

CASE: Aligning Coarse-to-Fine Cognition and Affection for Empathetic Response Generation

Jinfeng Zhou^{1,2*†} Chujie Zheng^{1†} Bo Wang^{2‡} Zheng Zhang¹ Minlie Huang¹

¹The CoAI Group, DCST, Institute for Artificial Intelligence, State Key Lab of Intelligent Technology and Systems,

¹Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing 100084, China

²College of Intelligence and Computing, Tianjin University, Tianjin, China

{jfzhou.mail, chujiezhengchn, zhangz.goal}@gmail.com, bo_wang@tju.edu.cn, aihuang@tsinghua.edu.cn

Abstract

Empathetic conversation is psychologically supposed to be the result of conscious alignment and interaction between the cognition and affection of empathy. However, existing empathetic dialogue models usually consider only the affective aspect or treat cognition and affection in isolation, which limits the capability of empathetic response generation. In this work, we propose the CASE model for empathetic dialogue generation. It first builds upon a commonsense cognition graph and an emotional concept graph and then aligns the user’s cognition and affection at both the coarse-grained and fine-grained levels. Through automatic and manual evaluation, we demonstrate that CASE outperforms state-of-the-art baselines of empathetic dialogues and can generate more empathetic and informative responses.¹

1 Introduction

Human empathetic conversations allow both parties to understand each other’s experiences and feelings (Keskin, 2014), which is crucial for establishing seamless relationships (Zech and Rimé, 2005) and is also integral to building a trustful conversational AI (Huang et al., 2020; Wang et al., 2021).

In social psychology, empathy consists of two aspects: cognition and affection (Davis, 1983). The cognitive aspect corresponding to the understanding of the user’s *situation and experiences* (Cuff et al., 2016). The affective aspect requires the comprehension of the user’s *emotional state* and his/her potential *emotional reaction* (Elliott et al., 2018). Although existing work of empathetic dialogue involves both aspects of empathy, there are still issues that need to be addressed. **First**, most work (Rashkin et al., 2019; Lin et al., 2019; Majumder

*Work done during internship at the CoAI Group.

†Equal contribution.

‡Corresponding author.

¹The project repository is available at <https://github.com/jfzhouyoo/CASE>

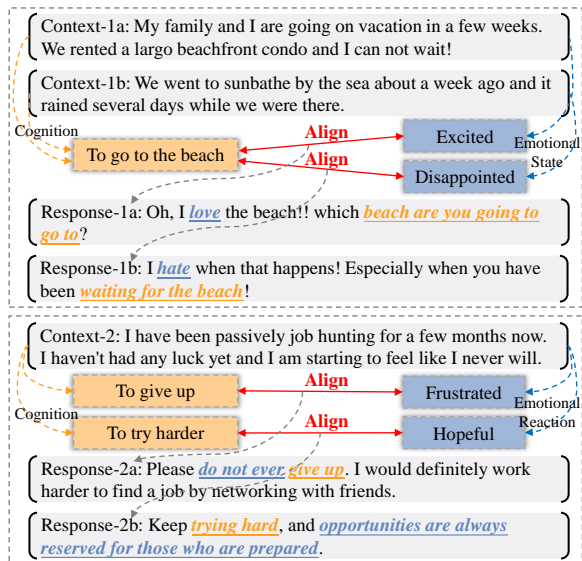


Figure 1: Examples from the EMPATHETICDIALOGUES dataset. The alignment of cognition and affection (i.e., emotional state and emotional reaction) leads to highly empathetic and informative expression in responses.

et al., 2020; Li et al., 2020, 2022) considers only the affective aspect, like detecting the user’s emotional state to enhance empathy expression. **Second**, although recent work explored both roles of cognition and affection in empathy expression (Zheng et al., 2021a; Sabour et al., 2022), they usually treat cognition and affection in isolation without considering their relationship.

However, human empathetic responses often result from conscious alignment and interaction between cognition and affection of empathy (Westbrook et al., 2011). For one thing, the user’s overall emotional state manifested in the context suggests the user’s attitude toward current situation (i.e., cognition). Thus, for the listener, aligning the user’s expressed cognition to the proper emotional state is essential for an appropriate empathetic response. As in case-1 of Figure 1, the alignment of cognition (i.e., *intent “to go to the beach”*) with different emo-

tional states (i.e., “*excited*” vs. “*disappointed*”) produces different appropriate empathetic expressions (i.e., “*love*” and “*which beach are you going to go to*” vs. “*hate*” and “*waiting for the beach*”), respectively. For another, the user’s situation drives the listener to infer the deeper specific cognitions and associate them with the underlying emotional reactions. In this way, the listener can produce a more actively empathetic response instead of only understanding and repeating the user’s expressed cognition. As in case-2 of Figure 1, building association between inferred cognitions and emotional reactions, i.e., “*to give up*” and “*frustrated*” vs. “*to try harder*” and “*hopeful*”, yields cognitively distinct but highly empathetic responses, i.e., *response-2a* vs. *response-2b*. The two cases highlight the necessity of aligning cognition and affection on both overall and specific (i.e., coarse and fine-grained) level for empathy modeling in response generation.

To this end, we align Cognition and Affection for reSponding Empathetically (CASE) on coarse and fine-grained levels by fusing sentence-level commonsense knowledge from COMET (Bosselut et al., 2019) and word-level concept knowledge from ConceptNet (Speer et al., 2017). Commonsense knowledge infers the user’s situation as cognition and infers emotional reactions to the situation, which are implied in the dialogue. Concept knowledge serves to extract the emotional state manifested in the dialogue. For encoding the two types of knowledge, we first construct commonsense cognition graph and emotional concept graph, where the initial independent representation of cognition and emotional concept is carefully adjusted by dialogue context adopting graph transformers. Then, we design a two-level strategy to align cognition and affection using mutual information maximization (MIM) (Appendix A) (Hjelm et al., 2019). The coarse-grained level considers overall cognition and affection manifested in the dialogue context to align contextual cognition and contextual emotional state, which are extracted with a knowledge discernment mechanism. The fine-grained level builds the fine-grained association between cognition and affection implied in the dialogue to align each specific cognition and corresponding emotional reaction. Further, an empathy-aware decoder is devised for generating empathetic expressions.

Our contributions are summarized as follows:

(1) We devise a unified framework to model the interaction between cognition and affection for in-

tegrated empathetic response generation.

(2) We construct two heterogeneous graphs involving commonsense and concept knowledge to aid in the modeling of cognition and affection.

(3) We propose a two-level strategy to align coarse-grained and fine-grained cognition and affection adopting mutual information maximization.

(4) Extensive experiments demonstrate the superiority of CASE in automatic and manual evaluation.

2 Related Work

2.1 Emotional & Empathetic Conversation

Emotional conversation gives the manually specified label preset as the emotion generated in the response (Zhou et al., 2018; Wei et al., 2019; Peng et al., 2022). Instead of giving a predefined emotion label, empathetic conversation (Chen et al., 2022; Kim et al., 2022) involves cognitive and affective empathy (Davis, 1983) and aims to fully understand the interlocutor’s situation and feelings and respond empathically (Keskin, 2014; Zheng et al., 2021b). For one thing, most existing works only focus on the affective aspect of empathy and make efforts to detect contextual emotion (Rashkin et al., 2019; Lin et al., 2019; Majumder et al., 2020; Li et al., 2020, 2022) while ignoring the cognitive aspect. For another, some research utilizes commonsense as cognition to refine empathetic considerations (Sabour et al., 2022). However, the relatively independent modeling between the two aspects (i.e., cognition and affection) violates their interrelated characteristics.

2.2 Commonsense & Concept Knowledge

As a commonsense knowledge base, ATOMIC (Sap et al., 2019) focuses on inferential knowledge organized as typed *if-then* relations. Six commonsense reasoning relations are defined for the person involved in an event, four of which are used to reason commonsense cognitions of a given event, i.e., PersonX’s intent before the event (*xIntent*), what PersonX need to do before the event (*xNeed*), what PersonX want after the event (*xWant*), and the effect of the event on PersonX (*xEffect*). Each commonsense cognition is aligned with user’s emotional reaction to the situation implied in the dialogue inferred by *xReact* (i.e., PersonX’s reaction to the event) in our approach. To obtain inferential commonsense knowledge, we use COMET (Bosselut et al., 2019), a pretrained generative model, to generate rich commonsense statements.

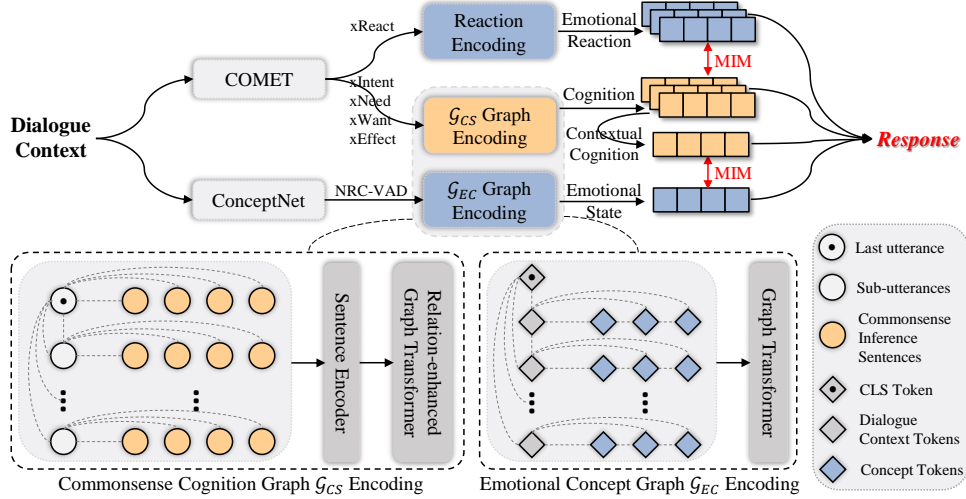


Figure 2: The architecture of the proposed CASE model.

Unlike commonsense knowledge that provides sentence-level commonsense expression, we adopt ConceptNet (Speer et al., 2017) as concept knowledge, which provides word-level human knowledge and is widely used in various NLP tasks (Zhang et al., 2020; Zhong et al., 2021; Zhou et al., 2021; Yang et al., 2022). Following Li et al. (2022), we use NRC_VAD (Mohammad, 2018) to assign emotion intensity to concepts in ConceptNet (processing details are in Li et al. (2022)) severed to extract the contextual emotional state manifested in the context, and align it with contextual cognition.

3 Approach

CASE framework is in Fig. 2. The dialogue context $X = [x_1, \dots, x_N]$ contains N utterances, where x_i denotes the i -th utterance. CASE contains three stages: (1) The graph encoding stage constructs and encodes heterogeneous commonsense cognition graph \mathcal{G}_{CS} and emotional concept graph \mathcal{G}_{EC} from the dialogue context X . (2) The coarse-to-fine alignment aligns coarse-grained (between contextual cognition and contextual emotional state) and fine-grained (between each specific cognition and corresponding emotional reaction) cognition and affection adopting MIM. (3) The empathy-aware decoder integrates the aligned cognition and affection to generate the response $Y = [y_1, y_2, \dots, y_M]$ with empathetic and informative expressions.

3.1 Graph Encoding

Commonsense Cognition Graph Construction Given the last utterance x_N of the dialogue context X , we segment it into the sub-utterances

$U = [u_0, u_1, u_2, \dots, u_t]$, where we prepend the whole x_N as u_0 for maintaining the global information of x_N . We use COMET to infer l commonsense cognition knowledge $K_i^r = [k_{i,1}^r, k_{i,2}^r, \dots, k_{i,l}^r]$ for each $u_i \in U$, where r is one of the four commonsense relations $\mathcal{R} = \{xIntent, xNeed, xWant, xEffect\}$, similar to Sabour et al. (2022). The idea is that human responses tend to inherit the above and transfer the topic. There are differences in the topic and connotation of different sub-utterances affecting the listeners’ concerns when responding empathetically.

For constructing the heterogeneous commonsense cognition graph \mathcal{G}_{CS} , we use the utterance set U and the commonsense cognition knowledge set $K_{CS} = \bigcup_{i=0}^t \bigcup_{r \in \mathcal{R}} K_i^r$ as vertices, i.e., vertex set $V_{CS} = U \cup K_{CS}$. There are seven relations of undirected edges that connect vertices. (1) The *self-loop* relation for each vertex $v_i^{CS} \in V_{CS}$. (2) The *global* relation between the whole x_N (i.e., u_0) and its sub-utterances $u_i (i \geq 1)$. (3) The *temporary* relation between any two successive sub-utterances u_j and u_{j+1} . (4) The four commonsense relations, i.e., $xIntent, xNeed, xWant, xEffect$, between the utterance $u_i \in U$ and the corresponding K_i^r .

We use a Transformer-based sentence encoder (cognition encoder) to first encode the vertices V_{CS} of the graph \mathcal{G}_{CS} . For each $v_i^{CS} \in V_{CS}$, we prepend with a special token $[CLS]$. Following Devlin et al. (2019), we collect the $[CLS]$ representation as the initial embedding matrix for \mathcal{G}_{CS} .

Emotional Concept Graph Construction We concatenate the utterances in the dialogue context X to obtain the token set T , i.e., $T = x_1 \oplus \dots \oplus$

$x_N = [w_1, \dots, w_n]$, where n is the number of all the tokens in the utterances in X . Following Li et al. (2022), we use ConceptNet to infer the related concepts for each token $w_i \in T$, among which only the top N' emotional concepts (according to the emotion intensity $\omega(c)$) are used for constructing \mathcal{G}_{EC} . Subsequently, the vertices V_{EC} in the heterogeneous emotional concept graph \mathcal{G}_{EC} contains a $[CLS]$ token, the dialogue context tokens T , and the above obtained emotional concepts. There are four relations of undirected edges that connect vertices. (1) The *self-loop* relation for each vertex $v_i^{EC} \in V_{EC}$. (2) The *global* relation between the $[CLS]$ token and other ones. (3) The *temporary* relation between any two successive tokens. (4) The *emotional concept* relation among a token and its related emotional concepts.

We initialize the vertex embedding for \mathcal{G}_{EC} by summing up the token embedding, the positional embedding, and the type embedding for each vertex (signaling whether it is a emotional concept or not).

Graph Encoder Given the commonsense cognition graph \mathcal{G}_{CS} , to capture the semantic relationship between vertices, we adopt the Relation-Enhanced Graph Transformer (Li et al., 2021) for graph encoding. It employs a relation-enhanced multi-head attention mechanism (MHA) to encode vertex embedding \hat{v}_{v_i} for vertex v_i (we omit the superscripts CS for simplicity) as:

$$\hat{v}_{v_i} = MHA_{v_k \in V_{CS}}(\mathbf{q}_{v_i}, \mathbf{k}_{v_k}, \mathbf{v}_{v_k}), \quad (1)$$

where the semantic relations between vertices are injected into the query and key vectors:

$$\mathbf{q}_{v_i} = \mathbf{v}_{v_i} + \mathbf{l}_{v_i \rightarrow v_k}, \quad \mathbf{k}_{v_k} = \mathbf{v}_{v_k} + \mathbf{l}_{v_k \rightarrow v_i}, \quad (2)$$

where $\mathbf{l}_{v_i \rightarrow v_k}$ and $\mathbf{l}_{v_k \rightarrow v_i}$ are learnable relation embeddings between vertices v_i and v_k . The self-attention is subsequently followed by a residual connection and a feed-forward layer, as done in the standard Transformer encoder (Vaswani et al., 2017). Finally, we obtain the commonsense cognition embedding \mathbf{cs}_i for each $v_i^{CS} \in V_{CS}$.

To encode the emotional concept graph \mathcal{G}_{EC} , we adopt a vanilla Graph Transformer (i.e., omitting the relation enhancement part in the above Graph Transformer). By superimposing the emotion intensity of each token, we obtain the emotional concept embedding \mathbf{ec}_i for each $v_i^{EC} \in V_{EC}$.

3.2 Coarse-to-Fine Alignment

Context Encoding Following previous works (Majumder et al., 2020; Sabour et al., 2022), we

concatenate all the utterances in the dialogue context X and prepend with a $[CLS]$ token: $[CLS] \oplus x_1 \oplus \dots \oplus x_N$. This sequence is fed into a standard Transformer encoder (context encoder) to obtain the representation \mathbf{S}_X of the dialogue context. We denote the representation of $[CLS]$ as \mathbf{s}_X .

Coarse-grained Alignment To reproduce the interaction of cognition and affection manifested in the dialogue context, we align contextual cognition and contextual emotional state at an overall level. They are separately acquired by cognitive and emotional knowledge discernment mechanisms, which select golden-like knowledge guided by response.

To obtain the contextual cognitive representation \mathbf{r}_{cog} , the knowledge discernment calculates the prior cognitive distribution $P_{CS}(\mathbf{cs}_i | X)$ over the commonsense cognition knowledge (that is, only K_{CS} rather than all the vertices V_{CS} in \mathcal{G}_{CS} , and we thus use $1 \leq i \leq |K_{CS}|$ for simplicity):

$$\mathbf{r}_{cog} = \sum_{i=1}^{|K_{CS}|} P_{CS}(\mathbf{cs}_i | X) \cdot \mathbf{cs}_i, \quad (3)$$

$$P_{CS}(\mathbf{cs}_i | X) = \text{softmax}_i(\mathbf{cs}_i^T \varphi_{CS}(\mathbf{s}_X)), \quad (4)$$

where $\varphi_{CS}(\cdot)$ is a MLP layer activated by tanh. Similarly, we calculate the prior emotional distribution $P_{EC}(\mathbf{ec}_i | X)$ ($1 \leq i \leq |V_{EC}|$) and obtain the contextual emotional representation \mathbf{r}_{emo} .

During training, we use the ground truth response Y to guide the learning of knowledge discernment mechanisms. We feed Y into the cognition encoder (used for initializing the embeddings of \mathcal{G}_{CS} above) and the context encoder to get the hidden states \mathbf{S}_Y^{cog} and \mathbf{S}_Y^{ctx} , where the $[CLS]$ representations are \mathbf{s}_Y^{cog} and \mathbf{s}_Y^{ctx} respectively. The posterior cognitive distribution $P_{CS}(\mathbf{cs}_i | Y)$ and the emotional one $P_{EC}(\mathbf{ec}_i | Y)$ are calculated as:

$$P_{CS}(\mathbf{cs}_i | Y) = \text{softmax}_i(\mathbf{cs}_i^T \mathbf{s}_Y^{cog}), \quad (5)$$

$$P_{EC}(\mathbf{ec}_i | Y) = \text{softmax}_i(\mathbf{ec}_i^T \mathbf{s}_Y^{ctx}). \quad (6)$$

We then optimize the KL divergence between the prior and posterior distributions during training:

$$L_{KL} = L_{KL}^{CS} + L_{KL}^{EC}, \quad (7)$$

$$L_{KL}^{CS} = \sum_{i=1}^{|K_{CS}|} P_{CS}(\mathbf{cs}_i | Y) \cdot \log \frac{P_{CS}(\mathbf{cs}_i | Y)}{P_{CS}(\mathbf{cs}_i | X)},$$

$$L_{KL}^{EC} = \sum_{i=1}^{|V_{EC}|} P_{EC}(\mathbf{ec}_i | Y) \cdot \log \frac{P_{EC}(\mathbf{ec}_i | Y)}{P_{EC}(\mathbf{ec}_i | X)}.$$

To further ensure the accuracy of discerned knowledge, similar to Bai et al. (2021), we employ the BOW loss to force the relevancy between cognitive / emotional knowledge and the target response. The BOW loss L_{BOW} is defined as:

$$L_{BOW} = -\frac{1}{|B|} \sum_{y_t \in B} \log \eta(y_t | \mathbf{r}'_{cog}, \mathbf{r}'_{emo}), \quad (8)$$

where $\eta(\cdot)$ is a MLP layer followed by softmax and the output dimension is the vocabulary size, B denotes the word bags of Y , $\mathbf{r}'_{cog} = \sum_{i=1}^{|K_{CS}|} P_{CS}(\mathbf{cs}_i | Y) \cdot \mathbf{cs}_i$, and $\mathbf{r}'_{emo} = \sum_{i=1}^{|V_{EC}|} P_{EC}(\mathbf{ec}_i | Y) \cdot \mathbf{ec}_i$.

Finally, we align the coarse-grained representations of the contextual cognition \mathbf{r}_{cog} and the contextual emotional state \mathbf{r}_{emo} using mutual information maximization (MIM). Specifically, we adopt the binary cross-entropy (BCE) loss L_{coarse} as the mutual information estimator that maximizes the mutual information between \mathbf{r}_{cog} and \mathbf{r}_{emo} :

$$\begin{aligned} L_{coarse} &= 2f_{coarse}(\mathbf{r}_{cog}, \mathbf{r}_{emo}) \\ &\quad - \log \sum_{\tilde{\mathbf{r}}_{emo}} \exp(f_{coarse}(\mathbf{r}_{cog}, \tilde{\mathbf{r}}_{emo})) \\ &\quad - \log \sum_{\tilde{\mathbf{r}}_{cog}} \exp(f_{coarse}(\tilde{\mathbf{r}}_{cog}, \mathbf{r}_{emo})), \quad (9) \end{aligned}$$

where $\tilde{\mathbf{r}}_{emo}$ and $\tilde{\mathbf{r}}_{cog}$ are the encoded negative samples. $f_{coarse}(\cdot, \cdot)$ is a scoring function implemented with a bilinear layer activated by sigmoid function:

$$f_{coarse}(\mathbf{a}, \mathbf{b}) = \sigma(\mathbf{a}^T \mathbf{W}_{coarse} \mathbf{b}). \quad (10)$$

Fine-grained Alignment To simulate the interaction of fine-grained cognition and affection implied in the dialogue during human express empathy, the fine-grained alignment builds the fine-grained association between each inferred specific cognition and corresponding emotional reaction.

For each $u_i \in U$, we infer the commonsense knowledge about emotional reaction $K_i^{xReact} = [k_{i,1}^{xReact}, \dots, k_{i,l}^{xReact}]$ using COMET, which is regarded as the user's possible emotional reaction to the current cognitive situation. Since $k_{i,j}^{xReact} \in K_i^{xReact}$ is usually an emotion word (e.g., happy, sad), we concatenate K_i^{xReact} and feed it into the Transformer-based encoder (reaction encoder) to get the representation of the emotional reaction \mathbf{H}_i^{er} . Similar to (Majumder et al., 2020) and (Sabour et al., 2022), we use the average pooling to represent the reaction sequence, i.e., $\mathbf{h}_i^{er} =$

Average(\mathbf{H}_i^{er}). To avoid over-alignment of out-of-context emotional reaction with cognition, we inject contextual information into the representation of reaction. We first connect \mathbf{h}_i^{er} with the context representation \mathbf{S}_X at the token level, i.e., $\mathbf{S}_i^{er}[j] = \mathbf{S}_X[j] \oplus \mathbf{h}_i^{er}$. Then another Transformer-based encoder takes \mathbf{S}_i^{er} as input and output the fused information $\mathbf{S}_i^{er'}$. We take the hidden representation of $[CLS]$ in $\mathbf{S}_i^{er'}$ as the emotional reaction representation \mathbf{er}_i of u_i .

Finally, we build the association between the inferred specific cognition $\{\bigcup_{j=1}^l \mathbf{cs}_{i,j}^r\}$ from u_i for $r \in \mathcal{R} = \{\text{xIntent, xNeed, xWant, xEffect}\}$ and the emotional reaction \mathbf{er}_i using MIM. Recall that $\{\bigcup_{i=0}^t \bigcup_{r \in \mathcal{R}} \bigcup_{j=1}^l \mathbf{cs}_{i,j}^r\}$ exactly correspond to the commonsense cognition knowledge set K_{CS} . The fine-grained BCE Loss L_{fine} is defined as:

$$\begin{aligned} L_{fine} &= \sum_{i=0}^t \sum_{r \in \mathcal{R}} \sum_{j=1}^l \left[2f_{fine}(\mathbf{cs}_{i,j}^r, \mathbf{er}_i) \right. \\ &\quad \left. - \log \sum_{\tilde{\mathbf{er}}_i} \exp(f_{fine}(\mathbf{cs}_{i,j}^r, \tilde{\mathbf{er}}_i)) \right. \\ &\quad \left. - \log \sum_{\tilde{\mathbf{cs}}_{i,j}^r} \exp(f_{fine}(\tilde{\mathbf{cs}}_{i,j}^r, \mathbf{er}_i)) \right], \quad (11) \end{aligned}$$

where $\tilde{\mathbf{er}}_i$ and $\tilde{\mathbf{cs}}_{i,j}^r$ are the encoded negative samples. $f_{fine}(\cdot, \cdot)$ is implemented as:

$$f_{fine}(\mathbf{a}, \mathbf{b}) = \sigma(\mathbf{a}^T \mathbf{W}_{fine} \mathbf{b}). \quad (12)$$

Altogether, the coarse-to-fine alignment module can be jointly optimized by L_{align} loss:

$$L_{align} = L_{BOW} + L_{KL} + L_{coarse} + \alpha L_{fine}, \quad (13)$$

where α is a hyper-parameter.

Emotion Prediction We fuse the contextual emotional state and emotional reaction to distill the affective representation, where we use \mathbf{er}_0 as the distillation signal of emotional reaction. This is because \mathbf{er}_0 is derived from the speaker's last utterance and represents the overall emotional reaction. A gating mechanism is designed to capture affective representation \mathbf{r}_{aff} :

$$\mathbf{r}_{aff} = \mu \cdot \mathbf{r}_{emo} + (1 - \mu) \cdot \mathbf{er}_0, \quad (14)$$

$$\mu = \sigma(\mathbf{w}_{aff}^T [\mathbf{r}_{emo}; \mathbf{er}_0]). \quad (15)$$

We project \mathbf{r}_{aff} to predict the user's emotion e :

$$P_{emo}(e) = \text{softmax}(\mathbf{W}_{emo} \mathbf{r}_{aff}), \quad (16)$$

which is supervised by the ground truth emotion label e^* using the cross-entropy loss:

$$L_{emo} = -\log P_{emo}(e^*). \quad (17)$$

3.3 Empathy-aware Response Generation

We employ a Transformer-based decoder to generate the response. To improve empathy perception in response generation, we devise two strategies to fuse the two parts of empathy (i.e., cognition and affection). First, we concatenate the cognitive and affective signals r_{cog} and r_{aff} with the dialogue context representation S_X at the token level, which is then processed by a MLP layer activated by *ReLU* to integrate cognition and affection into the dialogue context:

$$S'_X[i] = MLP(S_X[i] \oplus r_{cog} \oplus r_{aff}). \quad (18)$$

Second, we modify the original Transformer decoder layer by adding two new cross-attention to integrate commonsense cognition knowledge $K_{CS} = \{cs_i\}_{i=1}^{|K_{CS}|}$ and emotional concept knowledge $K_{EC} = \{ec_i\}_{i=1}^{|K_{EC}|}$, which are inserted between the self-attention and cross-attention for S'_X . The decoder then predicts the next token y_m given the previously decoded tokens $y_{<m}$, as done in the standard Transformer decoder. We use the negative log-likelihood loss L_{gen} to optimize the decoder:

$$L_{gen} = - \sum_{m=1}^M \log P(y_m | X, \mathcal{G}_{CS}, \mathcal{G}_{EC}, y_{<m}). \quad (19)$$

Finally, we jointly optimize the alignment loss, emotion prediction loss, generation loss, and diversity loss proposed by Sabour et al. (2022) as: $L = \gamma_1 L_{align} + \gamma_2 L_{emo} + \gamma_3 L_{gen} + \gamma_4 L_{div}$, where $\gamma_1, \gamma_2, \gamma_3$ and γ_4 are hyper-parameters.

4 Experiments

4.1 Experimental Setup

Dataset The experiments are conducted on the widely used EMPATHETICDIALOGUES (Rashkin et al., 2019) dataset, comprising 25k open domain conversations. In a conversation, the speaker confides personal experiences, and the listener infers the situation and emotion of the speaker and responds empathetically. Following Rashkin et al. (2019), we split the train/valid/test set by 8:1:1.

Baselines (1) *Transformer* (Vaswani et al., 2017): A vanilla Transformer-based response generation model. (2) *Multi-TRS* (Rashkin et al., 2019): A multi-task Transformer model that jointly optimizes response generation and emotion prediction. (3) *MoEL* (Lin et al., 2019): An empathy dialogue

Models	PPL	Dist-1	Dist-2	Acc
Transformer	37.65	0.47	2.05	-
Multi-TRS	37.45	0.51	2.12	0.347
MoEL	38.35	0.44	2.10	0.322
MIME	37.33	0.41	1.62	0.296
EmpDG	37.77	0.53	2.26	0.314
KEMP	37.32	0.55	2.31	0.341
CEM	36.86	0.64	2.84	0.373
CASE	35.37	0.74	4.01	0.402

Table 1: Results of automatic evaluation.

model that combines the output of multiple decoders for generating. (4) *MIME* (Majumder et al., 2020): An empathy dialogue model that mimics the user’s emotion for responding. (5) *EmpDG* (Li et al., 2020): An empathy dialogue generator that utilizes multi-resolution user emotions and feedback. (6) *KEMP* (Li et al., 2022): A knowledge-aware empathy dialogue model that only uses concept knowledge. (7) *CEM* (Sabour et al., 2022): A commonsense-aware empathetic chatting machine that only exploits commonsense knowledge.

Implementation Details We implemented all models with Pytorch. We initialize the word embeddings with pretrained GloVe word vectors (Pennington et al., 2014). The dimensionality of embeddings is set to 300 for all corresponding modules. We set hyper-parameters $l = 5$, $N' = 10$, $\alpha = 0.2$, $\gamma_1 = \gamma_2 = \gamma_3 = 1$ and $\gamma_4 = 1.5$. We use Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$ and $\beta_2 = 0.98$. The batch size is 16 and early stopping is adopted. The initial learning rate is set to 0.0001 and we varied it during training following Vaswani et al. (2017). The maximum decoding step is set to 30 during inference. All models are trained on a GPU-P100 machine. The training process of CASE is split into two phases. We first minimize L_{BOW} for pretraining knowledge discernment mechanisms, and then minimize L for training overall model.

4.2 Automatic Evaluation

In the model’s generation evaluation, we adopt the widely used Perplexity (PPL) and Distinct-1/2 (Dist-1/2) (Li et al., 2016). Perplexity evaluates the general generation quality of a model. Distinct-1/2 evaluates the generated diversity by measuring the ratio of unique unigrams/bigrams in the response. In the model’s emotion classification evaluation, we measure the accuracy (Acc) of emotion prediction. Following KEMP and CEM, we do not report

Models	PPL	Dist-1	Dist-2	Acc
CASE	35.37	0.74	4.01	0.402
w/o Graph	36.10	0.68	3.50	0.280
w/o Align	35.75	0.65	3.34	0.369
w/o CSGraph	35.51	0.64	3.18	0.375
w/o ECGraph	36.24	0.72	3.94	0.329
w/o CGAlign	35.67	0.68	3.60	0.378
w/o FGAlign	35.55	0.67	3.43	0.370

Table 2: Results of overall-to-part ablation study.

word overlap-based automatic metrics (Liu et al., 2016), e.g., BLEU (Papineni et al., 2002).

In Table 1, our model outperforms all baselines and achieves a significant improvement on all metrics. **First**, our model achieves about 4.0% reduction on PPL compared to the best baseline, which shows that CASE is more likely to generate ground truth responses. **Second**, our model achieves 15.6% and 41.2% improvement on Dist-1/2 compared to CEM, which indicates the superiority of CASE in generating informative responses at the unigrams and bigrams level. This is attributed to the coarse-to-fine alignment that allows CASE to inject more informative commonsense cognition on the premise of ensuring the perplexity of the generated response. **Third**, our model achieves about 17.9% and 7.8% improvement in prediction accuracy compared to KEMP and CEM, respectively. This verifies that CASE considers both aspects of affection (i.e., contextual emotional state and emotional reaction) more effectively than focusing only on a single aspect as KEMP and CEM.

4.3 Overall-to-Part Ablation Study

We conduct an overall-to-part ablation study in Table 2. In the overall ablation, **first**, we remove the commonsense cognition graph and emotional concept graph, called “w/o Graph”. The emotion prediction accuracy decreases significantly, which indicates that the two heterogeneous graphs make remarkable contribution to detecting emotion. **Second**, we remove the coarse-to-fine alignment, called “w/o Align”. The diversity of generation decreases significantly and emotion prediction accuracy drops distinctly. It supports our motivation that the alignment of cognition and affection leads to informative and highly empathetic expression.

In the part ablation, **first**, we remove two graphs, called “w/o CSGraph” and “w/o ECGraph”, respectively. From the results, we find that the contribution of the commonsense cognition graph is mainly

Comparisons	Aspects	Win	Lose	κ
CASE vs. EmpDG	Coh.	48.1 [‡]	39.2	0.54
	Emp.	51.9 [‡]	32.9	0.55
	Inf.	58.9 [‡]	31.6	0.50
CASE vs. KEMP	Coh.	44.4 [†]	41.8	0.45
	Emp.	50.0 [‡]	34.4	0.53
	Inf.	51.1 [‡]	34.0	0.53
CASE vs. CEM	Coh.	45.9 [‡]	42.2	0.51
	Emp.	53.2 [‡]	34.6	0.47
	Inf.	57.8 [‡]	29.8	0.56

Table 3: Human evaluation results (%) of CASE and baselines. The agreement ratio kappa $\kappa \in [0.41, 0.6]$ denotes the moderate agreement. ^{†,‡} represent significant improvement with p -value $< 0.1/0.05$, respectively.

to improve the diversity of generation (i.e., Dist-1/2), while the role of the emotional concept graph is mainly located in the recognition of emotion (i.e., Acc). This also supports our constructed motivation. **Second**, we remove coarse-grained and fine-grained alignments, called “w/o CGAlign” and “w/o FGAlign”, respectively. We observe that the alignment at the fine-grained level is more significant than the coarse-grained level in terms of overall contribution. This also matches our intuition that building the fine-grained association between cognition and affection is closer to the conscious interaction process during human express empathy.

4.4 Human Evaluation

Human Evaluation of CASE and Baselines

Here, 200 contexts are randomly sampled and each context is associated with two responses generated from our CASE and baseline. Following Sabour et al. (2022), three crowdsourcing workers are asked to choose the better one (**Win**) from two responses by considering three aspects, respectively, i.e., (1) Coherence (**Coh.**): which model’s response is more fluent and context-related? (2) Empathy (**Emp.**): which model’s response expresses a better understanding of the user’s situation and feelings? (3) Informativeness (**Inf.**): which model’s response incorporates more information related to the context? We use the Fleiss’ kappa (κ) (Fleiss, 1971) to measure the inter-annotator agreement. As in Table 3, the results show that CASE outperforms three more competitive baselines on all three aspects. Especially, CASE outperforms baselines significantly in terms of empathy and informativeness, which shows the superior of modeling the interaction between cognition and affection of empathy, and supports the observations from automatic evaluation.

Models	Coh.	Emp.	Inf.	Overall
CASE	3.88	3.48	3.62	3.58
w/o Graph	3.72	3.02	3.51	3.38
w/o Align	3.63	3.14	3.47	3.36
w/o CSGraph	3.74	3.25	3.42	3.33
w/o ECGraph	3.78	3.10	3.52	3.40
w/o CGAlign	3.72	3.27	3.53	3.41
w/o FGAlign	3.80	3.17	3.56	3.40

Table 4: Human evaluation results of CASE’s variants.

Models	PPL	Dist-1	Dist-2	Acc
Bart	15.17	2.77	16.41	0.419
BlenderBot	15.22	2.70	16.20	0.470
CASE-BlenderBot	15.40	2.92	17.66	0.492

Table 5: Analysis of integrating pre-trained model.

Human Evaluation on Variants of CASE To more intuitively verify the role of the key components of CASE in language expression, especially empathy ability, we conduct a scoring human evaluation for the variants of CASE. Besides the same settings as above, we require annotating the **Overall** preference score (1-5). As in Table 4, CASE achieves the highest scores in all aspects, indicating that all components contribute. The low empathy scores of “w/o Graph” and “w/o Align” as well as their variants further confirm the crucial role of graph structure and the effectiveness of alignment.

4.5 Applicability Analysis

To analyze the applicability of our method, we build it on the pre-trained model to explore whether it brings further benefits. We integrate BlenderBot (Roller et al., 2021) into CASE by replacing the encoder and decoder, and take the vanilla Bart (Lewis et al., 2020) and BlenderBot as baselines. All pre-trained models are small versions. As in Table 5, we found that CASE-BlenderBot integrating our method significantly outperforms finetune-only baselines. Although the overall performance of simple finetuning has achieved stage success, it is limited by the quality and scale of the dataset and lacks a more fine-grained design for the trait of human conversation. This also demonstrates the high-level applicability of our method for uncovering the underlying mechanisms of human conversation.

4.6 Case Study

Two cases from six models are selected in Table 6, among which CASE is more likely to express informative cognition in a highly empathetic tone.

This is due to two main advantages:

(1) Effective alignment between cognition and affection on two levels. For example, in the first case, on the fine-grained level, CASE associates the cognition “*to be safe*” with the affection “*good*” (i.e., emotional reaction) to appease the user’s “*Terrified*” experience, i.e., “*to stay safe*” and “*get a little better*”, in response. In the second case, on the coarse-grained level, in the user’s “*Embarrassed*” emotional state, CASE expresses empathetic affection “*it is not too bad*” with an informative cognitive statement, i.e., “*get it fixed*”, in response.

(2) Accurate identification of the conversational emotion integrating emotional concepts and reactions, being consistent with “Acc”. For instance, in the first case, the correct conversational emotion “*Terrified*” tends to be identified in the emotional concepts (“*frighten, terrify, etc.*”), while in the second case, the one “*Embarrassed*” tends to be identified in the emotional reactions (“*embarrassed, ashamed, etc.*”). Compared with baselines that cannot correctly perform two cases simultaneously, CASE identifies correct emotion in both cases by integrating emotional concepts and reactions.

5 Conclusion and Future Work

In this paper, for responding empathetically, we propose CASE to align cognition and affection by simulating their conscious interaction in human conversation. Extensive experiments verify the superiority of CASE on overall quality and empathy performance. Our work will also encourage future work to model the more complex interaction between cognition and affection in human conversation as well as other human language behaviors (Liu et al., 2021; Zheng et al., 2023).

Limitations

We discuss two limitations of this work as follows:

One limitation of our work is the lack of task-specific automatic metrics to evaluate the empathy of generated responses. Therefore, the evaluation of empathy relies more on human evaluation. Although human evaluation is a golden standard, automatic metrics help to conduct large-scale investigations. This is also a common limitation in current works on empathetic dialogue.

The second limitation is the passive response to the user’s cognition and affection. In many scenarios, empathy is used as a strategy for emotional support by responding to the user’s cognition and

Emotion	Terrified
Context	<i>Speaker</i> : As a kid, I was so frightened any time there was a tornado warning. Not sure if it was the wizard of oz effect or what i just remember being frozen with fear . <i>Listener</i> : Oh, wow! My son was terrified of tornadoes because of that darn movie, <i>Speaker</i> : My mom said the best thing once when I was babysitting and a tornado warning went off and I did not want my own fear to affect the little boy I was with. She was on the phone and just said you have big plans for the future, right? Nothing is going to get in the way of that. For some reason, It worked a charm .
MoEL	I am sorry to hear that. I hope you can get better soon!
MIME	I am sure it was a great feeling.
EmpDG	I am glad you are ok! I am sure you will have a great time.
KEMP	I would be too! I would have been terrified!
CEM	I am sure you will be fine!
Cognition	To make a plan, To be safe , To be helpful, Have a good life
Reaction	Scared, Hopeful, Nervous, Good, Happy
Concept	frighten, surprising, terrify, aghast, good , charming, protect
CASE	That is good , I am glad you are able to stay safe and get a little better!
Ground-Truth	Your mom sounds awesome! Wow, what a cool thing to say!
Emotion	Embarrassed
Context	<i>Speaker</i> : I burned my hair with my hair drier I am so embarrassed to go out in public. <i>Listener</i> : I am sorry to hear that, my wife has done the same thing. She wore a hat for almost a month when she left the house. <i>Speaker</i> : Yeah it is the worst , I look so weird with my hair like this it is so embarrassing.
MoEL	That is a pretty funny story.
MIME	I am sure it was a good idea.
EmpDG	Oh, that is a shame. I hope you are not too!
KEMP	I am sure you will not have to be able to do it again.
CEM	I am sorry to hear that.
Cognition	To look bad , To get a haircut, To fix it , To fix the hair , To get a wig
Reaction	Embarrassed, Bad , Ashamed, Sad
Concept	bad , ill, evil, freak, supernatural, exhausted, hope , sorrow
CASE	I hope you are able to get it fixed , or just let her know it is not too bad .
Ground-Truth	Well, like I told her, it will grow out, it will just take time.

Table 6: Generated responses from baselines and CASE. **Commonsense cognition, emotional reactions**, reasoned **emotional concepts** by **contextual words**, and corresponding **informative expressions** in responses are highlighted.

affection. However, besides passive response, emotional support also requires active emotion elicitation, which can be studied in future work.

Ethical Considerations

In this paper, our experiments adopt the widely used EMPATHETICDIALOGUES benchmark, an open-source dataset collected from Amazon Mechanical Turk (MTurk) that does not contain personal information. We also ensure the anonymization of the human evaluation. We believe that this work honors the ethical code of ACL.

Acknowledgements

This work was supported by the National Key Research and Development Program of China (No. 2021ZD0113304), the National Science Foundation for Distinguished Young Scholars (with No. 62125604) and the NSFC projects (Key project with No. 61936010).

This work was also supported by the National Natural Science Foundation of China (with No.

62272340, 61876128, 62276187).

References

- Jiaqi Bai, Ze Yang, Xinnian Liang, Wei Wang, and Zhoujun Li. 2021. **Learning to copy coherent knowledge for response generation**. 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 12535–12543. AAAI Press.
- 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.
- Mao Yan Chen, Siheng Li, and Yujiu Yang. 2022. **Emphi: Generating empathetic responses with human-like intents**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 1063–1074. Association for Computational Linguistics.
- Benjamin MP Cuff, Sarah J Brown, Laura Taylor, and Douglas J Howat. 2016. Empathy: A review of the concept. *Emotion review*, 8(2):144–153.
- Mark H Davis. 1983. Measuring individual differences in empathy: evidence for a multidimensional approach. *Journal of personality and social psychology*, 44(1):113.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Robert Elliott, Arthur C Bohart, Jeanne C Watson, and David Murphy. 2018. Therapist empathy and client outcome: An updated meta-analysis. *Psychotherapy*, 55(4):399.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. **Learning deep representations by mutual information estimation and maximization**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. **Challenges in building intelligent open-domain dialog systems**. *ACM Trans. Inf. Syst.*, 38(3):21:1–21:32.
- Sevgi Coşkun Keskin. 2014. From what isn't empathy to empathic learning process. *Procedia-Social and Behavioral Sciences*, 116:4932–4938.
- Wongyu Kim, Youbin Ahn, Donghyun Kim, and Kyong-Ho Lee. 2022. **Emp-rft: Empathetic response generation via recognizing feature transitions between utterances**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 4118–4128. Association for Computational Linguistics.
- 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*.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. **A mutual information maximization perspective of language representation learning**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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.
- Junyi Li, Wayne Xin Zhao, Zhicheng Wei, Nicholas Jing Yuan, and Ji-Rong Wen. 2021. **Knowledge-based review generation by coherence enhanced text planning**. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 183–192. ACM.
- Qintong Li, Hongshen Chen, Zhaochun Ren, Pengjie Ren, Zhaopeng Tu, and Zhumin Chen. 2020. **Empdgc: 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 *AAAI*.
- 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.
- Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. [Towards emotional support dialog systems](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3469–3483, Online. 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.
- Saif M. Mohammad. 2018. [Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words](#). 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 174–184. 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, Xingsheng Zhang, and Yajing Sun. 2022. [Modeling intention, emotion and external world in dialogue systems](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Virtual and Singapore, 23-27 May 2022*, pages 7042–7046. IEEE.
- 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.
- 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.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 300–325. Association for Computational Linguistics.
- 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.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. [Conceptnet 5.5: An open multilingual graph of general knowledge](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press.
- 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.
- Liuping Wang, Dakuo Wang, Feng Tian, Zhenhui Peng, Xiangmin Fan, Zhan Zhang, Mo Yu, Xiaojuan Ma, and Hongan Wang. 2021. [Cass: Towards building a social-support chatbot for online health community](#). *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–31.
- 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 *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 1401–1410. ACM.
- David Westbrook, Helen Kennerley, and Joan Kirk. 2011. [An introduction to cognitive behaviour therapy: Skills and applications](#). Sage.

- Zhitong Yang, Bo Wang, Jinfeng Zhou, Yue Tan, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. 2022. [Topkg: Target-oriented dialog via global planning on knowledge graph](#). In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, pages 745–755. International Committee on Computational Linguistics.
- 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.
- Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2020. [Grounded conversation generation as guided traverses in commonsense knowledge graphs](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2031–2043. Association for Computational Linguistics.
- Chujie Zheng, Yong Liu, Wei Chen, Yongcai Leng, and Minlie Huang. 2021a. [Comae: A multi-factor hierarchical framework for empathetic response generation](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 813–824. Association for Computational Linguistics.
- Chujie Zheng, Yong Liu, Wei Chen, Yongcai Leng, and Minlie Huang. 2021b. [CoMAE: A multi-factor hierarchical framework for empathetic response generation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 813–824, Online. Association for Computational Linguistics.
- Chujie Zheng, Sahand Sabour, Jiaxin Wen, Zheng Zhang, and Minlie Huang. 2023. Augesc: Dialogue augmentation with large language models for emotional support conversation. In *Findings of the Association for Computational Linguistics: ACL 2023*.
- Peixiang Zhong, Di Wang, Pengfei Li, Chen Zhang, Hao Wang, and Chunyan Miao. 2021. [CARE: commonsense-aware emotional response generation with latent concepts](#). 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 14577–14585. AAAI Press.
- 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.
- Jinfeng Zhou, Bo Wang, Ruifang He, and Yuexian Hou. 2021. [CRFR: improving conversational recommender systems via flexible fragments reasoning on knowledge graphs](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 4324–4334. Association for Computational Linguistics.
- Jinfeng Zhou, Bo Wang, Minlie Huang, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. 2022a. [Aligning recommendation and conversation via dual imitation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 549–561. Association for Computational Linguistics.
- Jinfeng Zhou, Bo Wang, Zhitong Yang, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. 2022b. [CR-GIS: improving conversational recommendation via goal-aware interest sequence modeling](#). In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, pages 400–411. International Committee on Computational Linguistics.
- Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. [S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization](#). In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 1893–1902. ACM.

A Mutual Information Maximization

Mutual information maximization (MIM) aims to measure the dependence between two random variables X and Y , and the mutual information (MI) between them is defined as: $MI(X, Y) = D_{KL}(P(X, Y) \| P(X)P(Y))$. However, maximizing MI directly is normally intractable. A successful practice to estimate MI with a lower bound is InfoNCE (Kong et al., 2020). Given two different views x and y of an input, InfoNCE is defined by:

$$\mathbb{E}_{P(X, Y)}[f_{\theta}(x, y)] - \mathbb{E}_{Q(\tilde{Y})}[\log \sum_{\tilde{y} \in \tilde{Y}} \exp f_{\theta}(x, \tilde{y})] + \log |\tilde{Y}|, \quad (20)$$

where f_{θ} is a learnable function with parameter θ . The set \tilde{Y} draws samples from a proposal distribution $Q(\tilde{Y})$, and it comprises $|\tilde{Y}| - 1$ negative samples and a positive sample y . One insight is that when \tilde{Y} always consists all values of Y and they are uniformly distributed, maximizing InfoNCE is analogous to maximize cross-entropy loss:

$$\mathbb{E}_{P(X, Y)}[f_{\theta}(x, y) - \log \sum_{\tilde{y} \in Y} \exp f_{\theta}(x, \tilde{y})]. \quad (21)$$

It shows InfoNCE is relevant to maximize $P_{\theta}(y|x)$ and approximates summation over elements in Y (i.e., partition function) by negative sampling (Zhou et al., 2020, 2022a,b). Upon the formula, we replace X and Y with specific cognition and affection to maximize MI between them.

ACL 2023 Responsible NLP Checklist

A For every submission:

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

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Sec.2 Related Work and Sec.4.1 Experimental Setup
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- 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)?
Sec.2 Related Work and Sec.4.1 Experimental Setup
- 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?
Sec. Ethical Considerations
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- 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.
Sec. 4.1 Experimental Setup

C Did you run computational experiments?

Left blank.

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

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?
Appendix B Implementation Details
- 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?
Sec.4 Experiments
- 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.)?
Not applicable. Left blank.
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Left blank.
- 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.?
Sec. 4 Experiments
- 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)?
Limited by the space. Crowdsourcing workers are from Amazon Mechanical Turk.
- 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?
Not applicable. Left blank.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. Left blank.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Not applicable. Left blank.