.
Gold	He will, I have faith.
Transformer	That is great! I am sure he will be fine.
EmoPrepend-1	I am sorry to hear that. I hope he gets it.
MoEL	That is a good news. I hope you get it for him.
EmpDG	I am <u>sorry</u> to hear that. I <u>hope</u> he gets <u>better</u> .
Emo.	Grateful
S_Cont.	X_1 : I joined a firm 6 months ago and then I got a promotion for junior manager. X_2 : Congratulations. That sounds like the fast track. X_3 : Yes, they thought that I have some extraordinary talent and that is the reason I got that position.
E_Cont.	joined, firm, ago, promotion, junior, manager, fast, track, thought, extraordinary, talent, reason, position.
Gold	Cool. Maybe next year you will be a partner!
Transformer	I am sorry to hear that.
EmoPrepend-1	That is great! i am glad you got it!
MoEL	That is great, I bet it is hard to do that!
EmpDG	<u>Wow</u> , that is <u>amazing</u> . I am sure they do <u>great</u> .

Table 5: Generated responses from Transformer, EmoPrepend-1, MoEL, and our EmpDG in two different speaker’s emotional states. S_Cont indicates the dialogue semantic context and E_Cont indicates the emotional context. X_i represents the i -th utterance in semantic context. Tokens in underline represent emotion-related words.

contains context-related words (“do”, “great”) and emotion-rated word (“amazing”).

6 Conclusion

In this paper, we propose an multi-resolution interactive empathetic dialogue model (EmpDG) to evoke more emotion perceptivity in dialogue generation. Two components are proposed to improve the performance of empathetic response generation. A multi-resolution empathetic generator combines coarse-grained dialogue-level and fine-grained token-level emotions to capture the nuanced emotional expressions in dialogue and evoke more emotion perceptivity. Two interactive discriminators utilize user feedback as additional context to interact with the generated response and dialogue context to optimize the long-term goal of empathetic conversation generation. Automatic and manual evaluation have shown that EmpDG can generate responses appropriate not only in content but also in empathy.

There are several future directions for this setting. First, one of the potential extension of EmpDG would be incorporating it with external knowledge (e.g., user profile or commonsense knowledge) to help emotion perceptivity. Second, in our setting, the emotional feedback and semantic feedback are separated. It is promising to model the interaction between semantic and emotional feedback. We leave these directions as future work.

Acknowledgements

We want to thank our anonymous reviewers for their feedback.

This work was supported by the National Key R&D Program of China with grant No. 2020YFB1406700, the Natural Science Foundation of China (61972234, 61902219, 61672324, 61672322, 62072279, 62072271), the Key Scientific and Technological Innovation Program of Shandong Province (2019JZZY010129), the Tencent AI Lab Rhino-Bird Focused Research Program (JR201932), the Fundamental Research Funds of Shandong University, the Foundation of State Key Laboratory of Cognitive Intelligence, iFLYTEK, P.R. China (COGOSC-20190003).

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein gan. *arXiv preprint arXiv:1701.07875*.
- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.

- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- 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.*, 93:309–317.
- Pierre Colombo, Wojciech Witon, Ashutosh Modi, James Kennedy, and Mubbasir Kapadia. 2019. Affect-driven dialog generation. In *NAACL*, pages 3734–3743.
- Shen Gao, Zhaochun Ren, Yihong Zhao, Dongyan Zhao, Dawei Yin, and Rui Yan. 2019. Product-aware answer generation in e-commerce question-answering. In *WSDM*, pages 429–437.
- Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. 2019. Dialoguecn: A graph convolutional neural network for emotion recognition in conversation. In *EMNLP-IJCNLP*, pages 154–164.
- Sayan Ghosh, Mathieu Chollet, Eugene Laksana, Louis-Philippe Morency, and Stefan Scherer. 2017. Affect-lm: A neural language model for customizable affective text generation. In *ACL*, pages 634–642.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *NeurIPS*, pages 2672–2680.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of wasserstein gans. In *NeurIPS*, pages 5767–5777.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazaré, and Jason Weston. 2019. Learning from dialogue after deployment: Feed yourself, chatbot! In *ACL*, pages 3667–3684.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Chao-Chun Hsu, Sheng-Yeh Chen, Chuan-Chun Kuo, Ting-Hao K. Huang, and Lun-Wei Ku. 2018. Emotionlines: An emotion corpus of multi-party conversations. In *LREC*.
- Chenyang Huang, Osmar Zaiane, Amine Trabelsi, and Nouha Dziri. 2018. Automatic dialogue generation with expressed emotions. In *NAACL*, pages 49–54.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014a. A convolutional neural network for modelling sentences. In *ACL*, pages 655–665.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014b. A convolutional neural network for modelling sentences. In *ACL*, pages 655–665.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Jingyuan Li and Xiao Sun. 2018. A syntactically constrained bidirectional-asynchronous approach for emotional conversation generation. In *EMNLP*, pages 678–683.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL*, pages 110–119.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016b. Deep reinforcement learning for dialogue generation. In *EMNLP*, pages 1192–1202.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017a. Adversarial learning for neural dialogue generation. In *EMNLP*, pages 2157–2169.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. Dailydialog: A manually labelled multi-turn dialogue dataset. In *IJCNLP*, pages 986–995.
- Xin Li, Piji Li, Wei Bi, Xiaojiang Liu, and Wai Lam. 2020. Relevance-promoting language model for short-text conversation. In *AAAI*, pages 8253–8260.
- Zhaojiang Lin, Andrea Madotto, Jamin Shin, Peng Xu, and Pascale Fung. 2019. Moel: Mixture of empathetic listeners. In *EMNLP-IJCNLP*.

- Zhaojiang Lin, Peng Xu, Genta Indra Winata, Farhad Bin Siddique, Zihan Liu, Jamin Shin, and Pascale Fung. 2020. Caire: An end-to-end empathetic chatbot. In *AAAI*, pages 13622–13623.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*, pages 2122–2132.
- Nurul Lubis, Sakriani Sakti, Koichiro Yoshino, and Satoshi Nakamura. 2018. Eliciting positive emotion through affect-sensitive dialogue response generation: A neural network approach. In *AAAI*.
- Saif M Mohammad and Peter D Turney. 2013. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Helmut Prendinger and Mitsuru Ishizuka. 2005. The empathic companion: A character-based interface that addresses users’ affective states. *Applied Artificial Intelligence*, 19(3-4):267–285.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2018. I know the feeling: Learning to converse with empathy. *CoRR*, abs/1811.00207.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *ACL*, pages 5370–5381.
- Alexey Romanov, Anna Rumshisky, Anna Rogers, and David Donahue. 2019. Adversarial decomposition of text representation. In *NAACL-HLT*, pages 815–825.
- Sashank Santhanam and Samira Shaikh. 2019. Emotional neural language generation grounded in situational contexts. *CoRR*, abs/1911.11161.
- João Sedoc, Sven Buechel, Yehonathan Nachmany, Anneke Buffone, and Lyle Ungar. 2020. Learning word ratings for empathy and distress from document-level user responses. In *LREC*, pages 1664–1673.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2015. Hierarchical neural network generative models for movie dialogues. *CoRR*, abs/1507.04808.
- Jamin Shin, Peng Xu, Andrea Madotto, and Pascale Fung. 2019. Happybot: Generating empathetic dialogue responses by improving user experience look-ahead. abs/1906.08487.
- Marcin Skowron, Mathias Theunis, Stefan Rank, and Arvid Kappas. 2013. Affect and social processes in on-line communication—experiments with an affective dialog system. *IEEE Transactions on Affective Computing*, 4(3):267–279.
- Zhenqiao Song, Xiaoqing Zheng, Lu Liu, Mu Xu, and Xuan-Jing Huang. 2019. Generating responses with a specific emotion in dialog. In *ACL*, pages 3685–3695.
- Haoyu Song, Yan Wang, Weinan Zhang, Xiaojiang Liu, and Ting Liu. 2020. Generate, delete and rewrite: A three-stage framework for improving persona consistency of dialogue generation. In *ACL*, pages 5821–5831. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NeurIPS*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*, pages 6000–6010.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Wei Wei, Jiayi Liu, Xianling Mao, Guibing Guo, Feida Zhu, Pan Zhou, and Yuchong Hu. 2019. Emotion-aware chat machine: Automatic emotional response generation for human-like emotional interaction. In *CIKM*, pages 1401–1410.

- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*.
- Emmanuelle Zech and Bernard Rimé. 2005. Is talking about an emotional experience helpful? effects on emotional recovery and perceived benefits. *Clinical Psychology & Psychotherapy: An International Journal of Theory & Practice*, 12(4):270–287.
- Wei-Nan Zhang, Lingzhi Li, Dongyan Cao, and Ting Liu. 2018a. Exploring implicit feedback for open domain conversation generation. In *AAAI*.
- Weinan Zhang, Yiming Cui, Yifa Wang, Qingfu Zhu, Lingzhi Li, Lianqiang Zhou, and Ting Liu. 2018b. Context-sensitive generation of open-domain conversational responses. In *COLING*, pages 2437–2447.
- Peixiang Zhong, Di Wang, and Chunyan Miao. 2019. An affect-rich neural conversational model with biased attention and weighted cross-entropy loss. In *AAAI*, volume 33, pages 7492–7500.
- Peixiang Zhong, Yan Zhu, Yong Liu, Chen Zhang, Hao Wang, Zaiqing Nie, and Chunyan Miao. 2020. Endowing empathetic conversational models with personas. *arXiv preprint arXiv:2004.12316*.
- Xianda Zhou and William Yang Wang. 2018. Mojtalk: Generating emotional responses at scale. In *ACL*, pages 1128–1137.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *AAAI*.