

Non-neural Models Matter: A Re-evaluation of Neural Referring Expression Generation Systems

Fahime Same^{♡*}, Guanyi Chen^{♠*}, and Kees van Deemter[♠]

[♡]Department of Linguistics, University of Cologne

[♠]Department of Information and Computing Sciences, Utrecht University

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

Abstract

In recent years, neural models have often outperformed rule-based and classic Machine Learning approaches in NLG. These classic approaches are now often disregarded, for example when new neural models are evaluated. We argue that they should not be overlooked, since for some tasks, well-designed non-neural approaches achieve better performance than neural ones. In this paper, the task of generating referring expressions in linguistic context is used as an example. We examined two very different English datasets (WEBNLG and WSJ), and evaluated each algorithm using both automatic and human evaluations. Overall, the results of these evaluations suggest that rule-based systems with simple rule sets achieve on-par or better performance on both datasets compared to state-of-the-art neural REG systems. In the case of the more realistic dataset, WSJ, a machine learning-based system with well-designed linguistic features performed best. We hope that our work can encourage researchers to consider non-neural models in future.

1 Introduction

Natural Language Generation (NLG) is concerned with the generation of natural language text from non-linguistic input (Gatt and Krahmer, 2018). One step in a classic generation pipeline (Reiter and Dale, 2000) is Referring Expression Generation (REG, Krahmer and van Deemter (2019) for an overview). REG has important practical value for commercial natural language generation (Reiter, 2017), computer vision (Mao et al., 2016), and robotics (Fang et al., 2015), for example. It has also been used as a tool to understand human language use (van Deemter, 2016). REG contains two different problems. One is to find a set of attributes to single out a referent from a set (also

called *one-shot REG*). The other is to generate referring expressions (REs) to refer to a referent at different points in a discourse (Belz and Vargas, 2007). We will focus on the latter task. We call this the *REG-in-context* task.

In earlier works, REG is often tackled in two steps (Henschel et al., 2000; Krahmer and Theune, 2002). The first step decides the form of an RE. For example, whether a reference should be a proper name (“*Marie Skłodowska-Curie*”), a description (“*the physicist*”), or a pronoun (“*she*”) at a given point in the context. The second step is concerned with content selection, i.e., the different ways in which a referential form can be realised. For example, to generate a description of *Marie Curie*, the REG system decides whether it is sufficient to mention her profession (i.e., “*the physicist*”) or whether it is better to mention her nationality as well (i.e., “*a Polish-French physicist*”).

Thanks to the rapid development of deep learning techniques, recent NLG models are able to generate RE in an End2End (E2E) manner, i.e., to tackle the selection of form and content simultaneously (Castro Ferreira et al., 2018a; Cao and Cheung, 2019; Cunha et al., 2020). The task of End2End (E2E) REG was proposed by Castro Ferreira et al. (2018a), who extracted a corresponding corpus from the WebNLG corpus (Castro Ferreira et al., 2018b)¹. Grounding on the WEBNLG dataset, they proposed a neural REG system built on a *sequence-to-sequence with attention* model. Their automatic and human evaluation results suggested that neural REG systems significantly outperform rule-based and feature-based machine learning (ML) baselines. However, it can be argued that Castro Ferreira et al. did not use very strong baselines for their comparison: OnlyName is a rule-based system that always generates a proper name given an entity, and Ferreira is a feature-based model

* Equal contribution. Order determined by swapping the order in Chen et al. (2021).

¹We refer to this extracted REG corpus as WEBNLG.

that uses Naive Bayes with only 3 simple features².

We present several rule-based and feature-based baselines to examine how neural models perform against “well-designed” non-neural alternatives. Note that a well-designed model is not necessarily complex. For example, it can be a rule-based system with one or two simple, “well-designed” rules. Since one of the advantages of neural E2E models is that they require little effort for feature engineering, we used two types of baselines, namely models that require minimal expert effort and models that use more demanding (but linguistically well-established) rules or features. Therefore, our main research question is: *Do state-of-the-art neural REG models always perform better than rule-based and machine learning-based models?*

To answer this question fairly, we consider the amount of resources used by each model. For example, the neural models require fewer human resources when it comes to linguistic expertise and annotation, but they require input from Deep Learning experts. Resources such as computing power and data needs should also be considered.

Another issue with previous studies concerns the datasets that were used: in WEBNLG, approximately 99.34% of entities in the test set also appear in the training set; consequently, evaluations using WEBNLG do not take unseen entities into consideration. Furthermore, since many sentences in WEBNLG are paraphrases of one another, evaluating neural models on WEBNLG alone may overestimate their performance. [Castro Ferreira et al. \(2019\)](#) recently extended WEBNLG to include unseen domains that contain many unseen entities³, and [Cunha et al. \(2020\)](#) have developed new models to handle them. Their test set has two subsets: one consists of documents 99.34% of whose entities are *seen*, while the other consists of documents 92.81% of whose entities are *unseen*. This arguably makes the data in WEBNLG unrealistic (see §2 for discussion). Therefore, we created what we believe to be a more realistic dataset based on the Wall Street Journal (WSJ) portion of the OntoNotes corpus ([Hovy et al., 2006](#); [Weischedel et al., 2013](#))⁴.

²The human evaluation in [Cunha et al. \(2020\)](#) showed a slightly different result: the `OnlyName` model performed as well as the Neural REG models in terms of fluency, grammaticality, and adequacy. However, since their human evaluation involved only two subjects, these outcomes need to be approached with caution.

³We used version 1.5 of the WEBNLG dataset in <https://github.com/ThiagoCF05/webnlg>.

⁴We used Ontonotes 5.0 licensed by the Linguis-

We evaluate all models on both WEBNLG and WSJ, using automatic and human evaluation experiments. The human experiments included a total of 240 participants and 16920 judgments.

This paper is structured as follows: in §2 and §3, we describe the datasets used and the REG models. In §4, we provide a detailed description of our automatic and human evaluations. In §5 and §6, we compare the results across different dimensions and make suggestions for future studies. The code for reproducing the results in this article can be found at: <https://github.com/a-quei/neuralreg-re-evaluation>.

2 Task and Datasets

This section explains the REG-in-context task and the two English datasets used to conduct the experiments.

2.1 The REG-in-context Task

Given a text whose REs have not yet been generated, and given the intended referent for each of these REs, the REG-in-context task is to build an algorithm that generates all these REs.

Consider the delexicalised text in Table 1. Given the entity “*AWH_Engineering_College*”, REG selects an RE based on the entity and its pre-context (“*AWH_Engineering_College is in “Kuttikkattoor” , India in the state of Kerala .*”), and its post-context (“*has 250 employees and Kerala is ruled by Kochi . The Ganges River is also found in India.*”).

2.2 The WEBNLG Dataset

[Gardent et al. \(2017\)](#) introduced the WEBNLG corpus for evaluating NLG systems. Using crowdsourcing, each crowdworker was asked to write a description for a given Resource Description Framework (RDF) triple (Table 1). The number of triples varied from 1 to 7. This corpus was later enriched and delexicalised ([Castro Ferreira et al., 2018a,b](#)) to fit the REG-in-context task. [Castro Ferreira et al. \(2019\)](#) further extended WEBNLG and divided the documents into test sets *seen* (where all data are from the same domains as the training data) and *unseen* (where all data are from different domains than the training data). This results in that almost all entities from the seen test set appear in the training set (9580 out of 9644), while only a

tic Data Consortium (LDC) <https://catalog.ldc.upenn.edu/LDC2013T19>.

Triples:

(AWH_Engineering_College, country, India)
(Kerala, leaderName, Kochi)
(AWH_Engineering_College, academicStaffSize, 250)
(AWH_Engineering_College, state, Kerala)
(AWH_Engineering_College, city, "Kuttikkattoor")
(India, river, Ganges)

Text: AWH Engineering College is in Kuttikkattoor, India in the state of Kerala. The school has 250 employees and Kerala is ruled by Kochi. The Ganges River is also found in India.

Delexicalised Text:

Pre-context: AWH_Engineering_College is in "Kuttikkattoor", India in the state of Kerala .

Target Entity: AWH_Engineering_College

Post-context: has 250 employees and Kerala is ruled by Kochi . The Ganges River is also found in India .

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

few entities from the unseen test set appear in the training set (688 out of 9644). Note that the maximum number of triples in the unseen set is five. So, one would expect the data in the unseen set to be less complex than the seen data.

We used version 1.5 of WEBNLG, which contains 67,027, 8278, and 19,210 REs in the training, development, and test sets. From the point of view of the present study, WEBNLG has some notable shortcomings. For a start, it consists of rather formal texts that may not reflect the everyday use of REs, and in which very simple syntactic structures dominate. The texts in WEBNLG also stand out for other reasons. For example, the texts are extremely short, with an average length of only 1.4 sentences. Consequently, as many as 85% of the REs are first-mentions, while 71% of the REs are proper names. Finally, in any given test sample, either more than 90% of the entities are seen or more than 90% are unseen. Realistic data should contain a reasonable amount of mixtures of seen and unseen entities. For all these reasons, we decided to test all algorithms on a second corpus as well.

2.3 The WSJ Dataset

Using the Wall Street Journal portion of the OntoNotes corpus, we constructed a new English REG dataset, following a similar approach as [Castro Ferreira et al. \(2018a\)](#). This corpus (WSJ) has very different characteristics from the WEBNLG. The WSJ consists of 582 newspaper articles containing 20,186, 2362 and 2781 REs in the training, development and test sets, respectively. The average length of the documents is 1189 words, and each document consists of 25 sentences on average. Furthermore, 23% of the instances are first-mention REs and the rest are subsequent mentions.

For each RE, we created its pre- and post-context at the local sentence-level and added K preceding and following sentences to the local context. We refer to K as the context length and set $K=2$ for this experiment. To create the dataset, we first delexicalised the REs. The dataset contains nearly 8000 coreferential chains. The REs in each chain were replaced with corresponding delexicalised expressions (similar to table 1). For delexicalisation, we used (1) the POS tag information, (2) the fine-grained annotation of the referential forms, and (3) the entity type of each referent. To delexicalise human REs, for example, we looked for concise but informative REs such as the combination of first and last names (e.g., "*Barack Obama*")⁵. When such an expression was found in a coreferential chain, its delexicalised version (tokens being separated by underscores, e.g., "*Barack_Obama*") was assigned to all REs in the chain. We then moved on to the next tag. Below is the order in which the human referents were searched and delexicalised: [firstname-lastname], [title-firstname-lastname], [modified firstname-lastname], [title-lastname], [lastname], [modified-lastname], [firstname]. For more details on the preparation of the WSJ documents and a delexicalised example, see Appendix A.

3 REG Models

In this section, we introduce the rule-based, ML-based, and the SOTA neural REG models. The term *ML-based* here refers to models that require feature engineering and follow a pipeline architecture.

⁵The reason for choosing concise and informative REs for delexicalisation is that these labels are also used in the realisation process.

3.1 Rule-based REG

Rule-based models have been widely used for generating REs in context (McCoy and Strube, 1999; Henschel et al., 2000). Here, we build rule-based systems for binary classification into two classes, namely pronominal and non-pronominal REs.

Simple Rule-based System (RREG-S). For the first rule-based system, we use 2 simple rules. The target entity r is realised as a pronominal RE if:

1. r is *discourse-old*;
2. r has no *competitor* in the current sentence and the previous sentence,

Otherwise, r is realised as a non-pronominal RE. An entity r is defined as discourse-old if it has been mentioned in the previous context. A competitor is an entity that can be referred to with the same pronoun as r .

We also build a dictionary that stores the pronouns associated with each entity. For seen entities, we extract pronouns from the training data. If an entity has multiple possible pronominal forms, we extract the most frequent one. For unseen entities, we determine their pronominal forms based on their meta-information, which is also used in E2E systems (Cunha et al., 2020). For example, if an entity in WEBNLG has the type PERSON and the gender FEMALE, we assign “*she*” to this entity.

For the surface realisation of each entity, we realise its non-pronominal form by replacing the underscores in the entity label with whitespaces (e.g., “*Adenan_Satem*” to “*Adenan Satem*”), as previously described by Castro Ferreira et al. (2018a). We realise the pronominal forms according to Castro Ferreira et al. (2016) by using the grammatical role of each entity (e.g., if the entity is in the object position, then we realise “*he*” as “*him*”).

Linguistically-informed Rule-based System (RREG-L). We build RREG-L by adopting a set of pronominalisation rules from Henschel et al. (2000). The fundamental concepts used by these rules are the idea of *local focus*, which is a simpler implementation of the Centering Theory (Grosz et al., 1995), and *parallelism*, i.e., whether r and its antecedent in the previous sentence have the same grammatical role (Henschel et al., 2000). The RREG-L is described in detail in Appendix B.

3.2 ML-based REG

The GREC Shared Task (Belz et al., 2010) triggered a plethora of ML-based models for building

REG-in-context (e.g., Greenbacker and McCoy, 2009; Hendrickx et al., 2008). These models differ from each other in the features and the ML algorithms they have used.

In this study, we build ML-based REG models using CatBoost (Prokhorenkova et al., 2018). It predicts whether a reference is realised as a pronoun, proper name, or description. Once the referential form is predicted, the next step is to select the content. The most frequent variant (with the same referential form as the predicted class) is selected in the training corpus given the referent and the full set of features. If no matching RE is found, a back-off method (Castro Ferreira et al., 2018a) is used, removing one feature at a time in order of importance. The order is calculated using the inherent feature importance method of the CatBoost algorithm. Depending on which features are used, we build two variants of ML-based models, namely ML-S and ML-L. The detailed list of the features used in these models can be found in Appendix C.

Features obtained by minimum effort (ML-S).

To find out what the upper bound is for a system that does not require any additional linguistic information or any additional annotation effort, we developed ML-S. In this model, we have relied only on the features that can be extracted directly from the corpus. Therefore, features such as grammatical role (which requires a syntactic parser) are not included in this model.

Linguistically Informed Features (ML-L).

To evaluate the upper bound performance of ML-based systems, we developed ML-L with the features that could affect the choice of referential form and could improve the overall accuracy of the REG systems suggested by the previous linguistic and computational studies (Ariel, 1990; Gundel et al., 1993; Brennan, 1995; Arnold and Griffin, 2007; Fukumura and Van Gompel, 2011; Kibrik et al., 2016; von Heusinger and Schumacher, 2019; Same and van Deemter, 2020). For example, we included features encoding grammatical role, recency, gender, and animacy in ML-L. Note that ML-L makes full use of the syntactic information⁶ and entity meta-information (e.g., GENDER and TYPE which are also used by both the rule-based systems and the neural models).

⁶The syntactic information is available in the annotations of the WSJ data set. We used spaCy to parse WEBNLG.

Model	RE Acc.↑	SED↓	BLEU↑	Text Acc.↑	Precision↑	Recall↑	F1↑
RREG-S	<u>54.60</u>	3.65	<u>72.05</u>	16.28	89.52	<u>77.57</u>	<u>82.28</u>
RREG-L	53.43	3.77	71.27	15.49	73.94	73.96	73.95
ML-S	54.35	3.70	70.89	15.43	71.70	63.52	66.39
ML-L	56.69	<u>3.66</u>	72.25	<u>16.36</u>	81.66	63.62	68.36
ATT+Copy	48.75	4.46	68.48	14.88	85.33	75.74	79.63
ATT+Meta	53.34	4.22	70.82	16.54	86.32	75.56	79.81
ProfileREG	40.96	7.40	61.04	11.39	<u>86.44</u>	87.40	86.91

Table 2: Automatic Evaluation Results on WEBNLG. Best results are **boldfaced**, whereas the second best are underlined. ↑ means the higher the metric, the better, while ↓ means the lower the better.

3.3 Neural REG

A limitation of the rule-based and ML-based models mentioned above is that they are not able to handle situations where an RE form (e.g., a proper name) can have multiple realisations, e.g., Lady Gaga/Stefani Germanotta. End2End NeuralREG can address this by generating REs from scratch. This study examines three NeuralREG systems that have been developed to deal with unseen entities as well. All of them were developed using the *sequence-to-sequence with attention* model (Bahdanau et al., 2014).

ATT+Copy. Cunha et al. (2020) proposed using three bidirectional LSTMs (Hochreiter and Schmidhuber, 1997) to encode a pre-context, a post-context, and the proper name of an entity (i.e., replacing underscores in entity labels with whitespaces) into three hidden vectors $h^{(pre)}$, $h^{(post)}$ and $h^{(r)}$, respectively. An auto-regressive LSTM-based decoder generates REs based on context vectors. To handle unseen entities, Cunha et al. used the copy mechanism, which allows the decoder to copy words from the contexts directly as output.

ATT+Meta. ATT+Meta (Cunha et al., 2020) used meta information of each entity to improve the quality of the generated REs. In each decoding step t , the context vector $v_t^{(c)}$ is concatenated with meta information embeddings before being fed to the decoder. In WEBNLG, meta information are the entity type $v^{(type)}$ and gender embeddings $v^{(gender)}$; while in WSJ, in addition to $v^{(type)}$ and $v^{(gender)}$, there is also plurality embedding $v^{(pl)}$.

ProfileREG. Cao and Cheung (2019) made ProfileREG to leverage the content of entity profiles extracted from Wikipedia. More specifically, instead of encoding the proper name of each entity, ProfileREG asks the entity encoder to encode the whole entity’s profile to obtain $h^{(r)}$. Note that since profiles of entities in WSJ are not accessible,

we evaluate ProfileREG only on WEBNLG.

4 Evaluation

We evaluated all the systems described in §3 on both WEBNLG and WSJ using automatic and human evaluations. We implemented the neural models based on the code of Cunha et al. (2020) and Cao and Cheung (2019)⁷. For WEBNLG, we used their original parameter setting, while for WSJ, we tuned the parameters on the development set and used the best parameter set.

To determine the optimal context length K of WSJ, we varied K from 1 to 5 sentences before and after the target sentence, then tested ATT+Meta on the development set with the different K contexts. It reaches the best performance when $K = 2$.

4.1 Automatic Evaluation

Metrics. Following Cunha et al. (2020), we evaluated REG systems from 3 angles. (1) **RE Accuracy** and **String Edit Distance** (SED, Levenshtein, 1966) were used to evaluate the quality of each generated RE. (2) After adding the REs to the original document, **BLEU** (Papineni et al., 2002) and **Text Accuracy** were used to evaluate the output text. (3) **Precision, recall,** and **F1** score were used to assess pronominalisation.

Results of WEBNLG. Table 2 depicts the results of WEBNLG⁸. Overall, the classic rule- and ML-based models performed better than neural models, while neural models did a better job on pronominalisation. For generating REs, ML-L had the best performance, as it obtained the highest RE

⁷ATT+Copy and ATT+Meta: github.com/rossanacunha/NeuralREG; and ProfileREG: github.com/mcao610/ProfileREG.

⁸Note that there is a discrepancy between our replication results and the results of Cunha et al. (2020). The reason for this difference is that we found a bug in the code for pre-processing provided by the original paper and fixed it after consultation with Cunha et al.

Model	RE Acc.↑	SED↓	BLEU↑	Text Acc.↑	Precision↑	Recall↑	F1↑
RREG-S	54.76/ 54.44	3.94/ 3.35	73.98/ 69.90	18.85/ 13.56	85.50/ 91.51	71.21/81.61	76.29/ 85.73
RREG-L	54.10/ <u>52.75</u>	3.86/3.68	73.56/68.72	18.49/12.41	69.66/77.65	73.34/74.43	71.30/75.90
ML-S	58.61/50.06	3.38/4.02	73.73/67.67	18.65/12.12	74.15/56.99	85.05/50.26	78.35/48.24
ML-L	63.01/50.32	3.30/4.03	76.10/67.91	20.30/12.33	83.68/73.90	85.00/50.21	84.32/47.98
ATT-Copy	<u>71.47</u> /25.84	<u>2.64</u> /6.28	80.46 /54.50	<u>26.39</u> /3.08	86.90 /83.66	87.75/68.12	<u>87.32</u> /72.97
ATT-Meta	71.64 /34.88	2.62 /5.82	<u>80.26</u> /60.00	27.88 /4.93	<u>86.48</u> / <u>86.66</u>	<u>87.97</u> /67.72	87.21/73.14
ProfileREG	68.26/13.43	3.27/11.55	78.24/41.13	21.82/0.7	<u>86.79</u> /86.04	94.80 / 82.66	90.33 / <u>84.24</u>

Table 3: Automatic Evaluation Results of WEBNLG for seen/unseen data respectively.

Model	RE Acc.↑	SED↓	BLEU↑	Text Acc.↑	Precision↑	Recall↑	F1↑
RREG-S	35.89	12.54	81.71	34.78	56.34	53.00	51.73
RREG-L	<u>37.22</u>	12.37	82.06	36.07	67.11	54.31	52.08
ML-S	37.18	12.56	<u>82.28</u>	36.07	<u>77.93</u>	56.70	55.52
ML-L	56.60	9.23	85.64	50.03	85.12	85.75	85.43
ATT+Copy	29.27	15.19	79.01	32.55	76.33	54.16	51.10
ATT+Meta	35.56	<u>12.11</u>	81.07	<u>36.83</u>	72.72	<u>70.50</u>	<u>71.42</u>

Table 4: Automatic Evaluation Results on WSJ. Note that since, for a WSJ document, it is extremely unlikely to generate every RE correctly, the text accuracy is always 0. Instead, we report sentence-level accuracy.

accuracy and BLEU scores and the second best SED and text accuracy score. For pronominalisation, ProfileREG yields the best performance, followed by RREG-S.

We were surprised to find that the simplest rule-based system, RREG-S, performs remarkably well. It not only defeats the linguistically informed, rule-based RREG-L, but also outperforms the SOTA neural models ATT+Copy and ATT+Meta on both RE generation and pronominalisation.

Table 3 shows the breakdown of the seen and unseen subsets. The SOTA neural models (i.e., ATT+Copy, ATT+Meta, and ProfileREG) have the top 3 performance on seen data, and the worst RE generation performance (i.e., RE Acc., SED, BLEU, and Text Acc.) on unseen data. The ML-based models achieve the fourth and fifth best performance on seen data, and lower performance (but not as low as the neural models) on unseen data. The nature of WEBNLG could explain this drop in performance on unseen data: the models may have limited ability to handle unseen entities, for instance, because they fail to conduct domain transfer (remember that unseen data comes from different domains than seen data). Since rule-based systems do not rely on training data, this explanation does not apply to them, which explains why they did not show the same drop in performance. In fact, they performed even better on unseen data, possibly because unseen data contained fewer triples than seen data (see §2). Concretely, rule-based systems have lower REG accuracy but higher pronominalisation

accuracy on unseen data compared to seen data. Additionally, ML-based models have low performance in the pronominalisation of unseen entities. The pronominalisation accuracy of the rule-based models is based on a 2-way distinction between a pronominal and a non-pronominal form, while the ML-based models make a 3-way distinction between a pronoun, a proper name and a description.

Another factor that might have lowered the performance of the ML models is the annotation practices in WEBNLG. Since these models are data-driven, the quality of the annotations directly affects their performance. It appears that whenever a (nominal) RE starts with a determiner, it is marked in WEBNLG as *description*; otherwise, it is marked as *proper name*. For instance, “*United States*” is marked as a proper name, while “*The United States*” is wrongly marked as a description. To allow comparison with previous work, we have not corrected the annotations, but it is important to keep in mind that this issue can cause ML-based models to underperform.

Results of WSJ. Table 4 shows the results of WSJ. Once again, ML-L performs best both in RE generation and in pronominalisation, outperforming the other models by a large margin. RREG-L outperforms RREG-S on WSJ on all evaluation metrics, which could be seen as confirmation of our hunch that WSJ contains different, and potentially more naturalistic texts than WEBNLG (see §2.2).

As for the neural models, the results suggest that meta-information can improve RE prediction

Model	All			Seen			Unseen		
	Fluency	Grammar	Clarity	Fluency	Grammar	Clarity	Fluency	Grammar	Clarity
RREG-S	5.76 ^A	5.73^A	5.71^A	5.73 ^A	5.62 ^A	5.68 ^A	5.79 ^A	5.83^A	5.75^A
RREG-L	5.68 ^A	5.52 ^A	5.67 ^A	5.63 ^A	5.45 ^A	5.64 ^A	5.74 ^A	5.60 ^{A,B}	5.70 ^A
ML-S	5.73 ^A	<u>5.65^A</u>	5.71^A	5.65 ^A	<u>5.64^A</u>	5.70 ^A	5.82^A	<u>5.65^{A,B}</u>	<u>5.72^A</u>
ML-L	5.78^A	5.63 ^A	5.67 ^A	5.73 ^A	5.63 ^A	5.62 ^A	5.82^A	5.62 ^{A,B}	<u>5.72^A</u>
ATT+Copy	5.65 ^A	5.62 ^A	<u>5.68^A</u>	5.71 ^A	<u>5.64^A</u>	<u>5.76^A</u>	5.58 ^A	5.60 ^{A,B}	5.59 ^A
ATT+Meta	5.68 ^A	5.56 ^A	5.66 ^A	5.69 ^A	5.68^A	5.65 ^A	5.67 ^A	5.43 ^B	5.66 ^A
ProfileREG	5.70 ^A	5.56 ^A	5.61 ^A	5.81^A	5.68^A	5.77^A	5.58 ^A	5.43 ^B	5.44 ^A
Human	5.81 ^A	5.69 ^A	5.82 ^A	5.77 ^A	5.69 ^A	5.83 ^A	5.84 ^A	5.70 ^A	5.80 ^A

Table 5: Human Evaluation Results on WEBNLG corpus. Rankings are determined by significance testing ($p < 0.01$). Per column, results that have *no* superscript letters in common are significantly different from each other.

Model	Fluency	Grammar	Clarity
RREG-S	76.13 ^{A,B}	75.74 ^{A,B,C}	78.03 ^A
RREG-L	72.56 ^{A,B}	73.38 ^C	74.76 ^A
ML-S	77.48 ^A	78.39 ^{A,B}	78.76 ^A
ML-L	77.52 ^A	78.45 ^A	79.45 ^A
ATT+Copy	74.43 ^{A,B}	74.57 ^{A,B,C}	75.63 ^A
ATT+Meta	71.94 ^B	72.95 ^{B,C}	73.95 ^A

Table 6: Human Evaluation Results on WSJ.

accuracy. Also, the inclusion of meta-information significantly boosts the recall of pronominalisation comparing ATT+Copy with ATT+Meta. Table 10 in Appendix D shows an original text and different outputs generated by the WSJ models.

4.2 Human Evaluation on WEBNLG

Materials. For WEBNLG seen entities, we randomly sampled 4 instances from each triple size group of 2-7 from the test set. In the case of the unseen data, we randomly chose 6 instances from size groups of 2-5. In this way, we obtained a total number of 48 reference instances (24 seen and 24 unseen). In addition to each reference instance, we selected its 7 different versions generated by the models (3 neural, 2 ML-based and 2 rule-based models). This yields a total of 384 items (48×8).

Design. The 384 items were randomly distributed into 12 lists of 32 items. Each list was rated by 10 participants. Participants were asked to rate each text for its fluency (“does the text flow in a natural, easy to read manner?”), grammaticality (“is the text grammatical (no spelling or grammatical errors)?”) and clarity (“does the text clearly express the data in the table?”) on a 7-point Likert scale anchored by 1 (very bad) and 7 (very good). The definition of each criterion was taken from Castro Ferreira et al. (2018a).

Participants. We used Amazon Mechanical Turk (MTurk) for human evaluation. We restricted MTurk workers to those located in the United States, with an approval rating of $\geq 95\%$ and 1,000 or more HITs approved. We rejected workers if they: (1) gave human-produced descriptions a score lower than 2 more than 3 times; or (2) gave scores with a standard deviation less than 0.5. 120 workers (12 lists \times 10 workers) participated, providing us with 11520 judgements (384 items \times 3 criteria \times 10 judgements/item). The participants were 80 males, 36 females, and 4 others/unanswered, with an average age of 37.

Results. Table 5 shows the results of the human evaluation WEBNLG. Few of the differences reach significance (using Wilcoxon’s signed-rank test with Bonferroni correction⁹), suggesting that WEBNLG may be ill-suited for differentiating between REG models¹⁰. The only two significant differences appear when comparing RREG-S with ATT+Meta and ProfileREG in terms of the grammaticality of unseen data. The results suggest that RREG-S is the best model for generating REs on WEBNLG, performing on a par with neural models on seen data and better than neural models on unseen data. Unlike our automatic evaluation, ATT+Meta does not outperform ATT+Copy in human evaluation.

4.3 Human Evaluation on WSJ

Materials. We randomly selected 30 documents from the test set of WSJ (reference text). We included the 6 different outputs generated by the 6 WSJ models (hereafter target texts). In this way we obtained a total of 180 reference-target pairs.

⁹P-values were multiplied by the number of comparisons.

¹⁰All non-significant differences in Table 5 and Table 6 are associated with p-values greater than 0.1.

Design. As mentioned in §2.3, the WSJ documents have an average length of 25 sentences. Since there are no input representations (e.g., in RDF) for WSJ, we decided to ask participants to compare texts using a Magnitude Estimation (ME) (Bard et al., 1996). The participants saw the reference and one of the target texts side by side, and they were asked to rate the target relative to the reference text. To make the task manageable for participants, texts were shortened to a maximum of the first 20 sentences. The 180 reference-target pairs were randomly distributed over 12 lists, each list having 15 items. Each list was rated by 10 participants. They were asked to rate the fluency, grammaticality and clarity of the target texts. The definition of fluency and grammaticality were as in the WEBNLG task, and clarity was defined as “how clearly does the target text allow you to understand the situation described in the standard text¹¹?”. The question asked for each of the 3 criteria was: assuming that standard text has a score of 100, how do you rate the fluency|grammaticality|clarity of target text? Participants were allowed to choose any positive number.

Participants. The MTurk worker restrictions were similar to the WEBNLG experiment. Workers with scores less than 5 standard deviations were rejected. The experiment included 120 participants, resulting in 5400 judgements (180 items×3 criteria×10 judgements/item). The participants were 65 males, 54 females, and 1 others/unanswered, with an average age of 38.

Results. Since typos are possible in ME (e.g., a worker might type 600 instead of 60), we excluded outliers, defined as a score that is lower than the median minus 3 standard deviations, or higher than the median plus 3 standard deviations of that item. The remaining scores were down-sampled for conducting significant testing. The results are shown in Table 6. Unlike WEBNLG, significant differences are frequent. For fluency, ML-S and ML-L perform the best while ATT+Meta performs the worst. For grammaticality, ML-L is still the best model, which significantly defeats RREG-L and ATT+Meta. A more detailed study is needed to investigate why RREG-L is the second worst in terms of grammaticality, which we found surprising. For clarity, no significant difference was found, perhaps because it was difficult for participants to compare long

documents. In sum, on WSJ, ML-L has the best performance, and the simpler ML-S and RREG-S also have considerably good performances.

5 Discussion

Why does Neural REG not defeat rule-based REG? Received wisdom has it that although neural models may be inferior to other models in terms of interpretability, they are nonetheless superior in terms of performance. Although it is possible that future neural models will perform better than the ones examined here, our results call into question whether this received wisdom is correct. One possible explanation is the observation that Neural NLG systems tend to perform very well on surface realisation tasks, but less well on tasks that focus on semantic content (see e.g., Reiter (2018) on hallucinations in the Data2Text generation tasks). REG, after all, is a task that focuses in large part on semantic content. There may be other reasons, which should be investigated in future work.

Role of Linguistically-informed Features. Rule-based models did particularly well on WEBNLG, outperforming other models. By contrast, on WSJ, the linguistically-informed feature-based model (ML-L) outperformed all other models. This suggests that the type of text, and consequently, the complexity of the REG task, might be a factor in choosing the REG method. Linguistically-informed features seem to have a more pivotal role in the case of more complex text types, whereas simpler texts can be handled at least as well by simpler rule-based models.

Resources Use. As mentioned before, different approaches require different amounts of **human** resources and annotation efforts. But we believe that other resource types should also be taken into consideration when models are compared, including the following: (1) **The amount of context:** the neural models access the whole pre-context and post-context for WEBNLG, while they access K preceding and K following sentences around the target entity for WSJ. The ML-based models extract features taking only the current sentence and the whole pre-context into account. The rule-based models only look at the current sentence and the previous one; (2) **External tools:** the neural models need no external tools, while the rule-based and ML-L models need a syntactic parser (which is also used for constructing datasets); (3)

¹¹We refer to the reference text as *standard text*.

External information: rule-based models, $ML-L$, and $ATT+Meta$ need entities’ meta-information. $ProfileREG$ requires the profile description of each entity, which, for most REG tasks, is hard to obtain; (4) **Computing resources:** the neural models need GPUs while other models can be constructed using merely personal computers; (5) **The amount of training data:** the rule-based models need no training data, while other models require training data (large-scale naturalistic versions of which, for the task of REG, is not available).

As we have seen, $RREG-S$ and $ML-S$ perform remarkably well on both WSJ and WEBNLG. Taking resources into consideration, the advantage of using a model such as $RREG-S$ and $ML-S$ becomes more pronounced. $RREG-S$ uses less human resources, less context, less computing resources, and no training data¹² compared to other models. $ML-S$ needs more context and training data; it probably also needs more human effort for feature engineering and selecting ML models, but it needs no external tools and no meta-information.

In aggregate, one’s choice of model may depend partially on what resources are available. For instance, for classic pipeline NLG systems, syntactic position and meta-information are often decided by earlier steps in the pipeline (Gatt and Krahmer, 2018). Therefore, if one’s aim is to rapidly construct a pipeline NLG system, then $RREG-S$ should probably be preferred.

Generalisability. We used neural REG to illustrate the importance of non-neural baselines. Our findings may not be generalisable to End2End NLG. However, if complex rule/template-based NLG systems are taken into account, Dušek et al. (2018) found that although these systems cannot defeat neural approaches, they still have competitive performance. It would be interesting to compare different types of models for other sub-tasks in the NLG pipeline (e.g., content determination, aggregation, and lexicalisation) in a similar way as has been done in the present paper¹³.

6 Conclusion

In this work, we have re-evaluated state-of-the-art Neural REG systems by considering four well-

¹²The pronominal form (e.g., he or she) of an entity can either be extracted from the training data or decided by the entity’s meta-information.

¹³Note that pipelined NLG systems are sometimes thought to yield better outputs than fully End2End NLG systems (Castro Ferreira et al., 2019)

designed rule- and ML-based baselines. In addition to the existing WEBNLG corpus, we built a new dataset for the task of *REG-in-context* on the basis of the WSJ corpus, arguing that this dataset may be more appropriate for the task. In the re-evaluation, we examined both our baselines and SOTA neural REG systems on both datasets, using automatic and human evaluations. The results suggest that the simplest rule-based baseline $RREG-S$ achieves equally good or better performance compared to SOTA neural models. Our results on the WSJ suggest that, on that corpus, the linguistically-informed ML-based model ($ML-L$) is best. We hope these results will encourage further research into the comparative strengths and weaknesses of neural, non-neural and hybrid methods in NLP.

In future, we have 4 items on our TODO list: (1) Investigate bottleneck features for Neural based models based on the feature set of $ML-L$; (2) Explore other neural architectures (e.g., testing models that leverage pre-trained language models) and construct larger realistic REG corpora; (3) Explore better human evaluation methods for longer documents that are better suited for evaluating the task of generating referring expressions in context; (4) Extend our research to other languages, especially in other language families, including languages that are morphological very rich or very poor and languages that frequently use zero pronouns (e.g., Chinese (Chen et al., 2018)).

Acknowledgements

We thank the anonymous reviewers for their helpful comments. Guanyi Chen is supported by China Scholarship Council (No.201907720022). Fahime Same is supported by the German Research Foundation (DFG)– Project-ID 281511265 – SFB 1252 “Prominence in Language” and the Junior Special Fund of the Cologne Center of Language Sciences (CCLS).

Ethical Considerations

We collected our human evaluations using Amazon Mechanical Turk. For the WEBNLG task, which used a 7-point Likert scale, the workers were paid 0.03\$ per item, in line with rates for similar tasks. For the more demanding WSJ task, we paid 0.10\$ per item. The payment for each task was set at \$7.5/hour (slightly above the US minimum wage, i.e., \$7.25/hour). We expected the amount to be a fair remuneration, but given the actual time some

participants needed, their remuneration turned out to be on the low side. In future crowd-sourcing experiments, we will base our remuneration on a more generous estimate of the duration per experimental task. We asked for demographic information, age, gender and English proficiency level, explicitly stating in the experiment that “Your information will be used for research purposes only. All your data will be held *anonymously*.” These fields were not marked as mandatory fields. The demographic information will not be made publicly available.

References

- Mira Ariel. 1990. *Accessing Noun-Phrase Antecedents*. Routledge.
- Jennifer E Arnold and Zenzi M Griffin. 2007. The effect of additional characters on choice of referring expression: Everyone counts. *Journal of memory and language*, 56(4):521–536.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Ellen Gurman Bard, Dan Robertson, and Antonella Sorace. 1996. Magnitude estimation of linguistic acceptability. *Language*, pages 32–68.
- Anja Belz, Eric Kow, Jette Viethen, and Albert Gatt. 2010. [Generating referring expressions in context: The GREC task evaluation challenges](#). In *Empirical Methods in Natural Language Generation: Data-oriented Methods and Empirical Evaluation*, volume 5790 of *Lecture Notes in Computer Science*, pages 294–327. Springer.
- Anja Belz and Sebastian Vargas. 2007. Generation of repeated references to discourse entities. In *Proceedings of the Eleventh European Workshop on Natural Language Generation*, pages 9–16. Association for Computational Linguistics.
- Susan E Brennan. 1995. Centering attention in discourse. *Language and Cognitive processes*, 10(2):137–167.
- Meng Cao and Jackie Chi Kit Cheung. 2019. [Referring 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, Emiel Kraemer, and Sander Wubben. 2016. [Towards more variation in text generation: Developing and evaluating variation models for choice of referential form](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 568–577, Berlin, Germany. Association for Computational Linguistics.
- Thiago Castro Ferreira, Diego Moussallem, Ákos Kádár, Sander Wubben, and Emiel Kraemer. 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 Kraemer, 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.
- Thiago Castro Ferreira, Chris van der Lee, Emiel van Miltenburg, and Emiel Kraemer. 2019. [Neural data-to-text generation: A comparison between pipeline and end-to-end architectures](#). 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 552–562, Hong Kong, China. Association for Computational Linguistics.
- 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.
- 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.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2018. [Findings of the E2E NLG challenge](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 322–328, Tilburg University, The Netherlands. Association for Computational Linguistics.

- Rui Fang, Malcolm Doering, and Joyce Y Chai. 2015. Embodied collaborative referring expression generation in situated human-robot interaction. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 271–278.
- Kumiko Fukumura and Roger PG Van Gompel. 2011. The effect of animacy on the choice of referring expression. *Language and cognitive processes*, 26(10):1472–1504.
- 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 Kraahmer. 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.
- Charles Greenbacker and Kathleen McCoy. 2009. UDel: generating referring expressions guided by psycholinguistic findings. In *Proceedings of the 2009 Workshop on Language Generation and Summarisation*, pages 101–102. Association for Computational Linguistics.
- Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. 1995. [Centering: A framework for modeling the local coherence of discourse](#). *Computational Linguistics*, 21(2):203–225.
- Jeanette K Gundel, Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language*, pages 274–307.
- Iris Hendrickx, Walter Daelemans, Kim Luyckx, Roser Morante, and Vincent Van Asch. 2008. [CNTS: Memory-based learning of generating repeated references](#). In *Proceedings of the Fifth International Natural Language Generation Conference*, pages 194–95, Salt Fork, Ohio, USA. Association for Computational Linguistics.
- Renate Henschel, Hua Cheng, and Massimo Poesio. 2000. Pronominalization revisited. In *Proceedings of the 18th conference on Computational linguistics-Volume 1*, pages 306–312. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. [OntoNotes: The 90% solution](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Andrej A. Kibrik, Mariya V. Khudyakova, Grigory B. Dobrov, Anastasia Linnik, and Dmitriy A. Zalmanov. 2016. [Referential choice: Predictability and its limits](#). *Frontiers in Psychology*, 7.
- Emiel Kraahmer and Mariët Theune. 2002. Efficient context-sensitive generation of referring expressions. *Information sharing: Reference and presupposition in language generation and interpretation*, 143:223–263.
- Emiel Kraahmer and Kees van Deemter. 2019. *Computational Generation of Referring Expressions: An Updated Survey*, chapter 19. Oxford University Press.
- Vladimir I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20.
- Kathleen E. McCoy and Michael Strube. 1999. [Generating anaphoric expressions: Pronoun or definite description?](#) In *The Relation of Discourse/Dialogue Structure and Reference*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Liudmila Ostroumova Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. 2018. [Catboost: unbiased boosting with categorical features](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6639–6649.
- Ehud Reiter. 2017. [A commercial perspective on reference](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 134–138, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Ehud Reiter. 2018. [Hallucination in Neural NLG](#).
- Ehud Reiter and Robert Dale. 2000. *Building Natural Language Generation Systems*. Studies in Natural Language Processing. Cambridge University Press.
- 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.

Kees van Deemter. 2016. *Computational models of referring: a study in cognitive science*. MIT Press.

Klaus von Heusinger and Petra B Schumacher. 2019. Discourse prominence: Definition and application. *Journal of Pragmatics*, 154:117–127.

Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, Mohammed El-Bachouti, Robert Belvin, and Ann Houston. 2013. [Ontonotes release 5.0](#).

A WSJ Dataset Construction

In this work, we used Ontonotes 5.0 licensed by the Linguistic Data Consortium (LDC) <https://catalog.ldc.upenn.edu/LDC2013T19>. We used the files in OntoNotes Normal Form (ONF) format. This format has combined all layers of the OntoNotes corpus, ranging from text (sentence and token segmented), parsing information, propositional content, and the coreference chains (Weischedel et al., 2013). These files were later rendered into XML format for full processing. As mentioned earlier, only the Wall Street Journal portion of this corpus has been used in the current study. The first and second person references were excluded from the dataset. Furthermore, we assumed that REs are presented in a linear order, therefore excluded cases such as union REs. As an example, in the sentence “[Mary], [John] and [David] received their booster shot last month”, the REs enclosed in brackets are included in the dataset, but their underlined union, “*Mary, John, and David*”, is excluded. We then delexicalised the REs as explained in §2.3. The delexicalised text in table 7 shows an example from the WSJ corpus.

B Details of RREG-L

Algorithm 1 describes the generation process of RREG-L. The system takes in the target entity r , the current sentence (u_2 , i.e., the one where the target RE is located), and the previous sentence (u_1). It starts with a rule checking whether an antecedent of r appears in u_1 (line 1). If the answer is no, then it realises r with its non-pronominal form. If such an antecedent exists, the system heads to check parallelism (Henschel et al., 2000). Concretely speaking, it checks whether or not r has the same grammatical role (i.e., subject or object) as its antecedent. If the parallelism holds, r is realised as a pronoun. Otherwise, we apply the “local focus” idea from Henschel et al. (2000),

Algorithm 1 The Linguistically Informed Rule-based REG Algorithm

Input: The target entity r , the sentence u_2 that r is in, and its previous sentence u_1 .

Output: The surface form of r .

```
1: if  $r$  has an antecedent in  $u_1$  then
2:   if  $r$  occurs in parallel context then
3:     RealiseProRE( $r$ )
4:   else
5:      $\mathcal{F} := \text{FocusSetConst}(u_1)$ 
6:     if  $r$ 's antecedent  $\in \mathcal{F}$  and  $r$  has no
       competitor  $r' \in \mathcal{F}$  then
7:       RealiseProRE( $r$ )
8:     else
9:       RealiseNONProRE( $r$ )
10:  else
11:    RealiseNONProRE( $r$ )
```

which builds upon the Centering Theory (Grosz et al., 1995). A referent is the local focus if it is (1) discourse-old, or (2) in the subject position. In line 5, the `FocusSetConst` function constructs a set \mathcal{F} , consisting of local focus entities in u_1 . If r 's antecedent is an element of \mathcal{F} , and r has no competitor being an element of \mathcal{F} , then we realise r as a pronoun (line 6-9). Both surface realisation functions (i.e., `RealiseProRE` and `RealiseNONProRE`) work similarly to RREG-S in realising pronominal and non-pronominal REs, respectively.

C Detailed list of features used in WEBNLG and WSJ feature-based ML models

Each feature is defined and calculated for each target entity r . Antecedent refers to the first co-referential RE preceding r . Tables 8 and 9 list the features used in the ML models.

Since we wanted to use the features in a back-off method for selecting the content of REs, we converted numerical features, such as recency, into categorical values. We tried different recency measurements on the WEBNLG and WSJ validation sets, and chose the ones which yielded the best results: (1) Word distance in all ML-based WEBNLG and WSJ models: 5 quantile groups; (2) Sentence distance in WEBNLG ML-R & ML-S, and WSJ ML-S models: 2 quantile groups; (3) Sentence distance in WSJ ML-L: 3 bins defined as whether the antecedent is in the

Text: Ronald B. Koenig , 55 years old , was named a senior managing director of the Gruntal & Co. brokerage subsidiary of this insurance and financial - services firm . Mr. Koenig will build the corporate - finance and investment - banking business of Gruntal , which has primarily been a retail - based firm . He was chairman and co-chief executive officer of Ladenburg , Thalmann & Co. until July , when he was named co-chairman of the investment-banking firm along with Howard L. Blum Jr. , who then became the sole chief executive . Yesterday , Mr. Blum , 41 , said he was n't aware of plans at Ladenburg to name a co-chairman to succeed Mr. Koenig and said the board would need to approve any appointments or title changes .

Delexicalized Text:

Pre-context: Mr. Koenig was named a senior managing director of the Gruntal brokerage subsidiary of this insurance and financial-services firm .

Target Entity: Mr. Koenig

Post-context: will build the corporate-finance and investment-banking business of Gruntal. Mr. Koenig was chairman and co-chief executive officer of Ladenburg until July , when Mr. Koenig was named co-chairman of Ladenburg along with Mr. Blum. Yesterday , Mr. Blum said Mr. Blum was n't aware of plans at Ladenburg to name a co-chairman to succeed Mr. Koenig and said the board would need to approve any appointments or title changes.

Table 7: An example data from the WSJ corpus. In the delexicalized text, every entity is underlined.

Feature Class	Definition
Referential status	Is r the first mention of the entity in the text?
Referential status	Is the antecedent of r in the same sentence?
Recency	Categorical distance between r and its antecedent in number of sentences
Recency	Categorical distance between r and its antecedent in number of words
Competition	Is there any other RE between r and its antecedent?
Position	A categorical feature marking whether r is the first, second, middle or last mention

Table 8: Features used in the ML-S models. Each feature is defined and calculated for each target entity r . Antecedent refers to the first co-referential referring expressions preceding r .

Feature Class	Definition
Grammatical Role	Grammatical Role of r
Grammatical Role	Grammatical Role of the antecedent
Meta information	Entity type (e.g. human, city, country, organization)
Meta information	Plurality: Is r plural or singular? (only WSJ)
Meta information	Gender
Recency	Categorical distance between r and its antecedent in number of words
Recency	Categorical distance between r and its antecedent in number of sentences
Recency	Categorical distance between r and its antecedent in number of paragraphs (only WSJ)

Table 9: Features used in the ML-L models.

same sentence, 1 sentence away, or more than 1 sentence away; (4) Paragraph distance in WSJ ML-L: 4 bins defined as whether r and its antecedent are in the same paragraph, 1 paragraph away, 2 paragraphs away, more than two paragraphs away. The paragraph information associated with the WSJ documents are taken from the PDTB parser at: github.com/WING-NUS/pdtb-parser/tree/master/external/aux_data/paragraphs.

D Sample texts generated by the WSJ models

Table 10 shows a reference text from the WSJ dataset in addition to the outputs generated by RREG-S (rule-based), ML-L (feature-based), and ATT-Meta (neural).

Original: **MGM Grand Inc.** said **it** filed a registration statement with the Securities and Exchange Commission for a public offering of six million common shares. **The Beverly Hills , Calif.-based company** said **it** would have 26.9 million common shares outstanding after the offering. **The hotel and Gaming company** said Merrill Lynch Capital Markets will lead the underwriters. Proceeds from the sale will be used for remodeling and refurbishing projects , as well as for the planned MGM Grand hotel / casino and theme park.

RREG-S: **MGM Grand Inc.** said **MGM Grand Inc .** filed a registration statement with the Securities and Exchange Commission for the offering. **MGM Grand Inc.** said **MGM Grand Inc.** would have 26.9 million common shares outstanding after the offering. **MGM Grand Inc.** said Merrill Lynch Capital Markets will lead the underwriters . Proceeds from the offering will be used for remodeling and refurbishing projects , as well as for the planned MGM Grand hotel / casino and theme park .

ML-L: **MGM Grand Inc.** said **it** filed a registration statement with the Securities and Exchange Commission for an offering of common shares. **MGM Grand Inc.** said **it** would have 26.9 million common shares outstanding after the offering. **MGM Grand Inc.** said Merrill Lynch Capital Markets will lead the underwriters . Proceeds from the offering will be used for remodeling and refurbishing projects , as well as for the planned MGM Grand hotel / casino and theme park.

ATT-Meta: **MGM Grand Inc.** said **MGM Grand Inc.** filed a registration statement with the Securities and Exchange Commission for the offering of the company of the market. **MGM Grand Inc.** said **it** would have 26.9 million common shares outstanding after the offering. **MGM Grand Inc.** said Merrill Lynch Capital Markets will lead the underwriters. Proceeds from the offering will be used for remodeling and refurbishing projects , as well as for the planned MGM Grand hotel / casino and theme park .

Table 10: Examples of an original text from the WSJ dataset together with the outputs generated by the models RREG-S, ML-L, and ATT-Meta. References to “MGM Grand Inc.” are boldfaced.