

Difference-aware Knowledge Selection for Knowledge-grounded Conversation Generation

Chujie Zheng[†], Yunbo Cao[‡], Daxin Jiang[§], Minlie Huang^{†*}

[†]Department of Computer Science and Technology, [†]Institute for Artificial Intelligence,

[†]State Key Lab of Intelligent Technology and Systems,

[†]Beijing National Research Center for Information Science and Technology,

[†]Tsinghua University, Beijing 100084, China

[‡]Smart Platform Product Department, Tencent, Beijing

[§]STCA NLP Group, Microsoft, Beijing

chujiezhengchn@gmail.com, yunbocao@tencent.com

djiang@microsoft.com, aihuang@tsinghua.edu.cn

Abstract

In a multi-turn knowledge-grounded dialog, the difference between the knowledge selected at different turns usually provides potential clues to knowledge selection, which has been largely neglected in previous research. In this paper, we propose a difference-aware knowledge selection method. It first computes the difference between the candidate knowledge sentences provided at the current turn and those chosen in the previous turns. Then, the differential information is fused with or disentangled from the contextual information to facilitate final knowledge selection. Automatic, human observational, and interactive evaluation shows that our method is able to select knowledge more accurately and generate more informative responses, significantly outperforming the state-of-the-art baselines. The codes are available at <https://github.com/chujiezheng/DiffKS>.

1 Introduction

Knowledge-grounded conversation generation aims at generating informative responses based on both discourse context and external knowledge (Ghazvininejad et al., 2018; Zhou et al., 2018a), where selecting appropriate knowledge is critical to the success of the task. Existing knowledge selection models generally fall into two types. One type is solely based on the context (Lian et al., 2019; Zhang et al., 2019; Meng et al., 2020; Ren et al., 2020), which we call **non-sequential selection** because knowledge selection at different turns is independent. The other type sequentially selects knowledge additionally conditioned on previously selected knowledge (Kim et al., 2020), which we

* Corresponding author.

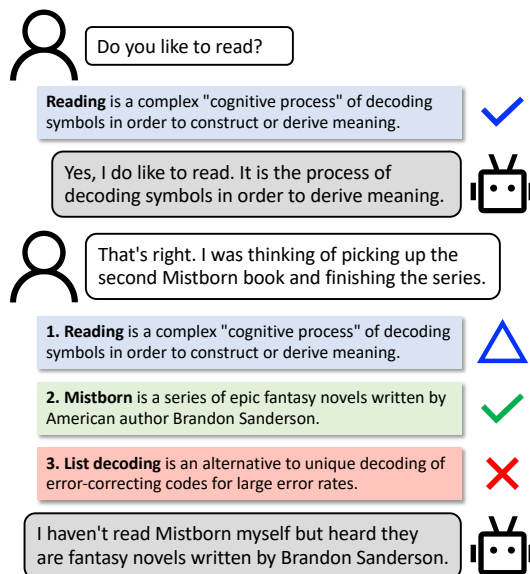


Figure 1: An example of difference-aware knowledge selection. The **blue** \triangle denotes that the corresponding knowledge has little difference from or is identical to the previously selected one, and selecting it may lead to *repetitive* responses. The **red** \times denotes that the difference is too large, and selecting it could make the response *incoherent* with the context.

call **sequential selection**. As shown in Kim et al. (2020), such a sequential way can better simulate a multi-turn dialog and facilitate knowledge selection in later turns.

However, the **difference** between selected knowledge at different turns has been largely neglected in prior studies, while it usually provides potential clues to knowledge selection. Figure 1 illustrates an example, where the dialog system selects one from candidate knowledge sentences (all relevant to the context) at the 2nd turn. Selecting the knowledge that has little difference from or even is identical to the previously selected one (like the

1st knowledge) may lead to generating *repetitive* responses, while too large difference (like the 3rd knowledge) would make the response *incoherent* with the context. As a result, the dialog system should avoid the knowledge which differs from the previously selected ones either too little or too largely, and instead select an appropriate knowledge sentence (the 2nd one) which can make the conversation flow smoothly and naturally.

We thus propose DiffKS, a novel **D**ifference-aware **K**nowledge **S**election method for knowledge-grounded conversation generation. It first computes the difference between the candidate knowledge sentences provided at the current turn and the previously selected knowledge. Then, in the two models we devise, the differential information is fused with or disentangled from the contextual information to facilitate final knowledge selection. Automatic and human evaluation on two widely-used benchmarks shows that our method is significantly superior over the state-of-the-art baselines and it can select knowledge more accurately and generate more informative responses.

Our contributions are summarized as follows:

- We propose to explicitly model and utilize the differential information between selected knowledge in multi-turn knowledge-grounded conversation for knowledge selection. We further devise two variants where the differential information is fused with or disentangled from the context information during knowledge selection.
- Automatic, human observational, and human interactive evaluations show that our method significantly outperforms strong baselines in terms of knowledge selection and can generate more informative responses.

2 Related Work

2.1 Knowledge-grounded Dialog Generation

Recently, a variety of neural models have been proposed to facilitate knowledge-grounded conversation generation (Zhu et al., 2017; Young et al., 2018; Zhou et al., 2018a; Liu et al., 2018). The research topic is also greatly advanced by many corpora (Zhou et al., 2018b; Moghe et al., 2018; Dinan et al., 2019; Gopalakrishnan et al., 2019; Moon et al., 2019; Tuan et al., 2019; Wu et al., 2019; Zhou et al., 2020). As surveyed in Huang et al. (2020),

existing studies have been mainly devoted to addressing two research problems: (1) **knowledge selection**: selecting appropriate knowledge given the dialog context and previously selected knowledge (Lian et al., 2019; Zhang et al., 2019; Meng et al., 2020; Ren et al., 2020; Kim et al., 2020); and (2) **knowledge-aware generation**: injecting the required knowledge to generate meaningful and informative responses (Ghazvininejad et al., 2018; Zhou et al., 2018a; Li et al., 2019; Qin et al., 2019; Yavuz et al., 2019; Zhao et al., 2020). Since selecting the appropriate knowledge is a precursor to the success of knowledge grounded dialog systems, we focus on the **knowledge selection** problem in this paper.

2.2 Non-sequential Knowledge Selection

The non-sequential selection models capture the relationship between the current context and background knowledge (Lian et al., 2019; Zhang et al., 2019; Meng et al., 2020; Ren et al., 2020). For instance, PostKS (Lian et al., 2019) estimates a posterior distribution over candidate knowledge sentences, which is based on both the context and the golden response, and only uses the context to estimate a prior distribution as an approximation of the posterior during inference.

Besides, Zhang et al. (2019); Meng et al. (2020); Ren et al. (2020) also belong to non-sequential selection models. Different from our work and Lian et al. (2019); Kim et al. (2020) that select knowledge from *candidate knowledge sentences*, their methods are devised for selecting important text spans or fragments from the *background knowledge document* that will be used in generation. Therefore these works have a different task setting from ours.

2.3 Sequential Knowledge Selection

The sequential selection models additionally make use of previously selected knowledge to facilitate knowledge selection (Kim et al., 2020). For instance, Kim et al. (2020) propose a Sequential Latent Knowledge Selection (SLKS) model. It keeps track of the hidden states of dialog history and previously selected knowledge sentences. Our method is parallel to SLKS because we also utilize the previously selected knowledge. However, we explicitly compute the difference between knowledge selected at different turns, while SLKS only encodes the already selected knowledge in an implicit way.

In addition, recently there emerge a number of works that propose RL-based models to select a path in *structured knowledge graph (KG)* (Xu et al., 2020a,b), which also select knowledge in a sequential way. While our method is designed to ground the conversation to *unstructured knowledge text*, we will leave as future work the application of our method to such KG-grounded dialog generation tasks (Wu et al., 2019; Moon et al., 2019; Zhou et al., 2020).

3 Methodology

3.1 Task Formulation

In a multi-turn dialogue, given a post and a sequence of knowledge sentences at each turn, our goal is to select appropriate knowledge and generate a proper response to the current context.

Formally, the post at the τ -th turn is a sequence of tokens $x^\tau = (x_1^\tau, \dots, x_{|x^\tau|}^\tau)$, and the response to be generated is $y^\tau = (y_1^\tau, \dots, y_{|y^\tau|}^\tau)$. The background knowledge $k^\tau = (k_1^\tau, \dots, k_{|k^\tau|}^\tau)$ contains a sequence of knowledge sentences provided at the τ -th turn. For each i , $k_i^\tau = (k_{i,1}^\tau, \dots, k_{i,|k_i^\tau|}^\tau)$ is a sequence of tokens in the i -th sentence.

Note that under the setting of multi-turn dialogue, we use $c^\tau \triangleq [x^{\tau-1}; y^{\tau-1}; x^\tau]$ as the given context at the τ -th turn, where $[\cdot; \cdot]$ denotes concatenation. In Section 3.2 and 3.4, we will omit the superscript τ for simplicity.

3.2 Encoders

The context is encoded with a bidirectional GRU (Cho et al., 2014):

$$(h_{c,1}, \dots, h_{c,|c|}) = \mathbf{BiGRU}_c(c), \quad (1)$$

where $h_{c,i} = [\vec{h}_{c,i}; \overleftarrow{h}_{c,i}]$. We use $h_c \triangleq [\vec{h}_{c,|c|}; \overleftarrow{h}_{c,1}]$ as the context representation. Similarly, the knowledge sentences are encoded with another BiGRU:

$$(h_{k,i,1}, \dots, h_{k,i,|k_i|}) = \mathbf{BiGRU}_k(k_i). \quad (2)$$

We use $h_{k,i} \triangleq [\vec{h}_{k,i,|k_i|}; \overleftarrow{h}_{k,i,1}]$ as the representation of k_i . Specifically, we add an empty sentence k_0 that indicates *no knowledge being used*.

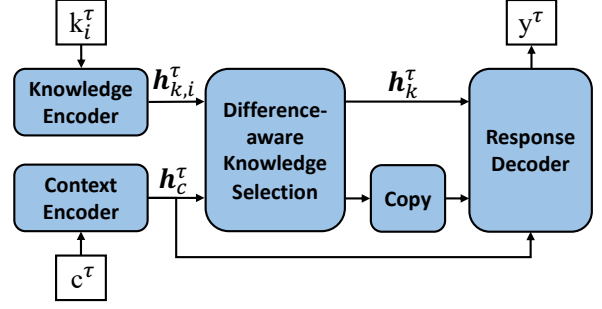


Figure 2: An overview of model structure.

3.3 Difference-aware Knowledge Selection

In order to select proper knowledge, our model gets aware of the difference between the current candidate knowledge sentences and the previously selected knowledge.

To make full use of the contextual dependency and relevance between the knowledge sentences¹, our model first compares candidate knowledge sentences to explore their correlations, where the comparison is conducted using BiGRU:

$$(r_0^\tau, \dots, r_{|k^\tau|}^\tau) = \mathbf{BiGRU}(h_{k,0}^\tau, \dots, h_{k,|k^\tau|}^\tau), \quad (3)$$

Then, the model computes the difference of each knowledge sentence r_i^τ from the knowledge selected in the previous M turns $\{h_k^{\tau-m}\}_{m=1}^M$:

$$o_i^\tau = \sum_{m=1}^M \lambda_m \mathbf{Diff}(h_k^{\tau-m}, r_i^\tau), \quad (4)$$

$$\sum_{m=1}^M \lambda_m = 1, \quad \forall m, \lambda_m \geq 0 \quad (5)$$

Inspired by Wang et al. (2018), we define the difference as follow:

$$\mathbf{Diff}(x, y) \triangleq \mathbf{F}([x - y; x \odot y]), \quad (6)$$

where \mathbf{F} is a fully connected layer activated with tanh. Note that at the first turn, we set o_i^1 to a zero vector because there is no differential information to be obtained.

For that intuitively the knowledge selected in the previous turn has the largest impact and most clues for the current selection, we studied the simplest case where $M = 1$, saying $o_i^\tau =$

¹For example, the knowledge sentences may be extracted from a document in order, or about the same topic like in Wizard of Wikipedia (Dinan et al., 2019).

$\text{Diff}(\mathbf{h}_k^{\tau-1}, \mathbf{r}_i^\tau)$, in the main experiments for simplicity.

Next, we introduce two variants where the differential information $\{\mathbf{o}_i^\tau\}_{i=0}^{|\mathbf{k}^\tau|}$ is fused with or disentangled from the contextual information during knowledge selection.

3.3.1 Fused Selection




Figure 3: Fused Selection module. The contextual information and the differential information are *fused* together to calculate the final knowledge selection distribution.

The Fused Selection module is shown in Figure 3. Directly taking \mathbf{o}_i^τ as an extra feature of \mathbf{k}_i^τ , it uses the context \mathbf{h}_c^τ to query the difference-enhanced knowledge sentences:

$$\beta_i^\tau = \mathbf{v}^\top \tanh(\mathbf{W}_{\text{que}} \mathbf{h}_c^\tau + \mathbf{W}_{\text{key}} [\mathbf{h}_{k,i}^\tau; \mathbf{o}_i^\tau]), \quad (7)$$

where \mathbf{v} , \mathbf{W}_{que} and \mathbf{W}_{key} are trainable parameters.

However, it is difficult to distinguish the respective contributions of contextual and differential information to knowledge selection in the above fused variant. We thus devise the disentangled variant as following, where the roles of two types of information are separated, which makes it feasible to conduct ablation study.

3.3.2 Disentangled Selection

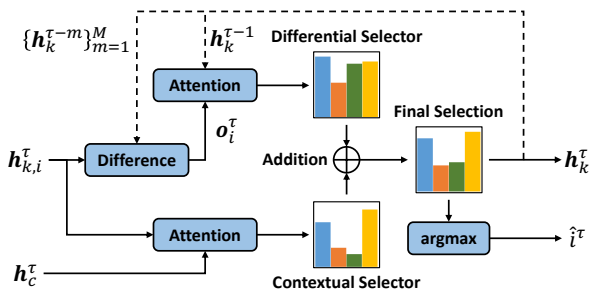


Figure 4: Disentangled Selection module. The contextual information and the differential information are *disentangled* to calculate two separate knowledge selection distributions in two independent selectors.

Figure 4 gives an overview of the Disentangled Selection module. It has two independent selectors. The **Contextual Selector** simply looks for

the knowledge sentence that has high relevance to the context, just like most existing knowledge selection models do. It only takes advantage of the context \mathbf{h}_c^τ to match each knowledge sentence itself $\mathbf{h}_{k,i}^\tau$, obtaining a context-aware selection distribution:

$$\beta_{\text{Ctx},i}^\tau = (\mathbf{h}_c^\tau)^\top \mathbf{h}_{k,i}^\tau. \quad (8)$$

In contrast, the **Differential Selector** focuses on predicting the next knowledge to be selected conditioned on the previously selected knowledge and differential information, which reveals the process of knowledge transition. Without the access to the contextual information, the Differential Selector views the previously selected knowledge $\mathbf{h}_k^{\tau-1}$ as query, and the knowledge sentence \mathbf{r}_i^τ with its differential information \mathbf{o}_i^τ as key, to estimate a difference-aware selection distribution:

$$\beta_{\text{Diff},i}^\tau = \mathbf{v}^\top \tanh(\mathbf{W}_{\text{que}} \mathbf{h}_k^{\tau-1} + \mathbf{W}_{\text{key}} [\mathbf{r}_i^\tau; \mathbf{o}_i^\tau]), \quad (9)$$

where \mathbf{v} , \mathbf{W}_{que} and \mathbf{W}_{key} are trainable parameters.

The final selection distribution is the summation of the distributions of two selectors:

$$\beta_i^\tau = \beta_{\text{Ctx},i}^\tau + \beta_{\text{Diff},i}^\tau. \quad (10)$$

Note that the Differential Selector relies on the previously selected knowledge, thus at the first turn, we set $\beta_{\text{Diff},i}^\tau$ to 0 for each i .

3.3.3 Selecting Knowledge

Finally, either adopting the Fused or Disentangled Selection module, the model selects the knowledge sentence with the highest attention score, and uses its representation for further generation²:

$$\alpha_i^\tau = \text{softmax}_i(\beta_i^\tau), \quad (11)$$

$$\hat{i}^\tau = \arg \max_i \alpha_i^\tau, \quad \mathbf{h}_k^\tau \triangleq \mathbf{h}_{k,\hat{i}^\tau}^\tau. \quad (12)$$

3.4 Decoder

The decoding state is updated by a GRU:

$$\mathbf{s}_t = \text{GRU}_D(\mathbf{s}_{t-1}, [e(y_{t-1}); \mathbf{h}_k]), \quad (13)$$

$$\mathbf{s}_0 = \mathbf{W}_D [\mathbf{h}_c; \mathbf{h}_k] + \mathbf{b}_D, \quad (14)$$

where \mathbf{W}_D and \mathbf{b}_D are trainable parameters, and $e(y_{t-1})$ denotes the embedding of the word y_{t-1} generated in the last time step. Then, the decoder

²The model is trained with teacher forcing, where the golden selected knowledge \mathbf{h}_{k,i^*}^τ is used during training.

outputs the generation probability over the vocabulary (without normalization):

$$\phi^G(y_t = w) = \mathbf{w}^T (\mathbf{W}_G \mathbf{s}_t + \mathbf{b}_G), \quad (15)$$

where \mathbf{W}_G and \mathbf{b}_G are trainable parameters, and \mathbf{w} is the one-hot vector of the word w . Meanwhile, a copy mechanism (Gu et al., 2016) is adopted to output additional copy probability of the words in the selected knowledge sentence $k_{\hat{i}}$ (without normalization):

$$\phi^C(y_t = w) = \sum_{j: k_{\hat{i},j}=w} (\mathbf{s}_t)^T \mathbf{H}(\mathbf{h}_{k_{\hat{i},j}}), \quad (16)$$

where \mathbf{H} is a fully connected layer activated with tanh. The final probability distribution is computed as follows:

$$\mathcal{P}(y_t = w) = \frac{1}{Z} \left(e^{\phi^G(y_t=w)} + e^{\phi^C(y_t=w)} \right), \quad (17)$$

where Z is the normalization term. Then we select the word from vocabulary with the highest probability, saying: $y_t = \arg \max_w \mathcal{P}(y_t = w)$.

3.5 Loss

The negative log likelihood loss is adopted:

$$\mathcal{L}_{\text{NLL}} = - \sum_{\tau=1}^T \sum_{t=1}^{|y^\tau|} \log \mathcal{P}(y_t^{\tau*}), \quad (18)$$

where $y_t^{\tau*}$ denotes the t -th word in the golden response at the τ -th turn and T is the length of turns in the whole dialogue. We also add supervision on the final knowledge selection distribution:

$$\mathcal{L}_{\text{KS}} = - \sum_{\tau=1}^T \log \alpha_{i^{\tau*}}^{\tau}, \quad (19)$$

where $i^{\tau*}$ denotes the index of the golden selected knowledge sentence at the τ -th turn. The total loss is their summation:

$$\mathcal{L} = \mathcal{L}_{\text{NLL}} + \lambda \mathcal{L}_{\text{KS}}. \quad (20)$$

where we set $\lambda = 1$ in our experiments.

4 Experiments

4.1 Datasets

We evaluated our method on two widely used benchmarks: Wizard of Wikipedia (WoW) (Dinan et al., 2019), and Holl-E (Moghe et al., 2018).

WoW (Dinan et al., 2019) contains multi-turn knowledge-grounded conversations, collected by wizard-apprentice mode. Each utterance of the wizard is grounded to a selected knowledge sentence, or indicated by that no knowledge is used. The dialogues are split into 18,430/1,948/965/968 for Train/Dev/Test Seen/Test Unseen respectively, with 4 turns per dialogue and 61 provided knowledge sentences per turn on average. Note that the test data is split into *Test Seen* (in-domain) and *Test Unseen* (out-of-domain), where Test Unseen contains topics that are never seen in Train or Dev.

Holl-E (Moghe et al., 2018) contains conversations in which one speaker is strictly instructed to give utterances by copying or modifying sentences from the given background document. Similarly, each utterance is annotated regarding the selected knowledge. Following Kim et al. (2020), we tokenized the background document into sentences, and meanwhile ensured that the annotated span is included in a whole sentence. The dialogues are split into 7,211/930/913 for Train/Dev/Test respectively, with 5 turns per dialogue and 60 provided knowledge sentences per turn on average.

4.2 Baselines

We compared our models with the following typical knowledge selection baselines:

MemNet (Ghazvininejad et al., 2018) stores knowledge sentences in its memory units, which are attentively read during decoding. We also evaluated a variant (**MemNet+**) where knowledge selection is supervised by the same \mathcal{L}_{KS} as our models do.

PostKS (Lian et al., 2019) estimates two knowledge selection distributions, where the prior distribution is based on only the context and the posterior one on both the context and the golden response, and their KL divergence is minimized during training. The knowledge selection of PostKS is supervised by a BOW loss. We also evaluated two variants, where one uses \mathcal{L}_{KS} instead of the BOW loss to supervise knowledge selection (**PostKS+**), and the other is further equipped with copy mechanism (**PostKS++**).

SLKS (Kim et al., 2020) improves PostKS by using two separate GRUs to update the states of dialog history and previously selected knowledge sentences respectively. For fair comparison, we replaced the pretrained BERT (Devlin et al., 2019) encoder and the Transformer (Vaswani et al., 2017) decoder in SLKS with BiGRU and GRU respec-

Models	ACC	BLEU-2/4	ROUGE-2	
WoW Seen				
MemNet	13.2**	6.6**	1.8**	3.2**
+ \mathcal{L}_{KS}	18.4**	7.2**	1.9**	3.3**
PostKS	13.8**	6.9**	1.8**	3.2**
+ \mathcal{L}_{KS}	22.5**	7.5**	2.3**	3.7**
+Copy	21.9**	9.9**	4.5**	5.6**
SLKS	23.4**	11.3	5.5	6.5
DiffKS _{Fus}	25.5	11.6	5.7	6.8
DiffKS _{Dis}	24.7	11.3	5.7	6.8
WoW Unseen				
MemNet	12.8**	5.7**	1.2**	2.3**
+ \mathcal{L}_{KS}	15.9**	5.9**	1.3**	2.3**
PostKS	13.6**	5.5**	1.2**	2.1**
+ \mathcal{L}_{KS}	15.8**	6.6**	1.5**	2.6**
+Copy	14.9**	7.9**	3.2**	3.9**
SLKS	14.7**	8.7**	3.7**	4.6**
DiffKS _{Fus}	19.7	10.0	4.7	5.6
DiffKS _{Dis}	18.3*	9.6	4.5	5.3
Holl-E				
MemNet	5.1**	8.0**	4.5**	8.9**
+ \mathcal{L}_{KS}	25.1**	7.7**	4.3**	9.0**
PostKS	6.1**	6.9**	3.9**	8.6**
+ \mathcal{L}_{KS}	29.5**	15.9**	8.2**	13.1**
+Copy	28.0**	26.5**	22.4**	23.1**
SLKS	28.6**	28.5**	24.5**	24.9*
DiffKS _{Fus}	33.0	29.5	25.5	25.9
DiffKS _{Dis}	33.5	29.9	25.9	26.4

Table 1: Automatic evaluation results. The best results are in **bold**. Significance tests were conducted between the best results and other competitors, with sign test for ACC, bootstrap resampling (Koehn, 2004) for BLEU, and Student’s t-test for ROUGE. */** indicate p -value < 0.05/0.005 respectively.

tively, and adopted the same copy mechanism in SLKS as in our models.

4.3 Implementation Details

All the models were implemented with PyTorch (Paszke et al., 2017). The sentences were tokenized with NLTK (Bird and Loper, 2004). We set the vocabulary size to 20K for WoW and 16K for Holl-E and used the 300-dimensional word embeddings initialized by GloVe (Pennington et al., 2014) or from a standard normal distribution $\mathcal{N}(0, 1)$. We applied a dropout rate of 0.5 on word embeddings. The hidden sizes were set to 200 for the encoders (totally 400 for two directions) and to 400 for the decoder. We adopted the ADAM (Kingma and Ba, 2015) optimizer with the initial learning rate set to 0.0005. The batch size was set to 8 dialogues. All the models share the same hyperparameter setting

and were trained for 20 epochs on one NVIDIA Titan Xp GPU. The checkpoints of our reported results were selected according to BLEU-4 on the Dev sets.

4.4 Automatic Evaluation

We used several automatic metrics: *ACC*, the accuracy of knowledge selection on the whole test set, corpus-level *BLEU-2/4* (Papineni et al., 2002), and *ROUGE-2* (Lin, 2004).

As shown in Table 1³, our method outperforms significantly all the baselines in all the metrics on three test sets (except BLEU and ROUGE on WoW Seen compared with SLKS), which indicates its superiority in selecting proper knowledge and generating informative responses. Compared to the baseline models, our models also demonstrate a stronger ability of generalization from in-domain (WoW Seen) to out-of-domain data (WoW Unseen). It is worth noting that on WoW Unseen, our DiffKS_{Fus} obtains a higher accuracy (19.7) of knowledge selection even than the BERT-enhanced SLKS in their original paper (18.3). We also observed that DiffKS_{Fus} performs a bit better on WoW while DiffKS_{Dis} on Holl-E. We conjecture that it is because in Holl-E, the golden selected knowledge among different turns usually has high contextual dependency (for example, they may be continuous sentences in the document), which makes it feasible to predict the next selected knowledge simply conditioned on the differential information.

4.5 Human Observational Evaluation

We conducted human observational evaluation with pair-wise comparison, where our two models were compared with PostKS++ and SLKS. 100 dialogues were respectively sampled from WoW Seen/Unseen. For each pair of dialogues generated from two models (suppose with T turns), annotators from Amazon Mechanical Turk were hired to give preferences (win, lose, or tie) for each response pair of all the T turns in terms of different metrics. Each pair-wise comparison of dialogues was judged by 3 curators. We adopted the following two metrics: *Naturalness* evaluates the fluency and readability of a response. *Appropriateness* evaluates the relevance to the context and whether

³We found in Kim et al. (2020) that BERT usually gives rise to a gain of 2-5 points in ACC, thus our results without using BERT as encoder are within a reasonable range comparing with those in the original reference.

A vs. B	Naturalness			Appropriateness		
	Win	Lose	κ	Win	Lose	κ
WoW Seen						
Fus / PostKS++	50.3*	42.5	.47	49.2*	43.1	.40
Fus / SLKS	44.5	43.3	.50	44.0*	38.7	.48
Dis / PostKS++	50.6*	44.9	.42	50.5*	44.4	.38
Dis / SLKS	42.7	43.8	.41	46.4	41.4	.47
Fus / Dis	43.2	42.8	.49	39.3	40.9	.57
WoW Unseen						
Fus / PostKS++	48.8*	43.2	.57	49.3**	40.5	.60
Fus / SLKS	47.9*	41.8	.44	47.3*	40.9	.47
Dis / PostKS++	52.0**	36.4	.46	46.8*	39.9	.49
Dis / SLKS	46.5*	39.7	.45	47.8*	42.3	.47
Fus / Dis	39.8	42.4	.52	41.5	37.8	.53

Table 2: Human observational evaluation results. Ties are not shown. Significance tests were conducted with sign test. κ denotes the Fleiss’ Kappa which measures annotation agreement.

Models	WoW Seen	WoW Unseen
Human[†]	4.13 (1.08)	4.34 (0.98)
PostKS++	2.30 (1.06)	2.13 (1.10)
SLKS	2.32 (1.11)	2.22 (1.15)
DiffKS_{Fus}	2.43 (0.96)	2.39 (1.16)
DiffKS_{Dis}	2.39 (1.17)	2.38 (1.19)

Table 3: Human interactive evaluation results. The standard deviation is marked in parentheses. The results of human[†] are from Dinan et al. (2019); Kim et al. (2020).

a response contains appropriate knowledge information to the context.

Results are shown in Table 2, where the Fleiss’ Kappa (Fleiss, 1971) values show almost moderate agreements ($0.4 < \kappa < 0.6$). Our models significantly outperform PostKS++ in both metrics, and also generally outperform SLKS in terms of Appropriateness. Again, the advantage of our models on WoW Unseen is more evident than on WoW Seen.

4.6 Human Interactive Evaluation

We further conducted human interactive evaluation where real humans converse with one model about a specific topic. We compared PostKS++ and SLKS with our two models. The workers from Amazon Mechanical Turk were asked to first select one topic from 2-3 provided candidate topics, and then converse with one of the models for 3-5 dialogue turns. After conversation, they were required to rate the dialog model with a 5-star scale in terms of the fluency and informativeness of the utterances and the coherence of the whole dialog. Following Dinan et al. (2019); Kim et al. (2020), the interactive evaluation was implemented with

Models	ACC	BLEU-2/4	ROUGE-2
WoW Seen			
DiffKS_{Dis}	24.7	11.3	5.7
w/o DiffSel	<u>22.3**</u>	<u>10.6**</u>	<u>4.9**</u>
w/o CtxSel	24.6	10.9	5.3*
WoW Unseen			
DiffKS_{Dis}	18.3	9.6	4.5
w/o DiffSel	<u>15.5**</u>	<u>8.8**</u>	<u>3.8**</u>
w/o CtxSel	18.4	9.1*	4.1*
Holl-E			
DiffKS_{Dis}	33.5	29.9	25.9
w/o DiffSel	<u>29.1**</u>	<u>27.9**</u>	<u>23.8**</u>
w/o CtxSel	31.6**	28.4**	24.7**

Table 4: Ablation tests. The larger performance drops between the two ablation models are underlined. The significance tests are conducted between the ablation models and the complete model DiffKS_{Dis}.

ParLAI (Miller et al., 2017). For each model, we averaged the scores from 150 collected conversations on each test set of WoW. We also reported the results of human-human dialog from Dinan et al. (2019); Kim et al. (2020), where each worker converses with another human and the latter has access to knowledge sentences just like the models do.

Results are shown in Table 3⁴, where DiffKS_{Fus} gains the highest scores and our models both outperform the other two state-of-the-art baselines, indicating that our models are favorably preferred by human annotators.

4.7 Ablation Test

In order to verify the effectiveness of the differential information in knowledge selection, we conducted ablation tests, which were specifically based on the disentangled variant DiffKS_{Dis}. In DiffKS_{Dis}, we removed either the Differential Selector (DiffSel) or the Contextual Selector (CtxSel), and trained the model with only one of the two selectors.

Results are shown in Table 4. Without the differential selector, the model performance is remarkably impaired in all the metrics on three test sets, indicating the importance of utilizing differential information. In comparison, removing the contextual selector is less influential (with less performance drop). We conjecture that this may result from the characteristics of datasets. For instance,

⁴We found in Dinan et al. (2019); Kim et al. (2020) that the stddev values of dialog models are usually between 1.0 and 1.4, thus our results are within a reasonable range.

Models	M	ACC	BLEU-2/4	ROUGE-2	
WoW Seen					
DiffKS_{Fus}	1	25.5	11.6	5.7	6.8
	2	26.3	11.7	5.8	7.0
	3	26.1	11.6	5.7	6.9
DiffKS_{Dis}	1	24.7	11.3	5.7	6.8
	2	26.1	11.7	6.0	7.1
	3	25.0	11.1	5.7	6.7
WoW Uneen					
DiffKS_{Fus}	1	19.7	10.0	4.7	5.6
	2	20.4	10.6	5.2	6.0
	3	19.5	9.8	4.8	5.6
DiffKS_{Dis}	1	18.3	9.6	4.5	5.3
	2	19.4	9.9	4.6	5.3
	3	19.1	9.9	4.5	5.2
Holl-E					
DiffKS_{Fus}	1	33.0	29.5	25.5	25.9
	2	33.2	30.1	26.1	26.2
	3	33.1	30.0	26.3	26.3
DiffKS_{Dis}	1	33.5	29.9	25.9	26.4
	2	33.9	31.2	27.2	26.9
	3	33.8	31.3	26.8	26.7

Table 5: Comparison between results with different M .

in WoW, the apprentice (without access to knowledge) usually reacts passively to the wizard (having access to knowledge). Thus the apprentice posts (contextual information) have limited influence in driving the conversation, which is instead affected or controlled by the wizard. In this case, our differential information that can predict the process of knowledge transition has more influence than the contextual information. In addition, same as Kim et al. (2020), the knowledge sentences in Holl-E are obtained by segmenting a long document into single sentences, which implies that there exists the relevance or contextual dependency between knowledge sentences. Consequently, the differential information is still able to provide considerable clues for knowledge selection even without access to the new user post (the context).

Furthermore, after removing DiffSel, DiffKS_{Dis} reduces to a vanilla knowledge selection model where the supervision \mathcal{L}_{KS} was directly applied on the ‘prior’ selection distribution. Nevertheless, the performance of the ablated model is sometimes competitive to the baselines (for instance, in terms of ACC, DiffKS_{Dis} w/o DiffSel obtains 22.3/15.5/29.1 vs. 21.9/14.9/28.0 of PostKS++). It may result from the gap between training and inference caused by the prior-posterior framework

Models	1 st	2 nd	3 rd	4 th	5 th
WoW Seen					
PostKS++	56.8	15.6	9.6	6.2	4.1
SLKS	57.4	18.4	10.1	8.9	5.4
DiffKS_{Fus}	57.4	22.5	12.8	9.8	7.4
DiffKS_{Dis}	56.6	21.5	11.2	10.2	7.9
WoW Uneen					
PostKS++	42.8	8.5	4.1	4.8	4.6
SLKS	43.0	6.1	5.2	4.9	5.0
DiffKS_{Fus}	40.9	21.2	10.5	7.7	4.6
DiffKS_{Dis}	40.2	16.1	10.3	7.7	6.1
Holl-E					
PostKS++	62.8	17.9	18.8	20.0	23.2
SLKS	65.2	18.4	19.2	21.3	19.6
DiffKS_{Fus}	65.8	22.3	22.1	25.5	25.8
DiffKS_{Dis}	63.9	23.0	23.4	26.0	28.3

Table 6: Knowledge selection accuracy over turns.

adopted in PostKS and SLKS, which may be not superior over directly training the prior selection distribution⁵.

5 Discussion

5.1 Difference From More Turns

To investigate the impact of increasing the turns of differential information (the M in Equ.4), we additionally experimented with $M = 2, 3$, and took the arithmetic average for simplicity in Equ.4, saying $\forall i, \lambda_i = 1/M$.

Results are shown in Table 5. We can find that $M = 2$ generally achieves the best performance compared with $M = 1, 3$ for both DiffKS_{Fus} and DiffKS_{Dis} (while $M = 3$ is still better than $M = 1$). It further turns out the effectiveness of explicitly modeling differential information. We also conjecture that the model performance would be further improved by assigning the nearest/farthest difference with the largest/smallest weight in Equ.4, saying $\lambda_1 > \lambda_2 > \dots > \lambda_M$, which is more reasonable than the simplified arithmetic average.

5.2 Accuracy Over Turns

To verify whether the sequential knowledge selection facilitates knowledge selection in later turns, we evaluated the accuracy of knowledge selection at different turns. The statistics are shown in Table 6. Our two models have the highest accuracy

⁵The prior-posterior framework is first proposed by PostKS without direct supervision \mathcal{L}_{KS} on knowledge selection. While in this paper and Kim et al. (2020) the supervision \mathcal{L}_{KS} is available, the prior-posterior framework may not be superior any more.

from the 2nd to 5th turns and outperform SLKS and PostKS++ (and SLKS also generally outperforms PostKS++). The results show that our models can select more accurate knowledge consistently over different turns.

5.3 Case Study

Topic: Georgia (U.S. state)	
k1:	Georgia is a state in the southeastern United States.
k2:	Georgia is known as the ‘Peach State’ and the ‘Empire State of the South’.
k3:	It began as a British colony in 1733, the last of the original Thirteen Colonies.
k4:	Named after King George II of Great Britain, the Province of Georgia covered the area from South Carolina down to Spanish Florida...
k5:	It was the last state to be restored to the Union, on July 15, 1870.
k6:	The history of Australia refers to the history of the area and people of the Commonwealth of Australia...
<i>Post 1: Hey... have you ever been to Georgia before?</i>	
PostKS++	(k1) Georgia is a state in the southeastern United States.
SLKS	(k1) Georgia is a state in the southeastern United States.
DiffKS_{Fus}	(k1) I have not, but I know it is a state in the southeastern United States.
DiffKS_{Dis}	(k1) I have never been to Georgia. It is a state in the southeastern United States.
<i>Post 2: I haven't been either, but I'd love to go. Peaches are delicious, especially from Georgia.</i>	
PostKS++	(k1) Yes, I know that Georgia is a state in the southeastern United States. △
SLKS	(k2) Georgia is known as the ‘Peach State’ and the ‘Empire State of the South’.
DiffKS_{Fus}	(k3) It began as a British colony in the last of the original thirteen colonies.
DiffKS_{Dis}	(k3) I've heard of the British colony in the US. It began as a British colony in the US.
<i>Post 3: Oh I didn't know that either! So it probably has a rich history too.</i>	
PostKS++	(k6) It refers to the history of the area and people of the Commonwealth of Australia and its preceding Indigenous and colonial societies. ×
SLKS	(k2) I don't know but I know that Georgia is known as the ‘Peach State’ and the ‘Empire State of the South’. △
DiffKS_{Fus}	(k4) It is named after King George II of Great Britain, the Province of Georgia covered the area from South Carolina.
DiffKS_{Dis}	(k5) It was the last state to be restored to the Union, on July 15, 1870.

Figure 5: Case study. We marked the selected knowledge sentence in parentheses before each response. The knowledge k1-k5 are about the topic *Georgia (U.S. state)*, while k6 is about *History of Australia*. The blue △ denotes duplicate responses resulting from repetitive knowledge selection. The red × denotes incoherent responses resulting from selecting a far different knowledge from previous turns.

We show a case from WoW Seen in Figure 5, which compares the responses generated by PostKS++, SLKS and our two models.

At the 2nd turn, PostKS++ generates almost the same responses as at the 1st turn due to the repetitive knowledge selection. Similar cases occur for SLKS at the 2nd and the 3rd turns. Moreover, PostKS++ selects a quite different knowledge sen-

tence at the 3rd turn from those at previous turns, which is about the topic *History of Australia* but not *Georgia (U.S. state)*. As a result, PostKS++ generates a response which is not coherent to the previous context at the 3rd turn. In contrast, our two models select both diverse and appropriate knowledge sentences at all the turns, thereby generating informative responses and making the dialog coherent and natural.

6 Conclusion

We present a novel difference-aware knowledge selection method for multi-turn knowledge-grounded conversation generation. Our method first compares the candidate knowledge provided at the current turn with the previously selected knowledge, and then selects appropriate knowledge to be used in generation. Experimental results show that our method is able to select knowledge more accurately and to generate more informative responses, outperforming significantly the state-of-the-art baselines.

Acknowledgments

This work was jointly supported by the NSFC projects (Key project with No. 61936010 and regular project with No. 61876096), and the Guoqiang Institute of Tsinghua University, with Grant No. 2019GQG1. We thank THUNUS NEXt Joint-Lab for the support.

References

- Steven Bird and Edward Loper. 2004. *NLTK: The natural language toolkit*. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. *Learning phrase representations using RNN encoder-decoder for statistical machine translation*. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. 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.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#). In *International Conference on Learning Representations*.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. [A knowledge-grounded neural conversation model](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qiang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. [Topical-Chat: Towards Knowledge-Grounded Open-Domain Conversations](#). In *Proc. Interspeech 2019*, pages 1891–1895.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Transactions on Information Systems*.
- Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2020. [Sequential latent knowledge selection for knowledge-grounded dialogue](#). In *International Conference on Learning Representations*.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations*.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Zekang Li, Cheng Niu, Fandong Meng, Yang Feng, Qian Li, and Jie Zhou. 2019. [Incremental transformer with deliberation decoder for document grounded conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 12–21, Florence, Italy. Association for Computational Linguistics.
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. [Learning to select knowledge for response generation in dialog systems](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5081–5087. International Joint Conferences on Artificial Intelligence Organization.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018. [Knowledge diffusion for neural dialogue generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1498, Melbourne, Australia. Association for Computational Linguistics.
- Chuan Meng, Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. Refnet: A reference-aware network for background based conversation. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- A. H. Miller, W. Feng, A. Fisch, J. Lu, D. Batra, A. Bordes, D. Parikh, and J. Weston. 2017. Parlai: A dialog research software platform. *arXiv preprint arXiv:1705.06476*.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. [Towards exploiting background knowledge for building conversation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2332, Brussels, Belgium. Association for Computational Linguistics.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. [OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy. 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.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. *Advances in Neural Information Processing Systems*.
- 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.

- Lianhui Qin, Michel Galley, Chris Brockett, Xiaodong Liu, Xiang Gao, Bill Dolan, Yejin Choi, and Jianfeng Gao. 2019. [Conversing by reading: Contentful neural conversation with on-demand machine reading](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5427–5436, Florence, Italy. Association for Computational Linguistics.
- Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. Thinking globally, acting locally: Distantly supervised global-to-local knowledge selection for background based conversation. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Yi-Lin Tuan, Yun-Nung Chen, and Hung-yi Lee. 2019. [DyKgChat: Benchmarking dialogue generation grounding on dynamic knowledge graphs](#). 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 1855–1865, Hong Kong, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Shuohang Wang, Mo Yu, Jing Jiang, and Shiyu Chang. 2018. [A co-matching model for multi-choice reading comprehension](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 746–751, Melbourne, Australia. Association for Computational Linguistics.
- Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. [Proactive human-machine conversation with explicit conversation goal](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3794–3804, Florence, Italy. Association for Computational Linguistics.
- Jun Xu, Haifeng Wang, Zhengyu Niu, Hua Wu, and Wanxiang Che. 2020a. Knowledge graph grounded goal planning for open-domain conversation generation. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Jun Xu, Haifeng Wang, Zhengyu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020b. Conversational graph grounded policy learning for open-domain conversation generation. In *Proceedings of the 58th Conference of the Association for Computational Linguistics*.
- Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao, and Dilek Hakkani-Tur. 2019. [DeepCopy: Grounded response generation with hierarchical pointer networks](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 122–132, Stockholm, Sweden. Association for Computational Linguistics.
- Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. [Augmenting end-to-end dialogue systems with commonsense knowledge](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Yangjun Zhang, Pengjie Ren, and Maarten de Rijke. 2019. Improving background based conversation with context-aware knowledge pre-selection. In *4th International Workshop on Search-Oriented Conversational AI (SCAI)*.
- Xueliang Zhao, Wei Wu, Chongyang Tao, Can Xu, Dongyan Zhao, and Rui Yan. 2020. [Low-resource knowledge-grounded dialogue generation](#). In *International Conference on Learning Representations*.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018a. [Commonsense knowledge aware conversation generation with graph attention](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4623–4629. International Joint Conferences on Artificial Intelligence Organization.
- Hao Zhou, Chujie Zheng, Kaili Huang, Minlie Huang, and Xiaoyan Zhu. 2020. KdConv: A chinese multi-domain dialogue dataset towards multi-turn knowledge-driven conversation. In *Proceedings of the 58th Conference of the Association for Computational Linguistics*.
- Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018b. [A dataset for document grounded conversations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.
- Wenya Zhu, Kaixiang Mo, Yu Zhang, Zhangbin Zhu, Xuezheng Peng, and Qiang Yang. 2017. [Flexible end-to-end dialogue system for knowledge grounded conversation](#). *CoRR*, abs/1709.04264.