

What can Neural Referential Form Selectors Learn?

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

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

[♡]Department of Linguistics, University of Cologne

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

Abstract

Despite achieving encouraging results, neural Referring Expression Generation models are often thought to lack transparency. We probed neural Referential Form Selection (RFS) models to find out to what extent the linguistic features influencing the RE form are learnt and captured by state-of-the-art RFS models. The results of 8 probing tasks show that all the defined features were learnt to some extent. The probing tasks pertaining to referential status and syntactic position exhibited the highest performance. The lowest performance was achieved by the probing models designed to predict discourse structure properties beyond the sentence level.

1 Introduction

Referring Expression Generation (REG) is one of the main stages of classic Natural Language Generation (NLG) pipeline (Reiter and Dale, 2000; Kraemer and van Deemter, 2012; van Deemter, 2016). REG studies are concerned with two different tasks. The goal of the classic REG task (also called one-shot REG), is to find a set of attributes to single out a referent from a set of competing referents. The second REG task (henceforth discourse REG) is concerned with the generation of referring expressions (RE) in discourse context. Belz and Varges (2007) phrase it as follows: *Given an intended referent and a discourse context, how do we generate appropriate referential expressions (REs) to refer to the referent at different points in the discourse?*

Classic discourse REG was usually understood as a two-step procedure. In the first step, the referential form (RF, i.e. the syntactic type) is determined. For instance, when referring to Joe Biden at a given point in a discourse, the first step is to decide whether to use a proper name (“Joe Biden”), a

description (“the president of the USA”), a demonstrative (“this person”) or a pronoun (“he”). The second step is to determine the RE content, that is, to choose between all the different ways in which a given form can be realised. For instance, to generate a description of Joe Biden, one needs to decide whether to only mention his job (e.g., *The president entered the Oval Office.*), or to mention the country as well (e.g., *The president of the United states arrived in Cornwall for the G7 Summit.*)

In earlier works, computational linguists linked REG to linguistic theories and built discourse REG systems on the basis of linguistic features. For example, Henschel et al. (2000) investigated the impact of 3 linguistic features namely recency, subjecthood, and discourse status on pronominalization, i.e. deciding whether the RE should be realised as a pronoun. Using these features, they used the notion of *local focus* as a criterion for detecting the set of referents that can be pronominalised. The same holds for feature-based models (see Belz et al. (2010) for an overview) where models are trained on linguistically encoded data.

More recently, a number of neural network-based REG models have been presented (Castro Ferreira et al., 2018a; Cao and Cheung, 2019; Cunha et al., 2020), where they propose to generate REs in an End2End manner without any feature engineering. They all used a benchmark dataset called WebNLG. These models generally follow the sequence-to-sequence framework (Sutskever et al., 2014), where there is an encoder for encoding the given discourse, and a decoder responsible for generating REs using the encoded information. The evaluation results suggested that these neural methods perform well not only for selecting the proper RFs, but also for producing fluent REs. However, it was unclear to what extent these neural models can encode linguistic features.

To conduct model inspection, we introduce a

* Equal contribution

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

Pos-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.

series of probing tasks. Using probing tasks is a well-established method to analyse whether a model’s latent representation encodes specific information. This approach has been widely used for analysing models in machine translation (Belinkov et al., 2017), language modelling (Giulianelli et al., 2018), relation extraction (Alt et al., 2020), and so on. Additionally, there had been various works on coreference resolution and bridging anaphora (Sorodoc et al., 2020; Pandit and Hou, 2021) which, similar to this paper, target the understanding of reference. More precisely, for a probing task, a diagnostic classifier is trained on representations from the model. Its performance embodies how well those representations encode the information associated with the probing task. The aim of this paper is to understand what linguistic features neural models encode when modelling REs.

Our main focus is on the encoding of linguistic features in the representations. In the linguistic tradition, the majority of RE production studies focus on Referential Form Selection (RFS), rather than RE content realisation. Our focus in the present work is likewise on RFS. To tackle RFS, we adopt the state-of-the-art neural REG model of (Castro Ferreira et al., 2018a). Additionally, to make comparison, we also propose (1) a strong baseline that uses only a single encoder (while Castro Ferreira et al. (2018a) used multiple encoders); and (2) to leverage pre-trained word embeddings (e.g., GloVe) or language models (e.g., BERT).

Therefore, in this paper, we first introduce the task of RFS on the basis of WebNLG corpus, and propose a number of neural models to tackle the task. Subsequently, we introduce 8 probing tasks, each of which is associated with a linguistic feature influencing the choice of RF. We examine our RFS models on these probing tasks in order to interpret and explain their behaviour. The code of each RFS model and the probing classifier is available at:

github.com/a-quei/probe-neuralreg.

2 Background

2.1 Discourse REG

Given a text whose REs have not yet been generated, and given the intended referent for each of these REs, the discourse REG task is to build an algorithm that generates all these REs. So far, this task has attracted many research efforts (e.g., Hendrickx et al. (2008); Greenbacker and McCoy (2009)) and it has been used in the GREC shared tasks (Belz et al., 2010).

More recently, this task was formulated into a format that goes together well with deep learning: Castro Ferreira et al. (2018a) introduced the End2End REG task, built a corresponding dataset based on WebNLG (Castro Ferreira et al., 2018b), and constructed NeuralREG models.

The WebNLG corpus was originally designed to assess the performance of NLG systems (Gardent et al., 2017). Each sample in this corpus contains a knowledge base described by a Resource Description Framework (RDF) triple (Table 1). Castro Ferreira et al. (2018a) and Castro Ferreira et al. (2018b) enriched and delexicalised the corpus to fit the discourse REG task. Table 1 shows a text created from a RDF, and its corresponding delexicalised version.

Taking the delexicalised text in Table 1 as an example, given the entity “AWH_Engineering_College”, REG chooses a RE based on that entity and its pre-context (“AWH_Engineering_College is in “Kuttikkattoor” , India in the state of Kerala .”) and its pos-context (“has 250 employees and Kerala is ruled by Kochi . The Ganges River is also found in India .”).

2.2 Factors that Influence RE Production

Languages display a large inventory of expressions for referring to entities (von Heusinger and Schumacher, 2019). In linguistics, the realisation choice

a speaker makes has been associated with the accessibility, i.e. activation of mental representations of a referent at a particular point in discourse: attenuated forms such as pronouns are often used to refer to highly accessible or highly activated referents, while richer forms such as descriptions and proper names are employed in referring to less accessible ones (Ariel, 1990; Gundel et al., 1993). Due to the central role of referring in communication, a wealth of research has tried to assess the influence of different features modulating the accessibility of a referent. von Heusinger and Schumacher (2019) refer to these features as *prominence-lending cues*, meaning that they increase the prominence status of their respective referents to some extent. In this section, we merely talk about the ones which will be taken up in our probing experiments, and will not further discuss cues such as animacy (Fukumura and van Gompel, 2011), competition (Arnold, 2010) and coherence relations (Kehler et al., 2008).

Referential status or *givenness* has been widely discussed in the literature (see Chafe (1976); Prince (1981)). When a new character is introduced into the discourse, the chance that this happens by means of a pronoun is slim (unless the referent is situationally given). Pronouns are reserved for referring to previously introduced (or given) referents.

Recency, another well-studied cue, is defined as the distance between the target referent and its antecedent. If a referent is not too far apart from its antecedent, then reduced forms are typically employed to refer to it.

There are also intra-clausal cues such as *grammatical role* (Brennan, 1995) and *thematic role* (Arnold, 2001) which impact the prominence status of referents. For instance, the subject of a sentence is perceived to be more prominent than the object.

Discourse-structural features affect the organisational aspects of discourse. Centering-based theories (Grosz et al., 1995) often use the notion of local focus to account for pronominalisation. *Local focus* takes the current and previous utterance into account. *Global focus*, on the other hand, situates a referent in a larger space, namely the whole text or a discourse segment (Hinterwimmer, 2019). Concepts such as the importance of a referent or familiarity are associated with the global prominence status of entities (Siddharthan et al., 2011).

Type	Classes
4-Way	Demonstrative, Description, Proper Name, Pronoun
3-Way	Description, Proper Name, Pronoun
2-Way	Non-pronominal, Pronominal

Table 2: 3 different types of RF classification.

3 Neural Referential Form Selection

In this section, we define the task of RFS built on the WebNLG dataset, and introduce a number of NeuralRFS models.

3.1 The RFS Task

Akin to REG, given the previous context $x^{(pre)} = \{w_1, w_2, \dots, w_{i-1}\}$ (where w is either a word or a delexicalised entity label), the target referent $w^{(r)} = \{w_i\}$, and the post context $w^{(pos)} = \{w_i, w_{i+1}, \dots, w_n\}$, a RFS algorithm aims at finding the proper RF \hat{f} from a set of K candidate RFs $\mathcal{F} = \{f_k\}_{k=1}^K$.

Regarding the possible RFs for the RFS task, we test 3 different classifications, depicted in Table 2. Due to the small number of demonstrative noun phrases in the dataset, we decided to also conduct a 3-way classification in which descriptions and demonstratives are merged. Also, most emphasis in the linguistic literature is on the pronominalisation issue. Therefore, we also included a 2-way classification task in the study.

As stated, the main goal of the paper is to understand which linguistic features are encoded by RFS neural models. Additionally, we were curious whether models trained solely for pronominalisation capture different contextual features in comparison with the other two classifications.

3.2 NeuralRFS Models

We build NeuralRFS models by (1) adopting the best NeuralREG model from Castro Ferreira et al. (2018a), and (2) proposing a new alternative which is simpler, and can easier incorporate pre-trained representations.

ConATT. We adopt the CATT model from Castro Ferreira et al. (2018a), which achieves the best performance on REG among the models they tested in their study. Given the inputs, we first use Bidirectional GRU (BiGRU, Cho et al., 2014) to encode $x^{(pre)}$ as well as $x^{(pos)}$. Formally, for each $k \in [pre, pos]$, we encode $x^{(k)}$ to $h^{(k)}$ with a BiGRU: $h^{(k)} = \text{BiGRU}(x^{(k)})$. Subsequently, differ-

Model	4-way			3-way			2-way		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
XGBoost	53.77	51.98	51.55	71.27	69.24	68.34	86.64	82.76	84.57
c-RNN	<u>68.79</u>	<u>62.95</u>	<u>64.96</u>	84.49	<u>82.52</u>	83.63	<u>90.31</u>	88.01	89.09
+GloVe	69.10	63.90	65.40	84.29	82.55	83.30	89.33	88.02	88.63
+BERT	62.63	61.80	62.15	83.02	81.44	82.15	90.98	88.00	89.42
ConATT	67.42	62.39	64.07	85.04	82.21	83.53	89.30	89.19	89.23
+GloVe	65.98	62.49	63.67	83.62	81.41	82.45	89.60	<u>88.06</u>	88.80

Table 3: Evaluation results of our RFS systems on WEBNLG. Best results are **boldfaced**, whereas the second best results are underlined.

ent from [Castro Ferreira et al. \(2018a\)](#), we encode $h^{(k)}$ into the context representation $c^{(k)}$ using self-attention ([Yang et al., 2016](#)). Concretely, given the total N steps in $h^{(k)}$, we first calculate the attention weight $\alpha_j^{(k)}$ at each step j by:

$$e_j^{(k)} = v_a^{(k)T} \tanh(W_a^{(k)} h_j^{(k)}), \quad (1)$$

$$\alpha_j^{(k)} = \frac{\exp(e_j^{(k)})}{\sum_{n=1}^N \exp(e_n^{(k)})}, \quad (2)$$

where 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 $h^{(k)}$:

$$c^{(k)} = \sum_{j=1}^N \alpha_j^{(k)} h_j^{(k)}. \quad (3)$$

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

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

where W_f is the weights in the feedforward layer. R is also used as the input of the probing classifiers (section 4). R is then fed for making the final prediction:

$$P(f|x^{(pre)}, x^{(r)}, x^{(pos)}) = \text{Softmax}(W_c R), \quad (5)$$

where W_c is the weight in the output layer.

c-RNN. In addition to ConATT, we also try a simpler yet effective structure, which uses only a single BiGRU. We name the framework it follows as the centred recurrent neural networks (henceforth c-RNN). Specifically, instead of using two separate BiGRUs to encode pre- and pos-contexts,

we first concatenate $x^{(pre)}$, $x^{(r)}$, and $x^{(pos)}$, and then encode them together:

$$h = \text{BiGRU}([x^{(pre)}, x^{(r)}, x^{(pos)}]). \quad (6)$$

Suppose that the target entity is in position i of the concatenated sequence, we extract the i -th representation from h_i for obtaining $R = \text{ReLU}(W_f h_i)$. After obtaining R , the rest of the procedure is the same as ConATT.

Pre-training. As a secondary objective of this study, we want to see whether RFS can benefit from pre-trained word embeddings and language models, whose effectiveness has not yet been explored in REG¹. For both c-RNN and ConATT, we try the GloVe embeddings ([Pennington et al., 2014](#)) to see how pre-trained word embeddings contribute to the choice of RF. For c-RNN, we try to stake it on the BERT ([Devlin et al., 2019](#)) model. In order to let BERT better encode the delexicalised entity labels, we first re-train BERT as a masked language model on the training data of WebNLG. We then freeze the parameters of BERT and use the model to encode the input, which is then fed into c-RNN².

Machine Learning (ML) based Model We used XGBoost ([Chen and Guestrin, 2016](#)) from the family of Gradient Boosting Decision Trees to train RFS classifiers. 5-fold-cross-validation was used to train the models. The classifiers were first trained on a wide range of features obtained from the WebNLG corpus (16 features). After running a variable importance analysis, we selected a subset of features for the final models. The detailed list of features are presented in Appendix A.

¹Previously, only [Cao and Cheung \(2019\)](#) used pre-trained embeddings, but no ablation study was done.

²We also explored other ways of using BERT, such as using only BERT plus a feed forward layer to obtain h , or not freezing parameters of BERT while training. The resulting models had low performance in all cases.




Figure 1: Confusion Matrices for 4-way classification results of XGBoost (left) and c-RNN+GloVe (right), where PRO, PN, DES, and DEM are pronoun, proper name, description and demonstrative respectively.

3.3 Evaluation

Implementation Details. We tuned hyperparameters of each of our models on the development set and chose the setting with the best macro F1 score. For the BERT model, we used the cased BERT-BASE³ and added all entity labels into the vocabulary to avoid tokenisation. When re-training BERT on WebNLG, we set the masking probability to 0.15 and trained it for 25 epochs.

For the 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 the macro averaged precision, recall, and F1 on the test set. We run each model for 5 times, and report the averaged performance. As for the dataset, we use the v1.5 of WebNLG (Castro Ferreira et al., 2019) and use only seen entities.

Results. Table 3 shows the results of different classification tasks. Generally, all neural variants outperform the machine learning baseline. The performance difference is small in the case of binary classification, while it is much bigger for 3- and 4-way classifications. This is because the 2-way classification (i.e., pronominalisation) has clearly less complexity than the other two alternatives, and, thus, the feature set used by the baseline results in almost similar outcomes to neural models.

Comparing neural variants to each other, the results show that the simpler c-RNN wins over ConATT in 4-way classification, and has on par performance with ConATT for 3- and 2-way classifications. One possible explanation is that ConATT first breaks down the input into three pieces (i.e., the target entity as well as pre- and pos-context),

encodes them separately, and merges the encoded representations back before being sent to make predictions. This “divide and merge” procedure might hinder the model from learning some useful information.

Regarding the effectiveness of incorporating pre-trained models, GloVe embeddings have positive impact on c-RNN only in case of 4-way predictions, and have no contribution to 2- and 3-way classifications. Moreover, it has negative effect on ConATT: the performance diminishes when GloVe is used. It is surprising to see that in case of c-RNN, BERT has negative effect on 4- and 3-way predictions (the F1 score reduced from 64.86 and 83.63 to 62.15 and 82.15 respectively). For pronominalisation, BERT slightly boosts the performance (from 89.09 to 89.42), but this boost is not as much as BERT’s boosting effect on other NLP tasks. This is probably because although BERT was re-trained on WebNLG delexicalised sentences, the entity labels still function as noise for BERT.

To obtain insights into the behaviours of the deep learning and classic ML-based models for RFS, we depict the confusion matrices of XGBoost and the best performing neural model c-RNN+GloVe in Figure 1 for the 4-way classification. The confusion matrices suggest that both models do a good job in selecting pronouns and proper names (that is why the performance difference in the 2-way classification is small), and both perform poorly in choosing demonstratives (probably due to the fact that demonstratives are extremely infrequent in WebNLG). The main difference between the two models is in distinguishing proper names from descriptions. The XGBoost model wrongly predicted the descriptions as proper names in 62.58% of the cases, while the neural c-RNN+GloVe

³huggingface.co/bert-base-cased

model did this wrong prediction in 20.18% of the times. This difference in the performance of the two models might be because the neural models learnt some useful features from the discourse which are not covered in our feature engineering procedure. Furthermore, after looking into the WebNLG dataset, we noticed that various RE cases are annotated incorrectly. For example, WebNLG annotates “United States” as a proper name, and “the United States” as a description. The incorrect annotations might increase the confusion between choosing description and proper name in both XGBoost and c-RNN+GloVe.

4 Probing RFS models

We use a logistic regression classifier as our probing classifier. Concretely, for each input, we first use a model discussed in section 3 to obtain its representation R . As mentioned in section 3, we ran each model five times and reported their averaged scores. For the probing tasks, we use the representations of the models with the best RFS performance on the development set.

4.1 Probing Tasks

Following our observations in section 2.2, we formulate the following probing tasks.

Referential Status. The referential status of the target entity influences the choice of RF in both linguistic (Chafe, 1976; Gundel et al., 1993) and computational studies (Castro Ferreira et al., 2016). In this study, we define referential status on two levels: discourse-level and sentence-level. The former (**DisStat**) has two possible values: (a) discourse-old (i.e., the entity has appeared in the previous discourse) and (b) discourse-new (i.e., the entity has not appeared in the previous discourse). Sentence-level referential status (**SenStat**) also consists of two values: (a) sentence-new (i.e., the RE is the first mention of the entity in the sentence), and (b) sentence-old (i.e., the RE is not the first mention).

Syntactic Position. Entities in subject position are more likely to be pronominalised than in object position (Brennan, 1995; Arnold, 2010). Therefore, in the syntax probing task (henceforth **Syn**), we do binary classification: subject or object.

Recency. Recency has been used as a vital feature in many of the previous REG or RFS systems (Greenbacker and McCoy, 2009; Kibrik et al., 2016). It measures the distance between the target

entity and its closest antecedent. There are various ways of estimating the recency of a target entity given its context. We hereby use two measures: (1) the number of sentences between the target entity and its antecedent (**DistAnt**), which consists of four possible values: the entity and its antecedent are (a) in the same sentence, (b) one sentence away, (c) more than one sentence away, and (d) the entity is a first mention (to distinguish first mentions from subsequent mentions). (2) whether there is an intervening referent between the target and its nearest antecedent (**IntRef**) (Greenbacker and McCoy, 2009). In other words, it checks whether the target and the preceding RE are coreferential. This feature has three 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. Note that the existence of intervening markables might signal the existence of a competition (if the intervening referent has the same animacy and gender values as the target RE).

Discourse Structure Prominence. As mentioned in section 2, the “organizational” properties of discourse may influence the prominence status of the entities. We introduce three probing tasks capturing different properties of the discourse. (1) *Local prominence* (**LocPro**): The idea of local prominence is coming from Centering Theory (Grosz et al., 1995). It is a hybrid feature of DisStat and Syn. Concretely, we use the implementation of Henschel et al. (2000): an entity is *locally prominent* if it is “discourse-old” and “realised as subject”. It is a binary feature with two possible values: (a) locally prominent, and (b) not locally prominent. (2) *Global prominence* (**GloPro**): This feature is based on the notion of global salience in Siddharthan et al. (2011), asking whether the entity is a minor or major referent in the text. According to them, “the frequency features are likely to give a good indication of the global salience of a referent in the document” (p. 820). We define a binary feature in which the most frequent entity in a text is marked as globally prominent. (3) *Meta-prominence* (**MetaPro**): In line with global prominence, we also want to explore to what extent prominence beyond a single text (e.g. on a text collection level) may impact the way people refer. In the context of the current circumstances, the sentence “I received *my vaccine* today” is unambiguous, and the RE *my vaccine* needs no extra modification (e.g. *my COVID-19 vaccine*); how-

Model	Type	DisStat	SenStat	Syn	DistAnt	IntRef	LocPro	GloPro	MetaPro
Random	-	49.57 (41.83)	33.11 (22.87)	49.65 (48.99)	25.19 (14.90)	33.30 (22.92)	50.05 (49.84)	49.75 (48.02)	25.24 (25.20)
Majority	-	86.91 (46.50)	86.91 (31.00)	61.27 (37.99)	86.91 (23.25)	86.91 (31.00)	56.28 (36.01)	68.49 (40.65)	28.12 (10.97)
c-RNN	4-way	85.16 (84.06)	93.28 (73.72)	94.16 (85.34)	92.84 (53.84)	91.71 (55.43)	83.37 (82.92)	70.62 (56.00)	44.76 (42.32)
	3-way	84.78 (83.72)	92.59 (72.60)	93.50 (83.60)	92.58 (54.78)	91.24 (53.21)	82.17 (81.67)	70.87 (56.70)	45.42 (41.79)
	2-way	88.84 (88.04)	92.77 (73.84)	93.49 (84.00)	92.53 (54.93)	91.01 (52.31)	86.08 (85.69)	71.24 (59.98)	44.32 (41.65)
c-RNN +GloVe	4-way	85.84 (84.85)	93.58 (74.59)	94.56 (87.04)	93.30 (55.67)	92.06 (55.93)	83.71 (83.20)	70.55 (53.53)	44.23 (41.71)
	3-way	85.09 (83.89)	91.89 (67.24)	93.23 (82.48)	91.72 (50.94)	90.92 (51.17)	82.08 (81.44)	70.20 (52.49)	45.58 (42.34)
	2-way	88.88 (88.02)	92.38 (71.25)	93.32 (82.67)	92.25 (53.67)	90.94 (51.43)	85.81 (85.22)	71.78 (63.17)	44.92 (41.03)
c-RNN +BERT	4-way	95.85 (90.64)	94.41 (78.04)	84.05 (82.71)	93.60 (56.91)	92.27 (54.30)	82.03 (81.67)	71.04 (54.24)	45.27 (43.07)
	3-way	94.00 (84.80)	92.74 (72.29)	85.12 (84.08)	92.57 (54.21)	91.28 (53.25)	82.92 (82.53)	71.69 (57.31)	43.64 (42.80)
	2-way	94.59 (87.28)	92.94 (69.69)	85.75 (84.74)	92.50 (54.19)	92.06 (54.88)	83.27 (82.77)	73.80 (63.07)	41.05 (40.75)
ConATT	4-way	94.86 (87.81)	94.12 (77.11)	88.64 (88.00)	93.69 (57.09)	92.11 (55.88)	86.93 (86.34)	72.22 (60.15)	48.37 (46.14)
	3-way	93.91 (84.39)	93.15 (74.19)	87.43 (86.66)	92.93 (55.26)	91.35 (54.09)	85.32 (84.56)	72.61 (60.61)	49.35 (47.47)
	2-way	93.74 (84.20)	92.78 (73.18)	89.01 (88.44)	92.50 (53.98)	91.19 (53.64)	87.05 (86.75)	70.65 (56.39)	44.24 (41.81)
ConATT +GloVe	4-way	94.86 (87.82)	94.10 (77.70)	87.98 (87.24)	93.66 (57.52)	92.10 (55.22)	86.06 (85.69)	71.94 (58.54)	53.19 (49.94)
	3-way	93.79 (84.35)	92.78 (72.83)	89.54 (88.91)	92.59 (54.23)	91.39 (51.96)	87.09 (86.80)	71.91 (59.05)	49.27 (46.36)
	2-way	93.81 (84.38)	92.86 (73.21)	87.69 (86.96)	92.84 (56.14)	91.50 (53.33)	85.61 (85.27)	72.48 (62.46)	44.47 (39.63)

Table 4: Results of each probing task. Results are reported in the format of A(B), where A is the accuracy and B is the macro F1.

ever, a couple of years from now, a richer RE may be needed to refer to the vaccine. The idea behind this exploratory feature is that people might use less semantic content to refer to the referents which are well known outside of the text. Based on the number of mentions of a target entity in the whole webNLG, four possible values, each of which representing an interval, are assigned to each RE: (a) $[0, 50)$, (b) $[50, 150)$, (c) $[150, 290)$, and (d) $[290, \infty)$. For example, the category $[0, 50)$ contains those entities that occur fewer than 50 times in the corpus.

4.2 Importance Analysis

We conducted a feature importance analysis to find out which features used in the probing tasks had the highest contributions to the feature-based ML models. This analysis functions as a sanity check to find out whether the representations have learnt




Figure 2: Feature importance of XGBoost classifiers for 4-way predictions. Higher loss shows greater importance of a feature. Results for 2-way and 3-way classification can be found in the Appendix B.

the features contributing the most to the RFS task.

To assess the importance of the features used in the probing tasks, we train XGBoost models, only using features from section 4.1, and calculate

the model-agnostic permutation-based variable importance of each model (Biecek and Burzykowski, 2021). Concretely, we measure the extent to which the performance changes if we remove one of the features. Figure 2 depicts the performance change for each feature. According to the figure, DisStat and Syn contribute the most. LocPro is the least important feature because it is a hybrid combination of DisStat and Syn. Removing it while keeping DisStat and Syn will not hurt the performance of the model a lot. Considering that DisStat and Syn are both highly vital features, LocPro is much more important than what the experiment suggests. In addition to DisStat and Syn probing tasks, we also expect high performance for the LocPro task.

4.3 Probing Results

We mentioned earlier that we conduct probing tasks to find out whether the RFS models’ latent representations encode the features mentioned in section 4.1. High performance in probing tasks would indicate that the features are encoded in the latent representations of the models.

We evaluate probing tasks using the accuracy and macro-averaged F1 scores. Each probing classifier was trained 5 times. We report the averaged value. Additionally, 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.

Results of Each Probing Task. Compared to the `random` baseline, all neural models have achieved higher performance on all tasks. (1) Referential status and syntactic position: all models exhibit consistently high performance on DisStat, SenStat, and Syn. This shows that, at least for the WebNLG corpus, all neural models can learn information of referential status and syntactic position; (2) Recency (i.e., DistAnt and IntRef): all models perform worse compared to the referential status and syntax probes. Although they do not have bad accuracy scores, their F1 scores are lower than that of DisStat, SenStat, and Syn, and are closer to the baselines. This finding is consistent with the results of section 4.2, where DistAnt and IntRef were found to be less important (comparing to DisStat and Syn). One possible explanation is that, in the WebNLG corpus, 67% of the documents contain only one sentence, making recency-related features play less role. As another possible explanation, in line with the previous probing works on corefer-

ence and bridging anaphora (Sorodoc et al., 2020; Pandit and Hou, 2021), models have more difficulty capturing long-distance properties; (3) Discourse structure prominence: since LocPro is a hybrid of DisStat and Syn, all models handled it to a large degree. Meanwhile, neural models appear to handle GloPro and MetaPro worse than other features since the performance of their corresponding probing tasks is closer to the baselines⁴. These results are in contrast with the importance analysis results, which suggested that both GloPro and MetaPro are important features (ranking 3 and 4 in Figure 2). Learning GloPro and MetaPro requires a model to have an overall understanding of the whole input document or the whole corpus, which the neural models might not be able to acquire.

Comparing c-RNN and ConATT. In section 3, we concluded that the c-RNN model works better than ConATT on 4-way RF classification. Nevertheless, when probing, we observed that ConATT does a better job in many tasks, including DisStat, LocPro, GloPro, and MetaPro. To understand why, we looked into the WebNLG dataset and found that 86.91% of the REs in WebNLG are first mentions, and 21% of the documents talk about the entity “*United States*”. This suggests that REs in WebNLG are not representative of the realistic use of REs. Therefore, although ConATT learns more contextual features, it still has a lower performance. ConATT’s better learning of referential status (i.e., DisStat) is probably a benefit of using self-attention, which helps the model capture longer dependencies than RNNs.

The Effect of Pre-training. As mentioned earlier, the secondary objective of this study is to find out whether RFS can benefit from pre-trained word embeddings and language models. The effect of incorporating the GloVe embeddings is not significant to c-RNN and ConATT. The major contribution of BERT is helping with learning DisStat (which is, again, probably a result of using self-attention). Akin to the above discussion, since the majority of the entities in WebNLG are first mentions, the increased accuracy boost in the DisStat task is not enough to boost the overall performance of RFS.

Comparing Different RF Classifications. It also appears that models learn different infor-

⁴Note that, for MetaPro, the `Majority` has low F1 score because the distribution of the values of MetaPro is balanced.

mation using different label sets (classes). For example, 2-way classification (i.e., pronominalisation) helps `c-RNN` learn more about referential status. But in case of models with attention mechanism (i.e., `ConATT`, `ConATT+GloVe` and `c-RNN+BERT` models), referential status is learnt better in 4-way classification models. Also, in case of `ConATT (+GloVe)`, we observe that more fine-grained classifications help the model learn more about meta prominence (i.e., `MetaPro`).

5 Conclusion

Our aim is to understand whether neural models capture the features associated with the task of RFS. To this end, we defined 8 probing tasks in which we focused on referential status, syntactic position, recency, and discourse structure. The probing results suggest that the probe classifiers always performed better than the `random` and the `majority` baselines. The performance was consistently good in the tasks associated with referential status, syntax and local prominence.

It is worth noting that probing has its own shortcomings. For instance, on the one hand, low probing performance does not always mean the feature is not encoded, but could also mean that such a feature does not matter to RFS. To mitigate this issue, we conducted a complementary ML-based variable importance analysis; in this analysis, discourse status and syntactic position came out as the factors with the highest contributions. These features were also predicted very well in the probing tasks. However, these results should still be taken with a pinch of salt: the variable importance has been conducted on the ML model and not on the neural models. We cannot be certain that the same features contribute to all the models similarly: a feature might be quite important in the machine learning model, but not as important in the neural models. On the other hand, some researches have questioned the validity of probing methods. They found out that it is difficult to distinguish between “learning the probing task” and “extracting the encoded linguistic information” (Hewitt and Liang, 2019; Kunz and Kuhlmann, 2020) for a probing classifier. This suggests that higher performance of a probing classifier does not necessarily mean more linguistic information has been encoded. This prevents us from directly quantifying *how well* the linguistic information has been learnt using the performance of probing classifiers and requires us to

make conclusions more carefully.

From our probing efforts, we conclude that: (1) All neural models have learnt some information about the features associated with the probing tasks, but how well they have learnt this information is yet to be assessed; (2) The `webNLG` corpus, which has often been used for the study of discourse REG, is not ideally suitable for studying discourse-related aspects of RFS, because the texts are too short and the majority of REs are first mentions. This leads to bias in the evaluation of RFS and REG algorithms; (3) When it comes to the question of how well a RFS feature can be learnt, it matters what neural architecture and label set are used, and whether the model is pre-trained or not. Using an attention mechanism and more fine-grained label sets help a model learn more information; (4) All models perform poorly in terms of learning those features, such as `GloPro` and `MetaPro`, that do not derive from the text itself but from the wider context in which it is written and read. We believe that future models should take these lessons into consideration.

In future, we plan to extend the current study from three angles. First, we plan to conduct experiments on different corpora. The `webNLG` corpus used in this study consists predominantly of extremely short documents with an average length of only 1.4 sentences/document; consequently the majority of REs are first mentions. We hope to find a more representative distribution of uses of referring expressions in other corpora such as `OntoNotes` (Hovy et al., 2006), which contain longer texts. Secondly, we plan to conduct experiments on other languages than English, in particular ones that favour zero pronouns (e.g., Chinese (Chen et al., 2018)), because these pose new challenges for the task of RFS. Thirdly, we plan to design new probing tasks on the basis of other factors that could influence RFS, such as animacy, competition and positional attributes (see Same and van Deemter (2020) for an overview).

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”.

References

- Christoph Alt, Aleksandra Gabryszak, and Leonhard Hennig. 2020. [Probing linguistic features of sentence-level representations in neural relation extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1534–1545, Online. Association for Computational Linguistics.
- Mira Ariel. 1990. *Accessing Noun-Phrase Antecedents*. Routledge.
- Jennifer E Arnold. 2001. The effect of thematic roles on pronoun use and frequency of reference continuation. *Discourse processes*, 31(2):137–162.
- 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.
- 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 Varges. 2007. [Generation of repeated references to discourse entities](#). In *Proceedings of the Eleventh European Workshop on Natural Language Generation (ENLG 07)*, pages 9–16, Saarbrücken, Germany. DFKI GmbH.
- Przemyslaw Biecek and Tomasz Burzykowski. 2021. *Explanatory model analysis: explore, explain, and examine predictive models*. CRC Press.
- 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, 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.
- 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.
- Wallace Chafe. 1976. Givenness, contrastiveness, definiteness, subjects, topics, and point of view. *Subject and topic*.
- 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 and Carlos Guestrin. 2016. [Xgboost: A scalable tree boosting system](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 785–794, New York, NY, USA. Association for Computing Machinery.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, 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.
- Kees van Deemter. 2016. *Computational models of referring: a study in cognitive science*. MIT Press.
- 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.
- 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.
- 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.
- Charles Greenbacker and Kathleen McCoy. 2009. UDeI: 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.
- Klaus von Heusinger and Petra B Schumacher. 2019. Discourse prominence: Definition and application. *Journal of Pragmatics*, 154:117–127.
- 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.
- Stefan Hinterwimmer. 2019. Prominent protagonists. *Journal of Pragmatics*, 154:79–91.
- 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.
- Andrew Kehler, Laura Kertz, Hannah Rohde, and Jeffrey L Elman. 2008. Coherence and coreference revisited. *Journal of semantics*, 25(1):1–44.
- 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*, 7:1429.
- Emiel Krahmer and Kees van Deemter. 2012. **Computational generation of referring expressions: A survey**. *Computational Linguistics*, 38(1):173–218.
- Jenny Kunz and Marco Kuhlmann. 2020. **Classifier probes may just learn from linear context features**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5136–5146, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Onkar Pandit and Yufang Hou. 2021. **Probing for bridging inference in transformer language models**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4153–4163, Online. Association for Computational Linguistics.
- 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.
- Ellen F Prince. 1981. Towards a taxonomy of given-new information. *Radical pragmatics*.

- 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.
- Advait Siddharthan, Ani Nenkova, and Kathleen McKeown. 2011. Information status distinctions and referring expressions: An empirical study of references to people in news summaries. *Computational Linguistics*, 37(4):811–842.
- Ionut-Teodor Sorodoc, Kristina Gulordava, and Gemma Boleda. 2020. [Probing for referential information in language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4177–4189, Online. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- 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.

A Further details on the XGBoost models

As mentioned earlier, for the RFS task, we firstly created the XGBoost models using a wide selection of features. Afterwards, we ran a variable importance analysis on the models, and chose a smaller subset of features for each classifier. The selected features are presented in Table 5.

B Further results of importance analysis

Figure 3 depicts the variable importance results for the 2-way and 3-way classification tasks. As mentioned in the paper, there is a high degree of agreement between the ordering of the variables in the 3 models.

To get a better idea about the contribution of each variable to the decisions made by the models, Figure 4 demonstrates the shapley values for 100 random orderings of explanatory variables in the 4-way classification model. The figure clearly shows that the model has failed to learn the demonstrative class. For other decisions, the model majorly uses 2 features, namely `DisStat` (referential status) and `Syn` (syntactic role).

Feature	Definition	2-way	3-way	4-way
Syn	Description is provided in the main text.	✓	✓	✓
Entity	Values: Person, Organisation, Location, Number, Other	✓	✓	✓
Gender	Values: male/female/other	✓	✓	✓
DisStat	Description is provided in the main text.	✓	✓	✓
SenStat	Description is provided in the main text.	-	✓	✓
DistAnt_S	Description is provided in the main text (DistAnt).	✓	✓	✓
DistAnt_W	Distance in number of words (5 quantiles)	✓	-	✓
Sent_1	Does RE appear in the first sentence?	✓	✓	✓
MetaPro	Description is provided in the main text.	✓	✓	✓
GloPro	Description is provided in the main text.	✓	✓	✓

Table 5: Features used in the XGBoost models.

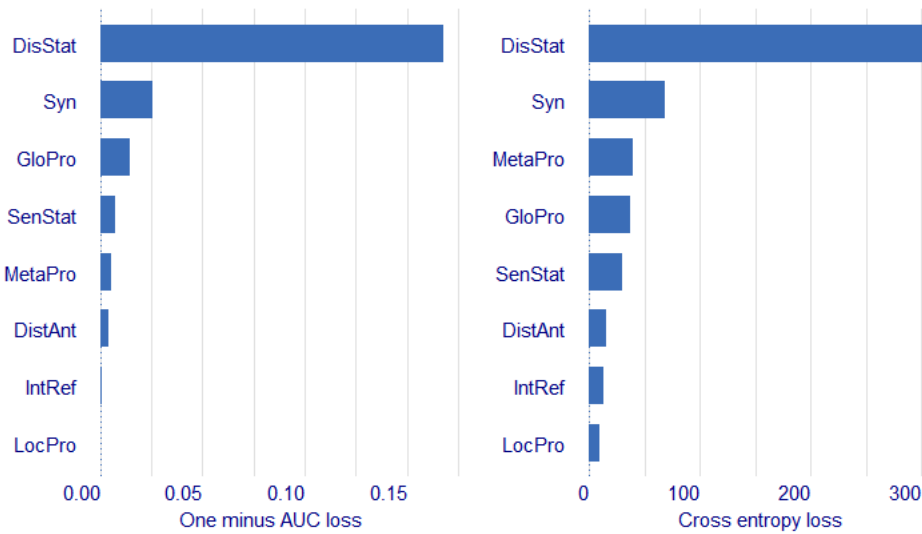


Figure 3: Feature Importance of the XGBoost 2-way (left figure) and 3-way (right figure) predictions.




Figure 4: Shapley values with box plots for 100 random orderings of explanatory variables in the XGBoost 4-class model.