

Transformer-GCRF: Recovering Chinese Dropped Pronouns with General Conditional Random Fields

Jingxuan Yang¹, Kerui Xu¹, Jun Xu^{2,3*}, Si Li^{1*}, Sheng Gao¹, Jun Guo¹,
Ji-Rong Wen^{2,3} and Nianwen Xue⁴

¹School of Artificial Intelligence, Beijing University of Posts and Telecommunications

²Gaoling School of Artificial Intelligence, Renmin University of China

³Beijing Key Laboratory of Big Data Management and Analysis Methods

⁴Department of Computer Science, Brandeis University

{yjx, xukerui, lisi, gaosheng, guojun}@bupt.edu.cn

junxu@ruc.edu.cn, jirong.wen@gmail.com, xuen@brandeis.edu

Abstract

Pronouns are often dropped in Chinese conversations and recovering the dropped pronouns is important for NLP applications such as Machine Translation. Existing approaches usually formulate this as a sequence labeling task of predicting whether there is a dropped pronoun before each token and its type. Each utterance is considered to be a sequence and labeled independently. Although these approaches have shown promise, labeling each utterance independently ignores the dependencies between pronouns in neighboring utterances. Modeling these dependencies is critical to improving the performance of dropped pronoun recovery. In this paper, we present a novel framework that combines the strength of Transformer network with General Conditional Random Fields (GCRF) to model the dependencies between pronouns in neighboring utterances. Results on three Chinese conversation datasets show that the Transformer-GCRF model outperforms the state-of-the-art dropped pronoun recovery models. Exploratory analysis also demonstrates that the GCRF did help to capture the dependencies between pronouns in neighboring utterances, thus contributes to the performance improvements.

1 Introduction

In pro-drop languages such as Chinese, pronouns can be dropped as the identity of the pronoun can be inferred from the context, and this happens more frequently in conversations (Yang et al., 2015). Recovering dropped pronouns (DPs) is a critical task for many NLP applications such as Machine Translation where the dropped pronouns need to be translated explicitly in the target language (Wang et al., 2016a,b, 2018). Recovering dropped pronoun is different from traditional pronoun resolution tasks (Zhao and Ng, 2007; Yin et al., 2017,

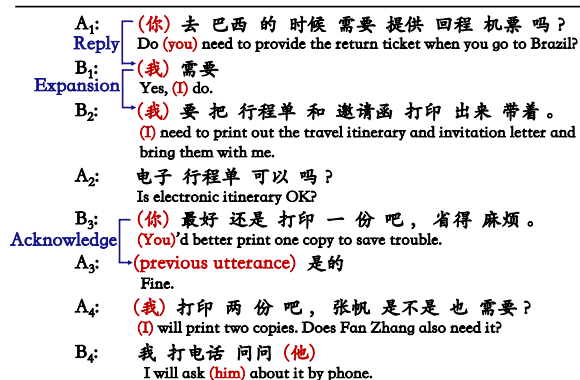


Figure 1: A conversation snippet between participant A and B. The dropped pronouns are shown in the brackets, and the dialogue patterns are marked with blue arrows.

2018), which aim to resolve the anaphoric pronouns to their antecedents. In dropped pronoun recovery, we consider both anaphoric and non-anaphoric pronouns, and we do not directly resolve the dropped pronoun to its antecedent, which is infeasible for non-anaphoric pronouns. We recover the dropped pronoun as one of 17 types pronouns pre-defined in (Yang et al., 2015), which include five types of abstract pronouns corresponding to the non-anaphoric pronouns. Thus traditional rule-based pronoun resolution methods are not suitable for recovering dropped pronouns.

Existing approaches formulate dropped pronoun recovery as a sequence labeling task of predicting whether a pronoun has been dropped before each token and the type of the dropped pronoun. For example, Yang et al. (2015) first studied this problem in SMS data and utilized a Maximum Entropy classifier to recover dropped pronouns. Deep neural networks such as Multi-Layer Perceptrons (MLPs) and structured attention networks have also been used to tackle this problem (Zhang et al., 2016; Yang et al., 2019). Giannella et al. (2017) used a linear-chain CRF to model the dependency be-

* Corresponding author

tween the sequence of predictions in a utterance.

Although these models have achieved various degrees of success, they all assume that each utterance in a conversation should be labeled independently. This practice overlooks the dependencies between dropped pronouns in neighboring utterances, and results in sequences of predicted dropped pronouns are incompatible with one another. We illustrate this problem through an example in Figure 1, in which the dropped pronouns are shown in brackets. The pronoun can be dropped as a subject at the beginning of a utterance, or as an object in the middle of a utterance. Pronouns dropped at the beginning of consecutive utterances usually have strong dependencies that pattern with three types of dialogue transitions (i.e., Reply, Expansion and Acknowledgment) presented in (Xue et al., 2016). For example, in Figure 1, the pronoun in the second utterance B_1 is “我 (I)”, the dropped pronoun in the third utterance B_2 should also be “我 (I)” since B_2 is an *expansion* of B_1 by the same speaker. Thus modeling the dependency between pronouns in adjacent sentences is helpful to recover pronoun dropped at utterance-initial positions. In contrast, the pronoun “他 (him)” dropped as an object in utterance B_4 should be recovered by capturing referent semantics from the context and modeling token dependencies in the same utterance.

To model the dependencies between predictions in the conversation snippet, we propose a novel framework called Transformer-GCRF that combines the strength of the Transformer model (Vaswani et al., 2017) in representation learning and the capacity of general Conditional Random Fields (GCRF) to model the dependencies between predictions. In the GCRF, a vertical chain is designed to capture the pronoun dependencies between the neighboring utterances, and horizontal chains are used for modeling the prediction dependencies inside each utterance. In this way, Transformer-GCRF successfully models the cross-utterance pronoun dependencies as well as the intra-utterance prediction dependencies simultaneously. Experimental results on three conversation datasets show that Transformer-GCRF significantly outperforms the state-of-the-art recovery models. We also conduct ablative experiments that demonstrate the improvement in performance of our Transformer-GCRF model derives both from the Transformer encoder and the ability of GCRF layer to model the dependencies

between dropped pronouns in neighboring utterances. All code is available at <https://github.com/ningningyang/Transformer-GCRF>.

The major contributions of the paper are summarized as follows:

- We conduct statistical study on pronouns dropped at the beginning of consecutive utterances in conversational corpus, and observe that modeling the dependencies between pronouns in neighboring utterances is important to improve the performance of dropped pronoun recovery.
- We propose a novel Transformer-GCRF approach to model both intra-utterance dependencies between predictions in a utterance and cross-utterance dependencies between dropped pronouns in neighboring utterance. The model jointly predicts all dropped pronouns in an entire conversation snippet.
- We apply the Transformer-GCRF model on three conversation datasets. Results show that our Transformer-GCRF outperforms the baseline models on all datasets. Exploratory experiments also show that the improvement is attributed to the capacity of the model to capture cross-utterance dependencies.

2 Related Work

2.1 Dropped pronoun recovery

As pronouns are frequently dropped in informal genres, Yang et al. (2015) first introduced dropped pronoun recovery as an independent task and used a Maximum Entropy classifier to recover DPs in text messages. Giannella et al. (2017) employed a linear-chain CRF to jointly predict the position, person, and number of the dropped pronouns in a single utterance, to exploit the sequential nature of this problem. With the powerful representation capability of neural network (Xu et al., 2020), Zhang et al. (2016) introduced a MLP neural network to recover the dropped pronouns based on the concatenation of word embeddings within a fixed-length window. Yang et al. (2019) proposed a neural network with structured attention to model the interaction between dropped pronouns and their referents using both sentence-level and word-level context, and again each dropped pronoun is predicted independently. Tong et al. (2019) further incorporated specific external knowledge to identify the referent more accurately. None of these methods consider

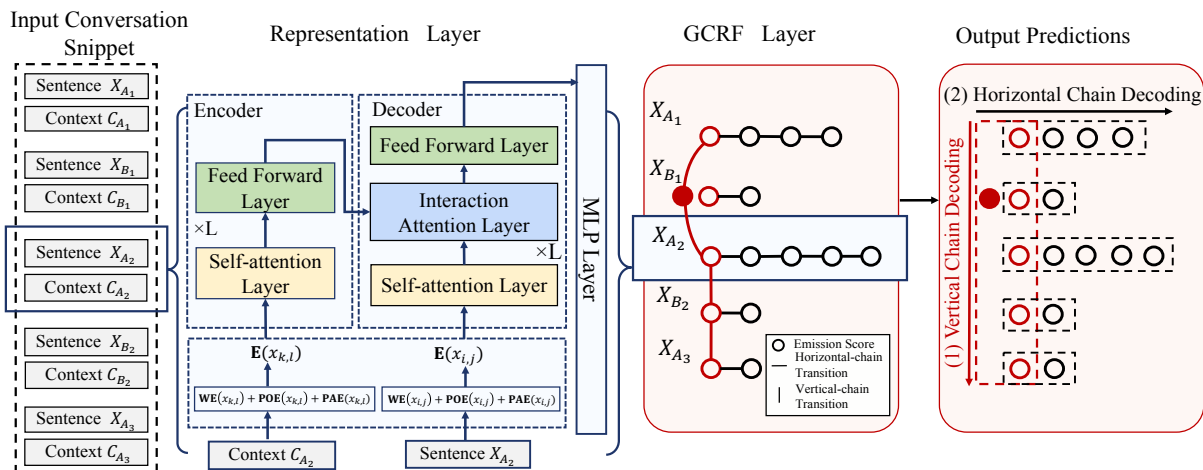


Figure 2: Overall architecture of our Transformer-GCRF model.

the dependencies between pronouns in neighboring utterances.

2.2 Zero pronoun resolution

Zero pronoun resolution (Zhao and Ng, 2007; Kong and Zhou, 2010; Chen and Ng, 2016; Yin et al., 2017, 2018) is a line of research closely related to dropped pronoun recovery. The difference between these two tasks is that zero pronoun resolution focuses on resolving anaphoric pronouns to their antecedents assuming the position of the dropped pronoun is already known. However, in dropped pronoun recovery, we consider both anaphoric and non-anaphoric pronouns, and attempt to recover the type of dropped pronoun but not its referent. Su et al. (2019) also presented a new utterance rewriting task which improves the multi-turn dialogue modeling through recovering missing information with coreference.

2.3 Conditional random fields

Conditional Random Fields (CRFs) are commonly used in sequence labeling. It models the conditional probability of a label sequence given a corresponding sequence of observations. Lafferty et al. (2001) made a first-order Markov assumption among labels and proposed a linear-chain structure that can be decoded efficiently with the Viterbi algorithm. Sutton et al. (2004) introduced dynamic CRFs to model the interactions between two tasks and jointly solve the two tasks when they are conditioned on the same observation. Zhu et al. (2005) introduced two-dimensional CRFs to model the dependency between neighborhoods on a 2D grid to extract object information from the web. Sut-

ton et al. (2012) also explored how to generalize linear-chain CRFs to general graphs. CRFs have also been combined with powerful neural networks to tackle sequence labeling problems in NLP tasks such as POS tagging and Named Entity Recognition (NER) (Lample et al., 2016; Ma and Hovy, 2016; Liu et al., 2018), but existing research has not explored how to combine deep neural networks with general CRFs.

3 Our Approach: Transformer-GCRF

We start by formalizing the dropped pronoun recovery task as follows. Given a Chinese conversation snippet $X = (x_1, \dots, x_n)$ which consists of n pro-drop utterances, where the i -th utterance $x_i = (x_{i1}, \dots, x_{im_i})$ is a sequence of m_i tokens, and additionally given a set of k possible labels $\mathcal{Y} = \{y_1, \dots, y_{k-1}\} \cup \{\text{None}\}$ where each y_j corresponds to a pre-defined pronoun (Yang et al., 2015) or ‘None’, which means no pronoun is dropped, the goal of our task is to assign a label $y \in \mathcal{Y}$ to each token in X to indicate whether a pronoun is dropped before this token and the type of pronoun. We model this task as the problem of maximizing the conditional probability $p(Y|X)$, where Y is the label sequence assigned to the tokens in X . The conditional probability of a label assignment Y given the whole conversation snippet X can be written as:

$$p(Y|X) = \frac{e^{s(X,Y)}}{\sum_{\tilde{Y} \in \mathcal{Y}_X} e^{s(X,\tilde{Y})}},$$

where $s(X, Y)$ denotes score of the sequences of predictions in the conversation snippet. The denominator is known as partition function, and \mathcal{Y}_X

contains all possible tag sequences for the conversation snippet X .

3.1 Overview of Transformer-GCRF

We score each pair of (X, Y) with our proposed Transformer-GCRF, as shown in Figure 2. When pre-processing the inputs, we attach a context to each pro-drop utterance \mathbf{x}_n in the snippet X . The context $C_n = \{\mathbf{x}_{n-5}, \dots, \mathbf{x}_{n-1}, \mathbf{x}_{n+1}, \mathbf{x}_{n+2}\}$ consists of the previous five utterances as well as the next two utterances following the practices in (Yang et al., 2019), and provides referent related contextual information to help recover the dropped pronouns. The representation layer uses the Transformer structure to encode the context C_n and generates representations for tokens in utterance \mathbf{x}_n from the decoder. The prediction layer then utilizes a generalized CRF to model the cross-utterance and inter-utterance dependencies between the predictions in the conversation snippet, and outputs the predicted sequence for tokens in the snippet.

3.2 Representation layer

We employ the encoder-decoder structure of Transformer (Vaswani et al., 2017) to generate the representations for the tokens in pro-drop utterance \mathbf{x}_i and context C_i separately.

3.2.1 Context encoder

The context encoder first unfolds all tokens in the context C_i into a linear sequence as: $(x_{i-5,1}, x_{i-5,2}, \dots, x_{i+2,m_{i+2}})$, and then inserts the delimiter ‘[SEP]’ between each pair of utterances. Following the Transformer model (Vaswani et al., 2017), the input embedding of each token $x_{k,l}$ is the sum of its word embedding $\mathbf{WE}(x_{k,l})$, position embedding $\mathbf{POE}(x_{k,l})$, and speaker embedding $\mathbf{PAE}(x_{k,l})$ as:

$$\mathbf{E}(x_{k,l}) = \mathbf{WE}(x_{k,l}) + \mathbf{POE}(x_{k,l}) + \mathbf{PAE}(x_{k,l}).$$

The token embeddings $\mathbf{E}(x_{k,l})$ are then fed into the encoder, which is a stack of L encoding blocks. Each block contains two sub-layers (i.e., a self-attention layer and a feed-forward layer) as:

$$\mathbf{H}^{(l)} = \text{FNN}(\text{SelfATT}(\mathbf{H}_Q^{(l-1)}, \mathbf{H}_K^{(l-1)}, \mathbf{H}_V^{(l-1)})), \quad (1)$$

for $l = 1, \dots, L$, where ‘FNN’ and ‘SelfATT’ denotes the feed-forward and self-attention networks respectively, and

$$\mathbf{H}^{(0)} = [\mathbf{E}(x_{i-2,1}), \mathbf{E}(x_{i-2,2}), \dots, \mathbf{E}(x_{i+1,m_{i+1}})].$$

In Equation 1, the self-attention layer first projects the input as a query matrix ($\mathbf{H}_Q^{(l-1)}$), a key matrix ($\mathbf{H}_K^{(l-1)}$), and a value matrix ($\mathbf{H}_V^{(l-1)}$). A multi-head attention mechanism is then applied to these three matrices to encode the input tokens in the context.

3.2.2 Utterance decoder

To generate the representations for tokens in the pro-drop utterance \mathbf{x}_i and exploit referent information from its context C_i , we utilize the decoder component of the Transformer to represent \mathbf{x}_i . Similar to the context encoder, the inputs to the utterance decoder are the embeddings of the tokens. Each embedding $\mathbf{E}(x_{i,j})$ is also a sum of its word embedding, position embedding, and speaker embedding. Then, the input to the decoder, denoted as $\mathbf{S}^{(0)}$, is a concatenation of all the token embeddings:

$$\mathbf{S}_i^{(0)} = [\mathbf{E}(x_{i,1}), \mathbf{E}(x_{i,2}), \dots, \mathbf{E}(x_{i,m_i})].$$

The decoder is still a stack of L decoding blocks. Each decoding block $\text{Dec}(\cdot)$ contains three sub-layers (i.e., a self-attention layer, an interaction attention layer, and a feed-forward layer) as:

$$\begin{aligned} \mathbf{S}_i^{(l)} &= \text{Dec}(\mathbf{S}_i^{(l-1)}, \mathbf{H}_i^{(L)}) \\ &= \text{FFN}(\text{InterATT}(\text{SelfATT}(\mathbf{S}_i^{(l-1)}), \mathbf{H}_i^{(L)})), \end{aligned}$$

for $l = 1, \dots, L$, where FFN is a feed-forward network, SelfATT is a self-attention network.

Finally, the output states of the decoder $\mathbf{S}^{(L)}$ are transformed into logits through a two-layer MLP network as:

$$\mathbf{P} = \mathbf{W}_1 \cdot \tanh(\mathbf{W}_2 \cdot \mathbf{S}^{(L)} + b_2) + b_1, \quad (2)$$

where the logits matrix \mathbf{P} of size $n \times m \times k$ will be fed into a subsequent prediction layer. k is the number of distinct tags, and each element $P_{i,j,l}$ refers to the emission score of the l -th tag of the j -th word in the i -th utterance.

3.3 GCRF layer

We utilize an elaborately designed general conditional random fields (GCRF) layer to recover dropped pronouns by modeling cross-utterance and intra-utterance dependencies between dropped pronouns.

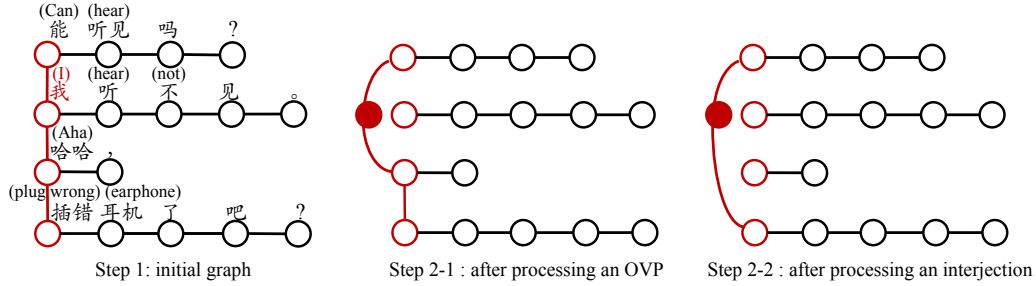


Figure 3: The GCRF graph construction. Step 1 constructs a initial graph. The tokens in each utterance are shown and the nodes corresponding to the first token in each utterance are highlighted in red; step 2-1 processes an OVP (in the second utterance) and adds an observed (shaded) node for token “我/(I)”; step 2-2 processes an interjection (in the third utterance) and skips the node corresponding to the token “哈哈/(Aha)”.

3.3.1 Graph construction in GCRF

Given a conversation snippet, a graph is constructed where each node, corresponding to a token, is a random variable y that represents the type of the pronoun defined in \mathcal{Y} . The edges in the graph are defined by the following two steps:

Step 1: Initial graph construction: We first split each compound utterance into several simple utterances by punctuation, and connect the nodes corresponding to the tokens in the same simple utterance with horizontal edges to model intra-utterance dependencies. Then we link the first tokens in consecutive utterances with a vertical chain to model the cross-utterance dependencies. Step 1 in Figure 3 shows an initial graph for a conversation snippet.

Step 2: Vertical edge refinement: Though the vertical chain constructed in Step 1 can capture most of the cross-utterance dependencies, they can be further refined considering the following two general cases in conversation:

- **Overt pronouns (OVP):** If an OVP appears as the first token in a utterance, it is clear that there is a dependency between the OVP and the dropped pronoun in neighboring utterances. To model this phenomenon, an observed node (with the value of its pronoun type) is inserted in the graph, and the vertical chain linked to the original node is moved to this new node. Step 2-1 in Figure 3 shows the refined graph after OVPs are processed.
- **Interjections:** If the first token in an utterance is an interjection (e.g., “嗯/ Well”, “哈哈/ Aha” etc.), it is better to skip the utterance in the vertical chain because the short utterance consisting of only interjections and punctuation does not provide useful information about the dependencies between pronouns.

Step 2-2 in Figure 3 shows the refined graph after interjections are processed.

3.3.2 Pronoun prediction

It is obvious that the GCRF is a special case of the 2D CRFs. To predict the labels of the nodes following the practices in (Zhu et al., 2005), we employ a modified Viterbi algorithm in which the nodes in the vertical chain are decoded first. Specifically, the constructed graph consists of two types of cliques: one from the horizontal chains and the other from the vertical chain. Given the emission score matrix \mathbf{P} outputted from the decoder layer (see Section 3.2.1), the joint score $s(X, Y)$ of the predictions can be computed by first computing the sum of horizontal chains and then summing up scores of the transitions in the vertical chain as:

$$s(X, Y) = \sum_{i=1}^n s_{hi} + \sum_{i=1}^{n-1} A_{T_i, T_{i+1}}^{(2)}, \quad (3)$$

$$s_{hi} = \sum_{j=1}^{m-1} A_{y_{i,j}, y_{i,j+1}}^{(1)} + \sum_{j=1}^m P_{i,j, y_{i,j}},$$

where $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$ are the transition matrices of the horizontal chains and the vertical chain, respectively; $A_{i,j}$ indicates the transition score from tag i to tag j ; and the node T_i is defined as,

$$T_i = \begin{cases} y_{\text{OVP}} & \text{if the node is an observed OVP} \\ y_{i,1} & \text{otherwise} \end{cases}$$

where $y_{\text{OVP}} \in \mathcal{Y}$ is the observed label corresponds to the specific OVP. The first term in Eq. (3) is the score corresponding to the horizontal chain cliques, and the second term corresponds to the vertical chain clique.

Algorithm 1 Transformer-GCRF Decoding.

Input: The emission score matrix \mathbf{P} ; Transition matrices $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$.

Output: The best path Y^*

```
1: for  $i = 1, \dots, n$  do
2:    $s_{hi}, \mathbf{bp}_i \leftarrow \text{ForwardScore}(\mathbf{P}_i, \mathbf{A}^{(1)})$ 
3: end for
4:  $\mathbf{P}_h = [s_{h1}, s_{h2}, \dots, s_{hn}]$ 
5:  $s(X, Y), \mathbf{bp}_v \leftarrow \text{ForwardScore}(\mathbf{P}_h, \mathbf{A}^{(2)})$ 
6:  $Y_{n,1}^* \leftarrow \arg \max (s(X, Y))$ 
7:  $\{Y_{1,1}^*, \dots, Y_{n-1,1}^*\} \leftarrow \text{TraceBack}(Y_{n,1}^*, \mathbf{bp}_v)$ 
8: for  $i = 1, \dots, n$  do
9:    $\{Y_{i,2}^*, \dots, Y_{i,m_i}^*\} \leftarrow \text{TraceBack}(Y_{i,1}^*, \mathbf{bp}_i)$ 
10: end for
11: function TRACEBACK( $y, \mathbf{bp}$ )
12:    $t = \text{length}(\mathbf{bp})$ 
13:    $z_1 \leftarrow y$ 
14:   for  $j = 2, \dots, t$  do
15:      $z_j \leftarrow \mathbf{bp}_{j,z_{j-1}}$ 
16:   end for
17:   return  $\{z_2, \dots, z_t\}$ 
18: end function
19: return  $Y^*$ 
```

3.4 Decoding the GCRF and Model training

The sequence that maximizes the conditional probability $p(Y|X)$ is outputted as the prediction:

$$Y^* = \arg \max_{Y \in \mathcal{Y}_X} p(Y|X).$$

A modified Viterbi algorithm is used to find the best labeling sequence. Specifically, we first apply the Viterbi algorithm to decode the vertical chain. Then, the vertical chain decoding results are used as the observed nodes in the graph, and the standard Viterbi algorithm is applied to each horizontal chain in parallel. Algorithm 1 shows the Transformer-GCRF decoding process.

Given a set of labeled conversation snippets \mathcal{D} , the model parameters are learned by jointly maximizing the overall log-probabilities of the ground-truth label sequences:

$$\max \sum_{(X,Y) \in \mathcal{D}} \log(p(Y|X)).$$

4 Datasets and Experimental Setup

Datasets: We evaluate the performance of Transformer-GCRF on three conversation benchmarks: Chinese text message dataset (SMS), OntoNotes Release 5.0, and BaiduZhidao. The

	Training		Test	
	#Sentences	#DPs	#Sentences	#DPs
SMS	35,933	28,052	4,346	3,539
TC	6,734	5,090	1,122	774
Zhidao	7,970	5,097	1,406	786

Table 1: Statistics of training and test sets on three conversational benchmarks.

SMS dataset is described in (Yang et al., 2015) and contains 684 text message documents generated by users via SMS or Chat. Following (Yang et al., 2015, 2019), we reserved 16.7% of the training set as the development set, and a separate test set was used to evaluate the models. The OntoNotes Release 5.0 was released in the CoNLL 2012 Shared Task. We used the TC section which consists of transcripts of Chinese telephone conversation speech. The BaiduZhidao dataset is a question answering dialogue corpus collected by (Zhang et al., 2016). Ten types of dropped pronouns are annotated according to the pronoun annotation guidelines. The statistics of these three benchmarks are reported in Table 1.

Baselines: State-of-the-art dropped pronoun recovery models are used as baselines: (1) MEPR (Yang et al., 2015) which leverages a set of elaborately designed features and trains a Maximum Entropy classifier to predict the type of dropped pronoun before each token; (2) NRM (Zhang et al., 2016) which employs two separate MLPs to predict the position and type of a dropped pronoun utilizing representation of words in a fixed-length window; (3) BiGRU which utilizes a bidirectional RNN to encode each token in a pro-drop sentence and makes prediction based on the encoded states; (4) NDPR (Yang et al., 2019) which models dropped pronoun referents by attending to the context and independently predicts the presence and type of DP for each token.

We also compare three variants of Transformer-GCRF as: (1) Transformer-GCRF(w/o refine) which removes Step 2 in Section 3.3.1 during the graph construction process, for exploring the effectiveness of processing OVP and interjections; (2) Transformer which removes the whole GCRF layer that globally optimizes the prediction sequences, and directly adds a MLP layer on the top of Transformer encoder to predict the dropped pronouns. It aims to explore the contribution of Transformer encoder among the total effectiveness of Transformer-

Model	Chinese SMS			TC of OntoNotes			BaiduZhida		
	P(%)	R(%)	F	P(%)	R(%)	F	P(%)	R(%)	F
MEPR (Yang et al., 2015)	37.27	45.57	38.76	-	-	-	-	-	-
NRM (Zhang et al., 2016)	37.11	44.07	39.03	23.12	26.09	22.80	26.87	49.44	34.54
BiGRU	40.18	45.32	42.67	25.64	36.82	30.93	29.35	42.38	35.83
NDPR (Yang et al., 2019)	49.39	44.89	46.39	39.63	43.09	39.77	41.04	46.55	42.29
NDPR-GCRF	51.27	45.45	47.73	39.45	43.55	40.53	39.60	49.54	43.39
Transformer	51.53	46.18	48.21	39.91	43.98	41.79	42.13	46.63	43.58
Transformer-GCRF(w/o refine)	52.32	47.50	49.23	40.18	44.02	42.01	42.41	47.76	43.62
Transformer-GCRF	52.51	48.12	49.81*	40.48	44.64	42.45*	43.30	46.54	43.92*

Table 2: Results in terms of precision, recall and F-score produced by the baseline systems and variants of our proposed Transformer-GCRF framework. ‘*’ indicates the improvement over the best baseline NDPR is significant (t-tests and p -value ≤ 0.05).

GCRF; (3) NDPR-GCRF which replaces the Transformer structure in the presentation layer with the NDPR model (Yang et al., 2019).

Training details: In all of the experiments, a vocabulary was first generated based on the entire dataset, and the out-of-vocabulary words are represented as “UNK”. The length of utterances in a conversation snippet is set as 8 in our work. In Transformer-GCRF, both the encoder and decoder in the Transformer have 512 units in each hidden layer. We augment each utterance with a context consisting of seven neighboring utterances according to the practice in (Yang et al., 2019). In each experiment, we trained the model for 30 epochs on one GPU, which took more than five hours, and the model with the highest F-score on the development set was selected for testing. Following (Glorot and Bengio, 2010), in all of the experiments the weight matrices were initialized with uniform samples from $[-\sqrt{\frac{6}{r+c}}, +\sqrt{\frac{6}{r+c}}]$, where r and c are the number of rows and columns in the corresponding matrix. Adam optimizer (Kingma and Ba, 2015) is utilized to conduct the optimization.

5 Results and Analysis

5.1 Performance Evaluation

We apply our Transformer-GCRF model to all three conversation datasets to demonstrate the effectiveness of the model. Table 2 reports the results of our Transformer-GCRF model as well as the baseline models in terms of precision (P), recall (R), and F-score (F).

From the results, we can see that our proposed model and its variants outperformed the baselines on all datasets. The best model Transformer-GCRF

achieves a gain of 2.58% average absolute improvement across all three datasets in terms of F-score. We also conducted significance tests on all three datasets in terms of F-score. The results show that our method significantly outperforms the best baseline NDPR ($p < 0.05$). The proposed Transformer-GCRF suffers from performance degradation when Step 2 is removed from the graph construction process (i.e., referring to the results of Transformer-GCRF(w/o refine) in Table 2), which demonstrates the important role of OVPs in modeling dependencies between different utterances, and the contribution of noise reduction resulting from skipping short utterances starting with interjections. Both our proposed Transformer-GCRF model and the variant Transformer-GCRF(w/o refine) model outperform the variant Transformer, which demonstrates that the effectiveness comes from not only the powerful Transformer encoder, but also the elaborately designed GCRF layer. Moreover, the variant NDPR-GCRF, which encodes the pro-drop utterances with BiGRU as NDPR (Yang et al., 2019), still outperforms the original NDPR. This shows that the proposed GCRF is effective in modeling cross-utterance dependencies regardless of the underlying representation.

5.2 Motivation and Effects of GCRF

5.2.1 Motivation by statistical results

The GCRF model is motivated with a quantitative analysis of our data, which shows that 79.6% of the dropped pronouns serve as the subject of a sentence, and occur at utterance-initial positions. The pronouns dropped at the beginning of consecutive utterances are strongly correlated with dialogue patterns and thus modeling conversational structures




Figure 4: Visualization of the transition weight between each pair of pronouns among 16 types of predefined pronouns (i.e., except the category ‘None’), obtained from the vertical chain transition matrix $\mathbf{A}^{(2)}$. Darker color indicates higher transition weight between these two types of pronouns.

helps improve recover dropped pronouns. Other pronouns dropped as objects in the middle of a utterance should be recovered by modeling intra-utterance dependencies.

To further explore the cross-utterance pronoun dependencies, we collected all pronoun pairs occurring at the beginning of consecutive utterances and classified the dependencies into one of the three dialogue transitions defined in (Xue et al., 2016). We found that 27.33% of the pairs correspond to *reply* transition, where the second utterance is a response to the first utterance, and 18.60% of pairs correspond to the *acknowledgment* transition, where the second utterance is an acknowledgment of the first utterance. In both cases, the utterances involve a shift of speaker, which is accompanied by a shift in the use of personal pronouns. Another 47.79% of the pairs correspond to the *expansion* transition, where the second utterance is an elaboration of the first utterance and the same pronoun is used.

5.2.2 Visualizing transition matrix of GCRF

To investigate whether our GCRF model actually learned the dependencies revealed by the quantitative analysis of our corpus, we visualize the transition matrix $\mathbf{A}^{(2)}$ of the vertical chain in Figure 4. We can see that the learned transition matrix matches well with the distribution of dialogue patterns. The matrix shows that the higher transition weights on diagonal correspond to the strong *expansion* transition in which the same pronoun is

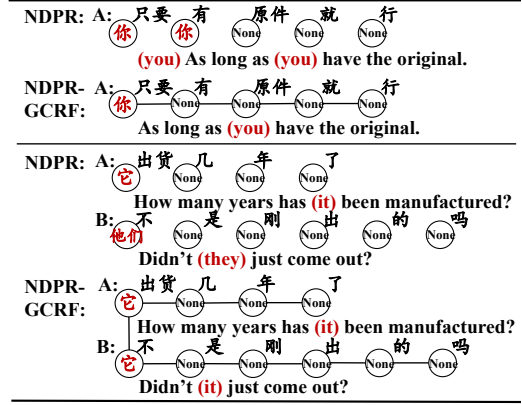


Figure 5: Example results of NDPR and NDPR-GCRF. The recovered pronouns are marked with red color and shown in brackets.

used in consecutive utterances and the transition weights between “我(I)” and “你(you)” (top-left corner) are high as well, indicating the strong *reply* transition. Moreover, the *acknowledgment* transition usually exists from the pronoun “previous utterance” to “我(I)” or “你(you)”.

5.2.3 Case studies

We demonstrate the effectiveness of GCRF by comparing the outputs of NDPR and NDPR-GCRF on the entire test set, and present some concrete cases in Figure 5. The examples show that the horizontal chains in GCRF contributes by preventing redundant predictions in the same utterance. For example, in the first case, the second pronoun “你(you)” is repeatedly recovered by NDPR since the dependency between the predictions of the first two tokens is ignored. The vertical chain contributes by predicting coherent dropped pronouns at the beginning of the utterances. For example, in the second case, the second utterance is a *reply* of the first one, and NDPR-GCRF recovers these two pronouns correctly by considering their dependency.

5.3 Effects of the Transformer architecture

We further study the effectiveness of multi-head attention in Transformer structure. Figure 6 shows an example conversation snippet with three utterances and the pronoun “它(it)” in the last utterance is dropped. The Transformer’s attention weights corresponding to three heads which are shown in blue, and the NDPR’s attention weights are shown in brown. From the results, we can see that “head 1” is responsible for associating “股票(stock)” with “它(it)” (in utterance A_1), “head 2” is responsible for associating “它(it)” with “它(it)”, and “head

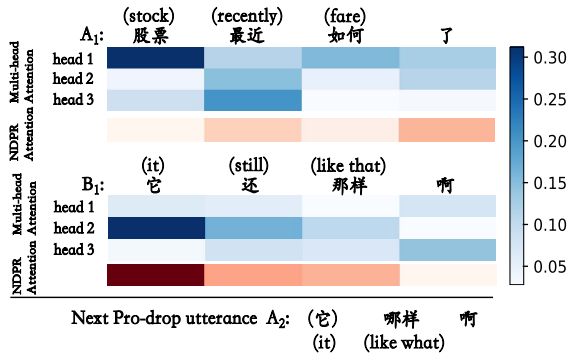


Figure 6: Visualization of multi-head attention in Transformer-GCRF and structured attention in NDPR.

3” is responsible for collecting noisy information, which is helpful for the training process (Michel et al., 2019; Correia et al., 2019). This is consistent with the observation in (Vig, 2019) that multi-head attention is powerful because it uses different heads to capture different relations. NDPR, on the other hand, captures all these the relations with a single attention structure. The results explain why Transformer is suitable for dropped pronoun recovery.

5.4 Error Analysis

Besides conducting the performance evaluation and analyzing the effects of different components, we also investigate some typical mistakes made by our Transformer-GCRF model. The task of recovering dropped pronouns consists of first identifying the referent of each dropped pronoun from the context and then recovering the referent as a concrete Chinese pronoun based on the referent semantics. Existing work has focused on modeling referent semantics of the dropped pronoun from context, and globally optimizing the prediction sequences by exploring label dependencies. However, there is also something need to do about how to recover the referent as a proper pronoun based on the referent semantics. For example, in two cases of Figure 7, the referents of the dropped pronouns are correctly identified, while the final pronoun was recovered as “(他们/they)” and “(它/it)” by mistake. We attribute this to that the model needs to be augmented with some common knowledge about how to recover a referent to the proper Chinese pronoun.

6 Conclusion and Future Work

In this paper, we presented a novel model for recovering the dropped pronouns in Chinese conversations. The model, referred to as Transformer-GCRF, formulates dropped pronoun recovery as

Context	A ₁ : 我给爷爷买的药他吃了吗? Did my grandfather take the medicine I bought for him?
Gold	B ₁ : (他) 吃了 (He) had taken the medicine.
Transformer-GCRF	B ₁ : (他们) 吃了 (They) had taken the medicine.
Context	A ₁ : 复合式听写是怎么做? How should I finish compound dictation questions ?
Gold	A ₂ : (它们) 就是听写句子吗? Do (they) require to write down the sentence you hear?
Transformer-GCRF	A ₂ : (它) 就是听写句子吗? Do (it) require to write down the sentence you hear?

Figure 7: Example errors made by Transformer-GCRF.

a sequence labeling problem. Transformer is employed to represent the utterances and GCRF is used to make the final predictions, through capturing both cross-utterance and intra-utterance dependencies between pronouns. Experimental results on three Chinese conversational datasets show that Transformer-GCRF consistently outperforms state-of-the-art baselines.

In the future, we will do some extrinsic evaluation by applying our proposed model in some downstream applications like pronoun resolution, to further explore the effectiveness of modeling cross-utterance dependencies in practical applications.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No.61702047), National Key R&D Program of China (2019YFE0198200), National Natural Science Foundation of China (No.61872338, No.61832017, No.62006234), Beijing Academy of Artificial Intelligence (No.BAAI2019ZD0305), Beijing Outstanding Young Scientist Program (No.BJJWZYJH012019100020098) and BUPT Excellent Ph.D. Students Foundation (No.CX2020305).

References

- Chen Chen and Vincent Ng. 2016. Chinese zero pronoun resolution with deep neural networks. In *Meeting of the Association for Computational Linguistic.*, pages 778–788.
- Goncalo M Correia, Vlad Niculae, and Andre F T Martins. 2019. Adaptively sparse transformers. *arXiv: Computation and Language*.
- Chris Giannella, Ransom K Winder, and Stacy Petersen. 2017. Dropped personal pronoun recovery in

- chinese sms. *Natural Language Engineering*, pages 905–927.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. *International Conference on Artificial Intelligence and Statistics.*, pages 249–256.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Fang Kong and Guodong Zhou. 2010. A tree kernel-based unified framework for chinese zero anaphora resolution. In *Conference on Empirical Methods in Natural Language Processing.*, pages 882–891.
- John D Lafferty, Andrew McCallum, and Fernando Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. *International Conference on Machine Learning*, pages 282–289.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *North American Chapter of the Association for Computational Linguistics.*, pages 260–270.
- Liyuan Liu, Jingbo Shang, Frank F Xu, Xiang Ren, Huan Gui, Jian Peng, and Jiawei Han. 2018. Empower sequence labeling with task-aware neural language model. *National Conference on Artificial Intelligence*, pages 5253–5260.
- Xuezhe Ma and Eduard H Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *Meeting of the Association for Computational Linguistics.*, pages 1064–1074.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one. *arXiv: Computation and Language*.
- Hui Su, Xiaoyu Shen, Rongzhi Zhang, Fei Sun, Pengwei Hu, Cheng Niu, and Jie Zhou. 2019. Improving multi-turn dialogue modelling with utterance rewriter. *arXiv: Computation and Language*.
- Charles Sutton, Andrew McCallum, et al. 2012. An introduction to conditional random fields. *Foundations and Trends® in Machine Learning*, 4(4):267–373.
- Charles A Sutton, Khashayar Rohanimanesh, and Andrew McCallum. 2004. Dynamic conditional random fields: factorized probabilistic models for labeling and segmenting sequence data. *International Conference on Machine Learning*, page 99.
- Jianzhuo Tong, Jingxuan Yang, Si Li, and Sheng Gao. 2019. Dropped pronoun recovery in chinese conversations with knowledge-enriched neural network. In *Chinese Computational Linguistics*, pages 545–557, Cham. Springer International Publishing.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Neural Information Processing Systems*, pages 5998–6008.
- Jesse Vig. 2019. A multiscale visualization of attention in the transformer model. *arXiv: Human-Computer Interaction*.
- Longyue Wang, Zhaopeng Tu, Shuming Shi, Tong Zhang, Yvette Graham, and Qun Liu. 2018. Translating pro-drop languages with reconstruction models. *AAAI Conference on Artificial Intelligence*.
- Longyue Wang, Zhaopeng Tu, Xiaojun Zhang, Hang Li, Andy Way, and Qun Liu. 2016a. A novel approach to dropped pronoun translation. *North American Chapter of the Association for Computational Linguistics*.
- Longyue Wang, Xiaojun Zhang, Zhaopeng Tu, Hang Li, and Qun Liu. 2016b. Dropped pronoun generation for dialogue machine translation. In *International Conference on Acoustics, Speech and Signal Processing*, pages 6110–6114.
- Jun Xu, Xiangnan He, and Hang Li. 2020. [Deep learning for matching in search and recommendation](#). *Foundations and Trends® in Information Retrieval*, 14(2–3):102–288.
- Nianwen Xue, Qishen Su, and Sooyoung Jeong. 2016. Annotating the discourse and dialogue structure of sms message conversations. *Linguistic Annotation Workshop on Meeting of the Association for Computational Linguistics*, pages 180–187.
- Jingxuan Yang, Jianzhuo Tong, Si Li, Sheng Gao, Jun Guo, and Nianwen Xue. 2019. Recovering dropped pronouns in chinese conversations via modeling their referents. In *North American Chapter of the Association for Computational Linguistics*.
- Yaqin Yang, Yalin Liu, and Nianwen Xue. 2015. Recovering dropped pronouns from chinese text messages. In *Meeting of the Association for Computational Linguistics.*, pages 309–313.
- Qingyu Yin, Yu Zhang, Weinan Zhang, and Ting Liu. 2017. Chinese zero pronoun resolution with deep memory network. In *Conference on Empirical Methods in Natural Language Processing.*, pages 1309–1318.
- Qingyu Yin, Yu Zhang, Weinan Zhang, Ting Liu, and William Yang Wang. 2018. Deep reinforcement learning for chinese zero pronoun resolution. *Meeting of the Association for Computational Linguistics*, pages 569–578.
- Weinan Zhang, Ting Liu, Qingyu Yin, and Yu Zhang. 2016. Neural recovery machine for chinese dropped pronoun. *arXiv: Computation and Language*.

Shanheng Zhao and Hwee Tou Ng. 2007. Identification and resolution of chinese zero pronouns: A machine learning approach. In *Conference on Empirical Methods in Natural Language Processing.*, pages 541–550.

Jun Zhu, Zaiqing Nie, Jirong Wen, Bo Zhang, and Weiyang Ma. 2005. 2d conditional random fields for web information extraction. *International Conference on Machine Learning*, pages 1044–1051.