

Improving the Robustness of Knowledge-Grounded Dialogue via Contrastive Learning

Jiaan Wang¹, Jianfeng Qu^{1*}, Kexin Wang¹, Zhixu Li^{2*}, Wen Hua³, Ximing Li⁴, An Liu¹

¹ School of Computer Science and Technology, Soochow University, Suzhou, China

² Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University, Shanghai, China

³ The Hong Kong Polytechnic University, Hong Kong SAR, China

⁴ College of Computer Science and Technology, Jilin University, China

jawang.nlp@gmail.com, jfqu@suda.edu.cn, kxwang1@stu.suda.edu.cn, zhixuli@fudan.edu.cn

Abstract

Knowledge-grounded dialogue (KGD) learns to generate an informative response based on a given dialogue context and external knowledge (*e.g.*, knowledge graphs; KGs). Recently, the emergence of large language models (LLMs) and pre-training techniques has brought great success to knowledge-grounded dialogue. However, when building KGD systems in real applications, there are various real-world noises that are inevitable to face. For example, the dialogue context might involve perturbations such as misspellings and abbreviations. In addition, KGs typically suffer from incompleteness and also might contain erroneous and outdated facts. Such real-world noises pose a challenge to the robustness of KGD systems and hinder their applications in the real world. In this paper, we propose an entity-based contrastive learning framework for improving the robustness of KGD. Specifically, we make use of the entity information in a KGD sample to create both its positive and negative samples which involve semantic-irrelevant and semantic-relevant perturbations, respectively. The contrastive learning framework ensures the KGD model is aware of these two types of perturbations, thus generating informative responses with the potentially noisy inputs in real applications. Experimental results on three benchmark datasets show that our method achieves new state-of-the-art performance in terms of automatic evaluation scores, verifying its effectiveness and potentiality. Furthermore, we show that our method can generate better responses than comparison models in both the noisy and the few-shot settings.

Introduction

Knowledge-Grounded Dialogue (KGD) aims to generate an informative response based on a given dialogue context and external knowledge to improve the usefulness and meaningfulness of the generated responses (Ghazvininejad et al. 2017; Zhou et al. 2018; Liu et al. 2019; Kim, Ahn, and Kim 2020; Li et al. 2021; Rashkin et al. 2021; Wu et al. 2022; Sun et al. 2023). As for the choice of knowledge source, structural knowledge graphs (KGs) are proven options (Moon et al. 2019; Jung, Son, and Lyu 2020; Wu et al. 2022), which consist of a lot of knowledge facts that are frequently used in daily life (Zheng et al. 2021; Zhang et al. 2022; Cao et al.




Figure 1: Illustrations of robustness in knowledge-grounded dialogue. The response in the first dialogue satisfies humans, while the second does not due to erroneous facts in the KG.

2023; Li et al. 2023). Despite the great success that has been achieved by these efforts, especially along with the rapid development of large language models (LLMs), existing work neglects the robustness of KGD methods facing real-world noises. As shown in Figure 1, the dialogue context might involve inevitable perturbations like misspellings and abbreviations. The KGs might also convey erroneous and outdated facts. These perturbations pose a challenge to the robustness of KGD methods.

As large language models (LLMs) become a promising way toward artificial general intelligence, many dialogue systems utilize LLMs as their backbone to achieve strong conversational ability. However, as pointed out by recent work (Wang et al. 2023d), there is still room for improving their robustness even for ChatGPT. Besides, some studies (Zhu, Song, and Liu 2023; Wang et al. 2023a,b) show the robustness of knowledge-enhanced LLM is also limited. Thus, improving the model robustness becomes one of the key challenges when deploying KGD models in real scenes.

In this paper, we propose an **entity-based contrastive learning** (EnCo) framework for improving the robustness of KGD models. Our key insight is to add perturbations in the given dialogue context as well as related knowledge, and let the KGD model learn how to generate an informative response with the potentially noisy inputs. Specifically, for a vanilla KGD sample, we create its positive samples and negative samples, where the positive samples should share faithful (similar or clipped) semantics with the vanilla sample, and the negative samples involve conflict semantics with the

*Corresponding authors.

vanilla ones. The object of EnCo ensures the model to represent similar KGD samples in a shared space by training the encoder to minimize the representation distance of them. For conflict samples, EnCo makes the encoder maximize their representations to let the model be aware of the perturbations in negative samples. Ideally, once the KGD model has the ability to distinguish perturbations in the potentially noisy dialogue context or KGs, its robustness could be ensured in real applications.

Therefore, the key in our EnCo framework is how to create positive and negative samples that should reflect real-world noises as much as possible. (i) In positive samples, there might involve semantic-irrelevant perturbations in the real scene like misspellings and abbreviations (c.f., the second dialogue in Figure 1). To this end, in view of paraphrasing which restates the same meaning in different lexical or syntactic expressions (Bhagat and Hovy 2013), we utilize a paraphrasing model to create such positive samples whose dialogue contexts ideally share similar semantics with that in the vanilla sample. However, in the preliminary experiments, we find that the paraphrasing model might change the entities in the vanilla context to other similar entities during paraphrasing. Though these paraphrased entities are relevant to the original ones, they also might introduce semantic-relevant perturbations to cause semantic gaps which should be avoided. Thus, we introduce a simple yet effective method (named entity-guided paraphrasing) to explicitly let the paraphrasing model be aware of the entities in the context and keep them unchanged during paraphrasing at both the training and inference stages. (ii) The negative samples should contain conflict semantics with the vanilla sample. In this manner, the KGD model can learn to distinguish the semantic-relevant perturbations in the real scene. Existing contrastive learning models typically adopt the rest of the samples in the same mini-batch as negative samples (Jaiswal et al. 2020; Poddar, Wang, and Reinspach 2022; Jiang et al. 2023) or retrieve negative samples from candidates (Karpukhin et al. 2020), which do not meet our requirements. To create such negative samples in our scene, we propose an entity-guided negative augmentation strategy that randomly deletes or replaces original entities in the vanilla context. To make the augmented negative samples hard to distinguish, we also edit the corresponding KG in the same way. For example, if we replace an entity e_a in the vanilla context with another one e_b , we also replace e_a with e_b in the KG to eliminate the potential semantic gap between the context and the KG in the negative samples. This is because the semantic gap might let the KGD model learn the shortcut from comparing the dialogue context and the KG (both of them are inputs to the model) to distinguish the semantic-relevant perturbations.

Experimental results on three KGD benchmark datasets show that our EnCo outperforms the previous state-of-the-art models in terms of widely-used automatic evaluation metrics, indicating its effectiveness and superiority. We also create a robustness test set based on existing benchmark data and show that our method could improve the robustness of KGD models when faced with noisy inputs. Furthermore, we also conduct few-shot evaluation and human evaluation

to suggest that our method could generate satisfactory responses compared to existing related methods.

Our contributions are concluded as follows:

- To the best of our knowledge, we are the first to study the robustness of knowledge-grounded dialogue models, and we introduce an entity-based contrastive learning framework for improving the robustness of KGD models.
- To let the KGD models be aware of semantic-relevant and semantic-irrelevant perturbations in real applications, we propose to utilize the entity information to guide the creation of positive and negative samples via paraphrasing and negative augmentation strategy, respectively.¹
- Experimental results on three KGD benchmark datasets show that our method outperforms the previous state-of-the-art models. Robustness evaluation further demonstrates that our method could generate informative responses with noisy inputs. In addition, we show the few-shot ability of our method.

Related Work

Knowledge-Grounded Dialogue

Knowledge-grounded dialogue (KGD) aims to incorporate external knowledge to generate informative responses. Structural knowledge graphs (KGs) are proven options to serve as the knowledge source, providing a large number of knowledge facts in the real world (Wang et al. 2022a,b, 2023c; Zhang et al. 2023). Many previous KGD work utilizes commonsense KGs (Zhou et al. 2018; Wu et al. 2020) or domain KGs (Wang et al. 2022c) to guide the generation.

Some researchers (Vougiouklis, Hare, and Simperl 2016; Ghazvininejad et al. 2017) attempt to leverage memory networks (Sukhbaatar et al. 2015) to store the relevant knowledge and then generate responses conditioned on both the dialogue context and the stored knowledge. To extract more relevant knowledge from the knowledge source, some work (Lian et al. 2019; Wu et al. 2020) selects the knowledge by approximating the prior distribution, *i.e.*, $P(\text{knowledge}|\text{context})$ based on the posterior distribution, *i.e.*, $P(\text{knowledge}|\text{context}, \text{response})$, leading to accurate knowledge selection. Besides, some researchers find that the response might involve words/tokens that existed in the knowledge resource (named knowledge words), and they employ various copy mechanisms to improve the generation probability of the knowledge words (Lin et al. 2020; Wu et al. 2020; Bai et al. 2021; Liang et al. 2021). To effectively excavate the structural information of KGs, previous literature also utilizes graph neural networks (GNNs) (Liu et al. 2019; Wu et al. 2021) or knowledge graph embeddings (KGE) (Zhou et al. 2018; Wang et al. 2022c) over KGs to obtain their structure-aware representation that is further incorporated into the dialogue generation process. Different from the previous KGD work which typically focuses on how to select relevant knowledge and how to generate high-quality responses conditioned on the relevant knowledge, we are the first to study the robustness of KGD models when faced with real-world noises.

¹<https://github.com/kxinwang2023/EnCo>

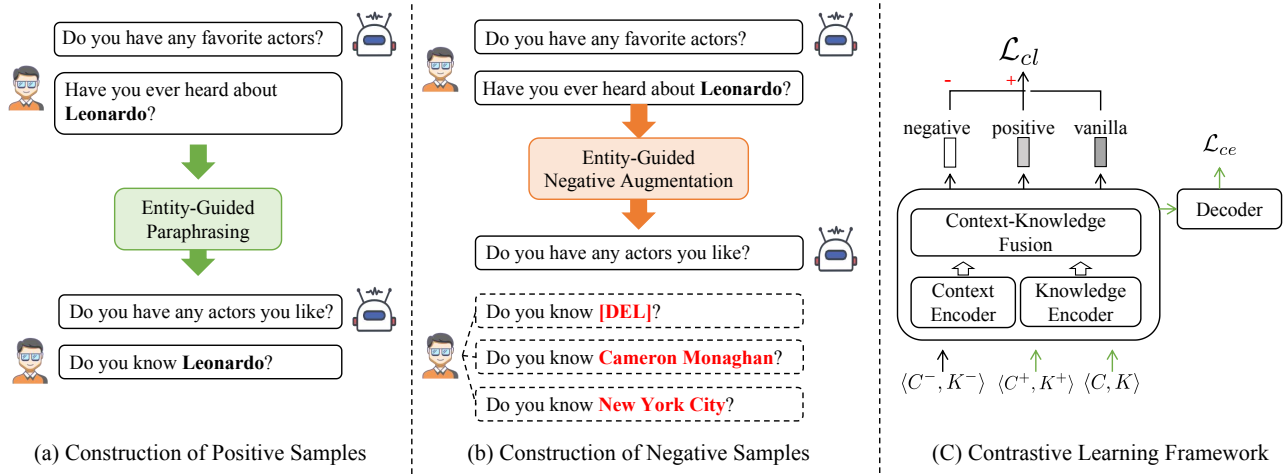


Figure 2: Illustrations of positive samples and negative samples in our entity-based contrastive learning framework.

Robustness of Dialogue Systems

There is a few work studies the robustness of dialogue systems. Sengupta, Krone, and Mansour (2021) focus on the robustness of task-oriented dialog systems. They show that common noise types (*e.g.*, misspellings) substantially degrade the accuracy of dialogue systems’ sub-models. Poddar, Wang, and Reinspach (2022) study the robustness of retrieval-based dialogue systems when faced with perturbations. They use contrastive learning as an auxiliary objective to learn robust dialogue context representations and make retrieval-based dialogue systems retrieve proper responses from candidates. Chen et al. (2023) find that the input orders of persona sentences significantly impact the quality and consistency of the personalized dialogue systems. They propose to learn robust representation under different persona orders and improve the consistency of response.

Existing robustness studies of dialogue systems generally focus on task-oriented dialogue, retrieval-based dialogue as well as personalized dialogue, whose settings are different from that of knowledge-grounded dialogue in our work.

Methodology

In this section, we first define the knowledge-ground dialogue (KGD) task and then elaborate on the details of our entity-based contrastive learning (EnCo) framework. As illustrated in Figure 2, for a given KGD sample, the EnCo framework utilizes an entity-guided paraphrasing model and entity-guided negative augmentation to construct positive and negative samples, respectively. To encode the KGD samples, EnCo involves a context encoder to encode the dialogue context, and a knowledge encoder to encode the relevant knowledge triples. Next, a context-knowledge fusion module is used to fuse the information of both dialogue context and external knowledge. Further, a decoder is employed to generate responses. For contrastive learning, EnCo makes the encoder minimize the representation distance between the given sample and the positive sample, and maximize that between the given sample and the negative sample. In

this way, the KGD model can learn to distinguish both the semantic-relevant and the semantic-irrelevant perturbations, improving its robustness when faced with real-world noises.

Task Definition

Given a context $C = \{u_1, u_2, \dots, u_{n-1}\}$ and the relevant knowledge triple set $K = \{(h_1, r_1, t_1), \dots, (h_m, r_m, t_m)\}$, where u_i represents the i -th utterance in the dialogue, (h_j, r_j, t_j) denotes the head entity h_j and tail entity t_j have the relation r_j . The goal of knowledge-grounded dialogue systems is to generate a proper response u_n based on the dialogue context C and the knowledge triples K .

The Construction of Positive Samples

To make the KGD model aware of the semantic-irrelevant perturbations in C , we paraphrase the context to construct the positive samples for a given KGD sample. In this way, the paraphrased samples ideally share similar semantics with the vanilla ones but with different expressions.

Considering the changing of entities during paraphrasing might also change the semantics involved in the context, we design a simple yet effective entity-guided paraphrasing that explicitly models the entity information in the following two steps: (1) we first use an off-the-shelf named entity recognition (NER) toolkit (*i.e.*, TexSmart²) to mine all the entities involved in the context C . Then, (2) we give the boundary information of entities by adding two special tokens: $[\text{Ent}]$ and $[\backslash\text{Ent}]$. For example, when paraphrasing the sentence “Do you know Leonardo?”, the NER toolkit first extracts the entity “Leonardo”, and then the special tokens are added to the entity boundary to form the input of the paraphrase model: “Do you know $[\text{Ent}]$ Leonardo $[\backslash\text{Ent}]$?”. Thus, the paraphrase model could be aware of the entity information lying in the input sentences, and not change them during paraphrasing.

The architecture of the paraphrase model. We use BART-large model (Lewis et al. 2019) as the backbone of the

²<https://texsmart.qq.com/en>

paraphrase model. This model contains standard transformer encoder-decoder architecture (Vaswani et al. 2017) (12 encoder layers, 12 decoder layers, 16 multi-head attention as well as 1024 hidden states). The BART model has been pre-trained on large-scale corpora with the self-supervised auto-denosing pre-training objectives (Lewis et al. 2019).

Training of the paraphrase model. ParaZh-22M (Hao et al. 2022) is a large-scale paraphrase dataset with about 22M sentence pairs. We utilize the NER toolkit (TextSmart) to pre-process the dataset. In detail, we first extract the entities from each sentence pair, and then add the special tokens to the entity boundaries. To let the paraphrase model not change the entities involved in the input sentences, we only reserve the sentence pairs whose source and target sentences contain the same entity set, resulting in about 6.5M training samples. The paraphrase model is trained to generate the golden paraphrased sentences in the text-generation style:

$$p_{\theta}(\hat{s}|s) = \prod_{t=1}^{|\hat{s}|} p_{\theta}(\hat{s}_t|s, \hat{s}_{1:t-1}) \quad (1)$$

where θ denotes the parameters of the paraphrase model, s and \hat{s} indicate the input sentence and the golden paraphrased sentence, respectively. $\hat{s}_{1:t-1}$ is the partial paraphrasing.

Inference of the paraphrase model. After training the paraphrase model, the model is used to paraphrase the dialogue context $C = \{u_1, u_2, \dots, u_{n-1}\}$. We input each utterance u_i to the paraphrase model, and use the top- k sampling strategy to decode the paraphrased utterance \hat{u}_i . All paraphrased utterances construct the context of the positive sample, which we denote as $C^+ = \{\hat{u}_1, \hat{u}_2, \dots, \hat{u}_{n-1}\}$.

Furthermore, to improve the robustness of the KGD model when faced with incomplete knowledge, the knowledge triple set used in the positive samples (denoted as K^+) is truncated from the vanilla ones (K). Specifically, we randomly remove $\lambda\%$ triples in K to obtain K^+ , and λ is randomly chosen from 0 to 15 for each sample. Note that removing information does not introduce conflict semantics, thus the truncated knowledge set can be used in the positive samples, making the KGD model learn to still generate informative responses with incomplete grounding knowledge.

The Construction of Negative Samples

Different from the positive sample $\langle C^+, K^+ \rangle$, a negative sample, denoted as $\langle C^-, K^- \rangle$, should introduce semantic-relevant perturbations to the given KGD sample $\langle C, K \rangle$. In view of existing contrastive learning models, they could be classified into two methods: (1) adopting the rest of the samples in the same mini-batch as negative samples (Jaiswal et al. 2020; Jiang et al. 2023) or (2) retrieving negative samples from candidates (Karpukhin et al. 2020). In KGD, the rest samples in the same mini-batch only contain irrelevant information and do not introduce semantic-relevant perturbations. Besides, we do not have negative candidates in the KGD scene. Thus, simply utilizing existing methods can not let the KGD model be aware of the semantic-relevant perturbations, making us decide to use the entity information lying in the context C to construct negative samples.

We first make use of the NER toolkit (TextSmart) to mine all the entities involved in C , and for each entity, it has a 30% probability to make one of the following changes (the remaining 70% probability does not change anything):

- *Randomly Deletion:* The entity will be removed from C , and we replace it with a special token “[DEL]”.
- *Relevant Replacement:* The entity will be replaced with another entity that has the same entity type as the entity.
- *Irrelevant Replacement:* The entity will be replaced with another one that has a different type from the entity.

The above three changes share 10%, 80% and 10% probabilities to employ for each of the changing entities (which are first selected with the 30% probability). In this manner, the changed context C^- might involve the different degrees of semantic-relevant perturbations. After that, the entities in the knowledge set K also conduct the same changes as the context C to obtain K^- . For example, if we delete an entity in C and mark it as “[DEL]”, the knowledge set K will also delete it and denote it as “[DEL]” (if has).

Contrastive Learning Framework

After constructing the positive and negative samples, we use a KGD model to encode all the samples and minimize the representation distance between the vanilla sample and the positive sample while maximizing that between the vanilla sample and the negative sample.

KGD model. As shown in Figure 2 (c), the KGD model receives dialogue context C and knowledge triple set K as inputs and generates responses. The model contains the following four modules: (1) a *context encoder* to calculate the context representation, which involves N_e stacked transformer encoder layers, where each layer consists of two sub-layers, a multi-head self-attention sublayer (SelfAttn) and a position-wise feed-forward network (FFN) sub-layer:

$$S_C^\ell = \text{SelfAttn}(H_C^{\ell-1}) + H_C^{\ell-1}, S_C^\ell \in \mathbb{R}^d \quad (2)$$

$$H_C^\ell = \text{FFN}(S_C^\ell) + S_C^\ell, H_C^\ell \in \mathbb{R}^d \quad (3)$$

where $H_C^{\ell-1}$ and H_C^ℓ denote the inputs and outputs of the ℓ -th layer, respectively, and H_C^0 is initialized as the embedding of input context C and d is the hidden dimension.

(2) A *knowledge encoder* is employed to calculate the knowledge representation. Following Wang et al. (2022c), we use TransR (Lin et al. 2015) to obtain the representation of each knowledge triple $(h_i, r_i, t_i) \in K$:

$$e_{h_i}, e_{r_i}, e_{t_i} \leftarrow \text{TransR}(h_i), \text{TransR}(r_i), \text{TransR}(t_i) \quad (4)$$

$$h_{k_i} = e_{h_i} \oplus e_{r_i} \oplus e_{t_i} \quad (5)$$

where e_{h_i} , e_{r_i} and e_{t_i} indicate the representations of h_i , r_i and t_i , respectively. h_{k_i} denotes the representation of the knowledge triple $k_i = (h_i, r_i, t_i)$ and \oplus means concatenation. In this way, we obtain all triple representations $\{h_{k_1}, h_{k_2}, \dots, h_{k_m}\}$ via the knowledge encoder.

(3) a *context-knowledge fusion module* to fuse the context and the knowledge representations. We adopt multi-head attention to fuse each triple representation h_{k_i} with the context information:

$$h_{k_i}^c = \|\|_{h=1}^m \text{Attn}_h(h_{k_i} \leftarrow H_C^{N_e}) \quad (6)$$

$$\text{Attn}_h = \sum_j \text{softmax}\left(\frac{Q_h(h_{k_i}) \cdot K_h(H_j^{N_e})}{\sqrt{d}}\right) \cdot V_h(H_j^{N_e}) \quad (7)$$

where $h_{k_i}^c$ indicates the context-enriched representation of triple k_i . $\prod_{h=1}^m$ indicates multi-head attention and m is the number of multi-head attention. Attn_h denotes the h -th head, whose query vector, key vector and value vector are denoted as $Q_h(\cdot)$, $K_h(\cdot)$ and $V_h(\cdot)$, respectively. $H_j^{N_e}$ means the j -th vector in $H_C^{N_e}$ that also indicates the representation of the j -th token in the context C .

(4) A *decoder* is used to generate responses, which consists of N_d stacked transformer decoder layers. To let the decoder be conditioned on both the context and the knowledge, we modified the vanilla transformer decoder layer by adding parallel cross-attention over context information $H_C^{N_e}$ and knowledge information $H_K^c = [h_{k_1}^c; h_{k_2}^c; \dots; h_{k_m}^c]$:

$$S_{dec}^\ell = \text{SelfAttn}(H_{dec}^{\ell-1}) + H_{dec}^{\ell-1} \quad (8)$$

$$C_{dec}^\ell = \text{CrossAttn}(S_{dec}^\ell, H_C^{N_e}) + \text{CrossAttn}(S_{dec}^\ell, H_K^c) + S_{dec}^\ell \quad (9)$$

$$H_{dec}^\ell = \text{FFN}(C_{dec}^\ell) + C_{dec}^\ell \quad (10)$$

where H_{dec}^ℓ denotes the state of the ℓ -th decoder layer. Then, at each decoding time step t , the top-layer (N_d -th) decoder hidden state $H_{dec,t}^{N_d}$ is fed into the softmax layer to produce the probability distribution of the next target token as:

$$p(u_n^t | C, K, u_n^{<t}) = \text{Softmax}(W_o H_{dec,t}^{N_d} + b_o) \quad (11)$$

where W_o and b_o are trainable parameters, and u_n^t denote the t -th token in the golden response u_n .

We use the vanilla sample and the positive samples to calculate cross-entropy loss:

$$\mathcal{L}_{van} = - \sum_{t=1}^{|u_n|} \log(p(u_n^t | C, K, u_n^{<t})) \quad (12)$$

$$\mathcal{L}_{pos} = - \sum_{(C^+, K^+) \in S^{pos}} \sum_{t=1}^{|u_n|} \log(p(u_n^t | C^+, K^+, u_n^{<t})) \quad (13)$$

$$\mathcal{L}_{ce} = \mathcal{L}_{van} + \mathcal{L}_{pos} \quad (14)$$

where S^{pos} denotes the set of positive samples.

Contrastive Loss. We also adopt the following contrastive loss over the vanilla, positive and negative KGD samples:

$$\mathcal{L}_{ctr} = - \sum_{C^+} \sum_{C^-} \log \frac{f(H_C^{N_e}, H_{C^+}^{N_e})}{f(H_C^{N_e}, H_{C^+}^{N_e}) + f(H_C^{N_e}, H_{C^-}^{N_e})} \quad (15)$$

where $f(a, b) = \exp(a^\top b)$. Therefore, our final loss is:

$$\mathcal{L} = \mathcal{L}_{ce} + \alpha \mathcal{L}_{ctr} \quad (16)$$

Experiments

Implementation Details

We implement our EnCo framework with PyTorch and Huggingface Transformers (Wolf et al. 2020) libraries. The context encoder and decoder are initialized by a pre-trained Chinese BART-large model, *i.e.*, the numbers (N_e and N_d) of

encoder layers as well as decoder layers are 12, and the hidden dimension d is 1,024. The number of multi-head attention in the context-knowledge fusion module (m) is set to 8. Following Wang et al. (2022c), the embedding size of entities and relations is set to 200. For each KGD sample, we create 5 positive and 5 negative samples. We train the KGD model on two 32GB Tesla V100 GPUs. We set the hyperparameters based on the preliminary experiments on the development set. We leverage the Adam optimizer with a default initial momentum and adopt linear warmup in the first 1,000 steps. The mini-batch size is set to 8, and the coefficient α in the final loss function is set to 1.0. We use minimal hyperparameter tuning using Learning Rates (LRs) in [1e-5, 2e-5, 3e-5, 5e-5] and epochs of 10 to 20. We find the model with LR of 5e-5 and 20 epochs to work best. All experimental results listed in this paper are the average of 3 runs.

Experimental Setups

Datasets. We conduct experiments on three widely-used public KGD datasets: (1) KdConv (Zhou et al. 2020) involves 4.5K dialogues and 86K utterances from music, travel and film domains. There are a total of 85K KGD samples in KdConv. The knowledge in KdConv contains 13.1K entities, 9.1K relations and 157.0K triples. (2) DuConv (Wu et al. 2019) contains 180K samples about film and entertainment (29K dialogues as well as 270K utterances). The knowledge in DuConv involves about 3.6M triples. (3) DuRecDial (Liu et al. 2020) has 145K samples with 10.2K dialogues. It also contains 21.8K entities, 454 relations and 222.2K triples. All samples in these three datasets are annotated with dialogue-level relevant knowledge triples.

Metrics. Following previous work (Zhou et al. 2020; Liu et al. 2020), we adopt perplexity (PPL), BLEU (Papineni et al. 2002) and DISTINCT (Li et al. 2016) to measure the fluency, relevance and diversity of the generated responses.

Baselines. We compare the proposed method with the following baselines: (1) **Seq2Seq** (Luong, Pham, and Manning 2015) is a stack RNN-based model with attention mechanisms; (2) **HRED** (Serban et al. 2015) is a hierarchical RNN-based generative model; (3) **Seq2Seq+Know** and (4) **HRED+Know** are introduced by Zhou et al. (2020), which fuse the context vector with the knowledge vector to form the initial state of the decoder; (5) **KIC** (Lin et al. 2020) uses recurrent knowledge interaction among response decoding steps to incorporate appropriate knowledge in dialogues; (6) **SDAN** (Cui et al. 2021) utilizes a knowledge-aware network to improve informativeness, and a syntactic latent variable network to generate syntactically diverse responses; (7) **MGCC_G** (Liu et al. 2020) consists of a goal-planning module and a goal-guided responding module to handle multi-type dialogs; (8) **KSPN** (Liu et al. 2022) proposes a knowledge selection-guided pointer network into the decoder to generate responses with the words from the captured knowledge; (9) **BART** (Lewis et al. 2019) is a transformer-based generative model which has been pre-trained with auto-denoising objectives; (10) **BART+Know** uses our knowledge encoder and context-knowledge fusion module to incorporate knowledge into BART; (11) **KRP-**

Model	KdConv (music)					KdConv (travel)					KdConv (film)				
	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)
Seq2Seq	16.17	28.89	10.63	2.52	12.69	10.44	29.61	14.91	3.75	19.01	23.88	26.97	8.53	2.51	13.62
HRED	16.82	29.92	11.17	2.71	14.07	10.90	30.92	15.61	4.15	20.52	24.74	27.03	8.30	2.55	14.12
Seq2Seq+K.	17.12	29.60	11.36	3.93	23.01	10.62	37.04	22.16	4.25	24.18	25.56	27.45	8.66	2.85	15.09
HRED+K.	17.69	29.73	11.59	3.80	22.00	11.15	36.87	21.31	3.98	24.06	26.27	27.94	8.73	2.86	15.81
KIC	13.06	30.41	13.87	3.31	23.49	8.46	37.21	23.30	3.45	25.47	11.29	28.12	9.53	2.63	26.74
SDAN	14.78	30.92	14.40	4.03	23.07	9.32	38.13	24.96	3.62	25.31	14.52	28.96	10.22	3.25	23.47
KSPN	3.59	35.55	23.77	3.32	29.22	2.51	44.39	33.85	2.89	25.96	3.80	29.44	14.92	3.03	30.27
BART	2.44	32.27	18.44	2.80	25.19	1.69	36.61	26.54	2.56	22.85	2.82	29.68	15.26	2.50	27.96
BART+K.	2.89	34.21	23.11	3.04	28.09	1.88	39.45	30.55	3.10	24.80	3.06	29.98	15.44	2.89	29.02
KRP-DS	3.05	36.88	23.89	3.55	29.47	2.08	45.00	34.69	2.71	26.59	3.17	30.14	15.80	3.29	30.84
EnCo (Our)	3.91	39.39	25.11	4.23	31.04	2.57	46.61	37.02	3.82	28.74	3.53	31.43	16.48	3.65	33.29

Table 1: Experimental results on the KdConv dataset. ↑ indicates higher is better. ↓ indicates lower is better. (K.: Knowledge)

	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)		
Seq2Seq	9.14	26.49	12.55	3.16	15.02
HRED	22.82	32.71	14.89	4.11	16.21
Seq2Seq+K.	10.96	28.32	13.96	4.30	16.48
HRED+K.	24.30	34.67	15.11	4.71	17.68
KIC	10.36	37.70	22.09	6.31	37.66
SDAN	8.42	36.47	17.31	5.53	18.60
MGCG_G	6.43	35.14	19.10	6.17	28.81
KSPN	4.02	38.92	17.01	6.02	37.02
BART	3.20	36.54	20.07	5.33	35.75
BART+K.	3.86	38.91	23.71	5.89	36.39
KRP-DS	5.06	39.06	23.96	5.83	36.91
EnCo (Our)	5.57	40.10	24.89	6.88	38.50

Table 2: Experimental results on the DuConv dataset.

	PPL (↓)	BLEU-1/3 (↑)	DIST-1/3 (↑)		
Seq2Seq	20.10	16.19	5.41	0.42	8.45
HRED	21.56	22.76	6.84	0.72	10.81
Seq2Seq+K.	22.82	18.81	6.34	0.66	9.92
HRED+K.	23.96	24.33	8.94	1.03	12.24
KIC	15.59	31.78	17.73	1.52	19.67
SDAN	19.72	26.11	10.17	2.17	14.01
MGCG_G	14.89	36.22	21.79	2.33	15.26
KSPN	5.50	38.01	23.19	2.13	18.67
BART	2.92	37.95	22.70	2.41	20.83
BART+K.	3.42	39.17	24.32	2.63	24.67
KRP-DS	3.50	39.02	24.41	2.51	25.03
EnCo (Our)	4.36	40.05	24.68	2.78	31.48

Table 3: Experimental results on the DuRecDial dataset.

DS (He et al. 2023) utilizes contextual information for path reasoning and guides knowledge prediction in KGD.

Main Results

Table 1, Table 2 and Table 3 show the results on KdConv, DuConv and DuRecDial, respectively. As we can see, the results on the three datasets indicate a similar trend. Specifically, EnCo achieves the best performance in terms of most metrics. Compared with the second-performance baseline (*i.e.*, KRP-DS), EnCo is significantly better than it with t-test $p < 0.01$ in terms of BLEU and DISTINC scores, showing its effectiveness. The positive samples created in our EnCo framework could also be regarded as a data augmentation approach for KGD, making the KGD model learn from diverse samples. In terms of PPL, BART achieves the best performance. Besides, we find that the PPL generally increases when incorporating knowledge into KGD models. For example, HRED achieves 16.82 PPL on KdConv (music) while the counterpart of HRED+Know is 17.69 (+0.87). We conjecture this is because the knowledge will influence the model’s output distributions during inferences, and thus increase PPL. This situation has also been found in previous work (Zhou et al. 2020; Wang et al. 2022c).

Robustness Study

To test the model’s performance when faced with real-world noises, we add the following noises into the test set of DuConv, respectively: (1) randomly deleting 30% words in the dialogue context; (2) randomly replacing 30% noun words with their synonyms³; (3) randomly deleting 30% utterances in the context except for the last utterance; (4) paraphrasing the dialogue context using our entity-guided paraphrase model; (5) introducing noises by ChatGPT with the prompt of “Dialogues in the real world can be noisy. For example, the original dialogue: $[D]$ in the real world might be like: $[\hat{D}]$. Please generate the noisy version of the following dialogue: $[D_i]$ ”. D and \hat{D} denote an original and noisy dialogue pair, which is written by our human annotators and serves as an in-context sample for ChatGPT to ground. There are a total of 20 in-context samples that involve different kinds of real-world noises like misspellings and abbreviations, and we randomly select one sample to prompt ChatGPT at each time. D_i is the input dialogue from the test set of DuConv. (6) randomly deleting 30% entities in KG; (7) randomly replacing 30% entities with other entities.

Table 4 lists the results of the robustness study. Our EnCo achieves the best performances among all noisy testing in

³<https://github.com/chatopera/Synonyms>

Model	Vanilla Test	Word Replacement	Synonym Replacement	Utterance Deletion	Paraphrasing	Noises (ChatGPT)	Entity Deletion	Entity Replacement
HRED+K.	34.67 / 19.81	25.44 / 15.01	28.57 / 18.16	30.12 / 19.72	28.86 / 18.48	26.78 / 17.03	29.11 / 18.54	27.52 / 17.30
BART+K.	38.91 / 28.50	32.37 / 21.94	33.50 / 23.12	34.03 / 23.62	33.29 / 22.95	31.58 / 19.94	32.91 / 22.15	31.77 / 21.09
EnCo	40.10 / 29.69	34.49 / 24.05	35.49 / 25.09	35.10 / 24.70	35.33 / 24.97	34.07 / 22.62	37.42 / 27.02	36.11 / 26.18

Table 4: Experimental results of robustness study in terms of BLEU-1 and BLEU-2.

	KdConv (music)	KdConv (travel)
EnCo	39.39 / 30.18 / 25.11	46.61 / 40.58 / 37.02
w/o \mathcal{L}_{pos}	36.59 / 28.01 / 23.54	43.29 / 37.51 / 34.05
w/o \mathcal{L}_{ctr}	36.02 / 27.19 / 23.21	42.80 / 36.95 / 33.47
BART+K.	34.21 / 26.14 / 23.11	39.45 / 33.72 / 30.55

Table 5: Results of ablations (BLEU-1/2/3) on DuConv.

terms of BLEU-1 and BLEU-2 (other metrics also show the same situations), verifying its effectiveness when faced with real-world noises. The positive and negative samples in EnCo let the KGD model be aware of both the semantic-irrelevant and the semantic-relevant perturbations, thus improving the model’s robustness in real applications.

Ablation Results

Compared with the BART+Know baseline, our EnCo adopts two improvements: (1) the positive samples are used as data augmentation (*i.e.*, the \mathcal{L}_{pos} in Eq. 14); (2) the contrastive learning framework (*i.e.*, the \mathcal{L}_{ctr} in Eq. 16). We conduct ablation studies to investigate the effect of \mathcal{L}_{pos} and \mathcal{L}_{ctr} by removing each of them. As shown in Table 5, when removing either \mathcal{L}_{pos} or \mathcal{L}_{ctr} , the model performance will decrease. Thus, the effectiveness of both improvements is proven.

Few-Shot Results

Since the positive samples in EnCo could also be regarded as a data augmentation approach. Following Poddar, Wang, and Reinspach (2022), we also attempt the following two trivial data augmentation methods to create positive samples: (1) *word deletion* randomly selects 70% of words in a dialogue context, and replaces them with a special token [DEL]; (2) *word reordering* randomly samples several pairs of words in a dialogue context (about 30% of context words), and switch them pairwise. We conduct experiments on DuConv using 1%, 10%, 50% and 100% training samples, respectively.

As shown in Figure 3, EnCo significantly surpasses all comparison models under each setting. Particularly, under the 1% setting, our model still achieves the best performances, indicating that our model works well in the few-shot setting as well. Besides, EnCo (WD) and EnCo (WR) denote using word deletion and word reordering to create positive samples in our EnCo framework, respectively. The vanilla EnCo outperforms these two variants, indicating the effectiveness of our entity-guided paraphrasing.

Human Study

Furthermore, we conduct human study on 200 random samples from DuConv. We compare the responses generated by

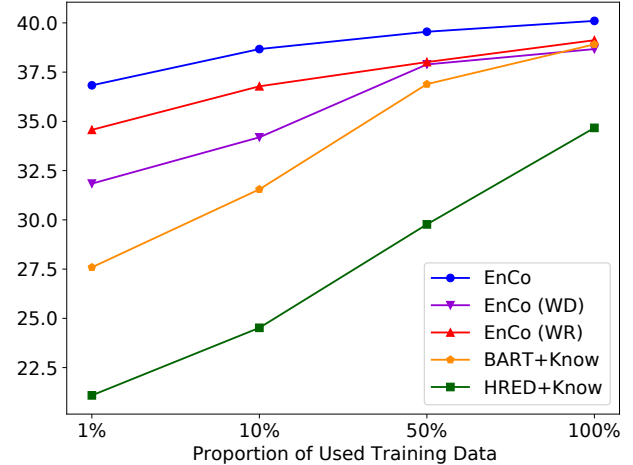


Figure 3: Few-shot result on DuConv (BLEU-1).




Figure 4: Results on human study.

EnCo with the responses generated by HRED+Know and BART+Know. Five master students are recruited to score the responses in terms of fluency, coherence and informativeness with a 3-point scale. As shown in Figure 4, EnCo outperforms HRED+Know and BART+Know on all three aspects, which verifies that our method performs better in generating fluent, coherent and informative responses.

Conclusion

In this paper, we study the robustness of knowledge-grounded dialogue (KGD) models when faced with real-world noises. We propose an entity-based contrastive learning (EnCo) framework to create positive and negative samples under the guidance of entity information. Then, we use contrastive learning to let the KGD model be aware of different perturbations. Experimental results on three datasets show that our method outperforms state-of-the-art baselines. The robustness study and few-shot study further indicate the superiority of our method under the various types of noises and the few-shot situations, respectively.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant No. 62102276, 62272334, 62072323), Shanghai Science and Technology Innovation Action Plan (No. 22511104700), the Zhejiang Lab Open Research Project (NO. K2022NB0AB04), the Natural Science Foundation of Jiangsu Province (Grant No. BK20210705, BK20211307), China Postdoctoral Science Foundation (Grant No. 2023M732563), the Natural Science Foundation of Educational Commission of Jiangsu Province, China (Grant No. 21KJD520005), Key Projects of Industrial Foresight and Key Core Technology Research and Development in Suzhou (SYC2022009), Engineering Lab of Bigdata and Intelligence of Jiangsu Province and Project Funded by the Priority Academic Program Development of Jiangsu Higher Education Institutions.

References

- Bai, J.; Yang, Z.; Liang, X.; Wang, W.; and Li, Z. 2021. Learning to copy coherent knowledge for response generation. In *Proceedings of AAAI 2021*, 14, 12535–12543.
- Bhagat, R.; and Hovy, E. 2013. What is a paraphrase? *Computational Linguistics*, 39(3): 463–472.
- Cao, Y.; Xu, J.; Yang, C.; Wang, J.; Zhang, Y.; Wang, C.; Chen, L.; and Yang, Y. 2023. When to Pre-Train Graph Neural Networks? From Data Generation Perspective! In *Proceedings of KDD 2023*.
- Chen, L.; Wang, H.; Deng, Y.; Kwan, W. C.; Wang, Z.; and Wong, K.-F. 2023. Towards Robust Personalized Dialogue Generation via Order-Insensitive Representation Regularization. In *Findings of ACL 2023*, 7337–7345.
- Cui, F.; Di, H.; Ren, H.; Ouchi, K.; Liu, Z.; and Xu, J. 2021. Syntactically Diverse Adversarial Network for Knowledge-Grounded Conversation Generation. In *Proceedings of EMNLP 2021*.
- Ghazvininejad, M.; Brockett, C.; Chang, M.-W.; Dolan, W. B.; Gao, J.; tau Yih, W.; and Galley, M. 2017. A Knowledge-Grounded Neural Conversation Model. In *Proceedings of AAAI 2017*.
- Hao, W.; Xu, H.; Xiong, D.; Zan, H.; and Mu, L. 2022. ParaZh-22M: A Large-Scale Chinese Parabank via Machine Translation. In *Proceedings of COLING 2022*.
- He, Q.; Xu, S.; Zhu, Z.; Wang, P.; Li, K.; Zheng, Q.; and Li, Y. 2023. KRP-DS: A Knowledge Graph-Based Dialogue System with Inference-Aided Prediction. *Sensors*, 23(15): 6805.
- Jaiswal, A.; Babu, A. R.; Zadeh, M. Z.; Banerjee, D.; and Makedon, F. 2020. A survey on contrastive self-supervised learning. *Technologies*, 9(1): 2.
- Jiang, C.; Ye, W.; Xu, H.; Huang, S.; Huang, F.; and Zhang, S. 2023. Vision Language Pre-training by Contrastive Learning with Cross-Modal Similarity Regulation. In *Proceedings of ACL 2023*.
- Jung, J.; Son, B.; and Lyu, S. 2020. AttnIO: Knowledge Graph Exploration with In-and-Out Attention Flow for Knowledge-Grounded Dialogue. In *Proceedings of EMNLP 2020*.
- Karpukhin, V.; Oguz, B.; Min, S.; Lewis, P.; Wu, L.; Edunov, S.; Chen, D.; and Yih, W.-t. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of EMNLP 2020*, 6769–6781.
- Kim, B.; Ahn, J.; and Kim, G. 2020. Sequential Latent Knowledge Selection for Knowledge-Grounded Dialogue. In *Proceedings of ICLR 2020*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; rahman Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2019. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of ACL 2019*.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, W. B. 2016. A Diversity-Promoting Objective Function for Neural Conversation Models. In *Proceedings of NAACL 2016*.
- Li, Q.; Guo, S.; Luo, Y.; Ji, C.; Wang, L.; Sheng, J.; and Li, J. 2023. Attribute-Consistent Knowledge Graph Representation Learning for Multi-Modal Entity Alignment. In *Proceedings of the ACM Web Conference 2023*, 2499–2508.
- Li, Y.; Peng, B.; Shen, Y.; Mao, Y.; Lidén, L.; Yu, Z.; and Gao, J. 2021. Knowledge-Grounded Dialogue Generation with a Unified Knowledge Representation. In *Proceedings of NAACL 2021*.
- Lian, R.; Xie, M.; Wang, F.; Peng, J.; and Wu, H. 2019. Learning to Select Knowledge for Response Generation in Dialog Systems. In *Proceedings of IJCAI 2019*.
- Liang, Y.; Meng, F.; Zhang, Y.; Xu, J.; Chen, Y.; and Zhou, J. 2021. Infusing Multi-Source Knowledge with Heterogeneous Graph Neural Network for Emotional Conversation Generation. In *Proceedings of AAAI 2021*.
- Lin, X.; Jian, W.; He, J.; Wang, T.; and Chu, W. 2020. Generating Informative Conversational Response using Recurrent Knowledge-Interaction and Knowledge-Copy. In *Proceedings of ACL 2020*, 41–52.
- Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; and Zhu, X. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of AAAI 2015*.
- Liu, M.; Zhao, P.; Liu, J.; Zhuang, Y.; and Yang, Y. 2022. Improving knowledge-based dialogue generation through two-stage knowledge selection and knowledge selection-guided pointer network. *Journal of Intelligent Information Systems*, 59(3): 591–611.
- Liu, Z.; Niu, Z.-Y.; Wu, H.; and Wang, H. 2019. Knowledge Aware Conversation Generation with Explainable Reasoning over Augmented Graphs. In *Proceedings of EMNLP 2019*.
- Liu, Z.; Wang, H.; Niu, Z.-Y.; Wu, H.; Che, W.; and Liu, T. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. In *Proceedings of ACL 2020*.
- Luong, T.; Pham, H.; and Manning, C. D. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of EMNLP 2015*, 1412–1421.
- Moon, S.; Shah, P.; Kumar, A.; and Subba, R. 2019. OpenDialog: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. In *Proceedings of ACL 2019*.

- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of ACL 2002*, 311–318.
- Poddar, L.; Wang, P.; and Reinspach, J. A. 2022. DialAug: Mixing up Dialogue Contexts in Contrastive Learning for Robust Conversational Modeling. In *Proceedings of COLING 2022*.
- Rashkin, H.; Reitter, D.; Tomar, G. S.; and Das, D. 2021. Increasing Faithfulness in Knowledge-Grounded Dialogue with Controllable Features. In *Proceedings of ACL 2021*.
- Sengupta, S.; Krone, J.; and Mansour, S. 2021. On the robustness of intent classification and slot labeling in goal-oriented dialog systems to real-world noise. *arXiv preprint arXiv:2104.07149*.
- Serban, I.; Sordani, A.; Bengio, Y.; Courville, A. C.; and Pineau, J. 2015. Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. In *Proceedings of AAAI 2015*.
- Sukhbaatar, S.; Weston, J.; Fergus, R.; et al. 2015. End-to-end memory networks. *Proceedings of NeurIPS 2015*.
- Sun, B.; Li, Y.; Mi, F.; Bie, F.; Li, Y.; and Li, K. 2023. Towards Fewer Hallucinations in Knowledge-Grounded Dialogue Generation via Augmentative and Contrastive Knowledge-Dialogue. In *Proceedings of ACL 2023*.
- Vaswani, A.; Shazeer, N. M.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proceedings of NeurIPS 2017*.
- Vougiouklis, P.; Hare, J.; and Simperl, E. 2016. A neural network approach for knowledge-driven response generation. In *Proceedings of COLING 2016*, 3370–3380.
- Wang, J.; Li, Z.; Zhang, T.; Zheng, D.; Qu, J.; Liu, A.; Zhao, L.; and Chen, Z. 2022a. Knowledge enhanced sports game summarization. In *Proceedings of WSDM 2022*, 1045–1053.
- Wang, J.; Liang, Y.; Meng, F.; Sun, Z.; Shi, H.; Li, Z.; Xu, J.; Qu, J.; and Zhou, J. 2023a. Is chatgpt a good nlg evaluator? a preliminary study. In *Proceedings of the 4th New Frontiers in Summarization Workshop*.
- Wang, J.; Liang, Y.; Meng, F.; Zou, B.; Li, Z.; Qu, J.; and Zhou, J. 2023b. Zero-shot cross-lingual summarization via large language models. In *Proceedings of the 4th New Frontiers in Summarization Workshop*, 12–23.
- Wang, J.; Qu, J.; Liang, Y.; Li, Z.; Liu, A.; Liu, G.; and Zheng, X. 2023c. Snowman: A Million-scale Chinese Commonsense Knowledge Graph Distilled from Foundation Model. *arXiv preprint arXiv:2306.10241*.
- Wang, J.; Xixu, H.; Hou, W.; Chen, H.; Zheng, R.; Wang, Y.; Yang, L.; Ye, W.; Huang, H.; Geng, X.; et al. 2023d. On the Robustness of ChatGPT: An Adversarial and Out-of-distribution Perspective. In *ICLR 2023 Workshop on Trustworthy and Reliable Large-Scale Machine Learning Models*.
- Wang, J.; Zou, B.; Li, Z.; Qu, J.; Zhao, P.; Liu, A.; and Zhao, L. 2022b. Incorporating commonsense knowledge into story ending generation via heterogeneous graph networks. In *Proceedings of DASFAA 2022*, 85–100. Springer.
- Wang, K.; Li, Z.; Wang, J.; Qu, J.; He, Y.; Liu, A.; and Zhao, L. 2022c. RT-KGD: relation transition aware knowledge-grounded dialogue generation. In *Proceedings of ISWC 2022*, 319–335. Springer.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of EMNLP 2020 (system demonstrations)*, 38–45.
- Wu, S.; Li, Y.; Xue, P.; Zhang, D.; and Wu, Z. 2022. Section-aware commonsense knowledge-grounded dialogue generation with pre-trained language model. In *Proceedings of COLING 2022*, 521–531.
- Wu, S.; Li, Y.; Zhang, D.; Zhou, Y.; and Wu, Z. 2020. Diverse and Informative Dialogue Generation with Context-Specific Commonsense Knowledge Awareness. In *Proceedings of ACL 2020*.
- Wu, S.; Wang, M.; Zhang, D.; Zhou, Y.; Li, Y.; and Wu, Z. 2021. Knowledge-Aware Dialogue Generation via Hierarchical Infobox Accessing and Infobox-Dialogue Interaction Graph Network. In *Proceedings of IJCAI 2021*.
- Wu, W.; Guo, Z.; Zhou, X.; Wu, H.; Zhang, X.; Lian, R.; and Wang, H. 2019. Proactive Human-Machine Conversation with Explicit Conversation Goal. In *Proceedings of ACL 2019*.
- Zhang, J.; Wang, J.; Wang, X.; Li, Z.; and Xiao, Y. 2023. AspectMMKG: A Multi-modal Knowledge Graph with Aspect-aware Entities. In *Proceedings of CIKM 2023*.
- Zhang, T.; Li, Z.; Wang, J.; Qu, J.; Yuan, L.; Liu, A.; Zhao, L.; and Chen, Z. 2022. Aligning internal regularity and external influence of multi-granularity for temporal knowledge graph embedding. In *Proceedings of DASFAA 2022*, 149–164. Springer.
- Zheng, D.; Xu, Z.; Meng, F.; Wang, X.; Wang, J.; and Zhou, J. 2021. Enhancing Visual Dialog Questioner with Entity-based Strategy Learning and Augmented Guesser. In *Findings of EMNLP 2021*, 1839–1851.
- Zhou, H.; Young, T.; Huang, M.; Zhao, H.; Xu, J.; and Zhu, X. 2018. Commonsense Knowledge Aware Conversation Generation with Graph Attention. In *Proceedings of IJCAI 2018*.
- Zhou, H.; Zheng, C.; Huang, K.; Huang, M.; and Zhu, X. 2020. KdConv: A Chinese Multi-domain Dialogue Dataset Towards Multi-turn Knowledge-driven Conversation. In *Proceedings of ACL 2020*.
- Zhu, H.; Song, Y.; and Liu, B. 2023. KTGAT: Improving the Robustness of Knowledge-enhanced Text Generation via Adversarial Training. *Proceedings of ICCEA 2023*.