

Recovering Dropped Pronouns in Chinese Conversations via Modeling Their Referents

Jingxuan Yang, Jianzhuo Tong, Si Li *, Sheng Gao, Jun Guo, Nianwen Xue †

Beijing University of Posts and Telecommunications

†Brandeis University

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

xuen@brandeis.edu

Abstract

Pronouns are often dropped in Chinese sentences, and this happens more frequently in conversational genres as their referents can be easily understood from context. Recovering dropped pronouns is essential to applications such as Information Extraction where the referents of these dropped pronouns need to be resolved, or Machine Translation when Chinese is the source language. In this work, we present a novel end-to-end neural network model to recover dropped pronouns in conversational data. Our model is based on a structured attention mechanism that models the referents of dropped pronouns utilizing both sentence-level and word-level information. Results on three different conversational genres show that our approach achieves a significant improvement over the current state of the art.

1 Introduction

Chinese is a pro-drop language, meaning that it is not always necessary to have an overt pronoun when the referent is clear from the context. This is in contrast with a non-pro-drop language like English, where an overt pronoun is always needed. For example, Kim (2000) shows that an overt subject is used only in 64% of the cases in Chinese while that percentage is over 96% in English.

Even in pro-drop languages like Chinese, pronouns are omitted to different degrees in different genres. Pronouns are dropped more often in informal conversational genres than in formal genres like newswire. Recovering these dropped pronouns (DPs) is important to applications such as Machine Translation where dropped pronouns need to be made explicit when Chinese is translated into a target language (Wang et al., 2016a,b, 2018) or Information Extraction where relations might involve entities referred to by these DPs.

* Corresponding author







	A: 这 电脑 你 帮 我 看看 接近 顶配 要 多少 钱 ? A : Can you help me figure out the cost of this computer with a near top-level configuration?
	B: 你 要 多少 寸 的 ? B : Which size of screen do you want?
	A: 我 要 大 一 点 的 吧 。 A : I want one with a bigger screen.
	B: 17 寸 的 ? (它) 很 重 的 。 B : A 17-inch one? It (the computer) is heavy.
	B: 17 寸 顶 级 配 置 是 5520 美 金 。 B : The 17-inch one with top-level configuration costs 5520 dollars .
	B: (它) 就 是 接 近 4 万 块 。 B : Or it (the price) is close to 40,000 RMB.

Figure 1: An example conversation between two people. The first dropped pronoun refers to “this computer” in the first sentence and is translated into “it” in English. It needs to be recovered from the wider context.

More concretely, recovering DPs involves i) first locating the position in a sentence where a pronoun is dropped, and ii) determining the type of the pronoun that is dropped. This is illustrated in Figure 1, where the recovered dropped pronouns are in parenthesis. Both instances of dropped pronouns can be replaced with the overt pronoun 它 (“it”), but they refer to different things in the conversation. Dropped pronoun recovery is different from zero pronoun (ZP) resolution (Chen and Ng, 2015; Yin et al., 2017, 2018) where the focus is on resolving an anaphoric pronoun to its antecedent, assuming the position of the zero pronoun is already determined. Here we do not attempt to resolve the dropped pronoun to its referent. Dropped pronoun recovery can thus be viewed as the first step of zero pronoun resolution, but this task is also important in its own right. In applications like machine translation, dropped pronouns only need to be identified and translated correctly but they do not need to be resolved.

Traditionally dropped pronoun recovery has

been formulated as a sequence labeling problem where each word in the sentence receives a tag that indicates whether the word has a dropped pronoun before it and which pronoun is dropped. Yang et al. (2015) leverages lexical, contextual and syntactic features to detect DPs before each word from a predefined list of pronouns. Giannella et al. (2017) utilizes a linear-chain conditional random field (CRF) classifier to simultaneously predict the position as well as the person number of a dropped pronoun based on lexical and syntactic information. These featured-based methods require labor-intensive feature engineering.

In contrast, Zhang et al. (2016) uses a multi-layer perceptron based neural network model to recover dropped pronouns, which eliminates the need for feature engineering. The input to their model is a context embedding constructed by concatenating the embeddings of word tokens in a fixed-length window. However, the referent of a dropped pronoun that provides crucial information to determine the identity of a pronoun is typically found outside the local context of the dropped pronoun, and thus cannot be effectively modeled with a window-based multi-layer perceptron model. For example, in Figure 1 the referent of the first dropped pronoun, the third person singular pronoun 它 (“it”) is 电脑 (“computer”), which appears in the very first utterance, several utterances before the pro-drop sentence. Therefore, long-distance contextual information needs to be captured in order to determine the type of the dropped pronoun.

In this paper, we describe a novel Neural Dropped Pronoun Recovery framework, named NDPR that can model the referent information in a much larger context. The model makes use of contextual information at two levels of granularity: the sentence level and the word level. An illustration of the NDPR framework is given in Figure 2. The referent modeling process is implemented with a structured attention mechanism. For each word token in the pro-drop sentence, the model attends to the utterance in which the referent of the dropped pronoun is mentioned by sentence attention and then zeros in to the referent by word attention. The resulting referent representation is combined with the representation of the dropped pronoun to make the final prediction on the location and type of the dropped pronoun. We demonstrate the effectiveness of our proposed model on three Chinese con-

versational datasets and results show that our approach outperforms the current state of the art by a fairly large margin.

We also perform ablation studies to explore the contribution of the different components of our model. We show that word-level attention is more effective in recovering *concrete* pronouns because it serves as a matching mechanism that matches the representation of the dropped pronoun against that of the DP referent. In contrast, sentence-level attention can be considered as an auxiliary tool for word-level attention to improve model accuracy by filtering out irrelevant referents. All code is available at <https://github.com/ningningyang/NDPR>.

Our contributions are summarized as follows:

- We propose a novel attention-based neural network to recover dropped pronouns in Chinese conversational data by modeling their referent.
- We evaluate our system on three conversational genres and show that our model consistently outperforms the state of the art by a large margin on all three datasets.
- We also present experimental results that demonstrate the effectiveness of various components in our framework and analyze some mistakes made by our model on each dataset.

2 Related Work

2.1 Dropped pronoun recovery.

Dropped pronoun detection originates from Empty Category (EC) detection and resolution, a task aimed to recover certain dropped elements in syntactic treebanks (Chung and Gildea, 2010; Cai et al., 2011; Xue and Yang, 2013). Dropped pronoun recovery was first proposed as an independent task in (Yang et al., 2015), which leveraged a set of specially designed features to recover DPs in Chinese text messages. Giannella et al. (2017) employed a linear-chain CRF classifier to jointly determine the position and person number of dropped pronouns in Chinese SMS messages using hand-crafted features. These traditional feature-based methods require heavy feature engineering. Zhang et al. (2016) for the first time utilized a multi-layer perceptron model to recover dropped pronouns. Each word is expressed as a concatenation of embeddings of word tokens in a fixed-length window. This method can not resolve the cases when

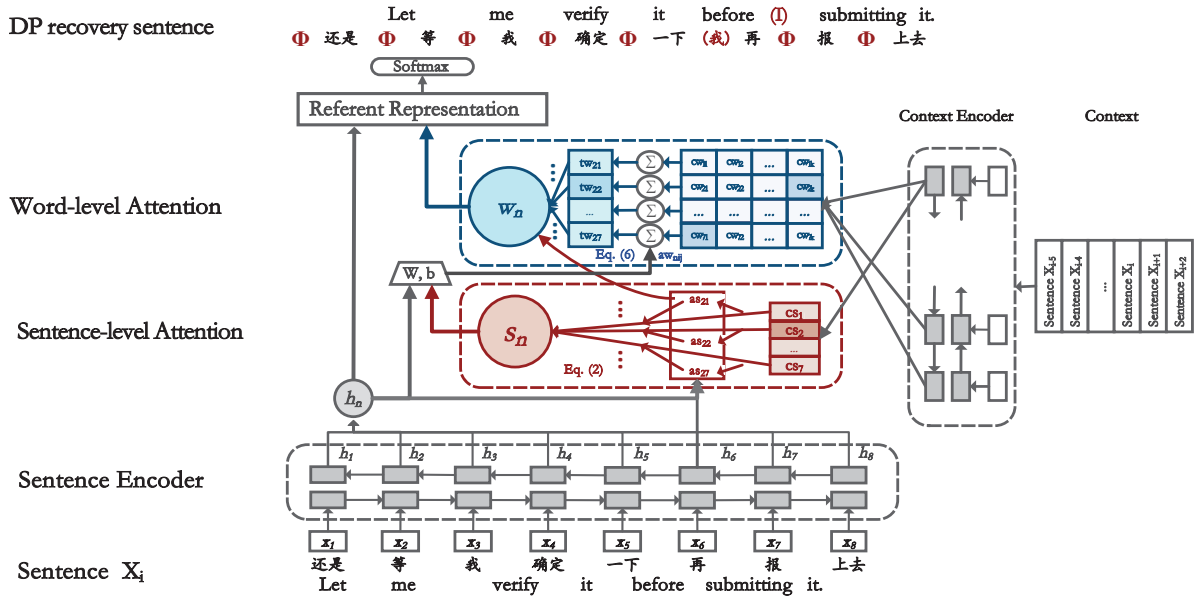


Figure 2: Neural Dropped Pronoun Recovery Framework.

referents are outside the local context, which is a common occurrence in conversational data. There has also been some work that attempts to recover dropped pronoun in pro-drop languages to help Machine Translation (Wang et al., 2016a,b, 2018). There are some inherent differences between their work and ours. In their work, they attempt to create training data by projecting pronouns from another language to Chinese, and the specific location of a dropped pronoun may not matter as much as long as it is translated correctly when Chinese is the source language. In contrast, we focus on recovering DPs from conversational context, and determining the correct position of a dropped pronoun is critical part of the task.

2.2 Zero pronoun resolution.

A line of research that is closely related to our task is zero pronoun (ZP) resolution, which aims to resolve Chinese pronouns to their antecedents. Converse and Palmer (2006) presented a rule-based method to resolve ZP by utilizing the Hobbs algorithm (Hobbs, 1978). Learning-based anaphoric resolution approaches have also been extensively explored. Zhao and Ng (2007) and Kong and Zhou (2010) proposed systems that perform ZP resolution by integrating syntactic features and position information in a system based on decision trees and context-sensitive convolution tree kernels. With the powerful learning capacity of neural networks, recent work focuses on learning representations to resolve zero pronouns or common

noun phrases (Ng, 2007; Yin et al., 2017, 2018). The DP recovery task explored in our work focuses on determining the position and type of dropped pronoun without attempting to resolve it to an antecedent.

2.3 Attention and memory network.

The attention mechanism is first proposed for neural machine translation (Bahdanau et al., 2014) and has been attempted in a variety of natural language processing tasks (Bahdanau et al., 2014; Hermann et al., 2015). Kim (2017) extended the basic attention mechanism to incorporate structural biases by combining graphical models with neural networks. Yang et al. (2016) proposed a word and sentence-level hierarchical attention network for document classification. Xing et al. (2018) also presented a hierarchical recurrent attention network (HRAN) to model the hierarchy of conversation context in order to generate multi-turn responses in chatbots. In contrast with these bottom-up mechanisms, our model adopts a top-down structured attention mechanism to construct the representation for a DP. Our model attempts to identify utterances which contain the referent of a DP, and then focus in on the representation of the referent itself. Our model draws inspiration from the memory network and its variants (Weston et al., 2015; Sukhbaatar et al., 2015; Henaff et al., 2017; Miller et al., 2016), where an external memory component is used to store and update knowledge.

3 Approach

Following (Yang et al., 2015), we formulate DP recovery as a sequential tagging problem. Given an input sentence $X = (x_1, x_2, \dots, x_s)$ and its contextual utterances $C = (X_1, \dots, X_m)$, our task is to model $P(Y|X, C)$ and predict a set of labels $Y = (y_1, y_2, \dots, y_s)$. Each element of Y indicates whether a pronoun is dropped and which type it should be before each word in the sentence X . The NDPR model consists of three components: (1) Input Layer; (2) Referent Modeling Layer; (3) Output Layer. We describe them in detail below.

3.1 Input Layer

This layer converts the raw inputs X and C into distributed representations.

3.1.1 Sentence X

During the encoding process, the input sentence of length s , is transformed into a sequence of d -dimensional word embeddings at first, and then fed into two different kinds of bidirectional recurrent neural networks (RNN) (Elman, 1991):

- **BiGRU** (Bahdanau et al., 2014): Each word $x_n, n \in \{1, \dots, s\}$ of the sentence is represented as a concatenation of forward and backward hidden states as: $h_n = [\overrightarrow{h}_n, \overleftarrow{h}_n]$. We aim to express DP by the hidden state of the first word after DP.
- **PC-BiGRU** (Yin et al., 2017): Pronoun-centered BiGRU (PC-BiGRU) also contains two independent GRU networks. For each word $x_n, n \in \{1, \dots, s\}$, the forward GRU_f encodes the preceding context of DP from left to right as \overrightarrow{h}_{n-1} , and the backward GRU_b models the succeeding context as \overleftarrow{h}_n in the reverse direction. Final hidden state of DP is also concatenated as $h_n = [\overrightarrow{h}_n, \overleftarrow{h}_{n-1}]$. The intuition of this structure is to express DP by the last word before DP and the first word after DP.

3.1.2 Context C

The context provides the necessary background information to recover DPs. In our work, we utilize five utterances preceding the one the DP is in and two utterances following current input sentence X as context C . The size of C is determined empirically since our statistics show that 97% of the dropped pronouns can be inferred based on information in the 7 sentences surrounding the sentence

in which the dropped pronoun occurs. The context C passes through the same encoder as X , yielding two kinds of memories: (1) sentence-level memory: concatenated final states of the forward and backward GRU for each contextual utterance $i = \{1, \dots, m\}$ as $cs_i = [\overleftarrow{cs}_i, \overrightarrow{cs}_i]$, holding sentence level background knowledge. (2) word-level memory: a set of concatenated hidden states at each time step $j = \{1, \dots, k\}$ as $cw_{i,j} = [\overleftarrow{cw}_{i,j}, \overrightarrow{cw}_{i,j}]$, expressing the contextual information of words in memory.

3.2 Referent Modeling layer

The referent modeling layer models the interaction between the input sentence X and its context C and outputs representation of DP referents. The relevant information for predicting the dropped pronoun is retrieved by a two-step attention-based inference mechanism in which the intermediate results of sentence-level attention is utilized to support subsequent word-level inference.

3.2.1 Identify utterances mentioning DP

This operation aims to figure out which utterances mention the intended referent of DP. We compute the relevance between the hidden states h_n representing the DP and its sentence level contextual representation cs_i , resulting in $rs_{n,i}$. After passing it through softmax, we obtain a sentence-level attention distribution $as_{n,i}$ as:

$$rs_{n,i} = h_n^T cs_i \quad (1)$$

$$as_{n,i} = \frac{e^{rs_{n,i}}}{\sum_{i=1}^m e^{rs_{n,i}}} \quad (2)$$

We then conclude sentence-level background information s_n of DP as a weighted sum of contextual utterance representations:

$$s_n = \sum_{i=1}^m as_{n,i} \cdot cs_i \quad (3)$$

3.2.2 Update DP representation

This operation updates representation of the DP by combining the original DP information h_n with sentence-level background knowledge s_n through a linear transformation as:

$$hs_n = W^{2d \times 4d} [h_n, s_n] + b^{2d}, \quad (4)$$

The updated DP state hs_n will be used in subsequent word level inference.




Figure 3: Counts of each type of dropped pronouns in different conversational data sets. Pronouns have a more balanced distribution in Chinese SMS and telephone conversation(TC). For BaiduZhidao, only concrete pronouns are annotated and the distribution of these pronouns is extremely uneven.

3.2.3 Identify DP referents

This operation aims to capture the word sequence that represents the intended referent of the DP. We cast the updated DP state hs_n onto the space of contextual words $cw_{i,j}$ by measuring the similarity $rw_{n,i,j}$ between them. The casting operation also serves as a regularization measure to restrict the inference space. The similarities are fed into the softmax function to give a word-level attention distribution $aw_{n,i,j}$ for each contextual word as:

$$rw_{n,i,j} = \mathbf{W}^{1 \times 2d} (\mathbf{h}s_n \odot cw_{i,j}) + b^1 \quad (5)$$

$$aw_{n,i,j} = \frac{e^{rw_{n,i,j}}}{\sum_{j=1}^k e^{rw_{n,i,j}}}, \quad (6)$$

The attention distribution $aw_{n,i,j}$ also provides an interpretation for DP recovery results based on the words it attends to.

Finally, we derive the word-level representation w_n from word states $cw_{i,j}$. Word-level attention $aw_{n,i,j}$ is used to compute a weighted sum of contextual word states $cw_{i,j}$ of each utterance, yielding the referent representation $tw_{n,i}$. Then, sentence-level attention as_i further filters out irrelevant referents, yielding the final referent representation w_n as:

$$tw_{n,i} = \sum_{j=1}^k aw_{n,i,j} \cdot cw_{i,j} \quad (7)$$

$$w_n = \sum_{i=1}^m as_{n,i} \cdot tw_{n,i} \quad (8)$$

3.3 Output Layer and Training Objective

The output layer predicts the recovery result based on DP state h_n and referent representation w_n . We feed the concatenation of these two parts into a 2-layer fully connected softmax classifier to give a categorical probability distribution over 17 candidate pronouns as:

$$\alpha_n = \tanh(\mathbf{W}_1 \cdot [h_n; w_n] + \mathbf{b}) \quad (9)$$

$$P(y_n|x_n, \mathbf{C}) = \text{softmax}(\mathbf{W}_2 \cdot \alpha_n + \mathbf{b}) \quad (10)$$

We train our model by minimizing cross-entropy between the predicted label distributions and the annotated labels for all sentences. The training objective is defined as:

$$\text{loss} = -\sum_{l \in N} \sum_{n=1}^s \delta(y_n|x_n, \mathbf{C}) \log(P(y_n|x_n, \mathbf{C})), \quad (11)$$

where N represents all training instances, s represents the number of words in each sentence; $\delta(y_n|x_n, c)$ represents the annotated label of x_n .

4 Datasets and Experimental Setup

4.1 Datasets

We perform experiments on three Chinese conversational datasets.

- Chinese text message (SMS) data, which consists of 684 SMS/Chat files. We use the same dataset split as (Yang et al., 2015), which reserves 16.7% of the training set as a held-out development set to tune the hyper-parameters and evaluate the model on a separate test set.

- OntoNotes Release 5.0, which is used in the CoNLL 2012 Shared Task. We use a portion of OntoNotes 5.0 that consists of transcripts of Chinese telephone conversation (TC) speech. The 9,507-sentence subset of OntoNotes only has coreference annotations for anaphoric zero pronouns. In order to formulate the task as an end-to-end sequence labeling task, we annotate all the dropped pronouns following annotation guidelines described in (Yang et al., 2015).
- BaiduZhidaao, which is a question answering dialogue dataset used in (Zhang et al., 2016). It contains 11,160 sentences that are annotated with 10 types of concrete dropped pronouns.

Figure 3 shows the statistics of each type of pronouns in the three data sets. According to (Yang et al., 2015), 5 out of 15 types of dropped pronouns in the SMS data set are *abstract pronouns* that do not correspond to any actual pronoun in Chinese. These can refer to an event (Event), the previous utterance (Previous Utterance), or an generic or unspecific entity (Generic). The other two abstract pronouns are used to indicate the subject of an existential construction (Existential) or a pleonastic subject (Pleonastic). The other ten types of Chinese pronouns are classified as concrete pronouns. The same 15 types of pronouns are used to annotate the TC data. We can see that the BaiDuZhidaao data set only has concrete pronouns.

4.2 Evaluation metrics

We use the same evaluation metrics as (Yang et al., 2015): precision (P), recall (R) and F-score (F).

4.3 Training details

Our model is implemented with Tensorflow. The vocabulary is generated from the training set, which contains 17,199 word types. Out-of-vocabulary (OOV) words are represented as UNK. The BiGRU encoder uses a hidden layer of 150 units. To train the model, we use the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.0003. We train the model for 8 epochs and select the model with the highest F-score on the development set for testing. Dropout rate is set at 0.2 on the fully connected layers. We use uniform initialization for the weight matrices and zero initialization for biases.

4.4 Baseline Methods and Model Variations

We compare our proposed NDPR framework with three baselines and implement two parallel experiments:

- **MEPR**: This model is provided by Yang (2015), which uses a maximum entropy (ME) classifier with hand-crafted features to recover dropped pronouns.
- **NRM**: This model is proposed by Zhang (2016), which uses two independent multi-layer perceptrons to locate the position of dropped pronouns and determine the type of dropped pronouns.
- **BiGRU**: This model encodes each word in the target sentence with a bidirectional GRU. The output representation of the BiGRU encoder is used to predict the type of the dropped pronoun. This method can be seen as a degenerate variant of our NDPR model without the attention mechanism.
- **NDPR**: This model uses the BiGRU encoder with both sentence-level and word-level attention. The word embeddings are initialized with pre-trained 300-D Zhihu QA vectors (Li et al., 2018) and fine-tuned when training the DP recovery model.
- **NDPR-rand**: Same as NDPR except that the word embeddings are randomly initialized.
- **NDPR-PC-BiGRU**: Same as NDPR but the encoder of utterance X is replaced with PC-BiGRU.

We also perform two ablation experiments to explore the effectiveness of sentence-level attention and word-level attention:

- **NDPR-S**: Same as NDPR but the referent modeling layer only utilizes sentence-level attention. The output layer makes prediction with only DP state h_n and sentence level information s_n .
- **NDPR-W**: Same as NDPR but the referent modeling layer only utilizes word-level attention. The output layer makes prediction with only DP state h_n and word-level information w_n .

5 Results and Discussion

5.1 Main Results

Table 1 shows experimental results of three baseline systems and variants of our proposed NDPR

Model	Chinese SMS			TC of OntoNotes			BaiduZhidao		
	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-rand	46.47	43.23	43.58	28.98	41.50	33.38	35.44	43.82	37.79
NDPR-PC-BiGRU	46.34	46.21	46.27	36.69	40.12	38.33	38.42	48.01	41.68
NDPR-W	46.78	46.61	45.76	38.67	41.56	39.64	38.60	50.12	43.36
NDPR-S	46.99	46.32	44.89	37.40	40.32	38.81	39.32	46.40	41.53
NDPR	49.39	44.89	46.39	39.63	43.09	39.77	41.04	46.55	42.94

Table 1: Results in terms of precision, recall and F-score on 16 types of pronouns produced by the baseline systems and variants of our proposed NDPR model. For NRM* (Zhang et al., 2016), we implement the proposed model as described in the paper.

Tag	NDPR-S	NDPR-W	NDPR
他们 (masculine they)	17.05	<u>23.28</u>	24.44
她 (she)	32.35	<u>33.72</u>	35.14
previous utterance	84.90	<u>86.08</u>	87.55
他 (he)	29.05	<u>31.20</u>	34.92
它 (it)	25.00	<u>26.67</u>	26.95
她们 (feminine they)	0	0	40.00
我 (I)	50.66	<u>50.90</u>	52.98
我们 (we)	31.49	<u>33.57</u>	34.81
你 (singular you)	42.88	<u>44.15</u>	44.31
pleonastic	25.89	22.29	28.46
generic	<u>11.61</u>	11.08	16.83
event	<u>6.15</u>	0	16.27
existential	<u>34.17</u>	30.84	38.71
你们 (plural you)	0	0	5.41
它们 (inanimate they)	16.00	19.15	13.89

Table 2: F-scores of our proposed model NDPR and its two variants (NDPR-S, NDPR-W) for concrete and abstract pronouns on the Chinese SMS test set.

model on Chinese SMS, TC section of OntoNotes, and BaiduZhidao. We can see that our proposed model and its variants outperform the baseline methods on all these datasets by different margins. Our best model, NDPR, outperforms MEPR by 7.63% in terms of F-score on the Chinese SMS dataset, and outperforms NRM by 16.97% and 8.40% on the OntoNotes and BaiduZhidao datasets respectively. Compared with the degenerate variant model BiGRU, our NDPR model also performs better on all three datasets, which demonstrate the effectiveness of referent modeling mechanism composed of sentence-level and word-level attention.

The experimental results also show that all components of our model have made a positive con-

tribution, which is evidenced by the fact that the full model NDPR outperforms the other variants by a small margin. The results also show that pre-trained word embeddings have had a positive impact on the model, as shown by the higher accuracy of the NDPR model over the NDPR-rand model, where the word embeddings are randomly initialized.

PC-BiGRU encoder seems to perform worse than vanilla BiGRU encoder, as shown by the slightly lower F-score of the NDPR-PC-BiGRU model. But it alleviates what we call the local pronoun repetition problem, the situation where a DP is recovered redundantly in a sentence before a verb and its adverbial modifier. We attribute this to that PC-BiGRU uses the context state h_{n-1} of the last word before DP instead of h_n , which makes the model pay more attention to the preceding modifier.

For the BaiduZhidao dataset, the NDPR-W model actually performs better than NDPR model, as indicated by the higher F-score for NDPR-W in the last column of Table 1. We attribute this to the fact that there are only concrete pronouns in this dataset. The combination of “我 (I)”, “你 (singular you)” and “它 (it)” accounts for 94.47% of the overall dropped pronoun population, for which the referent can be easily determined by word-level attention. Moreover, the fewer conversation turns in this data set mean there are few irrelevant referents that need to be filtered out by sentence-level attention.

5.2 Ablation Study

In this section, we dive a bit deeper and look at the impact of the attention mechanism on concrete and

Pro-drop sentence: A: (她)(she) 是用的 锤头 还是 扳手?
A: What did (she) use, a hammer or wrench?

A: 你 终于 出现 了 。
You finally showed up.
B: 你 别 乱说 。
Don't talk nonsense.
A: 啊 哈哈 。
Haha.
B: 最近 我的 电脑 它 让 我 女儿(daughter) 搞 坏 了
My daughter broke my computer recently.
A: 哈哈 她(she) 还 有 这 功 能 她(she) 不 简 单
Haha. She can even do that. She is not to be under-estimated.
B: 嗯 ， 破 坏 力 很 强 。
Yes, very destructive.
B: 她(she) 用 了 一 根 手 指 。
She just used a finger.




Figure 4: Visualization of attention in the NDPR model. The red pattern on the left shows the distribution of sentence-level attention and the blue pattern on the right shows the distribution of word-level attention. Darker color indicates higher attention weight.

abstract pronouns, respectively. Table 2 shows the F-scores of three variants of our proposed model for each type of pronouns on the entire Chinese SMS test set. The best results among these three variants are in boldface, and the better results between NDPR-S and NDPR-W are underlined.

We can see that for all types of concrete pronouns, the NDPR-W model outperforms the NDPR-S model by a significant margin. In general, the referent of concrete dropped pronoun is usually realized in the form of a phrase that consists of one or more word, and they can be accurately captured by word-level attention. The NDPR model that incorporates both word-level and sentence-level attention further improves upon NDPR-W, with the lone exception of 它们 (“they”). We believe the reason is that the sentence-level encoder cannot adequately represent the interaction between the multiple referents for plural pronouns like 它们 and only add noise to the representation.

However, this observation does not hold for abstract dropped pronouns. The NDPR-W performs comparably or slightly worse on four out of the five types of abstract pronouns. This is consistent with the fact that the referent of abstract pronouns

Chinese SMS	
Context :	A: 我还以为你大过年要搬家呢! A: I thought you would make a move during the Spring Festival. A: 我希望你能在报社附近 找个房子 。 A: I hope you can rent a house near the newspaper office.
	B: 嗯，我能 接上 最好了。 B: Well, it is best if the new house is ready .
Gold:	B: (event) 不行就只能搬回去了。 B: I have to move back if the new house is not accessible .
NDPR:	B: (我) 不行就只能搬回去了。 B: I have to move back if I can not.
Context :	A: 我可以从上海直飞到纽约吗? A: Can I take a direct flight from Shanghai to New York?
Gold:	B: (generic) 直飞都没有 B: There is no direct flight .
NDPR:	B: (我) 直飞都没有。 B: I can not flight directly.
TC of OntoNotes	
Gold:	(你) 有空， (你) 有空(你)可能还会到北京去，是吧? You may go to Beijing again when you are free, you are free, right?
NDPR:	(你) 有空， (我) 有空(我)可能还会到北京去，是吧? I may go to Beijing again when you are free, I am free, right?
BaiduZhidao	
Context :	A: 医院开的那些 退烧药 有用吗? A: Are those antipyretics prescribed by the hospital useful?
Gold:	B: (它们) 都有什么成分? B: What are the ingredients of them ?
NDPR:	B: (它) 都有什么成分? B: What are the ingredients of it ?

Figure 5: Example errors made by NDPR on different conversational datasets.

like “Event” is often an entire sentence. The full NDPR model still outperforms both NDPR-W and NDPR-S for all abstract pronouns.

5.3 Attention Visualization

Intuitively, sentence-level attention models the interaction between the dropped pronoun and each context utterance. The utterance containing the referent of the dropped pronoun should receive more attention. Similarly, word-level attention models the interaction between each word in the context and the dropped pronoun. Thus words that describe the referent of the pronoun should receive more attention if the model works as intended. Figure 4 shows a instance of attention distribution produced by our model.

In this example, the model correctly gives higher attention weights to three utterances that contain words that indicate the referent of the pronoun. At word level, the model also gives higher attention weights to the specific words that indicate the referent of the dropped pronoun. For example, words such as 她 (“she”) and 女儿 (“daughter”) have received higher weights, as indicated by the darker color. This suggests that in this case the model attends to the right utterances and words and works as we have expected.

5.4 Error Analysis

Figure 5 shows some typical mistakes made by our model on each genre. For the Chinese SMS dataset, the distribution of different types of pronouns is relatively balanced (See Table 3), and our model does a better job on concrete pronouns but stumbles on abstract pronouns like “event” and “generic” as shown in Table 2 as it is harder to attend to the correct parts of the context in these cases.

The TC data of OntoNotes is a transcription of telephone conversations where there are often repetitions, and our model struggles when people repeat what they said. If the same pronoun is mentioned repeatedly, our model can not capture this interaction, since each pronoun is recover independently. This suggests that one future improvement might involve using a sequence-based decoder.

In the BaiduZhida dataset, only concrete pronouns are annotated as shown in Figure 3. Pronouns like 它 (“it”) and 它们 (“they”) account for a large proportion of all dropped pronouns. For these two categories, the performance of our proposed method is hit-and-miss, which can be attributed to the absence of pronoun resolution common sense.

6 Conclusions

We have proposed an end-to-end neural network architecture that attempts to model the interaction between the dropped pronoun and its referent in order to recover dropped pronouns in Chinese conversational data. Our model is based on sentence-level and word-level attention and results show that our model consistently outperforms baseline methods when evaluated on three separate datasets. We further investigate the effectiveness of different components of our model by performing ablation experiments and demonstrate the interpretability of our model using attention visualization.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No.61702047, No.61300080) and Beijing Natural Science Foundation (No.4174098).

References

Dzmitry Bahdanau, Kyunghyun Cho, and Y Bengio. 2014. Neural machine translation by jointly

learning to align and translate. *Computer Science.*, 1409.

Shu Cai, David Chiang, and Yoav Goldberg. 2011. Language-independent parsing with empty elements. pages 212–216. Proceedings of the Meeting of the Association for Computational Linguistics.

Chen Chen and Vincent Ng. 2015. Chinese zero pronoun resolution: A joint unsupervised discourse-aware model rivaling state-of-the-art resolvers. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing.*, volume 2, pages 320–326.

Tagyoung Chung and Daniel Gildea. 2010. Effects of empty categories on machine translation. pages 636–645.

Susan P Converse and Martha Stone Palmer. 2006. *Pronominal anaphora resolution in Chinese.* Citeseer.

Jeffrey L. Elman. 1991. *Distributed Representations, Simple Recurrent Networks, And Grammatical Structure.* Springer US.

Chris Giannella, Ransom K Winder, and Stacy Petersen. 2017. Dropped personal pronoun recovery in chinese sms. *Natural Language Engineering*, 23(06):905–927.

Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, and Yann Lecun. 2017. Tracking the world state with recurrent entity networks. *International Conference on Learning Representations.*

Karl Moritz Hermann, Tomas Kociský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in Neural Information Processing Systems.*, pages 1693–1701.

Jerry R Hobbs. 1978. Resolving pronoun references. *Lingua*, 44(4):311–338.

Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. 2017. Structured attention networks. *International Conference on Learning Representations.*

- Young Joo Kim. 2000. Subject/object drop in the acquisition of Korean: A cross-linguistic comparison. *Journal of East Asian Linguistics*, 9(4):325–351.
- 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 *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 882–891. Association for Computational Linguistics.
- Shen Li, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du. 2018. Analogical reasoning on Chinese morphological and semantic relations. *Proceedings of the Meeting of the Association for Computational Linguistics*.
- Alexander H Miller, Adam Fisch, Jesse Dodge, Amirhossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-value memory networks for directly reading documents. *Conference on Empirical Methods in Natural Language Processing*, pages 1400–1409.
- Vincent Ng. 2007. Semantic class induction and coreference resolution. In *Proceedings of the Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*, pages 536–543.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in Neural Information Processing Systems*, pages 2440–2448.
- Longyue Wang, Zhaopeng Tu, Shuming Shi, Tong Zhang, Yvette Graham, and Qun Liu. 2018. Translating pro-drop languages with reconstruction models. *arXiv preprint arXiv:1801.03257*.
- Longyue Wang, Zhaopeng Tu, Xiaojun Zhang, Hang Li, Andy Way, and Qun Liu. 2016a. A novel approach to dropped pronoun translation. *arXiv preprint arXiv:1604.06285*.
- Longyue Wang, Xiaojun Zhang, Zhaopeng Tu, Hang Li, and Qun Liu. 2016b. Dropped pronoun generation for dialogue machine translation. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 6110–6114. IEEE.
- Jason Weston, Sumit Chopra, and Antoine Bordes. 2015. Memory networks. *International Conference on Learning Representations (ICLR)*.
- Chen Xing, Yu Wu, Wei Wu, Yalou Huang, and Ming Zhou. 2018. Hierarchical recurrent attention network for response generation. *National Conference on Artificial Intelligence*.
- Nianwen Xue and Yaqin Yang. 2013. Dependency-based empty category detection via phrase structure trees. pages 1051–1060.
- Yaqin Yang, Yalin Liu, and Nianwen Xue. 2015. Recovering dropped pronouns from Chinese text messages. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, volume 2, pages 309–313.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489.
- 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 *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*.