

Dialogue-oriented Pre-training

Yi Xu^{1,2,3}, Hai Zhao^{1,2,3,*}

¹ Department of Computer Science and Engineering, Shanghai Jiao Tong University

² Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University, Shanghai, China

³ MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

xuyi_2019@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn

Abstract

Pre-trained language models (PrLM) has been shown powerful in enhancing a broad range of downstream tasks including various dialogue related ones. However, PrLMs are usually trained on general plain text with common language model (LM) training objectives, which cannot sufficiently capture dialogue exclusive features due to the limitation of such training setting, so that there is an immediate need to fill the gap between a specific dialogue task and the LM task. As it is unlikely to collect huge dialogue data for dialogue-oriented pre-training, in this paper, we propose three strategies to simulate the conversation features on general plain text. Our proposed method differs from existing post-training methods that it may yield a general-purpose PrLM and does not individualize to any detailed task while keeping the capability of learning dialogue related features including speaker awareness, continuity and consistency. The resulted Dialog-PrLM is fine-tuned on three public multi-turn dialogue datasets and helps achieve significant and consistent improvement over the plain PrLMs.

1 Introduction

Recently, pre-trained language models (PrLMs) have shown impressive improvements for various downstream NLP tasks (Zhou et al., 2020; Ouyang et al., 2021; Zhang and Zhao, 2021; Radford et al., 2018; Yang et al., 2019; Zhang et al., 2020c; Clark et al., 2020; Li et al., 2021), including the response selection task for multi-turn dialogues, which takes a dialogue history as input and aims to select a

* Corresponding author. This paper was partially supported by National Key Research and Development Program of China (No. 2017YFB0304100), Key Projects of National Natural Science Foundation of China (U1836222 and 61733011), Huawei-SJTU long term AI project, Cutting-edge Machine Reading Comprehension and Language Model. This work was supported by Huawei Noah's Ark Lab.

Dialogue History:

A: *Could you please get me a train?*

B: *Sure I can help you find a train.*

B: *where are you coming from?*

B: *What time do you need to leave by?*

A: *I am leaving Leicester and I need to leave by 20:30.*

B: *What is your destination and day of travel?*

A: *Cambridge and on friday.*

A: *Can I just get the travel time for the train? Thanks!*

Response:

B: *The first train leaving after 20:30 is 142 21:09 and the travel time is 105 minutes.*

Table 1: A multi-turn dialogue example with interleaved or continuous utterances between two speakers.

most suitable response from a collection of answers (Zhou et al., 2016; Wu et al., 2017; Zhou et al., 2018b; Zhu et al., 2018; Zhang et al., 2018; Tao et al., 2019; Gu et al., 2019).

Pre-training tasks of all these PrLMs almost concentrate on two aspects: token prediction and sentence relation prediction. For example, the genetic BERT model (Devlin et al., 2019) uses masked language modeling (MLM) and next sentence prediction (NSP) objectives; ALBERT (Lan et al., 2020) predicts sentence order rather than NSP; ELECTRA (Clark et al., 2020) transfers MLM into a generating and then discriminating process like GAN (Goodfellow et al., 2014). However, these tasks are just devoted to incorporating token-level and sentence-level semantic information into embeddings, and cannot be sufficiently compatible with its dialogue-oriented characteristics.

Table 1 shows a multi-turn dialogue example. Compared with plain text, the utterance turn and speaker role keep shift as a conversation goes on, and the next utterance should keep continuous and consistent with the context. Besides, the two speakers may not follow strict shift rules, and one speaker may continuously shoot multiple utterances. Although some existing works have noticed such

nonlinear nature of multi-turn dialogues, they are limited to conducting post-training or pre-training in a specific domain and do not provide general-purpose dialogue-oriented PrLMs to fundamentally solve this problem (Xu et al., 2021a; Whang et al., 2021; Wolf et al., 2019; Zhang et al., 2020b; Henderson et al., 2020; Bao et al., 2020).

In this work, we make the first attempt to train a general-purpose dialogue-oriented PrLM. However, such a PrLM should be trained on huge dialogue data, which is hard to collect. Thus we propose three novel pre-training strategies (i.e., Insertion, Deletion, Replacement), so that we facilitate plain text originally for common PrLM training to simulate dialogue-like features. The resulted model, we denote as Dialog-PrLM, then is capable of effectively learning speaker awareness, continuity and consistency in a general way. Especially, for the convenient use of the downstream dialogue tasks, we introduce a special token [SOT] before each utterance to tell that it is a start of a turn and learn from these three strategies. These targeted pre-training tasks enable [SOT] to better represent each context utterance. We mimic dialogue-related features on conventional plain text, which can bring up the possibility that similar techniques could be adopted in other domains not only for dialogues.

Our pre-trained Dialog-PrLM is fine-tuned on three multi-turn dialogue response selection benchmarks, and obtains significant and consistent improvements over the plain PrLMs.

2 Related Work

For multi-turn dialogue response selection task, earlier works conduct single-turn match, which concatenates all the utterances in the history dialogue or just considers the last one to match with the candidate response (Lowe et al., 2015; Kadlec et al., 2015; Yan et al., 2016; Tan et al., 2016; Wan et al., 2016; Wang and Jiang, 2016). Recently, existing works tend to model the interaction between each dialogue utterance and the response, which usually adopt the encoding-matching-aggregation paradigm (Zhou et al., 2016; Wu et al., 2017; Zhang et al., 2018; Zhou et al., 2018a,b; Tao et al., 2019; Yuan et al., 2019). After encoding, distinct matching networks generate features for each utterance which are usually passed to GRU (Cho et al., 2014) for aggregating into a final matching score. Besides, some works adopt topic information (Xing et al., 2017; Wu et al., 2018; Xu et al., 2021b) or

conversation disentanglement to select the proper response (Jia et al., 2020; Wang et al., 2020).

There are more and more practice using powerful PrLMs as the model encoder (Zhang et al., 2021, 2020a; Zhu et al., 2020) like BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2020) and ELECTRA (Clark et al., 2020). Considering task domain difference from the general corpus for PrLM pre-training, recent studies start to conduct post-training on target multi-turn dialogue datasets to incorporate in-domain knowledge (Whang et al., 2020; Lu et al., 2020; Gu et al., 2020; Xu et al., 2021a; Whang et al., 2021). Whang et al. (2020) conduct post-training of MLM and NSP tasks as BERT. Rather than using the same tasks as PrLMs, Xu et al. (2021a) and Whang et al. (2021) both considers auxiliary tasks through post-training to enhance response selection.

Although the PrLMs which are trained on plain text have learned contextual semantic representation from token-level or sentence-level pre-training tasks like MLM, NSP, they all do not consider dialogue related features like speaker role, continuity and consistency. Despite some existing works (Xu et al., 2021a; Whang et al., 2021) considers that when conducting post-training, they are limited to a specific domain. (Wolf et al., 2019; Zhang et al., 2020b; Henderson et al., 2020; Bao et al., 2020) train on open-domain conversational data like Reddit for response selection or generation tasks, but they are limited to original pre-training tasks on plain text and ignore the dialogue related features. Besides, Wu et al. (2020) and Li et al. (2020) conduct task-specific training on collected dialogue corpora, but they also suffer from biased and limited amount of dialogue data.

Different from all the previous studies, we still make an attempt in obtaining a general-purpose PrLM but not aiming at any specific tasks like post-training methods. Meanwhile, our proposed dialogue-oriented pre-training enables the resulted PrLMs to especially capture dialogue related features in a general way.

3 Dialogue-oriented Pre-training

For dialogue-oriented pre-training, we split and sample sentences as "utterances" from the general Wikipedia corpus to simulate dialogue-like features. We design three dialogue-oriented pre-training strategies (i.e., Insertion, Deletion, Replacement) to jointly learn dialogue related char-

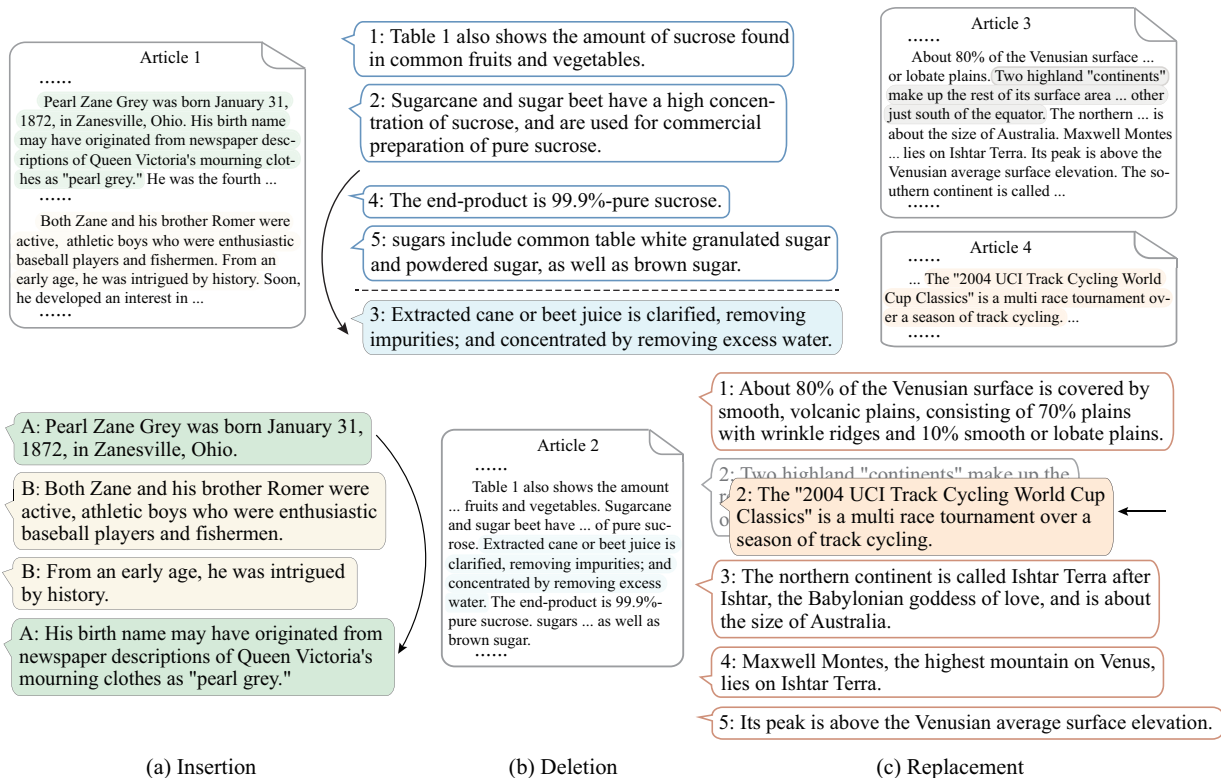


Figure 1: Three dialogue-oriented pre-training strategies.

acteristics based on the plain PrLMs. A special token $[SOT]$ is added before each "utterance", which tells that it is a start of a turn and matches the realistic scene of turn shift. The three tasks use the embedding of $[SOT]$ to represent each utterance and conduct targeted pre-training, which enables $[SOT]$ to learn dialogue related representation about speaker-awareness, continuity and consistency respectively. Figure 1 shows the overview of the three strategies.

Insertion In a real scenario, the speaker role might shift or not for each turn as a conversation goes on. Two speakers may carry out a conversation in turn or one speaker may continuously shoot multiple utterances. Considering a conversation session of four sentences between speaker A and B, we consider three possible cases: AABB, ABAB ABBA. The next time A speaks will happen after 0,1, or 2 turns from the last time. To enable Dialog-PrLM aware of the speaker role information, we should first simulate a two-party conversation on Wikipedia. We sample two continuous sentences $\{u_{A1}, u_{A2}\}$ in one paragraph of an article as the two utterances of A, and sample two continuous ones $\{u_{B1}, u_{B2}\}$ in the following paragraph of the same article as what B says. Sampling from the

same article is to ensure they are talking about one topic in general, which is in line with the realistic scenario and increases the difficulty of prediction. The "continuous" sentences simulate that A continues to express his opinion after being interrupted by B. We insert u_{A2} into $\{u_{B1}, u_{B2}\}$, and add $[SOT]$ before each utterance. We will not disrupt the utterance order inside one speaker, and also keep the overall order that A first then B. In this way, we can get three cases mentioned above. One possible input case is listed here:

$$X_{ins} = [CLS] [SOT] u_{A1} [SOT] u_{B1} [SOT] u_{A2} [SOT] u_{B2} [SEP]$$

The insertion task is to predict the next utterance of A. The whole sequence is encoded by PrLM, and we calculate the cosine similarity of the $[SOT]$ embedding of u_{A1} with the other three utterances as matching scores to predict the u_{A2} . These three scores are passed to a softmax layer and use cross entropy as the insertion loss \mathcal{L}_{ins} .

Deletion The plain pre-training tasks like MLM, NSP or SOP of PrLMs just enable the model to learn token-level or sentence-level semantic information, and they fail to catch dialogue related signals like continuity, which also helps to choose the answer that is coherent with context. We sample

continuous k sentences $\{u_1, u_2, \dots, u_k\}$ from one paragraph and randomly delete u_i from the first $k - 1$ (We do not choose u_k , as there is no [SOT] after u_{k-1}). The input sequence of the PrLM is:

$$X_{del} = [\text{CLS}] [\text{SOT}] u_1 \dots [\text{SOT}] u_{i-1} [\text{SOT}] u_{i+1} \dots [\text{SOT}] u_k [\text{SEP}] [\text{SOT}] u_i [\text{SEP}]$$

We append u_i at end and use [SEP] for separation. Similarly, we calculate the cosine similarity of the [SOT] embedding of u_i with the other remaining $k - 1$ utterances to predict the [SOT] of u_{i+1} , where u_i should be inserted back. These $k - 1$ scores are passed to a softmax layer and use cross entropy as the deletion loss \mathcal{L}_{del} .

Replacement The replacement task is to make Dialog-PrLM recognize the inconsistent utterance within a dialogue session, so that it will select the proper response which is consistent with the context in both style and context. Similar to deletion, we sample continuous k sentences $\{u_1, u_2, \dots, u_k\}$ from one paragraph, and then we sample one sentence u_r from another article, which is used to replace a randomly chosen u_i in $\{u_1, u_2, \dots, u_k\}$. The input sequence is:

$$X_{rep} = [\text{CLS}] [\text{SOT}] u_1 \dots [\text{SOT}] u_{i-1} [\text{SOT}] u_r [\text{SOT}] u_{i+1} \dots [\text{SOT}] u_k [\text{SEP}]$$

Each [SOT] is gathered after encoding, and passed to a linear layer to get a score:

$$\text{score}_{e_{u_j}} = W_r E_j + b_r$$

where $j = 1 \dots i - 1, r, i + 1, \dots, k$, and W_r, b_r are trainable parameters. E_j is the embedding of the j th [SOT]. These k scores are passed to a softmax layer and use cross entropy as the replacement loss \mathcal{L}_{rep} .

We adopt multi-task learning and define the final dialogue-oriented pre-training loss \mathcal{L}_{gen} as:

$$\mathcal{L}_{gen} = \mathcal{L}_{ins} + \mathcal{L}_{del} + \mathcal{L}_{rep}$$

4 Use of Dialogue-oriented Pre-training

Our Dialog-PrLM may be used in terms of domain fine-tuning or multi-task learning: (1) Domain fine-tuning: Our Dialog-PrLM is fine-tuned on the target response selection task. (2) Specific post-training: Our pre-training strategies are slightly adjusted and applied to specific multi-turn dialogue datasets. (3) Domain multi-task learning: the target response selection task jointly learns with the three auxiliary post-training tasks in (2) on Dialog-PrLM.

4.1 Domain Fine-tuning

After our dialogue-oriented pre-training, the target response selection task can be fine-tuned on our Dialog-PrLM. We denote the dataset as $D = \{(C, R, Y)_k\}_{k=1}^N$, where C is dialogue context, and R is the candidate response, and $Y \in \{0, 1\}$ is the label indicating whether R is a proper response for C . Besides, $C = \{U_1, \dots, U_n\}$ and $U_i, 1 \leq i \leq n$ is the i -th utterance in context C . We concatenate all utterances $\{U_i\}_{i=1}^n$ as well as the response R and add [SOT] before each to represent the following sequence:

$$X = [\text{CLS}] [\text{SOT}] U_1 [\text{SOT}] U_2 \dots [\text{SOT}] U_n [\text{SEP}] [\text{SOT}] R [\text{SEP}]$$

With the pre-training, Dialog-PrLM becomes effectively capable of representing each utterance. Therefore, rather than directly using [CLS], we pass all embeddings E of [SOT]s to GRU to model sequential interaction of the context and response, whose final hidden state H is used for generating matching score s for C and R :

$$H = \text{GRU}(E), \\ s = \text{sigmoid}(WH + b)$$

where $E \in \mathbb{R}^{(n+1) \times d}$, $H \in \mathbb{R}^d$, and $W \in \mathbb{R}^d, b \in \mathbb{R}$ are trainable parameters. The response selection loss $\mathcal{L}_{fine-tune}$ is:

$$p(Y | C, R) = sY + (1 - s)(1 - Y) \\ \mathcal{L}_{fine-tune} = -\frac{1}{N} \sum_{(C,R,Y) \in D} \log(p(Y | C, R))$$

4.2 Specific Post-training

Because the target dataset usually concentrates on a specific domain, existing works tend to introduce self-supervised post-training on the target domain in order to incorporate the in-domain knowledge. Here we can also apply the three strategies to the target multi-turn dialogue dataset as self-supervised auxiliary tasks, which jointly learn with the response selection task in Section 4.1.

To conduct the three auxiliary tasks, we should first sample from multi-turn dialogues to build post-training datasets. For the insertion task, there is a little difference from that in Wikipedia. We randomly choose k continuous utterances $\{u_1, u_2, \dots, u_k\}$ from a dialogue, and fix u_1 but randomly insert u_2 to any interval among $\{u_3, \dots, u_k\}$. The input sequence is:

$$X'_{ins} = [\text{CLS}] [\text{SOT}] u_1 [\text{SOT}] u_3 \dots [\text{SOT}] u_i [\text{SOT}] u_2 [\text{SOT}] u_{i+1} \dots [\text{SOT}] u_k [\text{SEP}],$$

We calculate the cosine similarity of the [SOT] embedding of u_1 with the other following utterances as matching scores to predict the u_2 . The loss is denoted as \mathcal{L}'_{ins} . Considering the following turn (u_2) tends to be more related to u_1 compared with the next utterance of u_1 's speaker (denoted as u_t), we do not predict u_t as what we do in Wikipedia. But they are both expected to recognize the most related utterance with u_1 , which helps select the proper response.

For the deletion and replacement tasks, we sample continuous k utterances from one dialogue, and conduct deletion or replacement in the same way as Wikipedia. The post-training losses for both tasks on the target domain are denoted as $\mathcal{L}'_{del}, \mathcal{L}'_{rep}$ respectively.

4.3 Domain Multi-task Learning

We apply the multi-task learning framework on the target domain on our Dialog-PrLM. Carried with dialogue related features from the general corpus, Dialog-PrLM is expected to learn from the target domain together with the target task. We train the response selection task in the same way as 4.1, and denote the loss as $\mathcal{L}'_{reselect}$. The final loss is to sum up response selection loss and the three auxiliary task losses:

$$\mathcal{L}_{final} = \mathcal{L}'_{ins} + \mathcal{L}'_{del} + \mathcal{L}'_{rep} + \mathcal{L}'_{reselect}$$

5 Implementation of Dialogue-oriented Pre-training

For dialogue-oriented pre-training, we sample train and valid datasets for the insertion, deletion and replacement tasks on both English and Chinese Wikipedia. To prevent information leakage (e.g. the model peeks at the correct utterance order from deletion samples when conducting replacement), we divide all the articles equally into three disjoint equal sets to sample from for the three tasks respectively. Data statistics are in Table 2.

Statistics	English		Chinese	
	Train	Valid	Train	Valid
#articles/task	330k	10k	67k	20k
Insertion	1.5M	47k	638k	189k
Deletion	1M	31k	508k	146k
Replacement	1M	32k	542k	160k

Table 2: Statistics for dialogue-oriented pre-training.

For English Wikipedia (totally 1,060,131 articles), we sample from 330,000/10,000 articles for training/evaluation respectively for each task. To

ensure data quality, we omit the "References" and "Literature" part. For insertion, we sample twice disjointly from one paragraph which has more than 4 sentences, and then sample twice from the next satisfactory one to construct two training samples, which goes on until the end of an article. For deletion and replacement, we sample k continuous sentences as a training sample from each paragraph with more than k sentences. We limit the maximum words of each sample to 400 to prevent overflows after tokenization.

For Chinese Wikipedia (totally 262,405 articles), we sample from 67,468/20,000 articles for training/evaluation respectively for each task. As Chinese corpus is much smaller than the English one, we sample disjoint k continuous sentences as much as we can from each paragraph with more than k sentences for the deletion and replacement tasks. As to insertion, we conduct sampling the same way as English. The maximum length of each sample is limited to 450.

6 Experiments

6.1 Datasets

For the target response selection task, our Dialog-PrLM is fine-tuned on three widely used benchmark datasets: (1) **E-commerce Corpus** (Zhang et al., 2018): includes conversations between customers and shopkeepers from the largest e-commerce platform Taobao in China. (2) **Douban Corpus** (Wu et al., 2017): consists of multi-turn conversations from the Douban group, which is a popular social networking service in China. (3) **Ubuntu Corpus (v1.0)** (Lowe et al., 2015) consists of English multi-turn conversations about technical support collected from chat logs of the Ubuntu forum.

As to the three auxiliary tasks for domain multi-task learning, we conduct sampling from the batch when training the response selection task. Dialogues will be neglected if they are less than 3 utterances. Different from general pre-training, we do not require every dialogue to have at least k sentences for all the three tasks.

For evaluation, we use the same metric $R_n@k$ as previous works, which selects k best matchable candidate responses among n and calculates the recall of the true ones. We also use MAP (Mean Average Precision), MRR (Mean Reciprocal Rank), and Precision-at-one P@1 as previous works.

	Model	E-commerce			Douban					Ubuntu			
		$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MAP	MRR	P@1	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$
	DialoGPT	-	-	-	-	-	-	-	-	-	79.0	88.5	97.1
	TOD-BERT	-	-	-	-	-	-	-	-	-	79.7	89.0	97.4
	BERT-[CLS]	62.7	82.2	96.2	58.7	62.7	45.1	27.6	45.8	82.7	81.9	90.4	97.8
◇	BERT-[SEP]	65.1	84.8	97.4	59.5	63.9	46.0	27.7	46.9	84.3	82.1	90.5	97.8
◇	Dialog-BERT	66.2	85.5	97.6	60.0	64.1	46.9	28.9	46.7	83.3	82.3	90.6	97.7
♣	BERT+multi-task	65.8	84.6	97.6	60.2	64.7	46.9	28.5	48.6	82.5	85.0	92.5	98.3
♣	Dialog-BERT+multi-task	68.0	85.3	97.7	60.9	64.9	48.0	30.0	47.9	82.9	85.4	92.8	98.5
	ELECTRA-[CLS]	58.2	79.6	96.9	59.0	63.2	44.8	27.6	47.3	82.8	82.5	90.7	97.8
	ELECTRA-[SEP]	60.4	80.6	96.3	58.8	62.5	44.2	26.9	46.3	84.1	82.2	90.7	97.8
♡	Dialog-ELECTRA	61.1	81.4	96.9	59.8	64.1	46.5	28.3	47.7	84.1	83.5	91.4	98.0
♠	ELECTRA+multi-task	68.1	86.8	97.9	61.4	65.3	47.5	29.6	50.6	83.8	86.6	93.4	98.5
♠	Dialog-ELECTRA+multi-task	68.3	86.3	98.0	61.6	65.6	48.3	30.0	49.8	84.7	86.8	93.6	98.6

Table 3: Response selection results on E-commerce, Douban and Ubuntu datasets. The symbols on the left indicate the corresponding comparison groups. The best results in each group are in boldface.

6.2 Experimental Settings

We work on two different PrLMs: BERT (*bert-base-uncased*, *bert-base-chinese*) and ELECTRA (*electra-base-discriminator*, *chinese-electra-180g-base-discriminator*)¹.

For our dialogue-oriented pre-training on Wikipedia, the max input sequence length is set to 512 after WordPiece tokenization. We set the learning rate as $2e-5$ with a warmup proportion of 10%. The plain PrLMs are continuously pre-trained with batch size of 8 per task for BERT and 16 per task for ELECTRA. It is trained for 1 epoch and evaluated every 10000 steps. The model with the best average accuracy of the three tasks is saved as Dialog-PrLM. The pre-training experiment needs 4 nVidia RTX 2080 GPUs.

For fine-tuning on our Dialog-PrLMs, the batch size is 32 and the max sequence length is 350. The model is trained for 5 epochs and evaluated after each epoch on the three datasets and both Dialog-PrLMs. Other settings are the same as dialogue-oriented pre-training. For domain multi-task learning on our Dialog-PrLMs, the batch size is 16, and the epoch is 3 for Douban on BERT, 4 for Ubuntu on ELECTRA and 5 for other cases. Other settings are the same as fine-tuning. The k value for both pre-training and domain multi-task learning is 5. The fine-tuning/multi-task learning experiments need 1/2 nVidia RTX 2080 GPUs.

6.3 Results

To verify the effectiveness of our method, we conduct extensive empirical studies on three multi-turn dialogue benchmarks. We are aware that applying complicated matching networks, speaker em-

beddings (Gu et al., 2020) or other various auxiliary tasks (Whang et al., 2020; Xu et al., 2021a; Whang et al., 2021) would achieve further improvement, but to fairly evaluate the general-purpose pre-training for dialogue tasks, we still follow the standard fine-tuning procedure on Dialog-PrLM by excluding those too advanced auxiliary techniques.

BERT-[CLS]: Each utterance U_i and the candidate response R are concatenated as $[\text{CLS}] U_1 U_2 \dots U_n [\text{SEP}] R [\text{SEP}]$ and then fed into the plain BERT model. The output embedding of $[\text{CLS}]$ is used for classification.

BERT-[SEP]: Rather than just use $[\text{CLS}]$, we append $[\text{SEP}]$ to each utterance or response to represent the previous sequence: $[\text{CLS}] U_1 [\text{SEP}] U_2 \dots [\text{SEP}] U_n [\text{SEP}] R [\text{SEP}]$, which is then fed into the original BERT model. The output embeddings of $[\text{SEP}]$ are gathered and fed into GRU for a matching vector like Section 4.1.

Dialog-BERT: This model conducts dialogue-oriented pre-training on the original BERT, we fine-tune on our Dialog-BERT through feeding $[\text{SOT}]$ embeddings to GRU.

BERT+multi-task: The response selection task is trained with the three auxiliary tasks on target datasets. We also add the special token $[\text{SOT}]$, and the only difference from Section 4.3 is that the joint learning is conducted on the original BERT.

Dialog-BERT+multi-task: As described in Section 4.3, we conduct domain multi-task learning on our pre-trained Dialog-BERT.

We also conduct fine-tuning on the English Ubuntu dataset with two dialogue related models: (1) DialoGPT (Zhang et al., 2020b) is an extension of GPT-2 that is pre-trained on Reddit data

¹Both datasets and code are available at <https://github.com/xyeas/Dialog-PrLM>

Model	E-commerce			Douban						Ubuntu		
	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MAP	MRR	P@1	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$
Dialog-BERT	66.2	85.5	97.6	60.0	64.1	46.9	28.9	46.7	83.3	82.3	90.6	97.7
w/o Insertion	64.9	83.6	97.7	59.0	63.2	45.1	27.8	46.5	83.2	82.1	90.5	97.8
w/o Deletion	64.2	83.7	97.5	59.0	63.1	44.7	27.5	45.9	83.8	82.1	90.6	97.8
w/o Replacement	64.4	84.7	97.6	59.8	63.6	45.4	28.4	46.8	83.6	82.1	90.4	97.8

Table 4: Ablation results for dialogue-oriented pre-training.

from scratch. (2) TOD-BERT (Wu et al., 2020) is trained on a combination of 9 task-oriented dialogue datasets over BERT and incorporates response selection objective. Experiments on ELECTRA are conducted in the same way with BERT. The results are in Table 3. Below PrLM- denotes BERT or ELECTRA.

Compared to the unsatisfactory results from both DialoGPT and TOD-BERT, it demonstrates the powerfulness and universality of our proposed dialogue-oriented pre-training. Compared with PrLM-[CLS], PrLM-[SEP] performs better in general except a little decrease in Ubuntu and Douban on ELECTRA, which shows that modelling the sequential interaction of the dialogue context and response helps improve performance.

After conducting the dialogue-oriented pre-training on Wikipedia, our Dialog-PrLM achieves further improvement on the three datasets and the two PrLMs, which shows that the three targeted training strategies enable the [SOT] token in Dialog-PrLM to grasp dialogue related nature (e.g. speaker-awareness, continuity, consistency) at the same time, so that it is more capable of representing an utterance compared with [SEP] in the plain PrLM (PrLM-[SEP]).

When we train the response selection task jointly with the three auxiliary tasks on target datasets, the domain multi-task learning on our Dialog-PrLM (Dialog-PrLM+multi-task) is still always performing better than on the plain PrLM (PrLM+multi-task). Having re-learned broader representation on general corpus, domain post-training further incorporates [SOT] with dialogue related feature from in-domain multi-dialogues and thus helps choose the correct response.

Compare PrLM-[SEP] with PrLM+multi-task, Dialog-PrLM with Dialog-PrLM+multi-task, domain multi-task learning indeed achieves improvements due to its incorporated in-domain dialogue related knowledge, which verifies the effectiveness of our proposed three strategies when applying to domain multi-turn dialogue datasets.

In conclusion, conducting dialogue related fea-

ture pre-training with our proposed three strategies on Wikipedia (Dialog-PrLM) helps achieve improvements when fine-tuning, and it will further improve when applying these strategies to domain multi-turn dialogues (Dialog-PrLM+multi-task).

7 Analysis

7.1 Ablation Study

In order to investigate the performance of each strategy, we conduct ablation experiments for both pre-training and domain multi-task learning. Results are shown in Tables 4 and 5 respectively.

The results in Table 4 indicate that the insertion, deletion and replacement tasks jointly contribute to the final increase. Influence on Ubuntu seems less than that on Douban and E-commerce, as the Ubuntu corpus contains many terminologies that do not usually appear in general corpora (e.g., apt-get, lsmod and grep) (Whang et al., 2020), so according to BERT+multi-task in Table 3, conducting domain post-training is more effective.

We also do ablation study for Douban on Dialog-BERT in Table 5 to explore the performance of three auxiliary tasks when applying to the target multi-turn dialogue datasets. Similarly, removing any part leads to worse performance, showing the necessity of each task.

Model	Douban				
	MAP	MRR	P@1	$R_{10}@1$	$R_{10}@2$
Dialog-BERT+multi-task	60.9	64.9	48.0	30.0	47.9
w/o Insertion	58.2	62.5	44.8	27.4	45.0
w/o Deletion	60.6	64.6	47.1	29.1	48.8
w/o Replacement	60.2	64.1	46.3	29.0	48.5

Table 5: Ablation for domain multi-task learning.

7.2 Utterance Representation Test

We have added a special token [SOT] before each utterance or response to represent the following sequence. After pre-training on Wikipedia on PrLMs, [SOT] of our Dialog-PrLM is expected to obtain the respective utterance representation through the three dialogue-oriented strategies.

Dialogue Context & Response	[CLS]	[SEP]	[SOT]
Douban			
1: how about going to Xitang during the Spring Festival?	0.8424	1	-0.4001
2: <i>I also want to go that time, but it will be more expensive for accommodation than weekend.</i>	0.9364	1	0.8666
3: have you contacted about where to live?	0.9694	1	-0.2902
4: if you haven't booked accommodation, it seems very difficult to book during the Spring Festival.	0.9495	1	-0.1643
5: will there be many people, or no shops open during the Spring Festival?	0.9024	1	-0.5064
6: it's okay if no stores open. i enjoy quiet and just want somewhere to drink and eat.	0.9737	1	-0.1644
7: but traffic is a hassle during the Spring Festival.	0.9684	1	-0.6139
8: <i>there are several people in this group xxx who go to Xitang, you can join them to cook together</i>	0.8865	1	0.8840
Response: is there anyone in Beijing, let's go together sometime.			
Ubuntu			
1: hello	0.9279	0.9956	-0.9791
2: what is the best server for update ubuntu	0.8027	0.9957	-0.9653
3: whichever is closest to you generally they all have the same content	0.9738	0.9949	-0.8846
4: have you read the documentation for the <i>script</i>	0.9643	0.9966	-0.2836
5: i know thats the error does the <i>script</i> s documentation indicate what modules or packages are queried to run it	0.8176	0.9964	0.5233
6: u can tell me how to install python modules and i will install the modules in the python <i>script</i>	0.8940	0.9977	0.5287
7: where did you get this <i>script</i>	0.9328	0.9980	0.5105
Response: can you be more specific what is this <i>script</i> supposed to do import logging			
E-commerce			
1: please help me issue an invoice with the title of xxx.	0.8864	1	-0.9791
2: okay.	0.8726	1	-0.8137
3: ok, is it difficult to package?	0.9249	1	-0.6777
4: not at all.	0.7843	1	-0.9359
5: please send me a picture tutorial later, i will do it myself.	0.8648	1	-0.3668
6: okay, if you have any question, please feel free to consult us.	0.9381	1	0.3211
7: fine, <i>can you deliver it today?</i>	0.9265	1	0.9899
8: yes, please check the <i>delivery</i> address.	0.9240	1	0.7117
9: correct.	0.6862	1	0.5704
10: ok, if pay before 16:00 today, <i>it will be delivered today, otherwise it will be delivered next day.</i>	0.9550	1	0.5866
11: i've already <i>paid</i> .	0.8663	1	0.0622
Response: ok, we will deliver the goods as soon as possible today, please wait at patience.			

Table 6: Examples from test sets of Douban, Ubuntu and E-commerce respectively. Similarity scores ≥ 0.5 of our Dialog-BERT([SOT]) are bold, indicating the corresponding utterances are considered relevant to the response.

To explore the semantic information of [SOT] of our Dialog-BERT, we calculate the cosine similarity of the correct response to each utterance in the dialogue context ([SOT]). Table 6 lists examples from the three target datasets respectively. For comparison, we use BERT to encode each utterance or response and use [CLS] for calculation ([CLS]). We also concatenate utterances and response with separation of [SEP] on BERT and then split to calculate similarity ([SEP]).

We observe that for all examples, both BERT([CLS]) and BERT([SEP]) can not discriminate which utterance is related with the correct answer. All the utterances are treated the same way including the irrelevant ones, which leads to much noise for response selection.

After conducting the dialogue-oriented pre-training on Wikipedia, Dialog-BERT learns a stark sense of "irrelevant" and "relevant". It is able to concentrate on the most critical utterances and distinguish from the irrelevant ones by a large margin.

For the example of Douban, Dialog-BERT realizes that the second and last utterance are most relevant. The response asks for someone to travel together, and is related to the second utterance which expresses a wish to go and the last one which gives a group to travel together. Dialog-BERT ignores the noise about accommodation and transportation, and is able to select related utterances among noise rather than just use the last one. For Ubuntu, Dialog-BERT concentrates on utterances about script and ignores the previous background information. For E-commerce, it recognizes the related last few utterances about delivery, and ignores the packaging information before. From examples from target datasets, our Dialog-BERT has absorbed related knowledge from our proposed dialog-oriented pre-training. The [SOT] could better represent each utterance, which can be utilized in tasks about representation like multi-party dialogue disentanglement.

8 Conclusion

This paper presents a novel general-purpose solution for dialogue tasks with pre-trained language models. To fill the gap between a detailed task and the LM task of PrLM, we propose dialogue-oriented pre-training on large scale of artificially built dialogue data which lets the resulted Dialog-PrLM enjoy both merits of general-purpose and capturing key dialogue related features including speak awareness, continuity and consistence. Our models are evaluated on three benchmark response selection datasets and achieve consistent performance improvement over the plain PrLMs.

References

- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. **PLATO: Pre-trained dialogue generation model with discrete latent variable**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pages 85–96.
- 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 2014)*, pages 1724–1734.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. **ELECTRA: Pre-training text encoders as discriminators rather than generators**. In *International Conference on Learning Representations (ICLR 2020)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2019)*, pages 4171–4186.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems*, pages 2672–2680.
- Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020. **Speaker-aware bert for multi-turn response selection in retrieval-based chatbots**. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2041–2044.
- Jia-Chen Gu, Zhen-Hua Ling, and Quan Liu. 2019. **Interactive matching network for multi-turn response selection in retrieval-based chatbots**. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2321–2324.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. **ConveRT: Efficient and accurate conversational representations from transformers**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174.
- Qi Jia, Yizhu Liu, Siyu Ren, Kenny Zhu, and Haifeng Tang. 2020. **Multi-turn response selection using dialogue dependency relations**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 1911–1920.
- Rudolf Kadlec, Martin Schmid, and Jan Kleindienst. 2015. **Improved deep learning baselines for ubuntu corpus dialogs**. *arXiv preprint arXiv:1510.03753*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. **Albert: A lite bert for self-supervised learning of language representations**. In *International Conference on Learning Representations (ICLR 2020)*.
- Junlong Li, Zhuosheng Zhang, Hai Zhao, Xi Zhou, and Xiang Zhou. 2020. **Task-specific objectives of pre-trained language models for dialogue adaptation**. *arXiv preprint arXiv:2009.04984*.
- Zuchao Li, Zhuosheng Zhang, Hai Zhao, Rui Wang, Kehai Chen, Masao Utiyama, and Eiichiro Sumita. 2021. **Text compression-aided transformer encoding**. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**. *arXiv preprint arXiv:1907.11692*.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. **The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems**. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294.
- Junyu Lu, Xiancong Ren, Yazhou Ren, Ao Liu, and Zenglin Xu. 2020. **Improving contextual language models for response retrieval in multi-turn conversation**. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1805–1808.
- Siru Ouyang, Zhuosheng Zhang, and Hai Zhao. 2021. **Dialogue graph modeling for conversational machine reading**. In *Findings of the Association for Computational Linguistics: ACL 2021*.

- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Ming Tan, Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2016. Lstm-based deep learning models for non-factoid answer selection. In *International Conference on Learning Representations (ICLR 2016)*.
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. Multi-representation fusion network for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 267–275.
- Shengxian Wan, Yanyan Lan, Jun Xu, Jiafeng Guo, Liang Pang, and Xueqi Cheng. 2016. Match-srnn: Modeling the recursive matching structure with spatial rnn. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI 2016)*, page 2922–2928.
- Shuohang Wang and Jing Jiang. 2016. [Learning natural language inference with LSTM](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2016)*, pages 1442–1451.
- Weishi Wang, Steven C.H. Hoi, and Shafiq Joty. 2020. [Response selection for multi-party conversations with dynamic topic tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 6581–6591.
- Taesun Whang, Dongyub Lee, Chanhee Lee, Kisu Yang, Dongsuk Oh, and HeuiSeok Lim. 2020. An effective domain adaptive post-training method for bert in response selection. In *Proc. Interspeech 2020*.
- Taesun Whang, Dongyub Lee, Dongsuk Oh, Chanhee Lee, Kijong Han, Dong-hun Lee, and Saebyeok Lee. 2021. Do response selection models really know what’s next? utterance manipulation strategies for multi-turn response selection. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2021)*.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Chien-Sheng Wu, Steven C.H. Hoi, Richard Socher, and Caiming Xiong. 2020. [TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 917–929.
- Yu Wu, Zhoujun Li, Wei Wu, and Ming Zhou. 2018. Response selection with topic clues for retrieval-based chatbots. *Neurocomputing*, pages 251–261.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. [Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 496–505.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2017)*.
- Ruijian Xu, Chongyang Tao, Daxin Jiang, Xueliang Zhao, Dongyan Zhao, and Rui Yan. 2021a. Learning an effective context-response matching model with self-supervised tasks for retrieval-based dialogues. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*.
- Yi Xu, Hai Zhao, and Zhuosheng Zhang. 2021b. Topic-aware multi-turn dialogue modeling. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*.
- Rui Yan, Yiping Song, and Hua Wu. 2016. [Learning to respond with deep neural networks for retrieval-based human-computer conversation system](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, page 55–64.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5753–5763.
- Chunyuan Yuan, Wei Zhou, Mingming Li, Shangwen Lv, Fuqing Zhu, Jizhong Han, and Songlin Hu. 2019. Multi-hop selector network for multi-turn response selection in retrieval-based chatbots. 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 111–120.
- Shuailiang Zhang, Hai Zhao, Yuwei Wu, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2020a. Dcmn+: Dual co-matching network for multi-choice reading comprehension. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. Dialogpt: Large-scale generative pre-training for conversational response generation. In *ACL, system demonstration*.
- Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. 2018. [Modeling multi-turn conversation with deep utterance aggregation](#).

- In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3740–3752.
- Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020c. Semantics-aware BERT for language understanding. In *the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhuosheng Zhang, Junjie Yang, and Hai Zhao. 2021. Retrospective reader for machine reading comprehension. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhuosheng Zhang and Hai Zhao. 2021. Advances in multi-turn dialogue comprehension: A survey. *arXiv preprint arXiv:2103.03125*.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018a. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*.
- Junru Zhou, Zhuosheng Zhang, Hai Zhao, and Shuailiang Zhang. 2020. **LIMIT-BERT : Linguistics informed multi-task BERT**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4450–4461.
- Xiangyang Zhou, Daxiang Dong, Hua Wu, Shiqi Zhao, Dianhai Yu, Hao Tian, Xuan Liu, and Rui Yan. 2016. Multi-view response selection for human-computer conversation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*, pages 372–381.
- Xiangyang Zhou, Lu Li, Daxiang Dong, Yi Liu, Ying Chen, Wayne Xin Zhao, Dianhai Yu, and Hua Wu. 2018b. Multi-turn response selection for chatbots with deep attention matching network. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, pages 1118–1127.
- Pengfei Zhu, Zhuosheng Zhang, Jiangtong Li, Yafang Huang, and Hai Zhao. 2018. **Lingke: a fine-grained multi-turn chatbot for customer service**. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations (COLING 2018)*, pages 108–112.
- Pengfei Zhu, Hai Zhao, and Xiaoguang Li. 2020. Dual multi-head co-attention for multi-choice reading comprehension. *arXiv preprint arXiv:2001.09415*.