

# Assessing Neural Referential Form Selectors on a Realistic Multilingual Dataset

Guanyi Chen<sup>♣</sup>, Fahime Same<sup>♡</sup>, and Kees van Deemter<sup>♣</sup>

<sup>♣</sup>Department of Information and Computing Sciences, Utrecht University

<sup>♡</sup>Department of Linguistics, University of Cologne

g.chen@uu.nl, f.same@uni-koeln.de, c.j.vandeemter@uu.nl

## Abstract

Previous work on Neural Referring Expression Generation (REG) all uses webNLG, an English dataset that has been shown to reflect a very limited range of referring expression (RE) use. To tackle this issue, we build a dataset based on the OntoNotes corpus that contains a broader range of RE use in both English and Chinese (a language that uses zero pronouns). We build neural Referential Form Selection (RFS) models accordingly, assess them on the dataset and conduct probing experiments. The experiments suggest that, compared to webNLG, OntoNotes is better for assessing REG/RFS models. We compare English and Chinese RFS and confirm that in both languages BERT has the highest performance. Also, our results suggest that in line with linguistic theories, Chinese RFS depends more on discourse context than English.

## 1 Introduction

Referring Expression Generation (REG) In Context is a key task in the classic Natural Language Generation pipeline (Reiter and Dale, 2000; Gatt and Krahmer, 2018). Given a discourse whose referring expressions (REs) have yet to be realised and given their intended referents, it aims to develop an algorithm that generates all these REs.

Traditionally, REG In Context (hereafter REG) is a two-step process. In the first step, the Referential Form (RF) is determined, e.g. whether to use a proper name, a description, a demonstrative or a pronoun. This step is the focus of this work and will be hereafter called Referential Form Selection (RFS). In the second step, the content of the RE is determined. For example, to refer to *Joe Biden*, one needs to choose from options such as “*the president*”, “*the 46th president of US*”.

In recent years, many works on REG have started to use neural networks. For example, Castro Ferreira et al. (2018a); Cao and Cheung (2019); Cunha et al. (2020) have proposed to generate REs in

an End2End manner, i.e., to tackle the selection of form and content simultaneously. Chen et al. (2021) used BERT (Devlin et al., 2019) to perform RFS. One commonality between these studies is that they were all tested on a benchmark dataset, namely webNLG (Gardent et al., 2017; Castro Ferreira et al., 2018b).

However, Chen et al. (2021) and Same et al. (2022) found that webNLG is not ideal for assessing REG/RFS algorithms because (1) it consists of rather formal texts that may not reflect everyday RE use; (2) its texts are very short and have a simple syntactic structure; and (3) most of its REs are first-mentions. These limitations led to some unexpected results when they tested their RFS models on webNLG. For example, advanced pre-trained models (i.e., BERT) performed worse than simpler models (i.e., single-layer GRU (Choi et al., 2014)) without any pre-training. By probing<sup>1</sup> various RFS models, they found that though BERT encodes more linguistic information, which is crucial for RFS, it still performs worse than GRU. In this study, we are interested in *how well each RFS model performs when tested on a dataset that addresses the above limitations* – in what follows, we call this a “realistic” dataset, for short.

Additionally, all the above studies were conducted on English only. It has been pointed out that speakers of East Asian languages (e.g. Chinese and Japanese) use REs differently from speakers of Western European languages (e.g. English and Dutch; Newnham (1971)). Theoretical linguists (Huang, 1984) have suggested that East Asian languages rely more heavily on context than Western European languages (see Chen (2022) for empirical testing and computational modelling). As a result, speakers of East Asian languages frequently use Zero Pronouns (ZPs), i.e. REs that contain no words and are resolved based merely

<sup>1</sup>Probing is an established method to analyse whether the latent representations of a model encode certain information.

<b>Text:</b> Amatriciana sauce is made with Tomato. It is a traditional Italian sauce. Amatriciana is a sauce containing Tomato that comes from Italy.
<b>Delexicalised Text:</b> Amatriciana_sauce is made with Tomato. Amatriciana_sauce is a traditional Italy sauce. Amatriciana_sauce is a sauce containing Tomato that comes from Italy.

Table 1: An example data from the webNLG corpus. In the delexicalised text, every entity is highlighted.

on their context.<sup>2</sup> This poses two challenges for REG/RFS models: (1) they need to be better able to encode contextual information; (2) they need to account for an additional RF (i.e. ZP). Therefore, we are curious to see *how well each RFS model performs when tested on a language that has more RFs and relies more on context than English*.

To answer the research questions above, we construct a “realistic” multilingual dataset of English and Chinese and try different model architectures, such as models with/without pre-trained word embeddings, and models incorporating BERT. We report the results and compare model behaviours on English and Chinese subsets. The code used in this study is available at: <https://github.com/a-quei/probe-neuralreg>.

## 2 Referential Form Selection (RFS)

Using webNLG, Castro Ferreira et al. (2018a) re-defined the REG task in order to accommodate deep learning techniques. Subsequently, Chen et al. (2021) adapted the definition to fit the RFS task. The first step is to remove from each RE all information about the RF of that RE. Concretely, as shown in Table 1, Castro Ferreira et al. (2018a) first “delexicalised” each text in webNLG by assigning a general entity tag to each entity and replacing all REs referring to that entity with that tag. In most cases, a tag is assigned to an entity by replacing whitespaces in its proper name with underscores, e.g. “Amatriciana sauce” to “Amatriciana\_sauce”.

For a target referent  $x^{(r)}$  (e.g. the second “Amatriciana\_sauce” in Table 1), given the referent, its pre-context in the discourse  $x^{(pre)}$  (e.g. “Amatriciana\_sauce is made with Tomato.”) and its post-context  $x^{(post)}$  (e.g. “is a traditional Italy

<sup>2</sup>For example, consider the question in Chinese: “你看见比尔了吗?” (Have you see Bill?). A Chinese speaker can reply “ $\emptyset$ 看见 $\emptyset$ 了。” ( $\emptyset$  saw  $\emptyset$ .) where the two  $\emptyset$  are ZPs that refer to the speaker himself/herself and “Bill” respectively.

EN	4-Way	Demonstrative, Description, Proper Name, Pronoun
	3-Way	Description, Proper Name, Pronoun
	2-Way	Non-pronominal, Pronominal
ZH	5-Way	Demonstrative, Description, Proper Name, Pronoun, ZP
	4-Way	Description, Proper Name, Pronoun, ZP
	3-Way	Non-pronominal, Pronoun, ZP
	2-Way	Overt Referring Expression, ZP

Table 2: Types of RF classification and possible classes. Demonstratives are grouped with descriptions in 3-way EN and 4-way ZH classifications under the category *Description*. The category *Non-pronominal* contains proper names, descriptions, and demonstratives.

sauce. Amatriciana\_sauce is a sauce containing Tomato that comes from Italy.”), the RFS task is to decide the proper RF  $\hat{f}$  (e.g., pronoun).

## 3 Dataset Construction

To construct a realistic multilingual REG/RFS dataset, we used the Chinese and English portions of the OntoNotes dataset<sup>3</sup> whose contents come from six sources, namely broadcast news, newswires, broadcast conversations, telephone conversations, web blogs, and magazines. We call the resulting Chinese subset OntoNotes-ZH and the English subset OntoNotes-EN. In the following, we describe the construction process.

First, for each RE in OntoNotes, we used the 3 previous sentences as the pre-context and the 3 subsequent sentences as the post-context. Similar to Chen et al. (2021), we are interested in different RF classification tasks. For Chinese, for example, we not only have a 2-way classification task where models have to decide whether to use a ZP or an overt RE, but also a 5-way task where models have to choose from a more fine-grained list of possible RFs. Table 2 lists all categories in both OntoNotes-EN and OntoNotes-ZH. Using the constituency syntax tree of the sentence containing the target referent and the surface form of the target, we automatically annotated each RE with its RF category. For example, an RE is considered a demonstrative if it is annotated in the syntax tree as a noun phrase and its surface form contains a demonstrative determiner.

Second, we excluded all coreferential chains consisting only of pronouns and ZPs. The pronominal

<sup>3</sup>OntoNotes is licensed under the Linguistic Data Consortium: <https://catalog.ldc.upenn.edu/LDC2013T19>.

	WebNLG	O-EN	O-ZH
Percentage of First Mentions	85%	43%	43%
Percentage of Proper Names	71%	21%	15%
Average Number of Tokens	18.62	106.44	139.55

Table 3: Statistics of WebNLG and OntoNotes. O-EN and O-ZH stand for OntoNotes-EN and OntoNotes-ZH.

chains consist mainly of first/second-person referents, and we do not expect much variation in referential form in these cases. In other words, we only included the chains that have *at least* one overt non-pronominal RE.

Third, we delexicalised the corpus following [Castro Ferreira et al. \(2018a\)](#). Additionally, since we used the Chinese BERT as one of our RFS models and it only accepts input shorter than 512 characters, we removed all samples in OntoNotes-ZH whose total length (calculated by removing all underscores introduced during delexicalisation and summing the length of pre-contexts, post-contexts, and target referents) is longer than 512 characters. Experiments with models other than BERT on the original OntoNotes-ZH show that this does not bias the conclusions of this study (see Appendix A).

Last, we split the whole dataset into a training set and a test set in accordance with the CoNLL 2012 Shared Task ([Pradhan et al., 2012](#)). Since ZPs in Chinese are only annotated in the training and development sets, following [Chen and Ng \(2016\)](#), [Chen et al. \(2018\)](#), and [Yin et al. \(2018\)](#), we used the development set as the test set and sampled 10% of the documents from the training set as the development data. Thus, we obtained OntoNotes-EN, where the training, development, and test sets contain 71667, 8149, and 7619 samples, respectively, and OntoNotes-ZH, where the training, development, and test sets contain 70428, 9217, and 11607 samples, respectively.

**OntoNotes vs. WebNLG.** Based on the nature of OntoNotes and the statistics in Table 3, we observe that: (1) the WebNLG data is all from DBpedia, while the OntoNotes data is multi-genre; (2) OntoNotes has a much smaller proportion of first mentions and proper names; and (3) the documents in OntoNotes are on average much longer than those in WebNLG.

Another difference between WebNLG and OntoNotes is in the ratio of seen and unseen entities in their test sets. [Castro Ferreira et al. \(2018b\)](#) divided the documents in the WebNLG’s test set

into *seen* (where all the data come from the same domains as the training data) and *unseen* (where all the data come from different domains than the training data). Almost all referents from the seen test set appear in the training set (9580 out of 9644), while only a few referents from the unseen test set appear in the training set (688 out of 9644).<sup>4</sup> In OntoNotes, 38.44% and 41.45% of the referents in the test sets of OntoNotes-EN and OntoNotes-ZH also appear in the training sets.

Having said this, OntoNotes largely mitigates the problems of WebNLG discussed in §1. If OntoNotes is a “better” and more “representative” corpus for assessing REG/RFS models, we can expect more “expected” results: models with pre-training outperform those without, and models that learn more useful linguistic information outperform those that learn less. We will detail our expectations in §5.

## 4 Modelling RFS

We introduce how we represent entities and how we adapt the RFS models of [Chen et al. \(2021\)](#).

### 4.1 Entity Representation

Unlike WebNLG, whose 99.34% of referents in the test set appear in the training set, the majority of referents in OntoNotes do not appear in both training and test sets. This means that RFS models should be able to handle unseen referents, but mapping each entity to a general entity tag with underscores would prevent the models from doing so ([Cao and Cheung, 2019](#); [Cunha et al., 2020](#)) because entity tags of unseen entities are usually out-of-vocabulary (OOV) words. Additionally, when incorporating pre-trained word embeddings and language models, using entity tags prevents entity representations from benefiting from these pre-trained models (again since the entity tags of unseen entities are usually OOV words).

Similar to [Cunha et al. \(2020\)](#), we replaced underscores in general entity tags (e.g. “*Amatriciana\_sauce*”) with whitespaces (henceforth, lexical tags, e.g. “*Amatriciana sauce*”). Arguably, there is a trade-off between using entity tags and using lexical tags. In contrast to lexical tags, the use of entity tags helps models identify mentions of the same entity in discourse, which has been shown to be a crucial feature for RFS. However, using entity tags prevents models from dealing with

<sup>4</sup>[Chen et al. \(2021\)](#) used only seen entities because the size of the underlying triples of the unseen test set differs from both the training set and seen test set.

unseen entities and reduces the benefit of using pre-trained language models. In §6.3, we compare the performance of using entity tags and lexical tags.

## 4.2 RFS Models

To build the RFS models, we use the two neural models from Chen et al. (2021): c-RNN and ConATT. Given the task definition in §2, models take pre-context  $x^{(pre)}$ , target referent  $x^{(r)}$ , and post-context  $x^{(post)}$  as inputs. As a result of using lexical tags, each target referent is no longer a single tag, but a sequence of tokens. In other words, instead of being  $\{w_i\}$ ,  $x^{(r)}$  is  $\{w_i, w_{i+1}, \dots, w_j\}$ . The other two inputs are pre-context  $x^{(pre)} = \{w_1, w_2, \dots, w_{i-1}\}$  and post-context  $x^{(post)} = \{w_{j+1}, w_{j+2}, \dots, w_n\}$ . The architectures of the models are as follows:

**c-RNN.** c-RNN concatenates  $x^{(pre)}$ ,  $x^{(r)}$  and  $x^{(post)}$ , and uses a single bidirectional GRU to encode them all. Formally, we obtain a sequence of hidden representations by  $h = \text{BiGRU}([x^{(pre)}, x^{(r)}, x^{(post)}])$ . We then use the summation of the hidden representations at the beginning and the end of the target referent (i.e.,  $i$  and  $j$ ) for calculating the final representation:

$$R = \text{ReLU}(W_f[h_i + h_j]), \quad (1)$$

where  $W_f$  is the weight in the feed-forward layer.  $R$  is then used for predicting the RF:

$$P(\hat{f}|x^{(pre)}, x^{(r)}, x^{(post)}) = \text{Softmax}(W_c R), \quad (2)$$

where  $W_c$  is the weight in the output layer.  $x$  can be initialised randomly or initialised by pre-trained word embeddings or language models. We tested both the vanilla c-RNN and c-RNN, whose input layer is initialised by pre-trained word embeddings or by BERT.

**ConATT.** ConATT first encodes  $x^{(pre)}$ ,  $x^{(r)}$  and  $x^{(post)}$  separately using three bidirectional GRUs and three self-attention modules (Yang et al., 2016). For each input  $x^{(k)}$ , we first obtain  $h^{(k)}$  using a BiGRU:  $h^{(k)} = \text{BiGRU}(x^{(k)})$ . Subsequently, given the total  $M$  steps in  $h^{(k)}$ , we first calculate the attention weight  $\alpha_j^{(k)}$  at each step  $j$  by:

$$\alpha_j^{(k)} = \frac{\exp(e_j^{(k)})}{\sum_{m=1}^M \exp(e_m^{(k)})}, \quad (3)$$

where  $e_j^{(k)} = v_a^{(k)T} \tanh(W_a^{(k)} h_j^{(k)})$ ,  $v_a$  is the attention vector and  $W_a$  is the weight in the attention

layer. The context representation of  $x^{(k)}$  is then the weighted sum of  $h^{(k)}$ :  $c^{(k)} = \sum_{j=1}^N \alpha_j^{(k)} h^{(k)}$ .

After obtaining  $c^{(pre)}$ ,  $c^{(r)}$  and  $c^{(post)}$ , we concatenate them with the target entity embedding  $x^{(r)}$ , and pass it through a feed forward network to obtain the final representation:

$$R = \text{ReLU}(W_f[c^{(pre)}, c^{(r)}, c^{(post)}]), \quad (4)$$

where  $[\cdot, \cdot]$  represents a concatenation operation. The prediction is made using Equation 2. The input layer of ConATT is initialised either randomly or by pre-trained word embeddings.

## 5 Hypotheses

OntoNotes reflects a broader range of RE use and is, therefore, more appropriate as a source of insights into the human use of REs. Thus, it is plausible to expect that the ‘‘unexpected results’’ of §1 will not occur when assessing RFS models (see §4) on OntoNotes. More specifically, we expect:

- $\mathcal{H}_1$  models that incorporate pre-training (i.e., pre-trained word embeddings and BERT, which has been proved to be effective in many NLP tasks) work better than those that do not;
- $\mathcal{H}_2$  ConATT, which has been shown to perform well on both REG (Castro Ferreira et al., 2018a) and co-reference resolution (Yin et al., 2018), works better than c-RNN;
- $\mathcal{H}_3$  models that learn more useful linguistic information (confirmed by probing experiments) perform better than those that learn less.

Comparing Chinese and English, we can see in Table 2 that Chinese has an additional category compared to English, namely ZP. Given the theory that Chinese speakers process ZPs in the same way as pronouns (Yang et al., 1999), we expect:

- $\mathcal{H}_4$  RFS models that work well in English would also work well in Chinese.

Additionally, since Chinese relies more on context than English (see §1), it is plausible to expect:

- $\mathcal{H}_5$  Chinese RFS models would benefit more from the use of contextual representations (i.e., BERT) than English RFS models.

Model	4-way			3-way			2-way		
	P	R	F	P	R	F	P	R	F
XGBoost	48.96	49.69	49.12	67.78	65.78	66.44	79.11	78.01	78.42
c-RNN	65.45	60.59	62.38	68.19	69.19	68.55	76.66	75.23	75.70
+Glove	<u>66.06</u>	<u>63.39</u>	<u>64.56</u>	<u>69.94</u>	<u>70.14</u>	<u>70.01</u>	<u>77.61</u>	<u>76.31</u>	<u>76.67</u>
+BERT	<b>73.57</b>	<b>75.94</b>	<b>74.59</b>	<b>80.53</b>	<b>81.81</b>	<b>81.03</b>	<b>87.21</b>	<b>86.97</b>	<b>87.08</b>
			(+19.57%)			(+18.21%)			(+15.03%)
ConATT	61.29	62.21	61.58	66.34	65.87	66.01	73.19	73.21	73.19
+Glove	63.71	61.70	62.51	67.18	66.88	67.00	75.17	74.48	74.75

Table 4: Evaluation results of the English RFS systems on OntoNotes-EN with lexical tags. Best results are **boldfaced**, whereas the second best results are underlined. “P”, “R” and “F” stand for macro-averaged precision, recall and F1 score. Each percentage below the F-score of BERT indicates how much c-RNN gains from using BERT compared to not using BERT.

Model	5-way			4-way			3-way			2-way		
	P	R	F	P	R	F	P	R	F	P	R	F
XGBoost	38.17	40.06	34.59	46.16	44.12	41.29	56.19	54.64	51.98	64.5	79.56	63.67
c-RNN	52.42	48.49	49.62	54.60	54.65	54.19	56.78	53.50	54.68	67.66	62.89	64.59
+SGNS	54.54	51.27	51.56	<u>57.78</u>	<u>56.75</u>	<u>57.16</u>	<u>59.57</u>	<u>56.19</u>	<u>57.46</u>	<u>67.74</u>	<u>65.33</u>	<u>66.37</u>
+BERT	<b>64.99</b>	<b>63.60</b>	<b>63.85</b>	<b>68.22</b>	<b>69.48</b>	<b>68.17</b>	<b>70.36</b>	<b>68.60</b>	<b>69.13</b>	<b>78.35</b>	<b>73.51</b>	<b>75.59</b>
			(+28.68%)			(+25.80%)			(+26.43%)			(+17.03%)
ConATT	51.78	48.28	49.25	54.27	53.08	52.98	53.67	49.47	50.79	63.25	56.92	58.28
+SGNS	<u>55.44</u>	<u>52.13</u>	<u>53.09</u>	55.88	54.94	54.18	55.01	53.06	53.87	64.98	61.38	62.69

Table 5: Evaluation results of the Chinese RFS systems on OntoNotes-ZH.

## 6 Experiments

In what follows, we first provide an overview of the implementation details of the RFS models. To understand what linguistic information can be learnt by each model, we introduce a series of probing experiments. We then discuss the performance of these models and answer the hypotheses.

### 6.1 Baseline and Implementation Details

Following Chen et al. (2021), we used a feature-based model, XGBoost (Chen et al., 2015), as our baseline. For pre-trained word embeddings, we used Glove (Pennington et al., 2014) for English and SGNS (Li et al., 2018) for Chinese; for BERT, we used “bert-base-cased” for English and “bert-base-chinese” for Chinese.<sup>5</sup> Since Chinese BERT is a character-based model, we use all Chinese

<sup>5</sup>(1) English Glove: <https://nlp.stanford.edu/projects/glove/>; (2) Chinese SGNS: <https://github.com/Embedding/Chinese-Word-Vectors>; (3) English BERT: [huggingface.co/bert-base-cased](https://huggingface.co/bert-base-cased); and (4) Chinese BERT: <https://huggingface.co/bert-base-chinese>.

models character-based. The results of the word-based models can be found in Appendix B.

We tuned the hyper-parameters of each of our neural models on the development set and chose the setting with the best macro F1 score. For training, we used a single Tesla V100. For the baseline XGBoost models, we set the learning rate to 0.05, the minimum split loss to 0.01, the maximum depth of a tree to 5, and the sub-sample ratio of the training instances to 0.5. We report macro-averaged precision, recall, and F1 on the test set. We run each model 5x and report the average performance.

### 6.2 Probing RFS Models

To test the hypotheses in §5 (especially  $\mathcal{H}_3$ ), we probed each RFS model using probing classifiers. Specifically, after training an RFS model, we extracted its hidden representations and used them to train a probing classifier for a particular linguistic feature. The performance of the probing classifier indicates how well the RFS model learns the feature (Belinkov et al., 2017; Giulianelli et al., 2018).

**Probing Tasks.** We used the probing tasks defined in Chen et al. (2021). These tasks pertain

Model	Type	DisStat	SenStat	Syn	DistAnt	IntRef	LocPro	GloPro
Random	-	49.99 (49.77)	33.06 (32.27)	50.10 (50.10)	25.17 (23.75)	33.09 (32.40)	49.94 (48.21)	50.38 (49.53)
Majority	-	55.95 (35.88)	44.05 (20.39)	50.14 (33.39)	44.05 (15.29)	44.05 (20.39)	68.08 (40.50)	63.08 (38.68)
c-RNN	4-way	64.73 (63.39)	54.41 (50.76)	74.73 (74.67)	51.66 (36.31)	50.52 (44.81)	74.57 (67.86)	63.89 (50.32)
	3-way	64.24 (63.30)	53.94 (50.45)	75.57 (75.55)	52.02 (36.78)	49.76 (42.83)	74.96 (68.26)	64.00 (49.71)
	2-way	64.45 (63.31)	53.55 (49.72)	73.90 (73.82)	51.55 (35.75)	49.67 (43.03)	73.50 (65.72)	63.39 (45.76)
c-RNN +GloVe	4-way	65.00 (64.24)	54.40 (51.39)	76.75 (76.75)	51.95 (37.09)	50.65 (44.94)	74.25 (67.26)	64.14 (51.44)
	3-way	65.17 (64.44)	55.14 (52.69)	78.06 (78.06)	52.81 (37.55)	50.73 (45.89)	75.46 (70.66)	64.67 (53.28)
	2-way	65.07 (64.26)	53.55 (49.34)	75.22 (75.06)	51.20 (35.87)	50.78 (45.04)	73.91 (67.22)	63.26 (47.49)
c-RNN +BERT	4-way	86.00 (85.67)	72.17 (69.46)	79.83 (79.73)	66.53 (50.36)	69.85 (65.99)	82.32 (80.08)	68.47 (60.06)
	3-way	83.74 (83.42)	71.56 (68.90)	81.17 (81.15)	65.35 (49.10)	68.03 (63.62)	85.05 (82.38)	67.82 (61.93)
	2-way	81.82 (81.12)	69.33 (67.07)	78.05 (77.89)	63.46 (47.97)	65.11 (62.06)	81.85 (77.45)	66.35 (53.37)
ConATT	4-way	64.37 (62.95)	52.20 (46.63)	73.37 (73.34)	49.74 (33.33)	49.55 (43.52)	74.04 (66.30)	63.57 (48.89)
	3-way	64.28 (61.87)	51.96 (45.92)	74.79 (74.76)	49.25 (31.91)	49.21 (41.64)	73.89 (67.51)	63.25 (48.61)
	2-way	62.07 (59.46)	49.45 (41.73)	64.44 (63.72)	48.05 (30.18)	47.85 (40.85)	71.24 (59.96)	63.32 (47.51)
ConATT +GloVe	4-way	65.39 (63.41)	53.51 (50.49)	79.96 (79.95)	51.51 (36.03)	50.52 (43.17)	76.05 (70.27)	63.79 (49.86)
	3-way	63.72 (61.79)	52.13 (45.39)	79.03 (79.00)	49.48 (33.03)	49.43 (41.53)	74.86 (68.46)	63.31 (48.97)
	2-way	63.77 (61.56)	50.73 (44.35)	74.20 (73.97)	48.77 (31.53)	49.24 (42.81)	72.31 (63.31)	63.15 (48.39)

Table 6: Results of the English RFS models on each probing task on the OntoNotes-EN dataset. A in A(B) is the accuracy and B is the macro F1.

to four classes of features, namely referential status (DisStat and SenStat), syntactic position (Syn), recency (DistAnt and IntRef), and discourse structure prominence (LocPro, GloPro). These features have been shown to matter for RFS in linguistic literature (Ariel, 1990; Gundel et al., 1993; Arnold, 2010; von Heusinger and Schumacher, 2019). The definition of each probing task is as follows: (1) **DisStat**: This feature has 2 values: (a) `discourse-old` (the entity appeared in the previous context), and (b) `discourse-new` (it did not); (2) **SenStat**: The sentence-level referential status feature has 3 values: (a) `sentence-new` (the RE is the first mention of the entity in the sentence), (b) `sentence-old` (the RE is not the first mention of the entity in the sentence), and (c) `first-mention` (the RE is the first mention of the entity in the discourse); (3) **Syn**: The syntax probing task is a binary classification task with val-

ues (a) `subject` and (b) `object`; (4) **DistAnt**: It contains four values: the entity and its antecedent are (a) `in same sentence`, (b) `one sentence apart`, (c) `more than one sentence apart`, and (d) the entity is a `first-mention` (to distinguish first mentions from subsequent mentions); (5) **IntRef**: This feature asks whether there is an intervening referent between the target RE and its nearest antecedent. There are 3 possible values: (a) the target entity is a `first-mention`, (b) the previous RE refers to the same entity, and (c) the previous RE refers to a different entity; (6) **LocPro**: is a hybrid of DisStat and Syn. It has 2 values: (a) `locally prominent`, and (b) `not locally prominent`. An entity is said to be locally prominent if it is both “discourse-old” and “realised as a subject”; (7) **GloPro**: This is a binary feature with two possible values: (a) `globally prominent`, and (b) `not globally prominent`. The

Model	Type	DisStat	SenStat	Syn	DistAnt	IntRef	LocPro	GloPro
Random	-	50.20 (49.93)	33.18 (32.70)	50.11 (49.79)	25.02 (23.81)	33.56 (33.01)	50.12 (46.44)	50.00 (44.27)
Majority	-	57.30 (36.43)	42.70 (19.95)	57.79 (36.62)	42.70 (14.96)	42.70 (19.95)	76.27 (43.27)	81.13 (45.09)
c-RNN	5-way	65.14 (62.80)	48.85 (45.89)	76.79 (75.94)	46.50 (28.49)	48.72 (45.78)	79.12 (65.54)	82.57 (52.03)
	4-way	64.60 (61.80)	48.76 (43.39)	76.30 (74.74)	45.75 (27.73)	47.84 (44.65)	79.11 (63.44)	81.97 (46.64)
	3-way	63.55 (61.19)	47.52 (41.52)	77.13 (76.11)	45.69 (26.43)	46.60 (41.13)	78.11 (61.70)	82.02 (45.76)
	2-way	61.32 (58.06)	46.09 (36.30)	77.95 (76.96)	45.23 (24.11)	45.71 (36.49)	77.86 (58.82)	82.11 (45.54)
c-RNN +SGNS	5-way	65.75 (63.52)	50.24 (47.24)	78.36 (77.28)	47.48 (30.71)	49.66 (46.13)	79.33 (66.11)	82.21 (50.37)
	4-way	66.07 (62.90)	50.93 (46.96)	78.41 (77.18)	47.64 (30.78)	50.57 (47.81)	80.11 (66.16)	82.24 (48.20)
	3-way	64.70 (62.87)	48.24 (42.54)	79.02 (77.81)	46.27 (27.51)	47.48 (43.59)	79.35 (64.17)	82.01 (46.11)
	2-way	62.48 (60.45)	46.30 (38.24)	78.50 (77.12)	45.38 (24.27)	44.82 (37.61)	77.72 (64.09)	81.93 (46.12)
c-RNN +BERT	5-way	76.17 (75.20)	59.58 (57.07)	79.42 (78.68)	56.14 (39.54)	59.89 (57.69)	81.86 (70.93)	82.05 (55.17)
	4-way	75.32 (73.96)	59.69 (57.66)	78.86 (78.15)	56.66 (37.12)	60.27 (56.90)	81.95 (69.68)	81.96 (46.60)
	3-way	74.46 (73.77)	58.41 (56.29)	80.48 (79.67)	55.91 (35.77)	59.39 (55.96)	82.71 (73.24)	81.91 (45.59)
	2-way	69.20 (68.10)	55.16 (52.08)	80.68 (79.84)	51.74 (29.71)	51.73 (52.36)	81.43 (71.30)	82.05 (45.07)
ConATT	5-way	65.36 (62.33)	48.50 (43.17)	75.14 (73.94)	46.44 (28.92)	48.25 (45.05)	78.90 (63.99)	82.02 (47.16)
	4-way	65.07 (61.91)	48.40 (43.15)	70.38 (67.48)	45.95 (26.41)	48.16 (44.15)	77.89 (57.31)	82.22 (47.27)
	3-way	62.93 (59.54)	45.14 (39.55)	70.38 (68.78)	43.85 (24.47)	45.28 (39.13)	77.34 (55.27)	82.06 (45.73)
	2-way	60.55 (52.10)	44.21 (32.85)	68.33 (65.67)	43.75 (21.78)	44.36 (32.66)	76.37 (49.38)	82.07 (45.35)
ConATT +SGNS	5-way	66.65 (63.48)	49.57 (45.60)	78.18 (77.27)	46.34 (29.77)	49.76 (46.84)	79.35 (64.16)	81.65 (50.72)
	4-way	66.09 (61.97)	49.43 (44.63)	75.87 (74.65)	46.04 (28.19)	49.20 (46.61)	79.50 (64.49)	82.22 (47.27)
	3-way	62.84 (58.79)	46.51 (38.78)	75.15 (74.09)	44.99 (24.66)	45.76 (38.51)	78.12 (60.19)	82.06 (45.73)
	2-way	62.65 (60.09)	46.76 (39.53)	74.17 (72.90)	44.31 (22.13)	44.84 (34.88)	77.53 (61.43)	82.07 (45.35)

Table 7: Results of the Chinese RFS models on each probing task on the OntoNotes-ZH.

most frequent entity in a text is marked as globally prominent.

**Probing Classifiers.** Following [Chen et al. \(2021\)](#), we use a logistic regression classifier as our probing classifier. When probing, we use  $R$  (see Equation 1 and 4) of the models with the best RFS performance on the development set as input representations. We evaluate probing classifiers using the accuracy and macro-averaged F1 scores. We run each probing classifier 5 times and report the averaged value. We use 2 baselines: (1) random: it randomly assigns a label to each input; and (2)

majority: it assigns the most frequent label in the given probing task to the inputs.

### 6.3 Experimental Results

**Results on each Language.** Table 4 and 5 show the results of each model on OntoNotes-EN and OntoNotes-ZH. In both languages, all neural RFS models defeat the baseline in 4-way and 5-way classifications, while models that does not use BERT have on-par or worse performance in 3-way and 2-way classifications. This suggests that feature-based models with linguistically-informed features

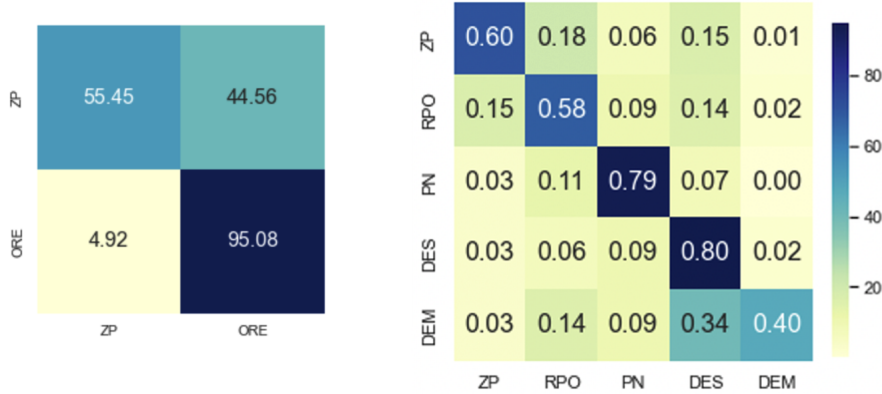


Figure 1: Confusion Matrix for Chinese 2-way c-RNN +BERT (left) and 5-way c-RNN +BERT (right) where ORE is overt RE, PRO is pronoun, PN is proper name, DES is description, and DEM is demonstrative.

Model	4-way			3-way			2-way		
	P	R	F	P	R	F	P	R	F
c-RNN	50.77	45.89	46.38	<u>60.83</u>	59.56	<u>59.94</u>	73.33	<u>72.58</u>	<u>72.84</u>
+Glove	<u>53.47</u>	<b>49.49</b>	<b>50.44</b>	<b>61.72</b>	<b>60.66</b>	<b>60.98</b>	<b>75.06</b>	<b>73.96</b>	<b>74.32</b>
ConATT	52.32	45.88	46.89	59.66	58.71	59.08	71.86	71.38	71.56
+Glove	<b>54.55</b>	<u>47.56</u>	<u>48.14</u>	59.75	<u>60.05</u>	59.85	<u>73.84</u>	72.32	72.66

Table 8: Evaluation results of RFS systems on OntoNotes-EN with entity tags.

can build remarkably good systems for RFS, but their performance decreases dramatically as the task becomes more fine-grained.

As for  $\mathcal{H}_1$ , word embeddings always improve RFS performance. The RFS tasks in both languages benefit strongly from using BERT. For instance, if we compare c-RNN +BERT to c-RNN for the full RFS tasks (i.e., 5-way classification in Chinese and 4-way classification in English), c-RNN +BERT improves the performance (F1 score) from 62.38 to 74.59 in English and from 49.62 to 63.85 in Chinese.

In both languages, contrary to our expectation  $\mathcal{H}_2$ , ConATT performs worse than c-RNN. Probing results presented in Tables 6 and 7 provide some explanations: in English, ConATT learns less information about referential status, syntactic position, and recency than c-RNN, and in Chinese, ConATT performs significantly worse than c-RNN in acquiring information about syntactic position.

Meanwhile, the results of the probing experiments suggest that expectation  $\mathcal{H}_3$ , that models that learn more useful information perform better, is true. Further evidence is provided by the observations that (1) BERT defeats all other models in almost all probing tasks and, therefore, defeats all

other models by a large margin; and (2) pre-trained word embeddings (GloVe and SGNS) help each model learn significantly more information about almost every feature except GloPro, and, therefore, improve RFS performance.

**English vs. Chinese.** In line with our expectation  $\mathcal{H}_4$ , models that work well for English also work well on modelling ZP in Chinese. However, deciding whether to use a ZP or an overt RE is generally harder than pronominalisation. For example, c-RNN achieves an F-score of 75.7 for the English 2-way task, while it is only 64.6 for Chinese.

Figure 1 shows the confusion matrices for the Chinese c-RNN +BERT 2-way and 5-way classifications. By comparing them, we find that fine-grained supervision helps with the choice between ZPs and overt REs. Focusing on 5-way classification, ZPs are quite often confused with pronouns. Linguistic theory suggests that attenuated forms such as pronouns and ZPs happen when the target referent is salient enough (Ariel, 2001). It is understandable that ZPs and pronouns are confused because it is hard for a model to make such a fine-grained decision about when the target referent is salient enough for pronominalisation but not for



pro-drop.

The results of both Chinese and English RFS tasks improve dramatically when using the contextual language model BERT. This is consistent with the probing results: in both languages, BERT helps a lot in acquiring all linguistic information except GloPro. To test our last hypothesis  $\mathcal{H}_5$ , we compute how much c-RNN gains from using BERT compared to not using BERT and report the numbers in Table 4 and 5. On average, c-RNN gains 17.60% from using BERT in English and 24.48% in Chinese. The results suggest that Chinese RFS benefits more from using BERT than English RFS. Nevertheless, we still cannot make conclusive statements about  $\mathcal{H}_5$ . Strictly speaking, these percentages are not directly comparable and the comparison cannot be fully controlled because for example: (1) the data is not fully parallel, and (2) the RFS tasks defined for the two languages differ from each other. For instance, unlike English RFS, Chinese RFS considers an extra category, namely ZP.

**Lexical Tags vs. Entity Tags.** To chart the benefits of lexical tags, we also ran models of Chen et al. (2021) on a version of OntoNotes-EN, in which entity tags are used instead of lexical tags. The results are presented in Table 8. Comparing this table to Table 4, we see that the performance of each model decreases significantly when the entity tags are used, especially in the 4-way and 3-way classifications. For example, the F-score of the 4-way c-RNN + GloVe model decreases from 64.56 to 50.44. As expected, these tags prevent the models from handling unseen entities.

## 7 Conclusion

To address the problem that all previous assessments of neural REG/RFS models were only tested on WebNLG, we built a realistic multilingual (English and Chinese) dataset based on the OntoNotes dataset, modified the RFS models accordingly and assessed them on this dataset. Although a few outcomes were against our expectations (e.g. ConATT performed worse than c-RNN), we found that our results are explainable using probing experiments. For example, models that use BERT, which performs best in the probing experiments, also beats all other models in RFS.

We also compared the English RFS to the Chinese RFS, which uses ZPs frequently and depends more on context than English. We found that RFS models that work for English can also model Chi-

nese ZPs. In line with the idea that Chinese relies more on context than English, the results suggest that Chinese RFS models benefited more from using contextualised language model BERT than those of English. However, as discussed, this needs to be further verified with more controlled experiments.

In future, we plan to extend our work from the following three perspectives: (1) testing other model explanation techniques, e.g., probing classifiers with control tasks (Hewitt and Liang, 2019) and attention analysis (Bibal et al., 2022); (2) assessing and probing RFS models on other languages (such as languages that are morphologically rich); and (3) trying more probing tasks based on factors that influence RFS, such as animacy, competition and positional attributes (see Same and van Deemter (2020) for more details).

## References

- Mira Ariel. 1990. *Accessing Noun-Phrase Antecedents*. Routledge.
- Mira Ariel. 2001. [Accessibility theory: An overview](#). In Ted Sanders, Joost Schilperoord, and Wilbert Spooren, editors, *Text Representation: Linguistic and psycholinguistic aspects*, volume 8, page 29. John Benjamins Publishing Company.
- Jennifer E Arnold. 2010. How speakers refer: The role of accessibility. *Language and Linguistics Compass*, 4(4):187–203.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaou Wang, Thomas François, and Patrick Watrin. 2022. [Is attention explanation? an introduction to the debate](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3889–3900, Dublin, Ireland. Association for Computational Linguistics.
- Meng Cao and Jackie Chi Kit Cheung. 2019. [Refering expression generation using entity profiles](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3163–3172, Hong Kong, China. Association for Computational Linguistics.

- Thiago Castro Ferreira, Diego Moussallem, Ákos Kádár, Sander Wubben, and Emiel Krahmer. 2018a. **NeuralREG: An end-to-end approach to referring expression generation**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1959–1969, Melbourne, Australia. Association for Computational Linguistics.
- Thiago Castro Ferreira, Diego Moussallem, Emiel Krahmer, and Sander Wubben. 2018b. **Enriching the WebNLG corpus**. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 171–176, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Chen Chen and Vincent Ng. 2016. **Chinese zero pronoun resolution with deep neural networks**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788, Berlin, Germany. Association for Computational Linguistics.
- Guanyi Chen. 2022. *Computational Generation of Chinese Noun Phrases*. Ph.D. thesis, Utrecht University.
- Guanyi Chen, Fahime Same, and Kees van Deemter. 2021. **What can neural referential form selectors learn?** In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 154–166, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Guanyi Chen, Kees van Deemter, and Chenghua Lin. 2018. **Modelling pro-drop with the rational speech acts model**. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 159–164, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, et al. 2015. Xgboost: extreme gradient boosting. *R package version 0.4-2*, 1(4).
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. **On the properties of neural machine translation: Encoder–decoder approaches**. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.
- Rossana Cunha, Thiago Castro Ferreira, Adriana Pagano, and Fabio Alves. 2020. **Referring to what you know and do not know: Making referring expression generation models generalize to unseen entities**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2261–2272, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. **Creating training corpora for NLG micro-planners**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 179–188, Vancouver, Canada. Association for Computational Linguistics.
- Albert Gatt and Emiel Krahmer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *J. Artif. Int. Res.*, 61(1):65–170.
- Mario Giulianelli, Jack Harding, Florian Mohnert, Dieuwke Hupkes, and Willem Zuidema. 2018. **Under the hood: Using diagnostic classifiers to investigate and improve how language models track agreement information**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 240–248, Brussels, Belgium. Association for Computational Linguistics.
- Jeanette K Gundel, Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language*, pages 274–307.
- John Hewitt and Percy Liang. 2019. **Designing and interpreting probes with control tasks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- C-T James Huang. 1984. On the distribution and reference of empty pronouns. *Linguistic inquiry*, pages 531–574.
- Shen Li, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du. 2018. **Analogical reasoning on Chinese morphological and semantic relations**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 138–143, Melbourne, Australia. Association for Computational Linguistics.
- Richard Newnham. 1971. *About Chinese*. Penguin Books Ltd.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. **GloVe: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. **CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes**. In *Joint Conference on*

- EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Ehud Reiter and Robert Dale. 2000. *Building Natural Language Generation Systems*. Studies in Natural Language Processing. Cambridge University Press.
- Fahime Same, Guanyi Chen, and Kees Van Deemter. 2022. [Non-neural models matter: a re-evaluation of neural referring expression generation systems](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5554–5567, Dublin, Ireland. Association for Computational Linguistics.
- Fahime Same and Kees van Deemter. 2020. [A linguistic perspective on reference: Choosing a feature set for generating referring expressions in context](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4575–4586, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Klaus von Heusinger and Petra B Schumacher. 2019. Discourse prominence: Definition and application. *Journal of Pragmatics*, 154:117–127.
- Chin Lung Yang, Peter C Gordon, Randall Hendrick, and Jei Tun Wu. 1999. Comprehension of referring expressions in chinese. *Language and Cognitive Processes*, 14(5-6):715–743.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California. Association for Computational Linguistics.
- Qingyu Yin, Yu Zhang, Weinan Zhang, Ting Liu, and William Yang Wang. 2018. [Zero pronoun resolution with attention-based neural network](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 13–23, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

## A Results on the Whole OntoNotes-ZH Dataset

The Chinese experiments in this paper were conducted on a subset of the original OntoNotes, each text of which contains less than 512 characters, since Chinese BERT can only accept texts shorter than 512 characters. For reference, we also tested models other than BERT on the whole OntoNotes-ZH dataset. In the whole OntoNotes-ZH dataset, there are 73607, 10008, and 12096 samples in the training, development, and test sets, respectively. Table 9 shows the results of the word-based Chinese RFS models on the whole OntoNotes-ZH dataset.

Comparing Table 9 with Table 10, we observe that the results are quite similar. The only exception is that the performance of  $c$ -RNN decreases from 55.16 to 53.86 in the 3-way classification, while the performance of ConATT does not change much.

## B Results of Using Word-based Models on OntoNotes-ZH

To conduct a fair comparison between BERT and other models, we built all our Chinese RFS models character-based. To justify this decision, we also test word-based models on OntoNotes-ZH. Table 5 shows the results of the word-based Chinese models.

Comparing the results in Table 10 and Table 5, there are slight differences, but these differences do not change our conclusions. For example, all models still perform worse than  $c$ -RNN +BERT by a large margin. ConATT can slightly defeat  $c$ -RNN in the 3-way and 2-way classifications but performs significantly worse in other settings.

Model	5-way			4-way			3-way			2-way		
	P	R	F	P	R	F	P	R	F	P	R	F
c-RNN	52.36	47.91	48.97	54.14	52.40	53.06	55.30	52.99	53.86	64.88	62.81	63.68
+SGNS	56.67	53.82	54.30	59.38	57.40	58.23	59.58	56.66	57.78	67.75	66.28	66.91
ConATT	50.41	45.45	46.86	51.27	49.80	50.35	59.06	54.43	56.11	63.71	63.75	63.73
+SGNS	52.33	48.60	49.37	53.48	51.64	52.38	60.53	56.18	57.69	67.86	64.97	65.95

Table 9: Evaluation results of our word-based Chinese RFS systems on the whole OntoNotes-ZH dataset.

Model	5-way			4-way			3-way			2-way		
	P	R	F	P	R	F	P	R	F	P	R	F
c-RNN	51.13	47.14	48.63	54.70	54.02	54.18	57.63	53.79	55.16	66.19	63.22	64.40
+SGNS	53.40	53.33	53.16	57.91	59.12	58.19	60.17	57.49	58.52	70.87	65.22	67.30
ConATT	48.52	45.15	46.26	56.34	49.92	49.26	56.24	55.70	55.94	65.33	64.28	64.75
+SGNS	50.58	47.04	48.31	54.68	51.85	52.62	59.93	55.79	57.32	67.15	65.29	66.11

Table 10: Evaluation results of our word-based Chinese RFS systems on a subset of the original OntoNotes-ZH dataset, each text of which contains less than 512 characters.