

Constructing Emotion Consensus and Utilizing Unpaired Data for Empathetic Dialogue Generation

Lei Shen^{1,2} Jinchao Zhang³ Jiao Ou^{1,2} Xiaofang Zhao^{1*} Jie Zhou³

¹Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

³Pattern Recognition Center, WeChat AI, Tencent Inc, China

{shenlei17z, oujiao17b, zhaoxf}@ict.ac.cn, {dayerzhang, withtomzhou}@tencent.com

Abstract

Researches on dialogue empathy aim to endow an agent with the capacity of accurate understanding and proper responding for emotions. Existing models for empathetic dialogue generation focus on the emotion flow in one direction, that is, from the context to response. We argue that conducting an empathetic conversation is a bidirectional process, where empathy occurs when the emotions of two interlocutors could converge on the same point, i.e., reaching an *emotion consensus*. Besides, we also find that the empathetic dialogue corpus is extremely limited, which further restricts the model performance. To address the above issues, we propose a dual-generative model, Dual-Emp, to simultaneously construct the emotion consensus and utilize some external unpaired data. Specifically, our model integrates a forward dialogue model, a backward dialogue model, and a discrete latent variable representing the emotion consensus into a unified architecture. Then, to alleviate the constraint of paired data, we extract unpaired emotional data from open-domain conversations and employ Dual-Emp to produce pseudo paired empathetic samples, which is more efficient and low-cost than the human annotation. Automatic and human evaluations demonstrate that our method outperforms competitive baselines in producing coherent and empathetic responses.

1 Introduction

Empathy, a fundamental trait of humans, describes the ability to place oneself in another person's position and share his/her feelings or emotions. Besides, it has been considered to be one of the most valuable affective phenomena for improving human-machine interactions (Zech and Rimé,

Joint work with Pattern Recognition Center, WeChat AI, Tencent Inc, China.*Xiaofang Zhao is the corresponding author.

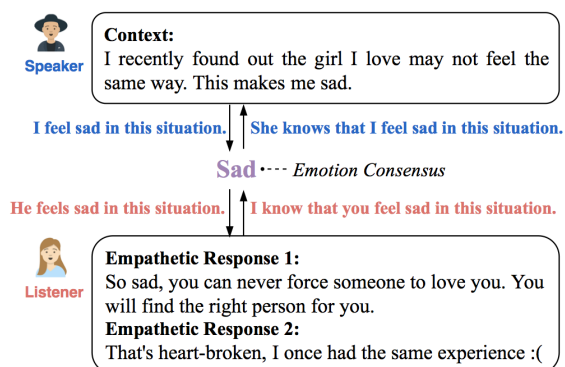


Figure 1: An example of conducting an empathetic conversation. Both responses show empathy to the speaker.

2005). The studies of empathy in natural language processing mainly include detecting empathy in spoken language or text (Buechel et al., 2018; Sharma et al., 2020), generating empathetic dialogue responses (Lin et al., 2019; Majumder et al., 2020), and constructing empathy lexicons (Sedoc et al., 2020) or datasets (Rashkin et al., 2019).

The empathetic dialogue generation task has been regarded as a unidirectional process from the context to response, and is modeled as a multi-task learning that combines the emotion understanding and the emotion-enhanced response generation. Therefore, existing work (Rashkin et al., 2019; Lin et al., 2019; Li et al., 2020; Majumder et al., 2020) mainly focuses on improving the accuracy of emotion classification or enhancing response generation via integrating the detected emotion factor.

Conducting an empathetic conversation is naturally a bidirectional process: the speaker conveys his/her emotion by describing a certain situation, then the listener receives that emotion and feeds his/her feeling back to the listener via a response. Then, the empathy is triggered when two interlocutors link similar experiences and their emotions could converge on the same point, i.e., reaching an *emotion consensus*. Take the case in Figure 1 as an

example. The emotion consensus “Sad” works as an intersection that connects both the speaker and listener, and it is a high-level abstraction behind the content, i.e., both two responses convey their acknowledgment of the sad feeling even with different expressions. Therefore, a unidirectional model is not enough to model the relationship between the context and response. Besides, previous models for this task only utilize paired data with limited capacity in a benchmark dataset, EMPATHETICDIALOGUES. Rather than manually annotating a larger empathetic dataset, we find that in open-domain conversations, there is large-scale emotional data that can be used to improve the performance. Compared with recognizing whether a context-response pair is empathetic, obtaining either an emotional context or response (named as *unpaired data* in this paper) can be easier with a well-trained classifier.

In this paper, we propose a **Dual-Generative** model for the **Empathetic** dialogue generation task (Dual-Emp), which simultaneously constructs emotion consensus and utilizes unpaired data. Dual-Emp combines a forward dialogue model (generating a response based on its context) and a backward dialogue model (generating a context based on its responses) with a discrete latent variable. Specifically, the forward and backward encoders convert the context and response into vectors at the same time, and then a discrete latent variable is used to capture the high-level emotion consensus shared in each context-response pair. Moreover, the latent variable and an emotion-enhanced attention mechanism are integrated into both forward and backward decoders to better express proper empathy. To utilize unpaired emotional data, we firstly extract them from open-domain conversations with emotions. Then we can get pseudo pairs by feeding either emotional responses or contexts to the backward or forward model. A joint training process is introduced to promote the semantic coherence between contexts and responses. Furthermore, two types of optimization methods are applied to better train the entire model with paired and unpaired data. Experimental results on a benchmark dataset EMPATHETICDIALOGUES show that Dual-Emp significantly outperforms competitive baselines in generating meaningful and related responses while expressing an appropriate empathy.

Our main contributions can be summarized as: (1) We point out that the empathetic dialogue generation contains bidirectional processes, and

highlight the importance of constructing emotion consensus. Besides, we propose a novel dual-generative model that couples a forward and a backward dialogue model with a discrete latent variable capturing the shared emotion consensus. (2) We utilize unpaired emotional data to break the constraint of paired empathetic data in the widely-used benchmark dataset EMPATHETICDIALOGUES. (3) Automatic and human evaluations show that our model outperforms competitive baselines in terms of fluency, coherence, and empathy.

2 Related Work

Emotion-Controllable Response Generation.

Infusing emotions into dialogue systems can make conversational agents more human-like and benefit the interactions between human and machine (Prendinger and Ishizuka, 2005). Emotion-controllable response generation aims to generate emotional responses conditioning on a manually-provided label. Existing work (Zhou et al., 2018; Zhou and Wang, 2018; Colombo et al., 2019; Song et al., 2019; Shen and Feng, 2020) focused on obtaining responses that are not only meaningful, but also in accordance with the desired emotion.

Empathetic Response Generation.

Rashkin et al. (2019) considered a richer and evenly distributed set of emotions, and released a dataset EMPATHETICDIALOGUES. Shin et al. (2020) formulated a reinforcement learning problem to maximize user’s sentimental feelings towards the generated responses. Lin et al. (2019) presented an encoder-decoder model with each emotion having a dedicated decoder. Majumder et al. (2020) introduced emotion grouping, emotion mimicry, and stochasticity to generate empathetic and various responses. Li et al. (2020) integrated knowledge to better understand dialogue contexts, and also designed an emotion-focused attention mechanism for emotional dependencies.

Dual Learning in NLP.

He et al. (2016) proposed Dual Learning (DL) for machine translation first, which considered the source to target language translation and target to source language translation as a dual task. After that, Tang et al. (2017) implemented a dual framework for the question-answering system. Both Zhang et al. (2018a) and Cui et al. (2019) used similar idea in dialogue generation task to produce coherent but not safe responses. Shen and Feng (2020) applied DL for emotion-controllable response generation with

three awards for emotions and semantics. Some researchers also exploited DL to relieve the need of paired data and make use of unpaired data in several areas, such as style transfer (Luo et al., 2019a,b), semantic understanding (Tseng et al., 2020), stylized response generation (Zheng et al., 2020a), and machine translation (Zheng et al., 2020b).

The differences between our model and previous methods are: (1) To improve the empathy understanding, we introduce a backward model to represent the response and a discrete latent variable to capture the emotion consensus shared by contexts and responses. (2) Our forward and backward models are connected by a latent variable, and *both of them can be updated at each iteration, while traditional DL can only fix one to update another.*

3 Proposed Method

For empathetic dialogue generation, a dialogue consists of utterances from a speaker and a listener. Given context $\mathbf{c} = \{S_1, L_1, S_2, L_2, \dots, S_t\}$, where $S_i = \{w_j^i\}_{j=1}^{|S_i|}$ denotes speaker and $L_i = \{w_j^i\}_{j=1}^{|L_i|}$ denotes listener, the goal is to track the speaker’s emotion state from \mathbf{c} , and generate a response $\mathbf{y} = L_t$ that is meaningful and empathetic.

3.1 Overview

The architecture of Dual-Emp is shown in Figure 2. Dual-Emp has five modules, the forward encoder f_{enc} , forward decoder f_{dec} , backward encoder b_{enc} , backward decoder b_{dec} , and z_e indicating a discrete latent variable. z_e can be inferred from both \mathbf{c} and \mathbf{y} and is used to capture emotion consensus shared in each $\langle \mathbf{c}, \mathbf{y} \rangle$ pair. Because of the existence of z_e , other modules are correlated and can better model both the semantic relation and the emotion connection between \mathbf{c} and \mathbf{y} .

3.2 Model Architecture

Since the backward dialogue model has the same architecture as the forward one, we specify the components of forward dialogue model below and omit those of backward model for space limitation. **Encoder.** Following the work of Lin et al. (2019), we firstly concatenate utterances in \mathbf{c} into a long sequence with length n and add a special token CTX to the beginning of \mathbf{c} inspired by BERT (Devlin et al., 2019). Then, each token w in \mathbf{c} is calculated as the sum of three embeddings:

$$\mathbf{E}_c(w) = \mathbf{E}_w(w) + \mathbf{E}_p(w) + \mathbf{E}_r(w), \quad (1)$$

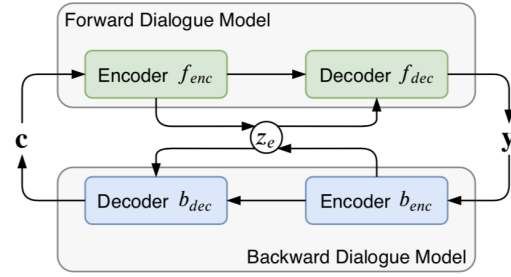


Figure 2: The architecture of Dual-Emp. It couples forward and backward dialogue models with a discrete latent variable z_e denoting emotion consensus.

where $\mathbf{E}_w(\cdot)$, $\mathbf{E}_p(\cdot)$, and $\mathbf{E}_r(\cdot) \in \mathbb{R}^{|V| \times d_{emb}}$ represent word embedding space, positional embedding space and role embedding space¹, respectively. Finally, a transformer encoder (Vaswani et al., 2017) f_{enc} is applied to get the context representation:

$$\mathbf{H} = f_{enc}(\mathbf{E}_c([\text{CTX}; \mathbf{c}])), \quad (2)$$

where “;” represents the concatenation operation, and $\mathbf{H} \in \mathbb{R}^{(n+1) \times d_{mod}}$. The contextualized encoding of CTX, i.e., $\mathbf{H}_0 \in \mathbb{R}^{d_{mod}}$, is used as the final representation of the entire context.

Emotion Consensus Construction. A K -way categorical latent variable $z_e \in [1, K]$ (Bao et al., 2020) is used to capture the emotion consensus shared by \mathbf{c} and \mathbf{y} . Inspired by Zhao et al. (2019), we define the prior distribution where we sample z_e from to be uniform², i.e., $p(z_e) = 1/K$. Correspondingly, the approximate posterior distribution is defined as follows:

$$q(z_e | \mathbf{c}) = \text{softmax}(\text{FFN}(\mathbf{H}_0)) \in \mathbb{R}^K, \quad (3)$$

where $\text{FFN}(\cdot)$ represents a feedforward network. This part can be considered as the emotion understanding on c . Here z_e has its own embedding space $\mathbf{E}_z \in \mathbb{R}^{K \times d_{mod}}$ to convert it into a vector, i.e., $\mathbf{E}_{[z]} = \mathbf{E}_z(z_e) \in \mathbb{R}^{d_{mod}}$. To supervise the emotion expression in $\mathbf{E}_{[z]}$, we train a classifier using the cross-entropy loss between $\mathbf{E}_{[z]}$ and ground-truth emotion label e^* :

$$p_e = \text{softmax}(\mathbf{W}_e \mathbf{E}_{[z]}), \quad (4)$$

$$\mathcal{L}_{emo} = -e^* \log p_e, \quad (5)$$

where \mathbf{W}_e is a trainable weight matrix.

¹The roles in \mathbf{c} is an alternating set of “speaker” and “listener”, while in \mathbf{y} , the role is “listener” only.

²Since emotion labels in EMPATHETICDIALOGUES are evenly-distributed, we set the prior distribution to be uniform.

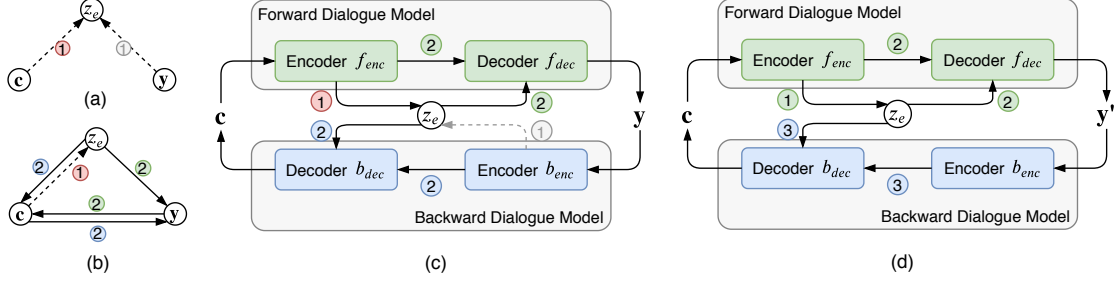


Figure 3: Illustration of the training process. (a) shows the inference of z_e . Both \mathbf{c} and \mathbf{y} are used to infer z_e that represents the emotion consensus shared by \mathbf{c} and \mathbf{y} . (b) shows the graphical model of (c), and (c) depicts procedures to compute \mathcal{L}_1 (Eq. 11)/ \mathcal{L}_2 (Eq. 12). (d) represents procedures to compute \mathcal{L}_3 (Eq. 15).

Decoder. Existing work (Lin et al., 2020; Li et al., 2020) mainly integrates the obtained emotion factor to either the first decoding position or all steps. To focus on emotion consensus dynamically, we apply an emotion-enhanced attention mechanism in the cross-attention layer of transformer decoder. We firstly concatenate $\mathbf{E}_{[z]}$ with token embeddings of the decoder input $\{y_i\}_{i=1}^{t-1}$ to get representations $\mathbf{Y} = \{y_i\}_{i=0}^{t-1}$ with $y_0 = \mathbf{E}_{[z]}$. Then we feed \mathbf{Y} into decoder f_{dec} .

Our decoder has similar structure to the transformer decoder. The input \mathbf{Y} is converted to \mathbf{D} by the self-attention layer. As \mathbf{H} and $\mathbf{E}_{[z]}$ serve different purposes, we design a cross-attention layer with two separate key-value matrices, and the encoder-decoder vectors are computed as follows:

$$\mathbf{C}_H = \text{MultiHead}(\mathbf{D}, \mathbf{H}, \mathbf{H}), \quad (6)$$

$$\mathbf{C}_Z = \text{MultiHead}(\mathbf{D}, \mathbf{E}_{[z]}, \mathbf{E}_{[z]}), \quad (7)$$

where $\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$ is a multi-head attention function taking a query matrix \mathbf{Q} , a key matrix \mathbf{K} , and a value matrix \mathbf{V} as inputs. The fully connected feedforward layer is defined as:

$$\hat{\mathbf{Y}} = \text{FFN}([\mathbf{C}_H; \mathbf{C}_Z]), \quad (8)$$

where $\hat{\mathbf{Y}} = \{\hat{y}_i\}_{i=1}^t$. Finally, the decoding distribution over the vocabulary of the next token is computed as:

$$p(y_t | y_{<t}, \mathbf{c}, z_e) = \text{softmax}(\mathbf{W}_o \hat{\mathbf{y}}_t), \quad (9)$$

where \mathbf{W}_o is a trainable weight matrix.

3.3 Training and Inference

We firstly describe how Dual-Emp can be trained with the paired data $\langle \mathbf{c}, \mathbf{y} \rangle$, and also the unpaired data \mathbf{c} or \mathbf{y} . Then a combined objective is derived to

optimize Dual-Emp using the paired and unpaired data at the same time.

Training with Paired Data. Given $\langle \mathbf{c}, \mathbf{y} \rangle$, we aim to maximize the log-likelihood of a joint probability $p(\mathbf{c}, \mathbf{y})$:

$$\log p(\mathbf{c}, \mathbf{y}) = \log \sum_{z_e} p(\mathbf{c}, \mathbf{y}, z_e). \quad (10)$$

Following the derivations from Zhao et al. (2018), Zhao et al. (2019), Tseng et al. (2020), and the variational inference (Kingma and Welling, 2014), an objective based on the evidence lower bound can be derived as:

$$\begin{aligned} \mathcal{L}_1 = & \mathbb{E}_{q(z_e|\mathbf{c})} \log p(\mathbf{y}|z_e, \mathbf{c}) + \mathbb{E}_{q(z_e|\mathbf{c})} \log p(\mathbf{c}|z_e, \mathbf{y}) \\ & - D_{\text{KL}}[q(z_e|\mathbf{c})||p(z_e)], \end{aligned} \quad (11)$$

where the first term denotes the forward dialogue model, $q(z_e|\mathbf{c})$ is the approximate posterior distribution of z_e , and is computed by the forward encoding process (red ① in Figure 3(c)). $p(\mathbf{y}|z_e, \mathbf{c})$ is the forward decoding process (green ② in Figure 3(c)); the second term denotes the reconstruction of \mathbf{c} , and $p(\mathbf{c}|z_e, \mathbf{y})$ is the backward decoding process (blue ② in Figure 3(c)); the third term is a Kullback-Leibler (KL) divergence between two distributions.

Analogously, the posterior distribution of z_e can be approximated by $q(z_e|\mathbf{y})$, and the objective can be converted as follows:

$$\begin{aligned} \mathcal{L}_2 = & \mathbb{E}_{q(z_e|\mathbf{y})} \log p(\mathbf{c}|z_e, \mathbf{y}) + \mathbb{E}_{q(z_e|\mathbf{y})} \log p(\mathbf{y}|z_e, \mathbf{c}) \\ & - D_{\text{KL}}[q(z_e|\mathbf{y})||p(z_e)], \end{aligned} \quad (12)$$

where terms have similar meanings to those in Eq. 11, and we only need to interchange “forward” and “backward”. Besides, the forward encoding process (red ① in Figure 3(c)) is replaced with the backward encoding process (gray ① in Figure 3). Detailed derivations can be found in Appendix. Therefore,

the final loss function for the paired data is:

$$\mathcal{L}_{cy} = \mathcal{L}_1 + \mathcal{L}_2 + \alpha \mathcal{L}_{emo}. \quad (13)$$

where α is a hyper-parameter.

Training with Unpaired Data. Given unpaired data \mathbf{c} (\mathbf{c} is an emotional context), we need to maximize the log-likelihood of a marginal probability $p(\mathbf{c})$:

$$\log p(\mathbf{c}) = \log \int_{\mathbf{y}} \sum_{z_e} p(\mathbf{c}, \mathbf{y}, z_e). \quad (14)$$

Then, we can get the evidence lower bound for the marginal probability:

$$\begin{aligned} \mathcal{L}_3 = & \mathbb{E}_{q(\mathbf{y}|z_e, \mathbf{c})} \mathbb{E}_{q(z_e|\mathbf{c})} \log p(\mathbf{c}|z_e, \mathbf{y}) \\ & - D_{\text{KL}}[q(z_e|\mathbf{c})||p(z_e)], \end{aligned} \quad (15)$$

where the first term is the reconstruction of \mathbf{c} , $q(z_e|\mathbf{c})$ is computed by the forward encoding process (① in Figure 3(d)), $q(\mathbf{y}|z_e, \mathbf{c})$ is the forward generation process (② in Figure 3(d)), and $p(\mathbf{c}|z_e, \mathbf{y})$ is the backward decoding process (③ in Figure 3(d)); the second term is a KL divergence.

The forward generation process $q(\mathbf{y}|z_e, \mathbf{c})$ is similar to the back-translation in machine translation (Zhang et al., 2018b), and we use f_{dec} to generate pseudo \mathbf{y}' given \mathbf{c} and z_e . Since the ground-truth \mathbf{y} is unobserved here, we apply reinforcement learning and policy gradient method (Williams, 1992) for training. The reward is designed as the probability of the model to reconstruct \mathbf{c} based on the generated $\hat{\mathbf{y}}$ and z_e :

$$r = p(\mathbf{c}|\hat{\mathbf{y}}, z_e). \quad (16)$$

Similarly, we can get an objective when utilizing unpaired \mathbf{y} (the emotional response):

$$\begin{aligned} \mathcal{L}_4 = & \mathbb{E}_{q(\mathbf{c}|z_e, \mathbf{y})} \mathbb{E}_{q(z_e|\mathbf{y})} \log p(\mathbf{y}|z_e, \mathbf{c}) \\ & - D_{\text{KL}}[q(z_e|\mathbf{y})||p(z_e)], \end{aligned} \quad (17)$$

where the first term is the reconstruction of \mathbf{y} , and the process is symmetrical to that of \mathcal{L}_3 . Detailed derivations can be found in Appendix. The final loss functions for unpaired data \mathbf{c} and \mathbf{y} are:

$$\mathcal{L}_c = \mathcal{L}_3 + \beta \mathcal{L}_{emo}, \quad (18)$$

$$\mathcal{L}_y = \mathcal{L}_4 + \gamma \mathcal{L}_{emo}, \quad (19)$$

where β and γ are two hyper-parameters.

Total Training Loss. During training, the paired empathetic data in EMPATHETICDIALOGUES and

the unpaired emotional data from open-domain conversations are used simultaneously. Then, the total loss can be summarized as:

$$\mathcal{L} = \mathcal{L}_{cy} + \mathcal{L}_c + \mathcal{L}_y. \quad (20)$$

Inference. During inference, given the input \mathbf{c} , only the forward dialogue model is applied. We use f_{enc} to encode \mathbf{c} and infer z_e , then employ f_{dec} to generate $\hat{\mathbf{y}}$ based on \mathbf{c} and z_e .

4 Experiments

In this section, we conduct experiments to evaluate our proposed method. We firstly introduce some empirical settings. Then we illustrate our results on both automatic and human evaluations. Finally, we show some cases generated by different models and do further analyses over our method.

4.1 Dataset

We conduct our experiments on the EMPATHETICDIALOGUES (Rashkin et al., 2019) dataset that consists of 24,850 conversations between two interlocutors. Each conversation in the dataset contains **one** emotion label, a situation where the speaker feels the exact emotion, and utterances about the speaker’s descriptions of the situation or the listener’s empathetic replies. There are 32 evenly-distributed emotion labels in the dataset. We apply the data provided by the original paper with the split ratio of 8:1:1 for training/validation/test set, and use the script released by Lin et al. (2019) to preprocess the data. Emotion labels are given as supervised signals in the training process, while during inference, they are predicted to evaluate the accuracy of emotion understanding.

4.2 Implementation Details

We optimize the models using Adam (Kingma and Ba, 2015) with a mini-batch size of 16. The learning rate is initialized to 1e-4 and we vary the learning rate following Vaswani et al. (2017). Similar to Lin et al. (2019), Li et al. (2020), and Majumder et al. (2020), we use pre-trained GloVe vectors (Pennington et al., 2014) to initialize the word embeddings. Besides, all common hyper-parameters are set the same as previous work, e.g., the hidden size d_{mod} and embedding size d_{emb} are set to 300. In order to alleviate the degeneration problem of variational framework, we apply KL annealing (Bowman et al., 2016) that is the same as in Zhou

et al. (2018). During inference, we use greedy decoding strategy and the maximum decoding step is set to 30. K equals to 32. α , β , and γ are simply set to 1. The running epoch is set to 30 with early stopping.

To get unpaired emotional data, we utilize two large-scale datasets of open-domain conversations provided by Zhong et al. (2020), namely Reddit and Twitter. Following Zhou et al. (2018) and Shen and Feng (2020), an emotion classifier is applied to obtain the ground-truth emotion label for each context and response. Here, we use the pre-trained classifier provided by Rashkin et al. (2019) to predict labels among 32 emotions. The classifier could return one emotion label with the highest probability. We firstly keep contexts or responses with the probability larger than a threshold $s = 0.60$. Then, we remove contexts or responses with length smaller than 3. Finally, 155,059 contexts and 149,672 responses are obtained as unpaired emotional data.

We use Pytorch³ to implement the codes, and our model is trained on a Titan Xp GPU with an average running time of 2 days.

4.3 Baselines

We compare our approach with five representative baselines: (1) **Multi-TRS** (Rashkin et al., 2019): A transformer-based model trained with emotion classification loss in addition to MLE loss, and the emotion label is classified from the encoder output; (2) **MoEL** (Lin et al., 2019): An extension to Multi-TRS, which softly combines the output states of the appropriate decoders and generates an empathetic response. Each decoder is optimized to focus on a specific emotion; (3) **EmpDG** (Li et al., 2020): A model that exploits coarse- and fine-grained emotions and introduces an interactive adversarial learning framework to use user feedbacks; (4) **DualVAE** (Tran and Nguyen, 2018): A model with two decoders: one is for CVAE, and the other is for response auto-encoding; (5) **MIME** (Majumder et al., 2020): A model that integrates emotion grouping, emotion mimicry, and stochasticity strategies to generate varied responses. MIME is also the state-of-the-art model for empathetic response generation. To make fair comparisons, we do not apply methods based on pre-trained models here, as both Dual-Emp and the above mentioned ones are not based on pre-trained models. Note that model (1) to (5) can only utilize the paired data.

³<https://pytorch.org/>

Additionally, we also design following models for ablation study: (6) **Sing-Emp-Paired**: A variation of Dual-Emp with only the forward model and paired empathetic data; (7) **Dual-Emp-Paired**: Dual-Emp with only paired empathetic data.

4.4 Evaluation Measures

Automatic Metrics. For automatic evaluation, we use followings metrics: (1) BLEU (Papineni et al., 2002); (2) Embedding-based scores (Average, Greedy, and Extrema)⁴ (Liu et al., 2016; Serban et al., 2017); (3) Perplexity (PPL) (Vinyals and Le, 2015); (4) Dist-1/2 (Li et al., 2016); (5) Emotion accuracy (the agreement between the ground-truth emotion labels and the predicted ones from Eq. 5). Emotion accuracy can be used to measure the ability of emotion understanding.

Human Evaluation. Firstly, we randomly sample 100 contexts and their corresponding responses from our model as well as the baselines. Next, we send pairs of the context and generated response from different models to three professional annotators without order. Annotators are asked to evaluate each pair independently based on three distinct metrics: Empathy, Relevance, and Fluency (Rashkin et al., 2019; Lin et al., 2019; Majumder et al., 2020). Empathy measures the degree of emotional understanding of context shown by the response; Relevance evaluates whether the generated responses are relevant on topic with the context; Fluency assesses the grammatical correctness and readability of the generated responses. Each metric is rated on five-scale with “5” represents the best performance.

Human A/B Test. In this part, we try to directly compare Dual-Emp with other baselines. We randomly sample 100 dialogues each for Dual-Emp vs. {Multi-TRS, MoEL, EmpDG, DualVAE, MIME}. Three annotators are given generated responses from either Dual-Emp or {Multi-TRS, MoEL, EmpDG, DualVAE, MIME} in random order, and are asked to choose the better response. They can either choose one of the responses or select “Tie” when the provided options are either both good or both bad. The result of each sample is determined by majority voting. Finally, we calculate the percentage of samples where the first or second model generates the better response and where these two models perform similarly.

⁴We employ a popular NLG evaluation project available at <https://github.com/Maluuba/nlg-eval>.

| Method | Automatic Evaluation | | | | | | | | Human Evaluation | | |
|-----------------|----------------------|--------------|--------------|--------------|--------------|-------------|-------------|--------------|------------------|-------------|-------------|
| | B | Avg | Gre | Ext | PPL | D1 (%) | D2 (%) | EA (%) | Emp | Rel | Flu |
| Multi-TRS | 2.56 | 0.938 | 0.786 | 0.541 | 33.82 | 0.68 | 2.62 | 35.17 | 3.08 | 3.21 | 3.14 |
| MoEL | 2.80 | 0.945 | 0.793 | 0.537 | 37.81 | 0.56 | 2.70 | 35.38 | 3.35 | 3.65 | 3.26 |
| Emp-DG | 2.79 | 0.935 | 0.788 | 0.532 | 34.31 | 0.47 | 2.10 | 34.35 | 3.27 | 3.54 | 3.38 |
| DualVAE | 2.76 | 0.941 | 0.791 | 0.540 | 33.46 | 0.77 | 3.21 | 35.36 | 3.36 | 3.62 | 3.45 |
| MIME | 2.82 | 0.946 | 0.794 | 0.536 | 37.53 | 0.51 | 2.68 | 34.88 | 3.40 | 3.79 | 3.41 |
| Sing-Emp-Paired | 2.77 | 0.944 | 0.790 | 0.533 | 32.71 | 0.75 | 2.91 | 28.75 | 3.57 | 3.77 | 3.49 |
| Dual-Emp-Paired | 2.86 | 0.950 | 0.792 | 0.542 | 32.56 | 0.80 | 3.09 | 36.82 | 3.62 | 3.86 | 3.57 |
| Dual-Emp | 2.91 | 0.957 | 0.796 | 0.545 | 31.01 | 1.08 | 3.23 | 37.53 | 3.82 | 4.08 | 3.62 |

Table 1: Automatic and human evaluation results. The metrics BLEU, Average, Greedy, Extrema, Dist-1, Dist-2, Emotion accuracy, Empathy, Relevance, and Fluency are abbreviated as B, Avg, Gre, Ext, D1, D2, EA, Emp, Rel, and Flu, respectively. Results show that Dual-Emp achieves the best performance on all metrics, especially a large improvement in Dist-1/2, Emotion accuracy, and Empathy.

| Dual-Emp vs. | Win | Loss | Tie | Kappa |
|--------------|-----|------|-----|-------|
| Multi-TRS | 43% | 27% | 30% | 0.563 |
| MoEL | 37% | 32% | 31% | 0.548 |
| Emp-DG | 40% | 28% | 32% | 0.506 |
| DualVAE | 39% | 30% | 31% | 0.527 |
| MIME | 36% | 32% | 32% | 0.569 |

Table 2: Results of human A/B test. Pairwise comparisons show that responses from Dual-Emp are more preferred by humans than those from baselines.

4.5 Experimental Results

Automatic Evaluation Results. The automatic evaluation results are shown in the left part of Table 1. The top part is the results of all baseline models, and we can see that Dual-Emp outperforms other methods on all metrics (t-test, p -value < 0.05). The improvements of Dual-Emp on *PPL*, *Dist-1/2*, and *Emotion accuracy* are significant, indicating that it can improve emotion understanding, and also enhance content fluency and diversity simultaneously. MoEL, Emp-DG, and MIME have similar performance, as they try to either improve the emotion understanding or intensify the emotion-based response generation.

The bottom part of Table 1 shows the results of our ablation study. Comparisons between Sing-Emp-Paired and Dual-Emp-Paired show the effectiveness of capturing emotion consensus with the assistance of both backward model and discrete latent variable. Especially, the noticeable improvement of *Emotion accuracy* indicates the discrete latent variable used for emotion prediction can help better model the emotion consensus by taking contexts and responses into consideration. In addition, we can find that with the support of unpaired emotional data, Dual-Emp achieves better results than Dual-Emp-Paired.

Human Evaluation Results. Human evaluation in Table 1 illustrates that Dual-Emp obtains the best

performance (t-test, p -value < 0.05) on all scores. This suggests that our bidirectional model with latent variable helps construct emotion consensus shared by contexts and responses, thus improving the topic consistency and evoking more empathetic expressions. Besides, as more unpaired emotional data is utilized, Dual-Emp can achieve better *Fluency*. Additionally, we carry out pairwise comparisons to directly compare the response quality in Table 2. The results confirm that responses from Dual-Emp are more preferred by humans. Agreements to measure the consistency among three annotators are calculated with Fleiss’ kappa (Fleiss and Cohen, 1973), and the kappa values indicate “moderate agreement” in our cases.

4.6 Case Study

Table 3 shows two examples generated by Dual-Emp and other baselines. In the first case, Dual-Emp generates the most context-consistent response with a proper “apprehensive” emotion by replying with words “scary” and “what happened”, whereas baselines fail to understand the negative emotion or express inappropriate contents. In the second case, Dual-Emp generates a coherent and informative response, which corresponds to a subtle emotion change of the context from “lost a job” to “hoping he can find a full time job soon”. The response is not only emotion-related, but also contains the correct personal pronoun “he” and keyword “job”.

4.7 Further Analysis

Effects of Backward Model and z_e . To gain an insight into the effectiveness of backward dialogue model and the latent variable z_e , we plot the *Emotion accuracy* score of each emotion label based on Sing-Emp-Paired and Dual-Emp-Paired in Figure 4.

| | | |
|---|-----------------------------------|---|
| 1 | Emotion | Apprehensive |
| | Context | I went skydiving last summer with my partner. It was so scary ! |
| | Ground-truth Response | Wow, that is scary. Were you nervous? |
| | MoEL | I am so sorry to hear that. I am glad you are okay! |
| | Emp-DG MIME Dual-Emp | Did you go check it out? Oh no! I am sorry to hear that. I hope you get it! That is scary! What happened to you? |
| 2 | Emotion | Hopeful |
| | Context | My husband lost a job, but I am hoping he can find a full time job soon. |
| | Ground-truth Response | He will, I have faith. |
| | MoEL | I am sorry to hear that. I am sure he will be fine. |
| | Emp-DG MIME Dual-Emp | I am sorry to hear that. I hope you have a good time! That is a good thing to do. I hope you have a good job! That is unfortunate. I hope he gets a new job. |

Table 3: Generated responses from MoEL, Emp-DG, MIME, and Dual-Emp. In Case 1, Dual-Emp generates the most context-consistent response with a proper ‘‘apprehensive’’ emotion by replying with words ‘‘scary’’ and ‘‘what happened’’. In Case 2, Dual-Emp captures a subtle emotion change of the context from ‘‘lost a job’’ to ‘‘hoping he can find a full time job soon’’. Besides, it contains the correct personal pronoun ‘‘he’’ and keyword ‘‘job’’.




Figure 4: Emotion accuracy over 32 emotions of Sing-Emp-Paired and Dual-Emp-Paired. The accuracy of Sing-Emp-Paired is unbalanced among all emotions, while Dual-Emp-Paired can not only improve the overall accuracy, but also exhibit a relatively even performance.

| s | #context | #response | B | Avg | Gre | Ext | PPL | D1 (%) | D2 (%) | EA (%) |
|------|----------|-----------|------|-------|-------|-------|-------|--------|--------|--------|
| 0.50 | 324,243 | 314,070 | 2.25 | 0.937 | 0.791 | 0.549 | 31.22 | 0.90 | 2.90 | 35.62 |
| 0.55 | 224,324 | 216,839 | 2.69 | 0.933 | 0.784 | 0.537 | 32.63 | 1.87 | 4.32 | 36.06 |
| 0.60 | 155,059 | 149,672 | 2.91 | 0.957 | 0.795 | 0.545 | 31.01 | 1.08 | 3.23 | 37.53 |
| 0.65 | 107,189 | 103,192 | 2.25 | 0.934 | 0.783 | 0.539 | 32.61 | 1.70 | 4.80 | 35.66 |
| 0.70 | 73,132 | 70,200 | 2.60 | 0.938 | 0.791 | 0.540 | 31.66 | 0.70 | 2.40 | 37.51 |

Table 4: Automatic evaluation results based on the number of unpaired data with different s values. Results show that more unpaired data does not lead to better results as some labels are not adequate with a low confidence.

As we can see, for Sing-Emp-Paired, some emotion categories can achieve pretty high accuracy, but in general, the accuracy is unbalanced among all emotions, which indicates that z_e cannot construct the emotion consensus well by only considering the contexts. In contrast, Dual-Emp-Paired not only improves the overall *Emotion accuracy*, but also exhibits a relatively even performance over all 32 emotions. Therefore, z_e can better understand the emotion via capturing emotion consensus with both forward and backward dialogue models.

Choices of the Unpaired Data. The threshold s we use in previous experiments equals to 0.60. Here, we choose different options to show their influence on the empathetic dialogue generation task.

Table 4 shows that more unpaired emotional data does not lead to better results as some labels are not adequate with a low confidence. The emotion classifier we applied to label the utterances from Reddit and Twitter is based on a 32-category classification task, thus it is hard to get very accurate results. Though the predicted emotion labels are noisy, these samples are good enough to train our model in practice.

5 Conclusion and Future Work

In this paper, we propose a dual-generative model, Dual-Emp, to generate the empathetic response given a context. We point out that conducting an

empathetic conversation is a bidirectional process, and empathy is mainly reflected by emotion consensus between the context and response. Then we couple forward and backward dialogue models with a discrete latent variable denoting the emotion consensus. Moreover, we integrate unpaired emotional data from open-domain conversations into Dual-Emp to relieve the need of paired data. Experimental results on a benchmark dataset show that Dual-Emp can generate fluent, related, informative, and empathetic responses. As the future work, we will prove the effectiveness of our method based on pre-trained models, and analyze how classification errors in unpaired data affect the generation.

Acknowledgements

We would like to thank all the reviewers for their insightful and valuable comments and suggestions.

References

- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. [PLATO: Pre-trained dialogue generation model with discrete latent variable](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 85–96, Online. Association for Computational Linguistics.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21, Berlin, Germany. Association for Computational Linguistics.
- Sven Buechel, Anneke Buffone, Barry Slaff, Lyle Ungar, and João Sedoc. 2018. [Modeling empathy and distress in reaction to news stories](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4758–4765, Brussels, Belgium. Association for Computational Linguistics.
- Pierre Colombo, Wojciech Witon, Ashutosh Modi, James Kennedy, and Mubbasir Kapadia. 2019. [Affect-driven dialog generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3734–3743, Minneapolis, Minnesota. Association for Computational Linguistics.
- Shaobo Cui, Rongzhong Lian, Di Jiang, Yuanfeng Song, Siqi Bao, and Yong Jiang. 2019. [DAL: Dual adversarial learning for dialogue generation](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 11–20, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Joseph L Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. [Dual learning for machine translation](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 820–828.
- 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*.
- Diederik P. Kingma and Max Welling. 2014. [Auto-encoding variational bayes](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. 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*, pages 4454–4466, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- 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)*, pages 121–132, Hong Kong, China. Association for Computational Linguistics.

- Zhaojiang Lin, Genta Indra Winata, Peng Xu, Zihan Liu, and Pascale Fung. 2020. [Variational transformers for diverse response generation](#). *ArXiv preprint*, abs/2003.12738.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike 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*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Fuli Luo, Peng Li, Pengcheng Yang, Jie Zhou, Yutong Tan, Baobao Chang, Zhifang Sui, and Xu Sun. 2019a. [Towards fine-grained text sentiment transfer](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2013–2022, Florence, Italy. Association for Computational Linguistics.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Xu Sun, and Zhifang Sui. 2019b. [A dual reinforcement learning framework for unsupervised text style transfer](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5116–5122. ijcai.org.
- Navonil Majumder, Pengfei Hong, Shanshan Peng, Jiankun Lu, Deepanway Ghosal, Alexander 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)*, pages 8968–8979, Online. 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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Helmut Prendinger and Mitsuru Ishizuka. 2005. The empathic companion: A character-based interface that addresses users’ affective states. *Applied Artificial Intelligence*, 19(3-4):267–285.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: A new benchmark and dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- João Sedoc, Sven Buechel, Yehonathan Nachmany, Anneke Buffone, and Lyle Ungar. 2020. [Learning word ratings for empathy and distress from document-level user responses](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1664–1673, Marseille, France. European Language Resources Association.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3295–3301. AAAI Press.
- Ashish Sharma, Adam Miner, David Atkins, and Tim Althoff. 2020. [A computational approach to understanding empathy expressed in text-based mental health support](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5263–5276, Online. Association for Computational Linguistics.
- 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*, pages 556–566, Online. Association for Computational Linguistics.
- Jamin Shin, Peng Xu, Andrea Madotto, and Pascale Fung. 2020. [Generating empathetic responses by looking ahead the user’s sentiment](#). In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 7989–7993. IEEE.
- Zhenqiao Song, Xiaoqing Zheng, Lu Liu, Mu Xu, and Xuanjing Huang. 2019. [Generating responses with a specific emotion in dialog](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3685–3695, Florence, Italy. Association for Computational Linguistics.
- Duyu Tang, Nan Duan, Tao Qin, Zhao Yan, and Ming Zhou. 2017. [Question answering and question generation as dual tasks](#). *ArXiv preprint*, abs/1706.02027.
- Van-Khanh Tran and Le-Minh Nguyen. 2018. [Dual latent variable model for low-resource natural language generation in dialogue systems](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 21–30, Brussels, Belgium. Association for Computational Linguistics.
- Bo-Hsiang Tseng, Jianpeng Cheng, Yimai Fang, and David Vandyke. 2020. [A generative model for joint natural language understanding and generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1795–1807, Online. 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.
- Oriol Vinyals and Quoc Le. 2015. [A neural conversational model](#). *ArXiv preprint*, abs/1506.05869.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- 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.
- Hainan Zhang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. 2018a. [Reinforcing coherence for sequence to sequence model in dialogue generation](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4567–4573. ijcai.org.
- Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou, and Enhong Chen. 2018b. [Joint training for neural machine translation models with monolingual data](#). 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 555–562. AAAI Press.
- Tiancheng Zhao, Kyusong Lee, and Maxine Eskenazi. 2018. [Unsupervised discrete sentence representation learning for interpretable neural dialog generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1098–1107, Melbourne, Australia. Association for Computational Linguistics.
- Tiancheng Zhao, Kaige Xie, and Maxine Eskenazi. 2019. [Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1208–1218, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yinhe Zheng, Zikai Chen, Rongsheng Zhang, Shilei Huang, Xiaoxi Mao, and Minlie Huang. 2020a. [Stylized dialogue response generation using stylized unpaired texts](#). *ArXiv preprint*, abs/2009.12719.
- Zaixiang Zheng, Hao Zhou, Shujian Huang, Lei Li, Xin-Yu Dai, and Jiajun Chen. 2020b. [Mirror-generative neural machine translation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Peixiang Zhong, Di Wang, Pengfei Li, Chen Zhang, Hao Wang, and Chunyan Miao. 2020. [Care: Commonsense-aware emotional response generation with latent concepts](#). *ArXiv preprint*, abs/2012.08377.
- 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 (Volume 1: Long Papers)*, pages 1128–1137, Melbourne, Australia. Association for Computational Linguistics.