

Don't Lose Yourself!

Empathetic Response Generation via Explicit Self-Other Awareness

Weixiang Zhao, Yanyan Zhao*, Xin Lu, Bing Qin
Research Center for Social Computing and Information Retrieval
Harbin Institute of Technology, China
{wxzhao, yyzhao, xlu, qinb}@ir.hit.edu.cn

Abstract

As a critical step to achieve human-like chatbots, empathetic response generation has attained increasing interests. Previous attempts are incomplete and not sufficient enough to elicit empathy because they only stay on the initial stage of empathy to automatically sense and simulate the feelings and thoughts of others via other-awareness. However, they ignore to include self-awareness to consider the own views of the self in their responses, which is a crucial process to achieve the empathy. To this end, we propose to generate **Empathetic response with explicit Self-Other Awareness (EmpSOA)**. Specifically, three stages, self-other differentiation, self-other modulation and self-other generation, are devised to clearly maintain, regulate and inject the self-other aware information into the process of empathetic response generation. Both automatic and human evaluations on the benchmark dataset demonstrate the superiority of EmpSOA to generate more empathetic responses. Our source code is available at <https://github.com/circle-hit/EmpSOA>.

1 Introduction

Empathy is a desirable trait of the engaging human conversation and is considered as the key step to human-like chatbots. In this paper, we focus on the task of empathetic response generation (Rashkin et al., 2019), which understands the feelings and situations of the user and responses properly.

According to one of the most influential theories of empathy proposed by Rogers (1992), *empathy is the ability to sense the others' private world as if it were our own, but without losing the "as if" condition*. Previous works (Lin et al., 2019; Li et al., 2020; Sabour et al., 2022) mainly focus on the prior part of the definition, which is referred as emotional contagion (Hatfield et al., 1993), and

*Corresponding author

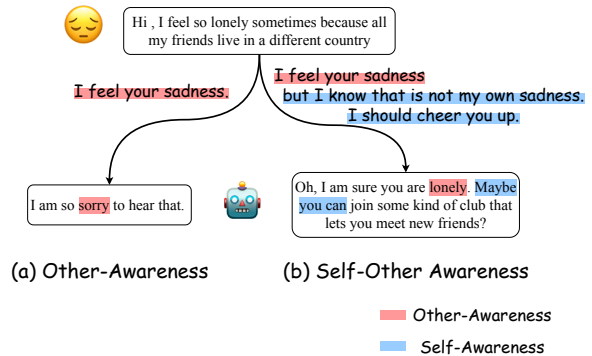


Figure 1: An example for the self-other awareness during an empathetic conversation from the EMPATHETIC-DIALOGUES (Rashkin et al., 2019) dataset. The user and the system are referred as the other and the self.

they automatically mimic the thoughts and feelings of the speaker to converge emotionally.

However, emotional contagion is only the initial component that precedes empathy and what makes empathy distinct from emotion contagion is a critical process called **self-other awareness** (Decety and Lamm, 2006). Thus, it is not sufficient enough for previous attempts to convey empathy because they only perform the single **other-awareness** to perceive the emotion and situation from the other and generate responses coupled with the same perceived emotion. As shown in the left part of Figure 1, with other-awareness, the self succeed to *feel the sadness* of the other. But the complete emotional overlap would induce the self to reinforce such sadness with the response of *so sorry to hear that*, which is not the goal of empathy.

And it is of great necessity to take the own view of the self into consideration to maintain the “as if” condition with a clear **self-awareness**, which is a conscious process to maintain and modulate the own views of the self during the empathetic interaction. In the right part of Figure 1, with the incorporation of the self-awareness, the self consciously holds in mind *that this it not my own sad-*

ness and I am responsible to *cheer you up*. Thus, the explicit self-other awareness plays pivotal roles in disentangling feelings and views of the self and the other, which constitutes a crucial perspective of empathy and largely contributes to generate more empathetic responses, especially when the other is in negative emotional states.

To this end, we propose to generate **Empathetic response with explicit Self-Other Awareness (EmpSOA)**. Inspired by the conceptual framework of information flow involved in human empathy (Decety and Lamm, 2006), we make such processes computable and abstract three stages in EmpSOA, named Self-Other Differentiation (SOD), Self-Other Modulation (SOM) and Self-Other Generation (SOG). Specifically, in SOD, we construct two heterogeneous graphs with four types of nodes to maintain the self-awareness representation and other-awareness representation, respectively. Among them, commonsense knowledge from COMET (Bosselut et al., 2019) is leveraged to manifest the fine-grained emotional and cognitive statuses of the self and the other. Further, we dynamically control the contributions of the self-other awareness representations in SOM and inject them into the process of empathetic response generation in SOG. Experimental results of both automatic and manual evaluations on the benchmark dataset demonstrate the superiority of EmpSOA to generate more empathetic responses.

The main contributions of this work are summarized as follows:

- We propose to generate empathetic responses via explicit self-other awareness, which constitutes a critical perspective of empathy.
- We devise a novel model EmpSOA to clearly maintain, modulate and inject the self-other aware information into the process of empathetic response generation.
- Results of extensive experiments on the benchmark dataset demonstrate the effectiveness of EmpSOA to identify the exact emotion of the other and generate more empathetic response.

2 Related Work

2.1 Empathetic Response Generation

Endowing empathy to the dialogue systems has gained more and more attentions recently. For previous attempts on empathetic response generation,

we divide them into two categories according to whether they incorporate both affection and cognition aspects of empathy. On the one hand, most existing works (Alam et al., 2018; Rashkin et al., 2019; Lin et al., 2019; Majumder et al., 2020; Li et al., 2020, 2022; Wang et al., 2022) only consider the affective aspect of empathy to understand the emotional state of the other and converge emotionally. On the other hand, Sabour et al. (2022) propose to comprehensively understand the emotional feelings and cognitive situations of the other with commonsense knowledge incorporated.

However, all previous methods only perceive the emotional or cognitive states of the other by the single other-awareness, ignoring to explicitly incorporate self-awareness to make an appropriate empathetic response with own views of the self.

2.2 Emotional Dialogue Generation

Emotion has been proven to be the key factor of achieving more engaging dialogue systems. Previous works explore two ways of incorporating emotion into dialogue generation. On the one hand, the generation-based methods (Zhou et al., 2018; Zhou and Wang, 2018; Shen and Feng, 2020) are proposed to generate emotional responses given a specified emotion label. On the other hand, retrieval-based (Qiu et al., 2020; Lu et al., 2021) methods aim to obtain emotional responses from candidates retrieved from the response repository. However, expressing the specified emotion in responses is merely the fundamental goal to achieve emotional dialogue systems, which is lack of the understanding for user’s feelings and situations required by the empathetic response generation.

3 Methodology

3.1 Task Definition

First, we define the task of empathetic response generation. Formally, let $D = [X_1, X_2, \dots, X_N]$ denotes a dialogue history with N utterances between the user (the other) and the system (the self), where the i -th utterance $X_i = [w_1^i, w_2^i, \dots, w_m^i]$ is a sequence of m words. Besides, each conversation is provided with an emotion label e from the total 32 available emotions to signal what the emotional tone that the other is grounded on. The goal is to generate the next utterance Y from the stand of the self that is coherent to the dialogue history D and empathetic to the other’s situation and feeling.

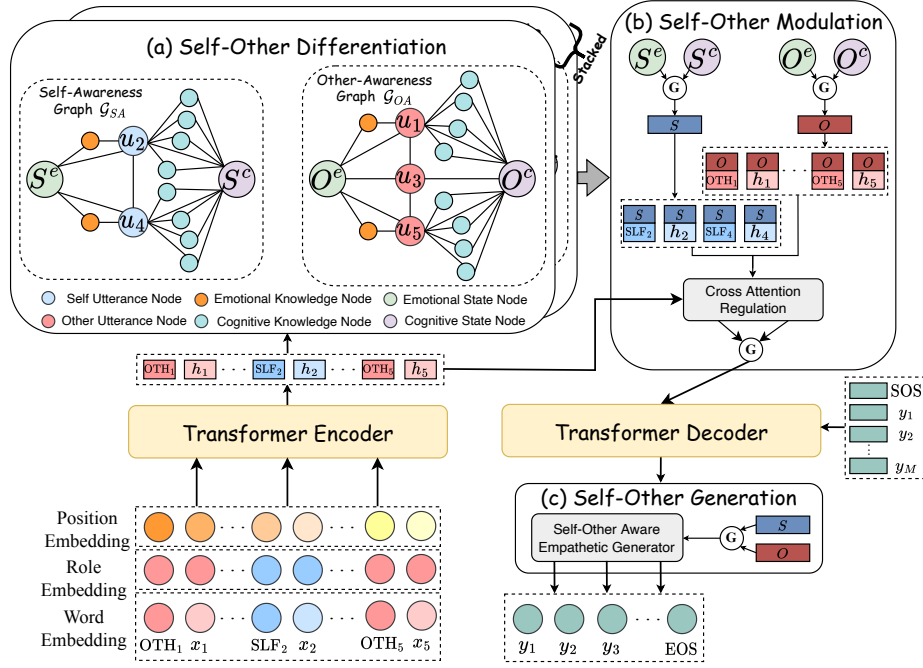


Figure 2: The overall architecture of our proposed EmpSOA model, which mainly consists of three modules: (a) Self-Other Differentiation; (b) Self-Other Modulation and (c) Self-Other Generation.

3.2 Overview of the Architecture

We display the overall architecture of EmpSOA in Figure 2. We abstract three main stages from the conceptual framework¹ of information flow involved in human empathy (Decety and Lamm, 2006) and make them computable in EmpSOA, including (a) Self-Other Differentiation (SOD), (b) Self-Other Modulation (SOM), and (c) Self-Other Generation (SOG). We first clearly disentangle the emotional and cognitive states of the self and the other to maintain the self- and other-awareness representations individually in SOD. Then in SOM, they are dynamically modulated and controlled to make different contributions to the self-other aware contextual information obtained from the context encoder. Finally, such self- and other-awareness representations are explicitly injected into the generation process in SOG to obtain the empathetic responses from views of both the self and the other.

3.3 Self-Other Aware Context Encoder

We adopt Transformer encoder (Vaswani et al., 2017) to obtain the contextual representations of the dialogue history. Following previous works (Li et al., 2022; Sabour et al., 2022), the dialogue is flattened into a word sequence. To make the encoder aware of the self-other distinction in the encoding

phase, we append two special tokens, [SLF] and [OTH], to the beginning of each utterance from the self and the other, respectively. Further, role embedding is added to supplement extra self-other aware information. The final input of the self-other aware context encoder are the sum of word embedding, role embedding and position embedding:

$$H^{so} = \text{Encoder}(E_w + E_r + E_p) \quad (1)$$

where $H^{so} \in \mathbb{R}^{N \times d_h}$ and d_h is the hidden dimension of the self-other aware context encoder.

3.4 Self-Other Differentiation

As mentioned above, the clear self-other awareness constitutes a crucial perspective of genuine empathy. To achieve this, we first devise self-other differentiation (SOD). Specifically, we construct two heterogeneous graphs, named self-awareness graph \mathcal{G}_{SA} and other-awareness graph \mathcal{G}_{OA} , to disentangle and maintain self- and other-awareness representations separately. Inspired by Sabour et al. (2022), both awareness representations consist of emotional and cognitive aspects. And we leverage commonsense knowledge from the external knowledge base ATOMIC (Sap et al., 2019) to imply the fine-grained emotional and cognitive knowledge of the self and the other at each dialogue turn. Such knowledge is highly related to the personal mental

¹More details about the conceptual framework of human empathy are provided in appendix file.

states and it has been widely used in many emotional dialogue-related tasks (Ghosal et al., 2020; Zhao et al., 2022a,b; Tu et al., 2022; Sabour et al., 2022; Peng et al., 2022).

Graph Construction. Since the way of constructing \mathcal{G}_{SA} and \mathcal{G}_{OA} is symmetrical, we only elaborate the formation of \mathcal{G}_{SA} for simplicity.

Nodes: There are four types of heterogeneous nodes in \mathcal{G}_{SA} to form the node set \mathcal{V}_{SA} , including (1) utterance nodes u_i , which are utterances in the dialogue history from the turn of the self; (2) external knowledge nodes e_i and c_i , which are the commonsense knowledge to imply emotional feelings and cognitive situations of the self at the fine-grained level; (3) emotional state node S^e and (4) cognitive state node S^c of the self.

Edges: We build edges \mathcal{E}_{SA} among these nodes to connect (1) adjacent utterance nodes; (2) each utterance node with its corresponding two external knowledge nodes; (3) the emotional state node with all utterance nodes and emotional knowledge nodes and (4) the cognitive state node with all utterance nodes and cognitive knowledge nodes.

Graph Initialization. We also take the self-awareness graph \mathcal{G}_{SA} as an example to describe how to initialize the four types of nodes. And the initialization ways and types of commonsense knowledge are same in both \mathcal{G}_{SA} and \mathcal{G}_{OA} .

For **utterance nodes** u_i , we obtain the corresponding hidden states of special tokens SLF_i from the self-other aware contextual representation H^{so} .

For **external knowledge nodes**, the generative commonsense transformer model COMET (Bosse-lut et al., 2019) is adopted to obtain the emotional and cognitive knowledge. We select relation types $xReact$ to manifest the emotional feelings and $\{xIntent, xNeed, xWant, xEffect\}$ ² to infer the cognitive situations at each dialogue turn of the self, which are consistent with those used in Sabour et al. (2022). Specifically, we adopt the BART-based (Lewis et al., 2020) variation of COMET, which is trained on the ATOMIC-2020 dataset (Hwang et al., 2021). And given each utterance X_i belonging to the self to form the input format $(X_i, r, [GEN])$, COMET would generate descriptions of inferential content under the relation r . Then hidden state representations from the last layer of COMET are obtained to initialize c_i and e_i .

²Please refer to the appendix file for the detailed definitions of all the relations.

For **emotional state node** and **cognitive state node** of the self, we randomly initialize them.

Self-Other Aware Graph Attention. Based on the self-awareness graph \mathcal{G}_{SA} and the other-awareness graph \mathcal{G}_{OA} , we apply the multi-head graph attention mechanism to update the node representations in each graph. Concretely, the graph attention operated on the node representation to update it from the information of other neighbourhoods can be written as:

$$\hat{v}_i = \parallel_{n=1}^H \left(\sum_{j \in N_i} \alpha_{ij} W_v^n v_j \right) \quad (2)$$

where \parallel denotes the concatenation of H attention heads, N_i is the neighbors of node i , and W_v^n is the linear transformation.

The attention weight of n -th head α_{ij}^n is utilised to measure the importance and relevance between the current node and its neighbours:

$$\alpha_{ij}^n = \frac{\exp((W_q^n v_i)^\top (W_k^n v_j))}{\sum_{j' \in N_i} \exp((W_q^n v_i)^\top (W_k^n v_{j'}))} \quad (3)$$

where W_q^n and W_k^n are both linear transformation.

After L stacked layers of multi-head self-other aware graph attention, both emotional state node and cognitive state node of the self and the other would aggregate the fine-grained self-other aware information, achieving the clear self-other differentiation for the following parts of our model.

Emotion Perception. Since we are provided with the golden emotion label of each conversation, an emotion classifier is devised to accurately comprehend the emotional state of the other. Unlike previous attempts that perform emotion classification without a clear differentiation between the emotional state of the self and the other (Li et al., 2020, 2022; Sabour et al., 2022), we exactly focus on the emotional state of the other O^e and the average of corresponding [OTH] tokens from H^{so} :

$$h_e = \text{Average}(OTHs) + O^e \quad (4)$$

where $OTHs$ is the sequence of OTH_i derived from H^{so} . Then, we pass h_e through a linear layer followed by the softmax operation to generate the emotion category distribution P_{emo} :

$$P_{emo} = \text{softmax}(W^e h_e) \quad (5)$$

where $P_{emo} \in \mathbb{R}^{n_e}$, $W^e \in \mathbb{R}^{d_h \times n_e}$ and n_e is the number of total available emotion categories.

During training, we perform the parameter learning by minimizing the Cross-Entropy (CE) loss between the emotion category distribution P_{emo} and the ground truth label e :

$$\mathcal{L}_{emo} = -\log(P_{emo}(e)) \quad (6)$$

3.5 Self-Other Modulation

Through SOD, we differentiate and maintain the self- and other-awareness representations. And what followed is Self-Other Modulation (SOM) module, a conscious and controlled process to determine to what extent we pay attention to them.

First, the emotional state and the cognitive state are dynamically fused by a gate mechanism to obtain the joint self-awareness representation S :

$$\begin{aligned} S &= g^s \odot S^e + (1 - g^s) \odot S^c \\ g^s &= \sigma([S^e; S^c]W^s + b^s) \end{aligned} \quad (7)$$

Similarly, the fused other-awareness representation O can be obtained by:

$$\begin{aligned} O &= g^o \odot O^e + (1 - g^o) \odot O^c \\ g^o &= \sigma([O^e; O^c]W^o + b^o) \end{aligned} \quad (8)$$

where \odot is the element-wise multiplication, σ is the sigmoid activation function and $W^s \in \mathbb{R}^{2d_h \times d_h}$, $W^o \in \mathbb{R}^{2d_h \times d_h}$, $b^s \in \mathbb{R}^{d_h}$ and $b^o \in \mathbb{R}^{d_h}$ are all trainable parameters.

Then, to refine the context with the self-other aware information, we respectively concatenate S and O to their corresponding self-other aware contextual representation H^{so} at the token level:

$$\hat{H}^s[i] = S \oplus H^s[i] \quad (9)$$

$$\hat{H}^o[i] = O \oplus H^o[i] \quad (10)$$

where H^s and H^o are the slices of H^{so} and belongs to the self and the other, respectively.

And Feed Forward Neural Network (Vaswani et al., 2017) is applied to perform the self-other aware context refinement in the point-wise way:

$$\tilde{H}^s = \text{FFN}^s(\hat{H}^s) \quad (11)$$

$$\tilde{H}^o = \text{FFN}^o(\hat{H}^o) \quad (12)$$

Finally, we adopt the cross attention mechanism to control and modulate the contribution of self-awareness context and other-awareness context:

$$C^s = \text{CROSS-ATT}^s(H^{so}, \tilde{H}^s) \quad (13)$$

$$C^o = \text{CROSS-ATT}^o(H^{so}, \tilde{H}^o) \quad (14)$$

And a gate is applied to obtain the modulated self-other aware contextual representation:

$$\begin{aligned} C^{so} &= g \odot C^s + (1 - g) \odot C^o \\ g &= \sigma([C^s; C^o]W^m + b^m) \end{aligned} \quad (15)$$

where $W^m \in \mathbb{R}^{2d_h \times d_h}$ and $b^m \in \mathbb{R}^{d_h}$ are trainable parameters.

3.6 Self-Other Generation

Finally, we devise the self-other generation (SOG) to inject the self-other aware information into the process of empathetic response generation.

For the target response $Y = [y_1, y_2, \dots, y_M]$, to generate the t -th word y_t , we firstly feed the previous generated words $y_{1:t-1}$ into the vanilla Transformer decoder (Vaswani et al., 2017). It is worth to mention that the input of cross attention is modified to the self-other aware contextual representation C^{so} derived from the SOM module:

$$h_t = \text{Decoder}(E_{y_{<t}}, C^{so}) \quad (16)$$

where $E_{y_{<t}}$ denotes the embeddings of the generated words before the time step t .

Then to make the generation process grounded on both views of the self and the other, we dynamically inject self- and other- awareness representations via the fusion of them and the hidden representation h_t of the t -th token:

$$\begin{aligned} h &= h_t + g^f \odot S + (1 - g^f) \odot O \\ g^f &= \sigma([h_t; S; O]W^f + b^f) \end{aligned} \quad (17)$$

where $W^f \in \mathbb{R}^{3d_h \times d_h}$ and $b^f \in \mathbb{R}^{d_h}$ are trainable parameters.

The distribution over the vocabulary for the t -th token can be obtained by a softmax layer:

$$P(y_t | y_{<t}, D) = \text{softmax}(Wh + b) \quad (18)$$

where D is the input dialogue history, $W \in \mathbb{R}^{|V| \times d_h}$ and V is the vocabulary size.

We utilise the standard negative log-likelihood as the generation loss function:

$$L_{gen} = -\sum_{t=1}^M \log P(y_t | D, y_{<t}). \quad (19)$$

A multi-task learning framework is adopted to jointly minimize the emotion perception loss, the generation loss and the diversity loss proposed by Sabour et al. (2022):

$$L = \gamma_1 L_{emo} + \gamma_2 L_{gen} + \gamma_3 L_{div} \quad (20)$$

where γ_1 , γ_2 and γ_3 are hyper-parameters.

4 Experiments

4.1 Dataset

We conduct our experiments on EMPATHETICDIALOGUES dataset (Rashkin et al., 2019). It is a large-scale multi-turn dataset with 25k empathetic conversations collected on the Amazon Mechanical Turk. In this dataset, empathetic conversations are carried out between a speaker and a listener (which is referred as the other and the self in this paper). In addition, 32 evenly distributed emotion labels are provided to signal the personal emotional feelings of the other. We use the same dataset split of 8:1:1 train/valid/test with that in Rashkin et al. (2019).

4.2 Baselines

We compare our proposed EmpSOA with the following competitive baselines.

- **Transformer** (Vaswani et al., 2017): The vanilla Transformer-based encoder-decoder generation model.
- **Multi-Task Transformer (Multi-TRS)** (Rashkin et al., 2019): A variation of the vanilla Transformer with an additional structure to perform emotion perception.
- **MoEL** (Lin et al., 2019): A Transformer-based model that captures emotions of the other and outputs an emotion distribution with multi decoders. By softly combining the output emotion distribution, each decoder is optimized to react to certain emotions, and generate an empathetic response.
- **MIME** (Majumder et al., 2020): Another Transformer-based model that mimics the emotion of the other to a varying degree by grouping emotions into two clusters. Stochasticity is introduced to yield emotionally more varied empathetic responses.
- **EmpDG** (Li et al., 2020): An adversarial empathetic response generation model that exploits both the coarse-grained dialogue-level and fine-grained token-level emotions, and the interactive user feedback.
- **KEMP** (Li et al., 2022): An encoder-decoder model that leverages external knowledge, including commonsense knowledge and emotional lexical knowledge, to explicitly understand and express emotions in empathetic dialogue generation.

- **CEM** (Sabour et al., 2022): For the first time to focus on both affection and cognition of empathy and leverage commonsense to draw more information about the user’s situation. It uses this additional information to enhance the empathy expression in generated responses.

4.3 Implementation Details

We implemented all the baselines and our model with 5 random runs. 300-dimensional pre-trained GloVe vectors (Pennington et al., 2014) are adopted to initialize the word embeddings and shared between the encoder and the decoder. The hidden dimension d_h is also set to 300 and the number of attention heads in SOD graph attention and SOM cross attention are 6. Loss weights γ_1 , γ_2 and γ_3 are set to 1, 1, and 1.5, respectively. Adam (Kingma and Ba, 2015) optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.98$ is used for training. Following Vaswani et al. (2017), we vary the learning rate during the training process with the initial learning rate of 0.0001. Early stopping is applied when training. And the training process is performed on one single Tesla V100 GPU with a mini-batch size of 16. For inference, we use a batch size of 1 and a maximum of 30 decoding steps for all models.

4.4 Evaluation Metrics

Automatic Evaluation. We apply three kinds of automatic metrics for evaluation: (1) Perplexity (**PPL**) measures the general quality of the generated responses; (2) Distinct- n (**Dist- n**) (Li et al., 2016) evaluates the diversity of the generated responses by measuring the ratio of unique n -grams; (3) Accuracy (**Acc**) of the emotion perception is utilised to evaluate the model capability for understanding the emotional state of the other. Following CEM, we do not report the word overlap-based automatic metrics such as BLEU (Papineni et al., 2002) because they are not appropriate for evaluating dialogue systems (Liu et al., 2016).

Human Evaluation. Following Sabour et al. (2022), we conduct the aspect-based pairwise preference test for human evaluation. Specifically, 100 response pairs generated by EmpSOA and baselines are randomly sampled. Then we ask 5 professional annotators to choose the better response following three aspects: (1) Coherence (**Coh.**): which response is more coherent and relevant to the dialogue history; (2) Empathy (**Emp.**): which response is more empathetic to show a better under-

Model	PPL	Acc	Dist-1	Dist-2
Transformer	37.62	-	0.43	1.98
Multi-TRS	37.73	32.86	0.43	1.92
MoEL	36.73	31.28	0.56	2.82
MIME	37.37	29.86	0.40	1.66
EmpDG	37.38	30.79	0.42	1.87
KEMP	36.39	36.57	0.61	2.65
CEM	36.49	37.34	0.60	2.85
EmpSOA (Ours)	35.02	48.32	0.71	3.96

Table 1: Comparison of our model against state-of-the-art baselines in terms of the automatic evaluation. The best results among all models are highlighted in bold.

standing of the other’s feelings and situations; (3) Informativeness (**Inf.**): which response contains more information related to the dialogue history.

5 Results and Analysis

5.1 Overall Results

Automatic Evaluation. Illustrated in Table 1, EmpSOA achieves the new state-of-the-art automatic evaluation results. Benefiting from the clear self-other awareness to maintain, modulate and inject the self- and other-awareness representations, EmpSOA is capable of generating empathetic responses of higher quality with the lowest PPL compared to all the baselines. In addition, the improvement on Dist-1 and Dist-2 indicates the superiority of EmpSOA in terms of generating more informative responses at the unigrams and bigrams level. Finally, although the similar external emotional knowledge is explored in CEM and KEMP, the prominent performance on emotion perception of the other can be ascribed to the explicit disentanglement of the emotional states between the self and the other.

Human Evaluation. As shown in Table 2, EmpSOA significantly outperforms three competitive baselines in terms of all three aspects, which demonstrates the superiority of EmpSOA to generate more empathetic and informative responses with explicit self-other awareness. In addition, we adopt the Fleiss’s kappa (Fleiss, 1971) to measure the overall inter-rater agreement. And the agreement ratio falls in the range of [0.41, 0.6], which denotes the moderate agreement.

5.2 Ablation Study

We conduct ablation studies to verify the effectiveness of the three key modules, SOD, SOM and

Comparisons	Aspects	Win	Lose
EmpSOA vs. CEM	Coh.	47.0 [‡]	38.0
	Emp.	47.0 [‡]	37.4
	Inf.	47.6 [‡]	36.6
EmpSOA vs. KEMP	Coh.	45.6 [‡]	34.8
	Emp.	45.8 [‡]	33.8
	Inf.	46.4 [‡]	34.6
EmpSOA vs. EmpDG	Coh.	48.2 [‡]	33.4
	Emp.	49.8 [‡]	35.6
	Inf.	47.8 [‡]	33.2

Table 2: Human evaluation results (%). Ties are not shown. [‡] represent significant improvement with p -value < 0.05 .

Model	PPL	Acc	Dist-1	Dist-2
EmpSOA	35.02	48.32	0.71	3.96
-SOG	36.65	46.36	0.57	2.59
-SOG & SOM	37.43	45.59	0.57	2.51
-SOG & SOM & SOD	37.18	34.71	0.52	2.19

Table 3: Results of ablation study. SOG, SOM, SOD refer to the self-other generation, self-other modulation and self-other differentiation module, respectively.

SOG, proposed in EmpSOA to achieve self-other awareness. Since they are highly correlated with each other, we remove each one of them according to the order of SOG, SOM and SOD individually.

Effect of Self-Other Generation. To investigate the impact of SOG module in generating self-other aware empathetic response, we discard the dynamic fusion of self- and other-awareness representations in each decoding step. Results are displayed in the second row in Table 3. Without the explicit injection of self-other aware information, the general quality and diversity of generation drops significantly. It manifests that it is of vital importance to explicitly offer the model the self-other aware information to generate more empathetic generation from the perspectives of both the self and the other.

Effect of Self-Other Modulation. Subsequently, we remove the SOM module to study the effectiveness of the modulation for the self-other aware information. The dropped results shown in the third row of Table 3 prove the importance to dynamically control and modulate different contributions of self- and other-awareness representations. Further, it reminds us that it is not sufficient enough to generate empathetic response just by the clear self-other differentiation without any self-other aware information incorporating into the generation process.

Emotion Context	Lonely Hi, I feel so lonely sometimes because all my friends live in a different country.
MIME	I am sorry to hear that. Do you have any friends?
EmpDG	That is sad. I am sorry to hear that.
KEMP	I am so sorry to hear that.
CEM	I am so sorry to hear that.
EmpSOA	I am sorry to hear that. Maybe you should try to talk to someone in the city and meet some new people.
Ground-Truth	Oh, I am sure you are lonely. Maybe you can join some kind of club that lets you meet new friends?

Table 4: Case study of the generated empathetic responses by our proposed EmpSOA and the baselines.

Model	PPL	Acc	Dist-1	Dist-2
EmpSOA	35.02	48.32	0.71	3.96
EmpNA	35.88	47.44	0.69	3.34
EmpOA	35.78	47.10	0.63	3.18
EmpSA	38.44	35.07	0.51	2.03

Table 5: Results of deeper analysis on self-other awareness with three variants of EmpSOA.

Effect of Self-Other Differentiation. Finally, the SOD module is discarded and results are shown in the last row. The significant decrease of emotion perception accuracy indicates that SOD make remarkable contribution to disentangle and perceive the emotion of the other. In addition, it is worth to mention that there is no any self-other aware related module in the current model. And compared to the complete EmpSOA, all results of the automatic evaluation decrease significantly, which supports our motivation that the clear self-other awareness contributes the crucial aspect of empathy.

5.3 Deeper Analysis on Self-Other Awareness

In this section, we demonstrate the in-depth analysis on how the explicit self-other awareness leads to more empathetic responses. Results are shown in Table 5. Concretely, three variants of EmpSOA are implemented, including **Empathetic** response generation with **No Awareness** (EmpNA), with **Other Awareness** (EmpOA) and with **Self Awareness** (EmpSA). We will elaborate each one of them.

For **EmpNA**, we merge the self- and other-awareness graphs to construct a single heterogeneous graph without the explicit differentiation between the emotional and cognitive state of the self and the other in the SOD module. Thus, the original self-awareness S and the other-awareness O are replaced by a joint representation and it is fed into the following SOM and SOG. Through this, empathetic response would be generated without any self-other aware information. The decreased perfor-

mance on generation quality and diversity confirms our constructed motivation to perform empathetic responses with self-other awareness.

For **EmpOA** and **EmpSA**, although we still differentiate the self-other awareness in the SOD module, only one of the self-awareness S or the other-awareness O is applied to SOM and SOG. Automatic evaluation results of both EmpOA and EmpSA decrease to a certain degree, which further verifies our claim to consider self-other awareness simultaneously when generating empathetic responses. Interestingly, the performance of EmpSA is much worsen than the other three models, which indicates that being selfish and over-focused on ourselves would neglect the feelings of the other, resulting in an improper way to elicit empathy.

5.4 Case Study

In Table 4, we show a case with responses generated by EmpSOA and the four baselines. It can be observed that all the models succeed to perceive the exact emotional state of the other and express *sorry* to achieve the emotional consensus. However, it is not sufficient enough to elicit empathy only in this way. Through the explicit self-other awareness, EmpSOA not only accurately reaches the emotional state of the other via other-awareness, but also attempts to stay conscious to avoid being overwhelmed by the feelings of the other and provide the valuable suggestions on *meeting new friends* via self-awareness, which is highly consistent with the empathy expressed in the ground-truth.

6 Conclusion and Future Work

In this paper, we propose EmpSOA to generate empathetic responses via explicit self-other awareness. Three stages including Self-Other Differentiation (SOD), Self-Other Modulation (SOM) and Self-Other Generation (SOG) are devised to achieve this goal. Experimental results on both automatic and

human evaluation demonstrate the superiority of EmpSOA to generate more empathetic responses.

In the future, we will explore the theory of self-other awareness in tasks that specified to elicit the positive emotion of the other.

7 Limitations

There are three points to discuss and they may inspire further investigation. First, since the length of empathetic conversations in the current benchmark dataset EMPATHETICDIALOGUES (Rashkin et al., 2019) is relatively short, the theory of self-other awareness could be explored under the circumstance of long conversations to maintain the self-awareness of chatbots for the long run. Second, for the better comprehension of self-other awareness, it is helpful to introduce more commonsense knowledge of higher quality. Finally, current automatic evaluation metrics are still not rational and proper to measure the ability of empathy. It is desirable to build better evaluation metrics for empathetic responses.

8 Ethics Statement

The open-source benchmark dataset EMPATHETICDIALOGUES (Rashkin et al., 2019) used in our experiments is collected by employed crowd-sourced workers, with user privacy protected and no personal information involved. Besides, the participants in our human evaluation are volunteered transparently informed of our research intent, with reasonable wages paid.

Acknowledgements

We thank the anonymous reviewers for their insightful comments and suggestions. This work was supported by the National Key RD Program of China via grant 2021YFF0901602 and the National Natural Science Foundation of China (NSFC) via grant 62176078.

References

Firoj Alam, Morena Danieli, and Giuseppe Riccardi. 2018. Annotating and modeling empathy in spoken conversations. *Computer Speech & Language*, 50:40–61.

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 Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers, pages 4762–4779. Association for Computational Linguistics.

Jean Decety and Claus Lamm. 2006. Human empathy through the lens of social neuroscience. *TheScientificWorldJOURNAL*, 6:1146–1163.

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

Deepanway Ghosal, Navonil Majumder, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. Cosmic: Commonsense knowledge for emotion identification in conversations. *arXiv preprint arXiv:2010.02795*.

Elaine Hatfield, John T Cacioppo, and Richard L Rapson. 1993. Emotional contagion. *Current directions in psychological science*, 2(3):96–100.

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*, pages 6384–6392. AAAI Press.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

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, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119. The Association for Computational Linguistics.

Qintong Li, Hongshen Chen, Zhaochun Ren, Pengjie Ren, Zhaopeng Tu, and Zhumin Chen. 2020. Empdg: Multi-resolution interactive empathetic dialogue generation. In *Proceedings of the 28th International*

- Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 4454–4466. International Committee on Computational Linguistics.
- Qintong Li, Piji Li, Zhaochun Ren, Pengjie Ren, and Zhumin Chen. 2022. **Knowledge bridging for empathetic dialogue generation**. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 10993–11001. AAAI Press.
- Zhaojiang Lin, Andrea Madotto, Jamin Shin, Peng Xu, and Pascale Fung. 2019. **Moel: Mixture of empathetic listeners**. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 121–132. Association for Computational Linguistics.
- 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 *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2122–2132. The Association for Computational Linguistics.
- Xin Lu, Yijian Tian, Yanyan Zhao, and Bing Qin. 2021. **Retrieve, discriminate and rewrite: A simple and effective framework for obtaining affective response in retrieval-based chatbots**. In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 1956–1969. Association for Computational Linguistics.
- Navonil Majumder, Pengfei Hong, Shanshan Peng, Jiankun Lu, Deepanway Ghosal, Alexander F. Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. **MIME: mimicking emotions for empathetic response generation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 8968–8979. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Wei Peng, Yue Hu, Luxi Xing, Yuqiang Xie, Yajing Sun, and Yunpeng Li. 2022. **Control globally, understand locally: A global-to-local hierarchical graph network for emotional support conversation**. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 4324–4330. ijcai.org.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. **Glove: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Lisong Qiu, Yingwai Shiu, Pingping Lin, Ruihua Song, Yue Liu, Dongyan Zhao, and Rui Yan. 2020. **What if bots feel moods?** In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 1161–1170. ACM.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. **Towards empathetic open-domain conversation models: A new benchmark and dataset**. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5370–5381. Association for Computational Linguistics.
- Carl R Rogers. 1992. The necessary and sufficient conditions of therapeutic personality change. *Journal of consulting and clinical psychology*, 60(6):827.
- Sahand Sabour, Chujie Zheng, and Minlie Huang. 2022. **CEM: commonsense-aware empathetic response generation**. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 11229–11237. AAAI Press.
- 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 *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*, pages 3027–3035. AAAI Press.
- Lei Shen and Yang Feng. 2020. **CDL: curriculum dual learning for emotion-controllable response generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 556–566. Association for Computational Linguistics.
- Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen, and Rui Yan. 2022. **MISC: A mixed strategy-aware model integrating COMET for emotional support conversation**. In *Proceedings of the 60th Annual*

Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 308–319, Dublin, Ireland. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Lanrui Wang, Jiangnan Li, Zheng Lin, Fandong Meng, Chenxu Yang, Weiping Wang, and Jie Zhou. 2022. Empathetic dialogue generation via sensitive emotion recognition and sensible knowledge selection. *arXiv preprint arXiv:2210.11715*.

Weixiang Zhao, Yanyan Zhao, Zhuojun Li, and Bing Qin. 2022a. Knowledge-bridged causal interaction network for causal emotion entailment. *arXiv preprint arXiv:2212.02995*.

Weixiang Zhao, Yanyan Zhao, and Xin Lu. 2022b. [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*, pages 4524–4530. ijcai.org.

Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. [Emotional chatting machine: Emotional conversation generation with internal and external memory](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 730–739. AAAI Press.

Xianda Zhou and William Yang Wang. 2018. [Mojitalk: Generating emotional responses at scale](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1128–1137. Association for Computational Linguistics.

A Conceptual Framework of Information Flow in Human Empathy

This framework (Decety and Lamm, 2006) shown in Figure 3 considers empathy as a construct that accounts for a sense of similarity in the feelings experienced by self and other (such translations go both ways, from other-to-self and from self-to-other), without confusion between the two agents. It involves both bottom-up and top-down information processing. Furthermore, it combines representational aspects, i.e., memories that are localized in distributed neural networks that encode information and, when temporarily activated, enable access to this stored information, as well as processes, i.e., computational procedures that are localized and are independent of the nature or modality of the stimulus that is being processed.

Inspired by this framework, we regard the above bottom-up information flow, which performs a sense of similarity in the feelings experienced by self and other, as the Self-Other Differentiation (SOD) to maintain clear self-other awareness in EmpSOA. Moreover, the top-down regulation process is abstracted as the Self-Other Modulation (SOM) to control the weighted contributions of self-other awareness. Finally, the Self-Other Generation (SOG) is similar to the meta-cognitive feedback, taking into account the self- and other-awareness information to elicit empathetic responses.

B Description of ATOMIC Relations

ATOMIC (Sap et al., 2019) is an atlas of everyday commonsense reasoning and organized through textual descriptions of inferential knowledge, where nine if-then relation types are proposed to distinguish causes vs. effects, agents vs. themes, voluntary vs. involuntary events, and actions vs. mental states. We give the brief definition of each relation.

xIntent Why does PersonX cause the event?

xNeed What does PersonX need to do before the event?

xAttr How would PersonX be described?

xEffect What effects does the event have on PersonX?

xWant What would PersonX likely want to do after the event?

xReact How does PersonX feel after the event?

oReact How does others' feel after the event?

oWant What would others likely want to do after the event?

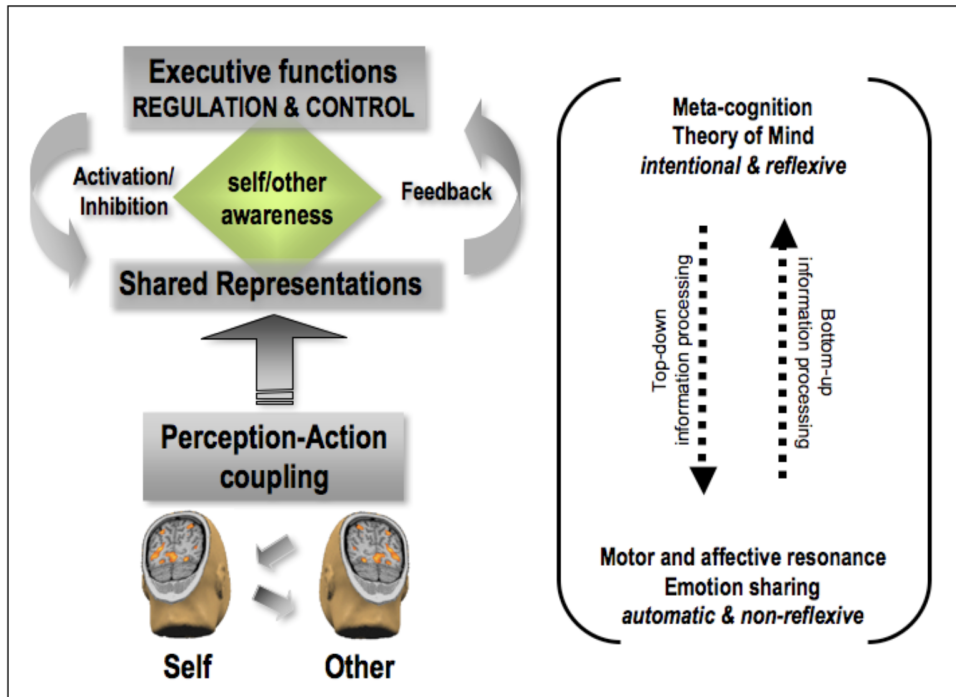


Figure 3: Schematic representation of the bottom-up (i.e., direct matching between perception and action) and top-down (i.e., regulation and control) information processes involved in human empathy. These two levels of processing are interrelated. Top-down regulation, through executive functions, modulates low levels and adds flexibility, making the individual less dependent on external cues. The meta-cognitive feedback plays a crucial role in taking into account one's own mental competence in order to react (or not) to the affective states of others. (Decety and Lamm, 2006)

oEffect What effects does the event have on others?

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 7
- A2. Did you discuss any potential risks of your work?
Section 8
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3,4,5

- B1. Did you cite the creators of artifacts you used?
Section 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section 4
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 4
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We perform experiments on public datasets doing naive incremental modeling works.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 4
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4

C Did you run computational experiments?

Section 4&5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- ✓ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4
 - ✓ C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 4
 - ✓ C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 3.4
- D** ✓ **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 4.4
- ✓ D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Section 4.4
 - ✓ D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Section 4.4 and Section 8
 - ✓ D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Section 5.1
 - ✓ D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Section 8
 - ✓ D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Section 4.4