

# Cross Copy Network for Dialogue Generation

Changzhen Ji<sup>1</sup>, Xin Zhou<sup>2</sup>, Yating Zhang<sup>2</sup>  
Xiaozhong Liu<sup>3</sup>, Changlong Sun<sup>2</sup>, Conghui Zhu<sup>1</sup> and Tiejun Zhao<sup>1</sup>

<sup>1</sup>Harbin Institute of Technology, Harbin, China

<sup>2</sup>Alibaba Group, Hangzhou, China

<sup>3</sup>Indiana University Bloomington, USA

czji\_hit@outlook.com, {eric.zx,ranran.zyt}@alibaba-inc.com

liu237@indiana.edu, changlong.scl@taobao.com, {conghui,tjzhao}@hit.edu.cn

## Abstract

In the past few years, audiences from different fields witness the achievements of sequence-to-sequence models (e.g., LSTM+attention, Pointer Generator Networks and Transformer) to enhance dialogue content generation. While content fluency and accuracy often serve as the major indicators for model training, dialogue logics, carrying critical information for some particular domains, are often ignored. Take customer service and court debate dialogue as examples, compatible logics can be observed across different dialogue instances, and this information can provide vital evidence for utterance generation. In this paper, we propose a novel network architecture - Cross Copy Networks (CCN) to explore the current dialog context and similar dialogue instances' logical structure simultaneously. Experiments with two tasks, court debate and customer service content generation, proved that the proposed algorithm is superior to existing state-of-art content generation models.

## 1 Introduction

As an important task in Natural Language Generation (NLG), dialogue generation empowers a wide spectrum of applications, such as chatbot and customer service automation. In the past few years, breakthroughs in dialogue generation technology focused on a series of sequence-to-sequence models (Sutskever et al., 2014). More recently, external knowledge is employed to enhance model performance. Wu et al. (2019); Liu et al. (2018) can assist dialogue generation by using knowledge triples. Similarly, Li et al. (2019); Rajpurkar et al. (2018); Huang et al. (2018); Reddy et al. (2019) explore document as knowledge discovery for dialogue generation, and Xia et al. (2017); Ye et al. (2020); Ghazvininejad et al. (2018); Parthasarathi and Pineau (2018) utilize unstructured knowledge

to explore in the open-domain dialogue generation. However, unaffordable knowledge construction and defective domain adaptation restrict their utilization.

Copy-based generation models (Vinyals et al., 2015; Gu et al., 2016) have been widely adopted in content generation tasks and show better results compared to sequence-to-sequence models when faced with out-of-vocabulary problem. Thanks to their nature of leveraging vocabulary and context distributions for content copy, it enables to copy the aforementioned named entities (e.g., person names, locations, company names) appeared in the above context) from the upper context to improve the specificity of the generated text.

In the task of dialogue generation, we can often observe the phrases/utterance patterns across different "similar dialogue" instances. For example, in customer service, the similar inquiries from the customers will get similar responses from the staff. It motivates us to build a model that can not only copy the content within the upper context of the target dialogue instance, but also learn the similar patterns across different similar cases of the target instance. Such external copy can be critical in some scenarios.

As shown in Figure.1, we propose two different kinds of copy mechanisms in this study: **vertical copy** context-dependent information within the target dialogue instance, and **horizontal copy** logic-dependent content across different 'Similar Cases' (SC). This framework is labeled as Cross-Copy Networks (CCN). As exemplar dialogue depicted, judges may repeat (horizontal copied) words, phrases or utterances from historical dialogues when those SCs sharing similar content, e.g., 'A sue B because of X and Y'.

In order to validate the proposed model, we employ two different dialogue datasets from two orthogonal domains - *court debate* and *customer ser-*

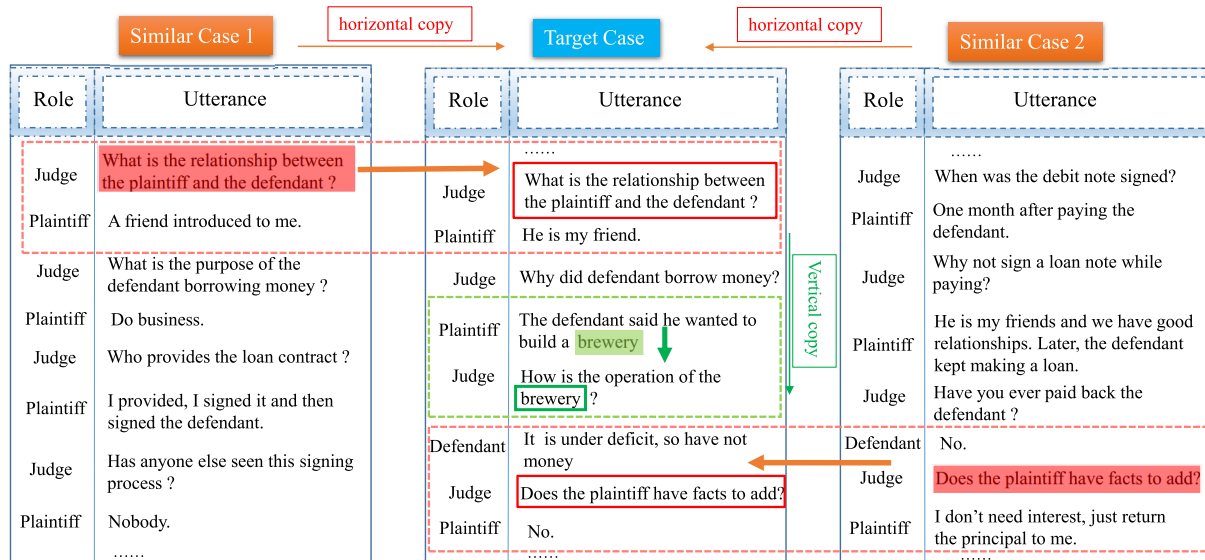


Figure 1: An example from the court debate showing the intuition of utterance generation by leveraging its context and the information from its similar cases. We name the copy process from the context as *vertical copy* and the one copied from its neighbor cases as *horizontal copy*.

*vice*. We apply proposed CCN to both datasets for dialogue generation. Experiments show that our model achieves the best results. To sum up, our contributions are as follows:

- We propose a new end-to-end model, the Cross Copy Networks (CCN), which enables internal (vertical) copy from the target dialogue and external (horizontal) copy from similar cases in the dataset without employing any external resources.
- We validate the proposed model by leveraging two different datasets - court debate and customer service datasets. Experiments show that our model has achieved State-of-the-art results in both domain datasets.
- To motivate other scholars to investigate this novel but an important problem, we make the experimental datasets publicly available<sup>1</sup>.

## 2 Model

In this section, we introduce the proposed model, the Cross Copy Network, which has three major components:

1. **Target Case Representation:** we obtain the target case representation with two attention distributions at the utterance layer and the dialogue layer, which contribute to the final attention distribution (Section 2.1);

2. **Similar Case Representation:** we fine-tune the pre-trained language model to obtain similar cases, and adopt the same method as the target case for encoding SC (Section 2.2);
3. **Cross Copy:** we learn two pointer distributions which are used to achieve internal (vertical) copy and external (horizontal) copy respectively (Section 2.3).

### 2.1 Target Case Representation

Given a dialogue  $\mathbf{D} = \{(\mathbf{U}, \mathbf{R})^L\}$  containing  $L$  utterances, the  $\mathbf{U}$  and  $\mathbf{R}$  stand for utterance and role of speaker, respectively, where each utterances in the dialogue is expressed as  $\mathbf{U}_i = \{w_{i1}, w_{i2}, \dots, w_{il}\}$ , and the  $l$  represents the length of the utterances. To distinguish SC and original context, we define the original context (historical dialogue) as **Target Case**.

Our encoder is shown in Figure 2. It is designed with hierarchical infrastructure consisting of three levels of components: utterance layer, dialogue layer and transformer layer.

#### 2.1.1 Utterance Layer

In the dialogue, role information can make critical contribution to the task of dialogue generation, and different roles may not share consistent lexical spaces. For role information  $\mathbf{R}_i$ , we utilize a 100-dimensional vector to represent different roles which is randomly initialized, and updated via back propagation.

<sup>1</sup><https://github.com/jichangzhen/CCN>

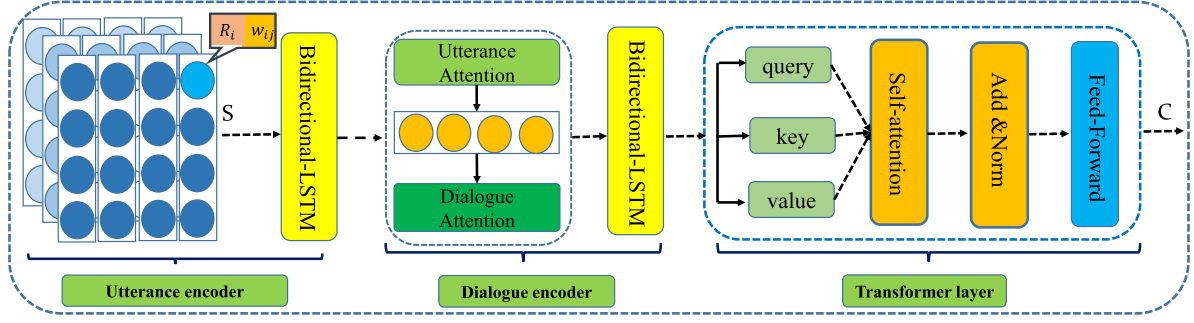


Figure 2: The encoder of CCN is divided into three levels: (1) **Utterance layer**: is used to encode role information and word level information; (2) **Dialogue layer**: is used to encode sentence level information; (3) **Transformer layer**: is used to capture long distance dependence for dialogue.

To take the role information into consideration for utterance representation learning, we concatenate the role information with each word of utterance expressed as  $\mathbf{S}_{ij}$ , and we use Bidirectional Long-Short Term Memory networks (**Bi-LSTM**) (Hochreiter and Schmidhuber, 1997) to encode the semantics of the utterance while maintaining its syntactic be expressed as  $\mathbf{h}^d$ .

In order to obtain the different importance of different historical dialogue information, we adopt the attention mechanism (Bahdanau et al., 2014) to obtain the utterance level’s attention distribution in historical dialogue  $a_j^u$  and utterance context  $\mathbf{h}_i^U$ :

$$\mathbf{h}_i^U = \sum_{j=1}^l a_j^u \mathbf{h}_{ij}^d \quad (1)$$

$$a_j^u = \frac{\exp(\tanh(W^u \mathbf{h}_{ij}^d + b^u)^T \mathbf{h}_{ij}^d)}{\sum_{j=1}^l \exp(\tanh(W^u \mathbf{h}_{ij}^d + b^u)^T \mathbf{h}_{ij}^d)} \quad (2)$$

The  $a_j^u$  represents the word probability distribution for the target utterance<sup>2</sup>.

### 2.1.2 Dialogue Layer

In order to represent the context information of the dialogue, we also use **Bi-LSTM** to encode the utterance dependency to obtain a global representation of an utterance as dialogue, denoted as  $\mathbf{h}^D$ .

We obtain the dialogue layer attention distribution  $a_i^d$ , which is a probability distribution over the prior utterances in the target dialogue. The  $a_i^d$  can be expressed as:

$$a_i^d = \frac{\exp(\tanh(W^d \mathbf{h}_i^D + b^d)^T \mathbf{h}_i^D)}{\sum_{i=1}^L \exp(\tanh(W^d \mathbf{h}_i^D + b^d)^T \mathbf{h}_i^D)} \quad (3)$$

<sup>2</sup>The  $W^u$  and  $b^u$  are learnable parameters, the  $\tanh$  is hyperbolic tangent function.

The final context attention distribution  $A^d$  of target case can be expressed as the product of  $a_j^u$  and  $a_i^d$ :

$$A^d = a_j^u * a_i^d \quad (4)$$

### 2.1.3 Transformer Layer

To expand the model’s ability to focus on different locations of long context, we adopt the self-attention with multi-heads (Vaswani et al., 2017) to explore an enhance representation, denoted as Transformer-Block. We feed  $\mathbf{h}^D$  to a N-layer Transformer-Block to suppress the long distance dependency for dialogue. Following this strategy, the final target case representation is:

$$\mathbf{C}^d = \text{Transformer}^N(\mathbf{h}^D) \quad (5)$$

## 2.2 Similar Case Representation

In this section, we introduce the approach of obtaining and representing similar cases.

### 2.2.1 Similar Case Finding

The similar cases (**SCs**) of the target case is discovered from the same dataset where the target case stays. To make it more efficient, we use ElasticSearch<sup>3</sup> to retrieve top 50 similar cases as candidates by leveraging the target case as a query and the all the other cases as documents. To make it more effective, we fine-tune the pre-trained **RoBERTa**<sup>4</sup> (Liu et al., 2019b) model. It uses a linear layer with sigmoid activation function on top of the pooled [CLS] representation from the

<sup>3</sup><https://www.elastic.co/products/elasticsearch>

<sup>4</sup>All the dialogues in the dataset are used to fine-tune RoBERTa. The average lengths of the candidate cases and the target case are both 106, so it can all fit within the RoBERTa model

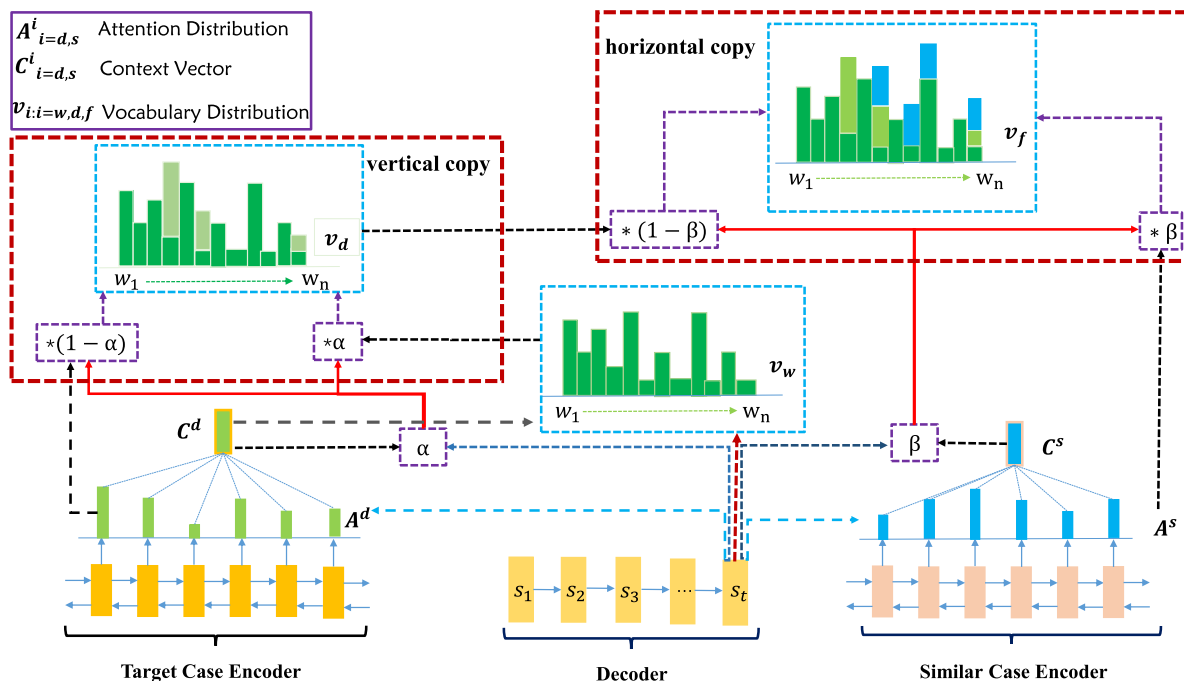


Figure 3: The decoder of CCN learns two pointer distribution to extend the original vocabulary twice. It learns the pointer distribution  $\alpha$  to obtain the content to be copied from the context (vertical copy) as well as the pointer distribution  $\beta$  to determine the content to be copied from its similar cases (horizontal copy).

concatenation of the target case and each candidate retrieved above as a binary classifier, to obtain similarity score.

### 2.2.2 Similar Case Encoding

For SC encoding, we adopt the same method as the target case<sup>5</sup>. We fuse role information with each word of utterance in the SC and then use **Bi LSTM** to obtain hidden state  $\mathbf{h}^s$ . Next, we adopt the attention mechanism to obtain the utterance layer distribution and dialogue layer distribution.

Therefore, we obtain attention distribution for different words  $a_j^{u*}$  of each utterance and different utterances  $a_i^{d*}$  of each SC. Then, we get final attention distribution  $A^s$  which can be expressed as the product of  $a_j^{u*}$  and  $a_i^{d*}$ .

Finally, we use the N-layer transformer-block to obtain the final SC representation  $\mathbf{C}^s$ .

### 2.3 Cross Copy

In this section, we learn two pointers distribution and to achieve internal (vertical) copy and external (horizontal) copy. The decoder's structure shown as Figure 3.

On the time step  $t$ , we concatenate target case context vector  $\mathbf{C}_t^d$  with decoder states  $s_t$  to get the

<sup>5</sup>We use two identical encoders to encode target case and similar case, while two encoders' parameters are not shared

distribution of the current vocabulary:

$$v_w = \text{softmax}(W^v(W^w[\mathbf{C}_t^d, \mathbf{s}_t] + b^w) + b^v) \quad (6)$$

where the  $W^w$ ,  $b^w$ ,  $W^v$  and  $b^v$  are learnable parameters.

For the cross copy, the algorithm execution process is divided into two stages.

At the first stage, we perform **vertical copy**. With the target case encoder hidden state  $\mathbf{h}^d$ , context vector  $\mathbf{C}^d$ , and decoder hidden states  $s$  as mentioned above, on the time step  $t$ , we can learn **vertical copy** probability distribution  $\alpha$ , it determines whether to copy the words from the historical dialogue. It can be expressed as Eq.7:

$$\alpha = \sigma(W^h * \mathbf{h}_t^d + W^c * \mathbf{C}_t^d + W^s * \mathbf{s}_t + b^d) \quad (7)$$

Combined with the attention distribution  $A^d$ , we can get the dynamic extended vocabulary  $v_d$  by pointer distribution  $\alpha$ :

$$v_d = \alpha * v_w + (1 - \alpha) * \sum_{i:w_i=w}^{l*L} A^d \quad (8)$$

In the second stage, we learn the **horizontal copy** probability distribution  $\beta$  of SC. For SC context vector  $\mathbf{C}^s$  and hidden state  $\mathbf{h}^s$ , on the time step

Table 1: Statistics of the CDD and JDDC

Dataset	CCD		JDDC	
	Dialogue	Utterance	Dialogue	Utterance
train	208,152	2,869,794	261,282	3,135,377
dev	26,018	364,345	32,660	391,983
test	26,020	371,554	32,661	391,480
<b>Total</b>	<b>260,190</b>	<b>3,605,693</b>	<b>326,603</b>	<b>3,918,840</b>

$t$ , we combine decoder hidden state  $\mathbf{s}$ , to get the **horizontal copy** pointer distribution  $\beta$ :

$$\beta = \sigma(W^{h*} * \mathbf{h}_t^s + W^{c*} * \mathbf{C}_t^s + W^{s*} * \mathbf{s}_t + b^s) \quad (9)$$

From the encoder, we get the attention distribution  $A^s$  of **SC**, we then perform a second expansion of the dynamic vocabulary to obtain the final vocabulary  $v_f$ :

$$v_f = (1 - \beta) * v_d + \beta * \sum_{i:w_i=w}^{l*L} A^s \quad (10)$$

It should be noted, if  $w$  is not in original vocabulary but in **SC**, the  $\beta$  is 1. Model will copy the  $w$  from **SC**. One of the advantages of our model is that it can produce out-of-vocabulary word.

In the formula above, the  $\sigma$  is sigmoid function. The  $W^h$ ,  $W^c$ ,  $W^s$ ,  $b^d$ ,  $W^{h*}$ ,  $W^{c*}$ ,  $W^{s*}$  and  $b^s$  are learnable parameters.

## 2.4 Loss function

In this dialogue generation task, for each dialogue **D**, the loss function is defined as:

$$\begin{aligned} loss &= -\log P(\mathbf{S} | \mathbf{D}) \\ &= -\sum_{j=1}^l \log P(w_{ij} | w_{i1:j-1}, \mathbf{D}) \end{aligned}$$

Denoting all the parameters in our model as  $\delta$ , then we obtain the following optimized objective function:

$$\min_{\theta} loss = loss + \lambda \|\delta\|_2^2 \quad (11)$$

To minimize the objective function, we use the diagonal variant of Adam (Zeiler, 2012).

## 3 Experimental Settings

### 3.1 Dataset

We employed two datasets for the experiment, Court Debate Dataset (**CDD**) from judicial field and Jing Dong Dialogue Corpus (**JDDC**) (Chen et al., 2019) from e-commerce field. The details

of the dataset are illustrated in Table 1. During training, the data is divided into the training set, development set and test set<sup>6</sup>.

#### 3.1.1 Court Debate Dataset

For **CDD**, we collected 121,016 court debate records of private lending dispute civil cases<sup>7</sup>. We take the judge and the historical conversation with the plaintiff and the defendant as the model input, and the judge’s utterance as the model output. These records are divided into 260,190 pairs of samples by experts with legal knowledge.

#### 3.1.2 Jing Dong Dialogue Corpus

Jing Dong Dialogue Corpus (**JDDC**)<sup>8</sup> contains 1,024,196 multi-turn dialogues, 20,451,337 utterances, and 150 million words. The average number of tokens contained in per sentence is about 7.4. In the experiments, we adopted the top 326,603 cases. The proposed algorithm and baselines are set to generate the utterances of the customer service, and the historical context between the customer service and the customer is set as input.

## 3.2 Evaluation Metrics

We adopt two evaluation methods to validate the proposed model: Automatic Evaluation and Human Evaluation.

### 3.2.1 Automatic Evaluation

To evaluate the effectiveness of the dialogue generated by **CCN**, we used ROUGE (Lin and Hovy, 2003) and BLEU (Papineni et al., 2002) scores to compare different models. We report ROUGE-1, ROUGE-L and BLEU to compare the advantages and disadvantages of each model.

### 3.2.2 Human Evaluation

In order to ensure the rationality/correctness of the generated utterance, we also conducted human evaluation. We randomly selected 300 samples from the test set. Then, we recruited five annotators<sup>9</sup> to judge the quality of generated utterance from two perspectives: (Ke et al., 2018; Zhu et al., 2019):

<sup>6</sup>The entire dataset is divided by a ratio of 8:1:1 for training, developing and testing, respectively.

<sup>7</sup>Private lending dispute cases are the most frequent cause of civil cases in China. This data set is provided by the High People’s Court of a province in China. All the court transcripts are manually recorded by the court clerk.

<sup>8</sup>[http://jddc.jd.com/auth\\_environment](http://jddc.jd.com/auth_environment)

<sup>9</sup>All annotators took basic annotation training before the experiment.



Table 2: Quantitative Evaluation. We report ROUGE-1, ROUGE-L and BLEU scores for each tested methods.

model	CCD			JDDC		
	ROUGE-1	ROUGE-L	BLEU	ROUGE-1	ROUGE-L	BLEU
LSTM(Hochreiter and Schmidhuber, 1997)	30.28	28.02	9.77	19.45	18.74	9.52
ByteNet(Gehring et al., 2017)	33.68	32.99	16.91	22.19	18.35	11.55
ConvS2S(Kalchbrenner et al., 2016)	35.92	31.48	16.34	26.53	21.08	11.64
S2S+attention(Nallapati et al., 2016)	36.91	33.12	18.52	28.44	22.34	13.42
PGN(See et al., 2017)	37.03	34.25	18.75	29.78	24.06	14.37
Transformer(Vaswani et al., 2017)	37.59	34.93	18.58	27.25	22.75	11.29
CCN(vertical-only)	37.15	34.51	19.65	30.27	25.08	15.75
CCN <sub>top-1</sub>	39.12	39.23	23.11	32.44	29.18	16.90
CCN <sub>top-2</sub>	40.43	38.16	23.24	33.56	31.17	18.52
CCN <sub>top-3</sub>	<b>41.10</b>	<b>39.82</b>	<b>24.75</b>	<b>34.17</b>	<b>32.37</b>	<b>19.53</b>

Table 3: Qualitative Evaluation. We report average score (Avg) and calculate the  $\kappa$  value in relevance and fluency. We recruited five annotators to evaluate the sentences generated by all the models. To be fair, for each input, we shuffled the output generated by all the models and then let the annotator to evaluate. The  $\kappa$  represents the consistency of evaluation by different annotators. And the  $\kappa$  coefficient between 0.48 and 0.82 means middle and upper agreement.

model	CCD				JDDC			
	Relevance		Fluency		Relevance		Fluency	
	Avg	$\kappa$	Avg	$\kappa$	Avg	$\kappa$	Avg	$\kappa$
LSTM(Hochreiter and Schmidhuber, 1997)	0.54	0.48	0.93	0.61	0.53	0.52	1.09	0.59
ByteNet(Gehring et al., 2017)	0.63	0.62	1.01	0.71	0.59	0.55	1.19	0.67
ConvS2S (Kalchbrenner et al., 2016)	0.64	0.51	1.05	0.82	0.67	0.71	1.13	0.56
S2S+attention(Nallapati et al., 2016)	0.89	0.55	1.32	0.69	0.88	0.48	1.26	0.57
PGN(See et al., 2017)	1.06	0.64	1.47	0.72	0.96	0.69	1.52	0.53
Transformer(Vaswani et al., 2017)	1.02	0.71	1.41	0.65	0.83	0.56	1.42	0.73
CCN(vertical-only)	1.04	0.58	1.56	0.68	0.95	0.59	1.54	0.66
CCN <sub>top-1</sub>	1.03	0.61	1.59	0.75	0.97	0.68	1.69	0.61
CCN <sub>top-2</sub>	1.10	0.59	1.64	0.72	<b>1.01</b>	0.72	<b>1.77</b>	0.70
CCN <sub>top-3</sub>	<b>1.12</b>	0.66	<b>1.69</b>	0.68	0.98	0.63	1.73	0.67

- **Relevance:** Generated utterance is logically relevant to the dialogue context and can provide meaningful information.
- **Fluency:** Generated utterance is fluent and grammatical.

The information on these two aspects are independently evaluated. For each aspect, we set three levels with scores: +2, +1, 0, in which 2 stands for excellent: for relevance, closely related to historical dialogue and be meaningful; for fluency, it has strong readability without grammatical error. 1 stands for good: for relevance, with some off-topic information; for fluency, sentence is readable, but with slight grammatical error. 0 means poor: for relevance, the sentence is off-topic or meaningless; for fluency, sentence has poor readability or serious grammatical errors. Finally, we obtain the weighted average score and kappa ( $\kappa$ ) of each model to compare the effect of the model.

### 3.3 Training Details

During the training process, we set the dimension of word embedding as 300 and use word2vec to build the initialization word vector. The dimension

of role embedding is set to 100 with random initialization. The hidden size is set to 300, we use 4 layer Transformer, where the number of heads equals to 8. The dropout probability is set to 0.8. Based on these settings, we optimize the objective function with a learning rate of  $5e-4$ . We perform the mini-batch gradient descent with a batch size of 64. We set maximum utterance length as 40 in decoder during generation (the generated sentence might contain sub-utterances).

## 4 Result discussion

### 4.1 Overall Performance

In the experiments, we select up to three similar cases to validate the effectiveness of the CCN, i.e., leveraging the most similar case (**top-1**), top two similar ones (**top-2**), and top three similar ones (**top-3**). In addition, we also test the variant of **CCN(vertical-only)** by only adopting vertical copy from the context, which is similar to the setting of the baseline **PGN** but with the proposed hierarchical dialogue encoders.

The performance of all the tested methods are reported in Table 2 and Table 3 in terms of quanti-

Case 1		Role	Utterance (context)	Role	Utterance (SC)
Input	Judge:	.....	You have just stated that the defendant has returned interest. Do you have a statement to the court?	Judge:	.....
	Plaintiff:	.....	At the beginning, the defendant <personname> paid monthly interest of 4,500 yuan on time, and the interest was paid until the end of 2014.	Plaintiff:	About 110,000 yuan of interest was paid.
Output	Groundtruth		<b>How did the defendant pay interest?</b>	S2S+attention	Does the plaintiff add ?
	CCN(top-1)		How did the defendant pay interest?	ConvS2S	Does the defendant have any additional comments ?
	Transformer		Does the plaintiff supplement the facts in this case ?	ByteNet	Proof by the defendant.
	PGN		How much interest paid ?	LSTM	Does the plaintiff add ?
	Judge:			Judge:	<b>How did the defendant pay interest ?</b>
Case 2		Role	Utterance (context)	Role	Utterance (SC)
Input	Judge:	.....	Does the RMB "50,000" stated on the debit note and receipt contain interest on the loan? Is it delivered ?	Judge:	.....
	Plaintiff:	.....	It is the principal of the loan and does not include interest. Full cash has been delivered to the defendant <personname> .	Plaintiff:	No
Output	Groundtruth		<b>Has the defendant &lt;personname&gt; returned the principal and interest after borrowing?</b>	S2S+attention	How much principal did the defendant borrow?
	CCN(top-1)		Whether the defendant <personname> return the principal and pay interest ?	ConvS2S	Does the defendant provide evidence ?
	Transformer		How interest is returned ?	ByteNet	The plaintiff continued to provide evidence ?
	PGN		How much interest did the defendant return?	LSTM	Does the plaintiff add ?
	Judge:			Judge:	After the loan expires, does the defendant <b>return the principal and pay interest ?</b>

Figure 4: Case Study. We show two examples to intuitively illustrate the performance of all the tested models. The content in SC is the top 1 similar case used by our proposed method CCN. Note that the color indicates the BLEU score of each generated output. The darker the color, the higher the BLEU score is.

tative and qualitative evaluation, respectively. As Table 2 shows, the proposed approach CCN with its variants outperform all the baselines in Rouge and Bleu metrics over the two datasets. We can also observe the increasing performance as the number of referred similar cases increases. As for the two qualitative criteria, CCN also shows better performance by a big margin compared to the baselines. Note that the Kappa value ( $\kappa$ ) indicates the agreement among the annotators.

As mentioned above, the increasing number of referred similar cases enables to bring about the improvement of performance, which demonstrates that the horizontal copy plays a critical role in dialogue generation without employing any external resources. However, in the training process, as the number of similar cases increases, the training speed is getting slower. Considering the time cost and memory limitation, only up to top three similar cases are utilized in this experiment to verify the proposed approach.

## 4.2 Case study

Figure 4 shows two examples to illustrate the performance of different tested methods. As depicted in case 1, comparing with the baselines, the CCN can learn dialogue logic from SCs and accurately

locate the sentence to complete the horizontal copy. Another important finding is that we can use SCs to obtain more accurate representation information. It can identify specific entities from the context for vertical copy while capturing the discourse patterns from the similar cases for horizontal copy to finally synthesize the sentence to be generated.

On the other hand, the baseline models are more inclined to generate general expressions which appear more frequent in the training data without much attention to the specific information in the context and the logical discourse patterns appearing in certain circumstances.

## 4.3 Error analysis

In order to explore the algorithm limitation and model capability boundary, we summarize the samples with high error rate. The following observations should be highlighted to scope the limitation of current model and enlighten future investigation for this track of research.

In the CDD, 53% of errors<sup>10</sup> occur when generated sentence contains the information that does not appear in context or in similar cases (e.g., “According to the provisions of Articles 44 and 45 of

<sup>10</sup>the error refers to the generated text whose either relevance or fluency score equals 0.

the Civil Procedure Law of the People’s Republic of China, if the parties find that the members of the collegiate bench and the clerk experienced any of the following circumstances, they have the right to apply for their evasion orally or in writing.”). Similarly in **JDDC**, such problem caused 47% of errors (e.g., “Sorry, we cannot refund you with the product [#price, #style, #brand, #specification, #color] you required.”). Such kind of law/product related information might need to leverage external expert knowledge. In addition, 23% and 36% of errors occur in **CDD** and **JDDC** respectively, when it comes to the long sentence to be generated (e.g., the sentence length is more than 30 words for the judge’s inquiry or for customer service response).

To address these problems in the future research, enhancing the long dependence of language models and establishing relations between different entities can be promising approaches.

## 5 Related work

### 5.1 Pointer Network

Pointer network (Vinyals et al., 2015) is a special network structure. It solved the problem of generating sequence depending on the input sequence. Based on this basis, CopyNet (Gu et al., 2016) and PGN (See et al., 2017) were proposed, which can copy the words in the context and form output sequence to cope with Out-Of-Vocabulary (OOV) problem. Nowadays, pointer networks are increasingly popular in NLP applications. In text summarization, Miao and Blunsom (2016) used it to select only suitable words from context instead of the entire dictionary for sentence compression; Nallapati et al. (2016) used it to speed up model convergence; Sun et al. (2018) used it to generate text title; Wang et al. (2019a) generated new conceptual words; Eric and Manning (2017) used it to develop a recurrent neural dialogue system; Shen et al. (2019) used it to enhance power of capturing richer latent alignment. It was also widely used in many other tasks, such as dependency parsing (Fernández-González, 2019; Liu et al., 2019a), question answering (Kadlec et al., 2016; Golchha et al., 2019), machine reading comprehension (Wang and Jiang; Wang et al., 2017), machine translation (Gulcehre et al., 2016), and language models (Merity et al., 2016).

Unlike previous studies, on the basis of internal copy, we introduce external copy, to establish a cross-copy structure, and achieve significant im-

provement.

### 5.2 Dialogue System

As an important task of NLP, dialogue system, has achieved great success and is widely used in practical applications, including customer service systems and chatbots. In recent years, with the development of deep learning technology, the neural network model has made significant progress: Su et al. (2019) solved the problem of information omitting and quoting in multiple rounds of dialogue by rewriting sentences; Lu et al. (2019) solved the problem of selecting the reply sentence in the dialogue system by adding features of time sequence and space; Du and Black (2019) solved the problem of lack of diversity in replies by using the dichotomy function to judge whether the two responses are similar.

There exist a number of prior studies to assist the task of dialogue generation through external knowledge: Wu et al. (2019); Liu et al. (2018) leveraged the structured knowledge triples to assist dialogue generation. Similarly, Li et al. (2019); Rajpurkar et al. (2018); Huang et al. (2018); Reddy et al. (2019) used documents as knowledge discovery for dialogue generation, and Xia et al. (2017); Ye et al. (2020); Ghazvininejad et al. (2018); Parthasarathi and Pineau (2018) utilized unstructured knowledge to explore in the open-domain dialogue generation.

With the deepening of dialogue generation, various new tasks have been proposed: Le et al. (2019) generated the most appropriate response by given video content, video title, and existing dialogue sentences; Tang et al. (2019) introduced how to lead the conversation to a specific goal in an open conversation; Wang et al. (2019b) introduced how to use different persuasion strategies in the dialogue to persuade people to donate to charities; Cao et al. (2019) were concerned about the application of dialogue analysis in the psychotherapy.

In our model, the CCN approaches to solve the problem of defective domain adaptation without any costly external knowledge.

## 6 Conclusion and Outlook

In this paper, we proposed a novel neural network structure-Cross Copy Networks, enabling both vertical copy (from dialogue context) and horizontal copy (from similar cases). Unlike prior models, the proposed CCN doesn’t need additional knowledge input, and it can be easily adopted to other



domains. We conduct experiments on two different datasets with both quantitative and human evaluation to validate the proposed model. Experimental results proved CCN’s superiority when comparing with a number of existing state-of-art text generation models, which tells the cross copy mechanism can successfully enhance the dialogue generation performance.

In future work, we will further investigate other content generation problems by leveraging multi-granularity copying mechanism. This study serves as the methodological foundation.

## 7 Acknowledgments

This work is supported by National Key R&D Program of China (2018YFC0830200; 2018YFC0830206; 2018YFC0830700).

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Jie Cao, Michael Tanana, Zac Imel, Eric Poitras, David Atkins, and Vivek Srikumar. 2019. *Observing dialogue in therapy: Categorizing and forecasting behavioral codes*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5599–5611, Florence, Italy. Association for Computational Linguistics.
- Meng Chen, Ruixue Liu, Lei Shen, Shaozu Yuan, Jingyan Zhou, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2019. The jddc corpus: A large-scale multi-turn chinese dialogue dataset for e-commerce customer service. *arXiv preprint arXiv:1911.09969*.
- Wenchao Du and Alan W Black. 2019. *Boosting dialog response generation*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 38–43, Florence, Italy. Association for Computational Linguistics.
- Mihail Eric and Christopher D Manning. 2017. A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue. *arXiv preprint arXiv:1701.04024*.
- Daniel Fernández-González. 2019. Left-to-right dependency parsing with pointer networks. In *Proceedings of NAACL-HLT*, pages 710–716.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. pages 1243–1252. JMLR. org.
- 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 *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Hitesh Golchha, Mauajama Firdaus, Asif Ekbal, and Pushpak Bhattacharyya. 2019. Courteously yours: Inducing courteous behavior in customer care responses using reinforced pointer generator network. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 851–860.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. *arXiv preprint arXiv:1603.06393*.
- Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the unknown words. *arXiv preprint arXiv:1603.08148*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih. 2018. Flowqa: Grasping flow in history for conversational machine comprehension. *arXiv preprint arXiv:1810.06683*.
- Rudolf Kadlec, Martin Schmid, Ondrej Bajgar, and Jan Kleindienst. 2016. Text understanding with the attention sum reader network. *arXiv preprint arXiv:1603.01547*.
- Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aaron van den Oord, Alex Graves, and Koray Kavukcuoglu. 2016. Neural machine translation in linear time. *arXiv preprint arXiv:1610.10099*.
- Pei Ke, Jian Guan, Minlie Huang, and Xiaoyan Zhu. 2018. Generating informative responses with controlled sentence function. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1499–1508.
- Hung Le, Doyen Sahoo, Nancy Chen, and Steven Hoi. 2019. *Multimodal transformer networks for end-to-end video-grounded dialogue systems*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5612–5623, Florence, Italy. 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. *arXiv preprint arXiv:1907.08854*.

- Chin-Yew Lin and Eduard Hovy. 2003. Automatic evaluation of summaries using n-gram co-occurrence statistics.
- Linlin Liu, Xiang Lin, Shafiq Joty, Simeng Han, and Lidong Bing. 2019a. [Hierarchical pointer net parsing](#). pages 1007–1017, Hong Kong, China. 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.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Junyu Lu, Chenbin Zhang, Zeying Xie, Guang Ling, Tom Chao Zhou, and Zenglin Xu. 2019. [Constructing interpretable spatio-temporal features for multi-turn responses selection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 44–50, Florence, Italy. Association for Computational Linguistics.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Yishu Miao and Phil Blunsom. 2016. Language as a latent variable: Discrete generative models for sentence compression. *arXiv preprint arXiv:1609.07317*.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. pages 311–318. Association for Computational Linguistics.
- Prasanna Parthasarathi and Joelle Pineau. 2018. Extending neural generative conversational model using external knowledge sources. *arXiv preprint arXiv:1809.05524*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Xiaoyu Shen, Yang Zhao, Hui Su, and Dietrich Klakow. 2019. [Improving latent alignment in text summarization by generalizing the pointer generator](#). pages 3760–3771, Hong Kong, China. Association for Computational Linguistics.
- Hui Su, Xiaoyu Shen, Rongzhi Zhang, Fei Sun, Pengwei Hu, Cheng Niu, and Jie Zhou. 2019. [Improving multi-turn dialogue modelling with utterance ReWriter](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 22–31, Florence, Italy. Association for Computational Linguistics.
- Fei Sun, Peng Jiang, Hanxiao Sun, Changhua Pei, Wenwu Ou, and Xiaobo Wang. 2018. Multi-source pointer network for product title summarization. pages 7–16. ACM.
- I Sutskever, O Vinyals, and QV Le. 2014. Sequence to sequence learning with neural networks. *Advances in NIPS*.
- Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric Xing, and Zhiting Hu. 2019. [Target-guided open-domain conversation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5624–5634, Florence, Italy. 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. pages 5998–6008.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. pages 2692–2700.
- Shuohang Wang and Jing Jiang. Machine comprehension using match- lstm and answer pointer.(2017). pages 1–15.
- Wenbo Wang, Yang Gao, Heyan Huang, and Yuxiang Zhou. 2019a. [Concept pointer network for abstractive summarization](#). pages 3074–3083, Hong Kong, China. Association for Computational Linguistics.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. pages 189–198.
- Xuwei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019b. [Persuasion for good: Towards a personalized persuasive dialogue system for social good](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. 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 goals. *arXiv preprint arXiv:1906.05572*.

- Yingce Xia, Fei Tian, Lijun Wu, Jianxin Lin, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2017. Deliberation networks: Sequence generation beyond one-pass decoding. In *Advances in Neural Information Processing Systems*, pages 1784–1794.
- Hao-Tong Ye, Kai-Lin Lo, Shang-Yu Su, and Yun-Nung Chen. 2020. Knowledge-grounded response generation with deep attentional latent-variable model. *Computer Speech & Language*, page 101069.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Qingfu Zhu, Lei Cui, Weinan Zhang, Furu Wei, and Ting Liu. 2019. Retrieval-enhanced adversarial training for neural response generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3763–3773.