

Structural Ambiguity and its Disambiguation in Language Model Based Parsers: the Case of Dutch Clause Relativization

Gijs Wijnholds

Leiden Institute of Advanced Computer Science
Leiden University
g.j.wijnholds@liacs.leidenuniv.nl

Michael Moortgat

Institute of Language Sciences
Utrecht University
m.j.moortgat@uu.nl

Abstract

This paper addresses structural ambiguity in Dutch relative clauses. By investigating the task of *disambiguation by grounding*, we study how the presence of a prior sentence can resolve relative clause ambiguities. We apply this method to two parsing architectures in an attempt to demystify the parsing and language model components of two present-day neural parsers. Results show that a neurosymbolic parser, based on proof nets, is more open to data bias correction than an approach based on universal dependencies, although both setups suffer from a comparable initial data bias.

1 Introduction

Ambiguity pervades natural language and as such forms one of the central challenges for natural language understanding (NLU) systems. Given the fact that most such systems rely on large-scale deep learning architectures, the presence of structural biases in the training data used may affect a system’s capacity for disambiguation. Specifically in the case of parsing, typical architectures rely on the assumption of just a single correct parse, although many may exist. This assumption then may force a bias into the training, both on the lexical and on the syntactic level.

In this paper, we study syntactic ambiguities in Dutch, where a structural ambiguity affects the interpretation of relative clauses. The preferred reading as subject or object relativisation will typically be determined by lexical choice. A running example in Dutch, with its two possible interpretations in English, is given below:

- (a) de dokter die de patiënt geneest
- (b) *the doctor who cured the patient*
- (c) *the doctor whom the patient cured*

In this example, the verb ‘cure’ displays a strong selectional preference for a doctor as its subject and a patient as its object, so one would expect

that the subject-relative interpretation (*b*) is the preferred reading. However, grounding the phrase in a prior sentence disambiguates the ambiguous relative clause: by prepending the phrase “The patient cured the doctor”, one is led to infer that the object-relative reading of (*c*) is in fact preferred. In extreme cases, one finds examples of lexical choice that *block* one of the readings, typically due to the semantic class of the subject and object being different. An example is “De man die de boterham eet” (*The man who eats the sandwich*), where it is semantically implausible for the verb’s arguments to be reversed. The situation is summarized in Table 1, where different possible orderings of the relative clause together with a prior sentence lead to a different expected readings of the relative clause.

Arguing that a parser typically exploits statistical properties of its training corpus, but should additionally rely on both lexical and syntactic cues inside of that corpus, we carry out an experiment test a parser’s capacity for disambiguation in context. Specifically, we extract a set of selectional preferences for Dutch transitive verbs, that we classify according to their *reversibility*, i.e. whether subject and object can be interchanged, additionally indicating whether there is a strong preference for a noun as subject or object. This leads to three classes of (*s, v, o*) triples that we then use to generate a test set of Dutch relative clauses together with prior sentences, to test a parser’s capability of disambiguation in context.

By evaluating two different parsing regimes that both are built on top of a language model, we investigate the encoding of structural bias in the parser training data, the possibility of mitigating structural bias, and attempt to pinpoint to what extent lexical knowledge or syntactic information is employed by the parsers in question.

Our contributions in this paper are therefore threefold: we (1) provide a dataset of selectional preferences for Dutch (*s, v, o*) triples with an addi-

| Prior sentence | Target phrase | Correct reading | Plausible |
|-------------------------------|----------------------------------|-----------------|-----------|
| De man eet de boterham. | De man die de boterham eet | subj. rel. | ✓ |
| De man eet de boterham. | De boterham die de man eet | obj. rel. | ✓ |
| De boterham eet de man. | De man die de boterham eet | n/a | ✗ |
| De dokter geneest de patient. | De dokter die de patient geneest | subj. rel. | ✓ |
| De patient geneest de dokter. | De dokter die de patient geneest | obj. rel. | ✓ |
| De dokter geneest de patient. | De patient die de dokter geneest | obj. rel. | ✓ |
| De patient geneest de dokter. | De patient die de dokter geneest | subj. rel. | ✓ |

Table 1: Different cases in our disambiguation experiment, where the provided prior sentence determines the interpretation of the target phrase. The top three rows are cases of an *irreversible* (s, v, o) triple where interchanging subject and object leads to an implausible case that is not included in our experiments. The bottom four rows are cases of a *reversible* (s, v, o), albeit with a strong selectional preference for one interpretation. By adding the prior sentence, the interpretation of the target phrase is disambiguated.

tional layer of classification according to semantic noun classes, and (2) create a novel test set targeting structural ambiguity in the interpretation of Dutch relative clauses. Finally, we (3) provide a number of experiments indicating that structural bias is easily encoded, but not so easily mitigated in a language model-based parser. The code and data for this work is distributed, but organized in a single repository¹.

2 Background

Probing and syntactic sensitivity Previous work has used probing, where a small neural network is attached to a large language model to extract task-specific information, to argue that large scale language models like BERT have internalized some linguistic knowledge during pretraining (Tenney et al., 2019), and there appears to be some consensus of the syntactic awareness of BERT models (Rogers et al., 2020). Specifically, studies have indicated the possibility of extracting parse trees from BERT representations successfully (Hewitt and Manning, 2019; Vilares et al., 2020).

Another line of research into syntactic sensitivity investigates the probabilities of language models in a masked language modelling environment, where studies typically define surprisal rates to measure the degree to which the language model’s predictions coincide with human-like behaviour in the face of syntactic ambiguities (Futrell et al., 2019; Hu et al., 2020; Arehalli et al., 2022; Aina and Linzen, 2021), typically focusing on garden-path effects. A related approach uses priming to investigate the language models’ response to structurally

similar sentences (Sinclair et al., 2022).

These studies differ from the current paper in that we explicitly target sentences that are not disambiguated without adding extrasentential context, which should for a parser to infer the intended syntactic analysis. Hence, our setup relies on a probing-like paradigm where we evaluate a parser on top of a language model.

Dutch NLP The rising interest in large scale language models in the NLP community has led to a number of investigations for Dutch specifically. Two dominant Dutch language models have been developed, based on the respective BERT (Devlin et al., 2019) (BERTje, de Vries et al. (2019)) and the RoBERTa (Liu et al., 2019) architecture (RobBERT, Delobelle et al. (2020)). On the evaluation side there have been several studies using Dutch-specific phenomena to evaluate the respective monolingual language models: the work of (Wijnholds and Moortgat, 2021) introduces a parallel Natural Language Inference (NLI) dataset for Dutch, showing that Dutch NLI is more difficult to tackle than its original English version. More recently, a range of probing studies has been performed investigating verb-subject dependencies in the face of syntactic constructions involving discontinuity (Kogkalidis and Wijnholds, 2022; Moortgat et al., 2023) and ellipsis (Haagen et al., 2022), aside from a more subtle Dutch NLI challenge (Wijnholds, 2023).

More closely related to constructions involving relative pronouns is the work of Allein et al. (2020), which introduces relative pronoun prediction for Dutch *die* and *dat* as a binary classification task, where the surrounding context determines the

¹https://github.com/gijswijnholds/relpron_disambiguation

choice of the neuter (*dat*) or non-neuter (*die*) pronoun. This can be modelled as a masked language modelling task (Delobelle et al., 2020), reaching high performance. A more complex variation of this task is defined by Bouma (2021), where the experiment investigates the language model’s capacity to predict relative pronoun attachment.

Against these studies, our approach differs in that we do not directly target the language model probabilities, but rather investigate the attached parsers, thereby indirectly assessing the contextualization power of the underlying language model. In that sense, the current work is more in line with prior, more theoretical work that takes parser ambiguity into account (Moortgat and Wijnholds, 2017; Moortgat et al., 2020; Wijnholds et al., 2020).

Selectional preferences The tendency of predicates to combine with certain arguments is known as selectional preference (Katz and Fodor, 1963) and is relevant to different NLP applications, such as word sense disambiguation (McCarthy and Carroll, 2003) and semantic role labelling (Gildea and Jurafsky, 2002). A plurality of datasets exist to evaluate automatic selectional preference acquisition (McRae et al., 1998; Padó, 2007; Zhang et al., 2019), and while different probabilistic and neural methods have been evaluated on the preference induction task (Resnik, 1997; Van de Cruys, 2014). More recently, some studies investigated how knowledge about selectional preference is encoded in pretrained embeddings, but with inconclusive results (Metheniti et al., 2020; Muthupari et al., 2022). These studies have largely focused on English. For Dutch, while previous work tried to exploit selectional preferences to improve parser accuracy (van Noord, 2010), in this work we rather use selectional preference to investigate the bias of existing parsers, as a precursor to potential architectural considerations, in a setting where language models are commonly employed in parser development. We make the code for the different components available online.

3 Data Generation

In order to generate a suitably large set of relative clause patterns, we proceed with a pipeline for extracting suitable subject-verb-object triples from a large corpus. First we perform a probability-based extraction of base triples. Then we use lexical information from a dictionary to filter out relevant nouns and classify them according to their seman-

tic class, after which in a final step we perform a manual filtering and classification of the obtained subject-verb-object triples to guarantee correctness.

Triple extraction For the first step, we need a way to extract subject-verb-object triples from a large corpus. To this end, we iterate over Lassy Large (van Noord et al., 2013) – a 700M word corpus with automatically assigned syntactic annotations – and extract all the cases of transitive verbs and their respective subjects and objects.

To compensate for parsing errors and infrequent observations, we rely on posterior probability (Resnik, 1997), the simplest measure that was shown to give best performance on a selectional preference acquisition task (Zhang et al., 2019). Posterior probability allows us to easily filter out those triples that occurred most frequently, allowing us to consider the most canonical triples only. For a given (s, v, o) triple, its posterior probability is defined as follows:

$$p(s, v, o) = \frac{f(s, v, o)}{\sum_{s', o'} f(s', v, o')}$$

Dictionary-based classification After gathering all (s, v, o) triples and removing stopwords, we perform a dictionary-based filtering and classification in two steps. First, we obtain an exhaustive list of Dutch nouns with their respective semantic class² from the *Algemeen Nederlands Woordenboek*³, a comprehensive online dictionary of Dutch.

We use these categories to classify all the subjects and objects in the extracted (s, v, o) triples, after which we organize triples according to the most frequent noun categories, preventing as much as possible infrequent or illicit combinations of a verb with a given pair of nouns. By organizing triples by whether subject and object fall into the same semantic category, we have an initial estimate of whether a triple is *irreversible*, and in cases where the semantic category of the subject and object coincide, we estimate how strong the preference of the verb for the particular subject and object is based on the frequency of the (s, v, o) triple relative to its inverted (o, v, s) triple.

Manual filtering This initial estimation gives us a comprehensive list of subject-verb-object triples,

²person, animal, plant, substance, object, abstract, mass noun

³<https://anw.ivdnt.org>

that we finally filter manually, making a per case decision on whether a triple is *irreversible*, reversible with *strong* preference for the given subject and object, or reversible with a *weak* such preference, meaning that object and subject could be swapped without leading to an implausible triple. An example of the latter is “De toerist herkent de reiziger” (*The tourist recognizes the traveller*).

Generating relative clauses After the two-step process described above, we have a robust set of (s, v, o) triples that we can use to generate the desired test cases for our experiment. In total, we obtained 3304 irreversible triples, 370 triples with a strong preference for the regular relation, and 724 where subject and object could be easily interchanged.

From these triples, we generate relative clauses following the pattern displayed in Table 1. First, for an irreversible (s, v, o) triple, we create a relative clause with the subject s as the head noun (**S Pron O V**) as well as the reversed relative clause (**O Pron S V**), which due to the irreversibility of the (s, v, o) triple must be analyzed with the subject-relative and object-relative reading respectively. We prepend in both cases a prior sentence in SVO order to be able to inspect the effect of adding this context to the parser’s input.

For the reversible triples, we generate four cases. Again, we generate two variations of the relative clause, but additionally vary the prior sentence so it will force a reading of the relative clause, that is ambiguous without this context. This allows us to compare parser performance on both lexical and syntactic cues.

In the generation process, we use the *Algemeen Nederlands Woordenboek* to extract the gender of each noun (*de* for gendered nouns, *het* for neuter nouns), and the gender of the relative pronoun (*die* for gendered head nouns, *dat* for neuter head nouns).

4 Parsing Regimes

In our experiments we want to evaluate parsers that were built on top of a large language model, in order to distinguish the effect of the language model from that of the specific parsing strategy employed. For comparison purposes we test two different parsing regimes.

Neural proof nets The first parser we examine is a neurosymbolic parser based on a multi-modal

type-logical grammar that simultaneously encodes function-argument structure and dependency roles (Kogkalidis et al., 2023, 2020). This parsing setup exists alongside other *neuralizations* of categorial grammar parsers (Clark, 2021), but was explicitly developed for Dutch.

The architecture of this parser follows the typical structure of a categorial parser, where a *supertagging* component assigns to words logical formulas that encode their intended combinatorial behaviour, followed by a process of *proof search* that combines the formulas into a proof representing the full parse history. In the system of (Kogkalidis et al., 2023), supertagging is implemented as a graph decoding network, that learns to construct the tree structures representing the formulas of the logical formalism from an underlying (Dutch) BERT model (de Vries et al., 2019). Proof search is implemented as *neural proof net* search, which amounts to linking atomic subformulas of opposite polarity in a way that determines the correct dependency and function-argument relations.

For a full exposition of this parser we refer the reader to (Kogkalidis et al., 2020); for the sake of our experiments it is enough to consider the two possible correct supertagging assignments for the Dutch relative clause, illustrated in Table 2.

Universal dependencies The second parser we evaluate approaches parsing as a sequence labelling task, following the work of Strzyz et al. (2019).

Specifically, sentences are encoded using a relative part-of-speech based encoding, with each word assigned a triple (i, p, d) where p refers to the part-of-speech of the word’s head, i indicating its relative location, and d the word’s dependency label. For example, the label $(+1, V, nsubj)$ says that the current word is in the *nsubj* dependency relation with respect to the first word to its right that carries the *V* part-of-speech tag. The root is encoded by labelling its dependent with $(-1, \text{ROOT}, \text{root})$. A full example is given in Table 3 which contains the encoding of the two possible parses for the Dutch relative clause. The relative part-of-speech based encoding was found to be the highest performing in the experiments of Strzyz et al. (2019) and so we use it in our experiments. Labels are considered atomically, and as such the parser is implemented as a standard token classification model, where a token-level classifier is fine-tuned along with a BERT model.

| | | | | | | |
|------------|-----------------------------------|---------|---|-----------------------------------|--------|--|
| | De | patiënt | die | de | dokter | geneest |
| Subj. rel. | $\square_{det}(N \rightarrow NP)$ | N | $\diamond_{relcl}(\diamond_{su}VNW \rightarrow S) \rightarrow \square_{mod}(NP \rightarrow NP)$ | $\square_{det}(N \rightarrow NP)$ | N | $\diamond_{obj1}NP \rightarrow \diamond_{su}VNW \rightarrow S$ |
| Obj. rel. | $\square_{det}(N \rightarrow NP)$ | N | $\diamond_{relcl}(\diamond_{obj1}VNW \rightarrow S) \rightarrow \square_{mod}(NP \rightarrow NP)$ | $\square_{det}(N \rightarrow NP)$ | N | $\diamond_{obj1}NP \rightarrow \diamond_{su}VNW \rightarrow S$ |

Table 2: Supertagging assignment of the neural proof net parser for the subject-relative and object-relative interpretation of the Dutch relative clause. Function-argument structure is encoded by the linear implication, where $A \rightarrow B$ denotes a function consuming a phrase of type A to produce a result of type B . The unary operations \diamond_d, \square_d encode dependency structure, where heads assign dependency role d to their *complements* by means of \diamond_d marking, and \square_d allows *adjuncts* to project their dependency role d . The parser assigns a *higher-order* formula to the relative pronoun, allowing the implicit gap for either the subject or object of the verb (“geneest”) in the body of the relative clause to be identified with the head noun (“de patiënt”). For the sake of our experiments, we can inspect the formula assigned to the relative pronoun ‘die’ to determine the interpretation of the relative clause, given the annotation of the gap type VNW with either the *su* or *obj1* dependency.

| | De | patiënt | die | de | dokter | geneest |
|------------|---------------------|-------------------------|-----------------------|---------------------|-----------------------|---------------------------|
| Subj. rel. | (+1,N, <i>det</i>) | (-1,ROOT, <i>root</i>) | (+1,V, <i>nsubj</i>) | (+1,N, <i>det</i>) | (+1,V, <i>obj</i>) | (-2,N, <i>acl:relcl</i>) |
| Obj. rel. | (+1,N, <i>det</i>) | (-1,ROOT, <i>root</i>) | (+1,V, <i>obj</i>) | (+1,N, <i>det</i>) | (+1,V, <i>nsubj</i>) | (-2,N, <i>acl:relcl</i>) |

Table 3: Relative part-of-speech based encoding of the subject-relative and object-relative interpretation of the Dutch relative clause. In both cases, the label assigned to the relative pronoun ‘die’ determines the interpretation of the relative clause.

5 Evaluation & Results

In our experiments, we test the parsers on three different scenarios: first, we inspect the existing structural bias present in the parsers due to training data statistics, in a setting where the parser only gets fed the relative clause. In such a setting, we would expect the parser to have a strong bias for irreversible cases, but an even distribution in accuracy on weakly reversible cases. Next, we observe the effect of parsing the relative clause when the underlying language model is allowed to contextualize against the prior sentence, expecting this to aid the parser in assigning the correct reading. In this setting, we would ideally hope for high accuracy across the board. Finally, we examine the effect of additionally finetuning the parser components on a small amount of training data to see if the parser can pick up on the task.

5.1 Experimental setup

We organize the test cases into a train/dev/test split in order to compare the parser baseline against a finetuning setting, where the parser is additionally trained to recognize examples of disambiguating context to learn how to choose the correct interpretation of the relative clause.

Data preparation In order to not allow overfitting of the model – we feed it data that is engineered to be task-specific – we select a small amount of training data against larger development and test

sets. We separate the verbs involved in the three different data sets as another measure to avoid overfitting. This leads to a training set of 2640 samples, against a development set of 4400 samples and a test set of 5556 samples.

Parser setup We initially train both parsers on the original examples from Lassy-Small, as the neural proof net parser was trained on this. However, plain evaluation on the test data is not an option given that we want to prepend a prior disambiguating sentence; the fact that BERT employs positional embeddings makes the parsers unsuitable for the task. Hence we train the parsers from scratch, prepending a random number of unattended tokens, ranging between 5-80 tokens. This ensures that the parser will be robust against the position shifting in later experiments. Due to the difference in parsing regime, baseline scores for the position-shifted parsers follow different evaluation metrics, displayed in Table 4.

| NPN | | UD | |
|-------|-------|-------|-------|
| 92.98 | 54.87 | 88.37 | 86.92 |

Table 4: Baseline parser accuracy scores.

For the neural proof net parser we report tagging accuracy (percentage of total supertags correctly predicted) and frame accuracy (percentage of sentences for which all supertags were correctly

predicted). These numbers are only slightly lower than the original parser.⁴ For the UD parser, we compute unlabelled and labelled attachment score, which are comparable to state of the art.⁵

5.2 First experiment

For the initial evaluation of the parsers, we assess their disambiguation performance in two scenarios: first, we ask the parsers to parse only the ambiguous phrase, inspecting the initial bias the parsers have obtained from their training data. In the next experiment, we allow the underlying language model to contextualize against the disambiguating prior sentence, and we let the parser component then parse the ambiguous phrase to see if it can successfully exploit the LM’s contextualization capabilities.

Table 5 displays the results for the first scenario, where we present the relative clause in two possible orders: one in which the regular order is presented (**S Pron O V**), and one in which subject and object are interchanged (**O Pron S V**). In the case of irreversible triples, this means that in the regular order we expect a high-performing parser to always assign the subject-relative interpretation, but for the reversed order we expect the object-relative interpretation. For the reversible cases, we expect a 50/50 accuracy for both presentations if there is a weak preference for either order, and a skew towards the subject-relative interpretation in the case of strong lexical preference.

The result displays a clear preference for the subject-relative interpretation. In the case of irreversible triples the parsers both pick up on the fact that presented a reversed order must obtain the object-relative interpretation. On the other hand, the results for the reversible triples are significantly below expectation, in the sense that regardless of the presented word order, they will almost always assign a subject-relative interpretation ($> 92.03\%$ on the left, $< 2.32\%$ on the right).

These results are somewhat to be expected: the subject-relative reading prevails in the training data that the parsers were trained on: a total of 306 cases of the subject-relative interpretation occur in Lassy Small, versus 32 cases of the object-relative interpretation. One could then argue that this is indeed the natural intended interpretation so it should have been picked up by any parser replicating its training

⁴Tagging accuracy: 93.21, frame accuracy: 56.36

⁵The Spacy Dutch UD parser `n1_core_news_lg` reports a UAS score of 87 and a LAS score of 83 (<https://spacy.io/models/nl>)

| <i>Neural Proof Nets</i> | S die O V (<i>subj-rel</i>) | O die S V (<i>obj-rel</i>) |
|---------------------------------|---|--|
| Irreversible | 98.76 | 61.86 |
| Reversible-strong | 95.37 | 2.32 |
| Reversible-weak | 97.64 | 1.63 |
| <i>Universal Dependencies</i> | S die O V (<i>subj-rel</i>) | O die S V (<i>obj-rel</i>) |
| Irreversible | 95.72 | 27.85 |
| Reversible, Strong pref. | 94.88 | 0.65 |
| Reversible, Weak pref. | 92.03 | 0.26 |

Table 5: Accuracy results for three different relative clauses without context. Left: presenting the relative clause in regular word order. Right: presenting the relative clause in reversed order. These results indicate the baseline parsing preference without any disambiguating prior sentence.

data, explaining the results in Table 5.

Contextualization However, if it were the case that the parser can easily exploit the information embedded in the language model, we would expect to see that setting the model such that the prior sentence is indeed attended to by the underlying BERT model, the performance would increase. Table 6 displays the results for this second scenario. Here, by introducing the prior sentence as disambiguating context, the ideal parser scores upward to 100% everywhere, thus indicating it can make the correct parse in context.

| First | SVO | SVO | OVS | OVS |
|---------------|---------------------|--------------------|--------------------|---------------------|
| Second | S die O V | O die S V | S die O V | O die S V |
| Reading | (<i>subj-rel</i>) | (<i>obj-rel</i>) | (<i>obj-rel</i>) | (<i>subj-rel</i>) |
| <i>NPN</i> | | | | |
| Irrev. | 98.88 | 69.77 | N/A | N/A |
| Strong | 98.98 | 10.13 | 1.81 | 93.05 |
| Weak | 98.83 | 4.76 | 2.14 | 97.19 |
| <i>UD</i> | | | | |
| Irrev. | 96.91 | 35.03 | N/A | N/A |
| Strong | 96.61 | 1.29 | 0.51 | 96.06 |
| Weak | 94.77 | 0.46 | 0.45 | 94.30 |

Table 6: Accuracy results for three different relative clauses with context, i.e. the prior sentence is attended to by the language model prior to parsing. These results indicate the effect of contextualizing on parsing disambiguation performance.

We observe a stable accuracy for the cases of

subject-relative readings, with significant accuracy gains for the object-relative reading that is less persistent in the original training data. This shows that, to some extent, the added information of the underlying language model gives the parser the incentive to pick up on the grammatical relations in the prior sentence. However, the results are not particularly encouraging: while both parsers do increase in accuracy overall, their strong bias towards a subject-relative interpretation remains.

5.3 Finetuning

After evaluation of the influence of the prior sentence through the contextualization of the BERT embeddings, we additionally finetune the parsers on our task, to see whether the parser could in principle assign the correct interpretation given that the underlying BERT embeddings have access to the prior sentence for contextualization. We explicitly do not update the language model itself, as we want to investigate the parser’s capacity for disambiguation, and allowing the language model to update would be too prone to overfitting (Rogers et al., 2020).

The finetuning scenario thus serves as a means to measure the extent to which the parsers’ strong structural bias can be mitigated. Given the fact that the relative clause has an unambiguous interpretation in the contextualized scenario, training is straightforward, and the results reflect whether the parser by itself can pick up on the contextualized lexical information provided by the language model. These results are displayed in Table 7.

| First | SVO | SVO | OVS | OVS |
|---------------|---------------------|--------------------|--------------------|---------------------|
| Second | S die O V | O die S V | S die O V | O die S V |
| Reading | (<i>subj-rel</i>) | (<i>obj-rel</i>) | (<i>obj-rel</i>) | (<i>subj-rel</i>) |
| <i>NPN</i> | | | | |
| Irrev. | 89.98 | 91.67 | N/A | N/A |
| Strong | 65.64 | 74.45 | 40.02 | 30.98 |
| Weak | 61.31 | 64.43 | 46.59 | 45.55 |
| <i>UD</i> | | | | |
| Irrev. | 89.95 | 71.89 | N/A | N/A |
| Strong | 72.92 | 24.43 | 19.71 | 65.56 |
| Weak | 67.95 | 17.63 | 14.43 | 64.52 |

Table 7: Accuracy results for three different relative clauses with context, i.e. the prior sentence is attended to by the language model prior to parsing, after finetuning different parts of the parser model on the task itself.

In these results we observe that it is indeed pos-

sible to leverage the training task to even out the parsers’ bias through grounding, leaving the neural proof net parser accuracy evenly distributed over cases of subject-relative and object-relative interpretations. On the other hand the UD parser does not adapt to the task that well and retains the strong bias toward a subject-relative interpretation. Overall we observe that the price one pays for the increased accuracy in cases of the object-relative interpretation, is a significant drop of performance in the subject-relative case, showing that developing a balanced parser is no easy task.

6 Discussion

In the setup of our experiments, we were careful to develop test cases that target the parsers in a few different settings. Rather than expecting high performance overall, the aim of the experiment is to both measure the prevalence of structural bias in the parser, as well as measuring to what extent such bias can be mitigated, if present. As such, the experimental results shouldn’t be taken as proof that the parsers in themselves are necessarily insufficient.

Rather, it should be taken as a point to argue that parsers should generally take ambiguity into account, and while lexical ambiguity can be addressed by means of a neural model iterating over a balanced training dataset, the inability to accommodate syntactic ambiguity, both in the training corpora used, as well as in the parser architectures involved, poses a problem that our experiments confirm.

Selectional preferences Readers may be critical about the choice of a ternary definition of *reversibility* since selectional preferences are considered to be graded, expressed in terms of degree of plausibility of a given combination of verbs and nouns. Indeed, the selectional preference dataset of (Zhang et al., 2019) considers degrees of plausibility. It is important to note that such a modelling is incompatible with our current experimental setup, and exactly for the reason we outline above, that the parsers by default assume a single output parse and the disambiguation experiment intends to pinpoint exactly one such parse. Allowing a parser to express a probability distribution over parse trees would allow one to more closely match the way selectional preferences are modelled, but that is (unfortunately) outside of the scope of this work.

Finetuning Aside from the methodological viewpoint above, one may argue that we need not care about the presented experiment as we could simply finetune underlying language model together with the parser on top and achieve high performance on the experiments. While it is true that finetuning the BERT model alongside the parser leads to peak performance, this is most likely due to the language model picking up on positional information quickly, and not due to the model becoming a better parser. To back this claim, we measure parser accuracy metrics on the original training corpus, for finetuned models that only adapt the parser or include the underlying language model. The metrics are displayed in Table 8.

| | NPN | | UD | |
|--------------------|-------|-------|-------|-------|
| | Tag | Frame | UAS | LAS |
| Parser only | 72.54 | 10.32 | 87.41 | 85.92 |
| LM+Parser | 73.13 | 9.77 | 80.39 | 78.58 |

Table 8: Parser accuracies on the test set of Lassy Small, after finetuning on the disambiguation experiment. The top row gives the results for the actual finetuned models that we evaluate, where the bottom row indicates accuracies for models where the language model is included in the finetuning process.

For the neural proof net parser, we observe that performance drastically declines in both cases, with larger decline for the case where the language model’s parameter were included in the finetuning process. The UD parser on the other hand does not decrease in performance so much, but has a strong decline once the language model is included. This highlights two points: first, given that the neural proof net parser adapted itself better to the task of our main experiment, we conclude that the price to pay for task adaptation is a decline in overall parser performance. Second, we argue that including the language model in finetuning leads to overfitting on the task, and reduces the parser’s overall accuracy.

7 Conclusion

By introducing a synthetic test set of naturalistic Dutch relative clauses, we carried out an experiment to investigate the sensitivity to structural bias of two parsing architectures that are both based on a BERT-style language model. The experiments show that both parsers pick up on a structural preference for a subject-relative reading of the relative clause, following a strong statistical bias coming

from the data they were trained on. Further experimentation shows that a more complex neurosymbolic parsing regime adapts more easily to a bias correcting finetuning setup than a universal dependencies parser implemented as a sequence labelling model. However, in both cases performance on the task is severely below expectation, and we hope that this work inspires further work on careful data augmentation and parser development.

8 Limitations

A main limitation of the present study is that we did not find a satisfying way to mitigate the bias of the parsers, leaving this as an open problem for future work. Additionally, it would be interesting to study the performance of language models that already encode syntactic cues as part of their pretraining; such a model for Dutch has been developed by Tziafas et al. (2023) with the goal to reduce training data and model size of a BERT model.

It could be viewed as a limitation the the current study investigates a phenomenon typical for Dutch; reproducing the exact same study in English is not possible, as the English relative clause is disambiguated on its own accord due to lexical or word order cues. However, further investigation in a multilingual setting could identify similar structural constructions that allow for further experimentation on a larger scale.

Acknowledgments

Both authors gratefully acknowledge support from the Dutch Research Council (NWO) under the scope of the project “A composition calculus for vector-based semantic modelling with a localization for Dutch” (360-89-070).

References

- Laura Aina and Tal Linzen. 2021. [The language model understood the prompt was ambiguous: Probing syntactic uncertainty through generation](#). In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 42–57, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Liesbeth Allein, Artuur Leeuwenberg, and Marie-Francine Moens. 2020. [Binary and multitask classification model for Dutch anaphora resolution: Die/dat prediction](#). *arXiv preprint arXiv:2001.02943*.
- Suhas Arehalli, Brian Dillon, and Tal Linzen. 2022. [Syntactic surprisal from neural models predicts, but](#)

- underestimates, human processing difficulty from syntactic ambiguities. In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 301–313, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Gosse Bouma. 2021. Probing for dutch relative pronoun choice. *Computational Linguistics in the Netherlands Journal*, 11:59–70.
- Stephen Clark. 2021. Something old, something new: Grammar-based CCG parsing with transformer models. *CoRR*, abs/2109.10044.
- Tim Van de Cruys. 2014. A neural network approach to selectional preference acquisition. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 26–35, Doha, Qatar. Association for Computational Linguistics.
- Pieter Delobelle, Thomas Winters, and Bettina Berendt. 2020. RobBERT: a Dutch RoBERTa-based Language Model. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3255–3265, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. 2019. Neural language models as psycholinguistic subjects: Representations of syntactic state. 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 32–42, Minneapolis, Minnesota. Association for Computational Linguistics.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational linguistics*, 28(3):245–288.
- Tessel Haagen, Loïs Dona, Sarah Bosscha, Beatriz Zamith, Richard Koetschryter, and Gijs Wijnholds. 2022. Noun phrase and verb phrase ellipsis in dutch: Identifying subject-verb dependencies with bertje. *Computational Linguistics in the Netherlands Journal*, 12:49–63.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. A systematic assessment of syntactic generalization in neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Jerrold J Katz and Jerry A Fodor. 1963. The structure of a semantic theory. *language*, 39(2):170–210.
- Konstantinos Kogkalidis, Michael Moortgat, and Richard Moot. 2020. Neural proof nets. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 26–40, Online. Association for Computational Linguistics.
- Konstantinos Kogkalidis, Michael Moortgat, and Richard Moot. 2023. SPINDLE: Spinning raw text into lambda terms with graph attention. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, Dubrovnik, Croatia. Association for Computational Linguistics.
- Konstantinos Kogkalidis and Gijs Wijnholds. 2022. Discontinuous constituency and BERT: A case study of Dutch. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3776–3785, Dublin, Ireland. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Diana McCarthy and John Carroll. 2003. Disambiguating nouns, verbs, and adjectives using automatically acquired selectional preferences. *Computational Linguistics*, 29(4):639–654.
- Ken McRae, Michael J Spivey-Knowlton, and Michael K Tanenhaus. 1998. Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, 38(3):283–312.
- Eleni Metheniti, Tim Van de Cruys, and Nabil Hathout. 2020. How relevant are selectional preferences for transformer-based language models? In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1266–1278, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Michael Moortgat, Konstantinos Kogkalidis, and Gijs Wijnholds. 2023. *Diamonds Are Forever: Theoretical and Empirical Support for a Dependency-Enhanced Type Logic*, pages 57–87. *Logic and Algorithms in Computational Linguistics 2021 (LACompLing2021)* Springer International Publishing, Cham.

- Michael Moortgat, Mehrnoosh Sadrzadeh, and Gijs Wijnholds. 2020. [A Frobenius Algebraic Analysis for Parasitic Gaps](#). *Journal of Applied Logics*, 2631(5):823.
- Michael Moortgat and Gijs Wijnholds. 2017. [Lexical and derivational meaning in vector-based models of relativisation](#). In *Proceedings of the 21st Amsterdam Colloquium*, page 55. ILLC, University of Amsterdam.
- Mughilan Muthupari, Samrat Halder, Asad Sayeed, and Yuval Marton. 2022. [Where’s the learning in representation learning for compositional semantics and the case of thematic fit](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 28–39, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Gertjan van Noord. 2010. Self-trained bilexical preferences to improve disambiguation accuracy. *Trends in Parsing Technology: Dependency Parsing, Domain Adaptation, and Deep Parsing*, pages 183–200.
- Gertjan van Noord, Gosse Bouma, Frank Van Eynde, Daniël de Kok, Jelmer van der Linde, Ineke Schuurman, Erik Tjong Kim Sang, and Vincent Vandeghinste. 2013. [Large Scale Syntactic Annotation of Written Dutch: Lassy](#), pages 147–164. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Ulrike Padó. 2007. The integration of syntax and semantic plausibility in a wide-coverage model of human sentence processing.
- Philip Resnik. 1997. [Selectional preference and sense disambiguation](#). In *Tagging Text with Lexical Semantics: Why, What, and How?*
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Arabella Sinclair, Jaap Jumelet, Willem Zuidema, and Raquel Fernández. 2022. [Structural Persistence in Language Models: Priming as a Window into Abstract Language Representations](#). *Transactions of the Association for Computational Linguistics*, 10:1031–1050.
- Michalina Strzyz, David Vilares, and Carlos Gómez-Rodríguez. 2019. [Viable dependency parsing as sequence labeling](#). 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 717–723, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Georgios Tziafas, Konstantinos Kogkalidis, Gijs Wijnholds, and Michael Moortgat. 2023. [Improving BERT pretraining with syntactic supervision](#). In *Proceedings of the 2023 CLASP Conference on Learning with Small Data (LSD)*, pages 176–184, Gothenburg, Sweden. Association for Computational Linguistics.
- David Vilares, Michalina Strzyz, Anders Søgaard, and Carlos Gómez-Rodríguez. 2020. [Parsing as pretraining](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9114–9121.
- Wietse de Vries, Andreas van Cranenburgh, Arianna Bisazza, Tommaso Caselli, Gertjan van Noord, and Malvina Nissim. 2019. [BERTje: A Dutch BERT model](#). *arXiv preprint arXiv:1912.09582*.
- Gijs Wijnholds. 2023. [Assessing monotonicity reasoning in Dutch through natural language inference](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1494–1500, Dubrovnik, Croatia. Association for Computational Linguistics.
- Gijs Wijnholds and Michael Moortgat. 2021. [SICK-NL: A dataset for Dutch natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1474–1479, Online. Association for Computational Linguistics.
- Gijs Wijnholds, Mehrnoosh Sadrzadeh, and Stephen Clark. 2020. [Representation learning for type-driven composition](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 313–324, Online. Association for Computational Linguistics.
- Hongming Zhang, Hantian Ding, and Yangqiu Song. 2019. [SP-10K: A large-scale evaluation set for selectional preference acquisition](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 722–731, Florence, Italy. Association for Computational Linguistics.