

A Logic-Based Framework for Natural Language Inference in Dutch

Lasha Abzianidze
Konstantinos Kogkalidis

L.ABZIANIDZE@UU.NL
K.KOGKALIDIS@UU.NL

UiL OTS, Utrecht University, Utrecht, the Netherlands

Abstract

We present a framework for deriving inference relations between Dutch sentence pairs. The proposed framework relies on logic-based reasoning to produce inspectable proofs leading up to inference labels; its judgements are therefore transparent and formally verifiable. At its core, the system is powered by two λ -calculi, used as syntactic and semantic theories, respectively. Sentences are first converted to *syntactic* proofs and terms of the linear λ -calculus using a choice of two parsers: an Alpino-based pipeline, or Neural Proof Nets. The syntactic terms are then converted to *semantic* terms of the simply typed λ -calculus, via a set of hand designed type- and term-level transformations. Pairs of semantic terms are then fed to an automated theorem prover for natural logic which reasons with them while using the lexical relations found in the Open Dutch WordNet. We evaluate the reasoning pipeline on the recently created Dutch natural language inference dataset, and achieve promising results, remaining only within a 1.1–3.2% performance margin to strong neural baselines. To the best of our knowledge, the reasoning pipeline is the first logic-based system for Dutch. The code is available at git.io/JzdGd.

1. Introduction

Among the many Natural Language Understanding tasks, Natural Language Inference (NLI) is of particular interest. An NLI task can be broadly summarised as follows: given two natural language utterances, a *premise* and a *hypothesis*, decide whether the former entails, contradicts, or is neutral with respect to the latter. An NLI system requires the capacity for manipulating syntactic structure as well as lexical meaning, necessitating a holistic approach to yield meaningful results.

In recent years, the advent of neural models has set new benchmarks for NLI tasks, with general-purpose language models based on the pre-train and fine-tune paradigm claiming the lion's share in the literature. Despite their indisputable performance, such models suffer from a variety of downsides, opaqueness and unpredictability being the most striking (Sanchez et al. 2018, Glockner et al. 2018, McCoy et al. 2019). In practical terms, a neural system might achieve high accuracy scores, but provides limited insight on how it arrived at a decision, thus prohibiting manual verification of the inference process and its outcome. At the same time, the high expressiveness of modern neural architectures makes them prone to detect and capitalise on subtle statistical patterns and annotation artifacts common in popular NLI datasets (Gururangan et al. 2018, Poliak et al. 2018, Tsuchiya 2018), artificially inflating their performance within the evaluation domain, but failing to generalise on out-of-distribution inputs. These issues are further pronounced in use cases where reliability, robustness, and interpretability are of major importance.

In stark contrast to neural models, logic-based methods for NLI boast transparency, reliability and formal rigor, often at the cost of a drop in performance. Logic-based systems offer not just a prediction, but rather the full explanation behind it, allowing a deep inspection of their inner workings that extends beyond mere quantitative comparisons. They are also much more reliable in their non-neutral (i.e. entailment or contradiction) predictions than neural NLI systems. Witness to this is the fact that disagreements between a logic-based model and a dataset's ground truth annotations are often due to noise or errors in the latter.

In this work, we utilise the SICK dataset (Marelli et al. 2014b) in its recent Dutch translation (Wijnholds and Moortgat 2021) as an experimental testbed for the first cross-lingual application of LangPro (Abzianidze 2017), a Natural Tableau-based theorem prover. Our inputs to the prover are semantic expressions in higher-order logic based on simple type theory. To obtain semantic expressions, we transform typological grammar derivations procured from two wide-coverage parsers: a pipeline based on Alpino (Bouma et al. 2001), and Neural Proof Nets (Kogkalidis et al. 2020a). Following careful tuning, the prover employs the above expressions, in combination with the lexical semantic relations found in the Open Dutch WordNet database (Postma et al. 2016), to generate logical inferences for the dataset, achieving a final accuracy of about 79%.

The structure of the paper is as follows. In §2 we initiate the unfamiliar reader to the formal systems we are utilising, namely typological grammars, λ -calculi and natural tableaux. In §3 we move to a more practical territory, describing the tools and processes we use and detailing each of our framework’s components. We then describe our experiments and expose our results on SICK-NL in §4, and make our assessments based on extensive qualitative error analysis in §5. We draw comparisons to related work in §6. The last section is reserved for some conclusive remarks and suggestions for future directions.

2. Background

In this section, we provide a brief expository note to the formal systems we employ at each step of our inference pipeline.

We begin by detailing our syntactic framework of choice in §2.1. Motivated by the need for a transparent syntax-semantics interface, we employ a semantically-g geared typological grammar (Morrill 2012, Moot and Retoré 2012). Typological grammars are rooted in the logical tradition of formal linguistics. One of their biggest appeals is their affinity to semantic expressions due to the propositions-as-types interpretation that equates proofs with programs and propositions with functional types. As such, they make for an ideal candidate in the application envisaged here.

We then describe the logic used for linguistic semantics in §2.2. The use of formal logic to model natural language inference is common practice in formal semantics. The usual suspect is first-order predicate logic, due to its well-behavedness, on the one hand, and the wide accessibility of out-of-the-box automated theorem provers and model builders, on the other (Blackburn and Bos 2005, Blackburn et al. 2001). Despite its attractiveness at a first glimpse, capturing meaning with first-order logic formulas can be notoriously difficult (common pain points include, among others, the representation of phenomena involving subsecutive adjectives and generalised quantifiers). We therefore opt for a higher-order logic in the form of a simply-typed λ -calculus. Aside from its expressive power, it boasts a clear syntax that resembles linguistic expressions; as such, its terms are easy to obtain from parse structures.

We conclude the section by describing the procedure of reasoning with natural language sentences in §2.3. Our inferential engine is powered by Natural Tableau (Muskens 2010, Abzianidze 2016), as inspired by the Natural Logic project (van Benthem 2008, Moss 2010), a study of reasoning with meaning representations close to natural language, and the semantic tableau method, one of the most popular proof procedures for formal logics (D’Agostino et al. 1999).

2.1 Syntax

Typological grammars use a *lexicon* to assign *types* to words and model parse structures as logical *proofs* built with the aid of a small set of inference rules. Our typological grammar’s logical backbone is the implication-only fragment of Intuitionistic Linear Logic (Wadler 1993). Its types \mathcal{T} are inductively defined as a set closed under a single binary operator, i.e. they form a magma (\mathcal{A}, \multimap) , where:

- $\mathcal{A} \subset \mathcal{T}$ a finite set of basic or *atomic* types

- \multimap the linear implication (or lollipop), an operator such that $\tau_1, \tau_2 \in \mathcal{T} \iff \tau_1 \multimap \tau_2 \in \mathcal{T}$

A type of the form $\tau_1 \multimap \tau_2$ (shorthand: $\tau_1(\tau_2)$) is called *complex*, and is used to denote a linear transformation that will consume an *argument* of type τ_1 to produce a *result* of type τ_2 . In the linguistic setup, words that can stand on their own are assigned atomic types, whereas words requiring complements are assigned complex types; a simplified but representative lexicon is depicted in Table 1.

ganzen, eenden, bessen	::	np	je, me	::	$pron$
zwemmen	::	$np(s)$	gaf	::	$pron(np(pron(s)))$
eten	::	$np(np(s))$	die	::	$(pron(s))(np(np))$
rode, blauwe	::	$np(np)$	en	::	$\forall \alpha : \alpha(\alpha)$

Table 1: A toy lexicon of simple linear types.

In its simplest form, the type logic provides three rules of inference, through which complex expressions can be built from simple ones. Owing to the remarkable equivalence between logics and λ -calculi known as the Curry-Howard correspondence (Sørensen and Urzyczyn 2006), each logical rule has an analogue in the *term* language of the linear λ -calculus:

$$\frac{\Gamma \vdash s^{\tau_1(\tau_2)} \quad \Delta \vdash t^{\tau_1}}{\Gamma, \Delta \vdash (s \ t)^{\tau_2}} \multimap E \quad \frac{}{x^\tau \vdash x^\tau} Ax \quad \frac{\Gamma, x^{\tau_1} \vdash s^{\tau_2}}{\Gamma \vdash (\lambda x.s)^{\tau_1(\tau_2)}} \multimap I$$

The implication elimination rule ($\multimap E$) posits that, given the derivability of an expression s of type $\tau_1(\tau_2)$ from some context Γ and the derivability of an expression t of type τ_1 from some context Δ , from the two contexts together we can derive the term $s \ t$ of type τ_2 , corresponding to the application of s to t . Together with our toy lexicon, this rule already suffices to derive terms for a few simple sentences:

- (1) a. eenden zwemmen
($zwemmen^{np(s)} \ eenden^{np}$)^s
- b. eenden eten rode bessen
($eten^{np(np(s))} \ (rode^{np(np)} \ bessen^{np})^{np} \ eenden^{np}$)^s
- c. je gaf me bessen
($gaf^{pron(np(pron(s)))} \ me^{pron} \ bessen^{np} \ je^{pron}$)^s

The next two rules are a crucial component of the type logic, giving us access to hypothetical reasoning, a tool required for the derivation of higher-order syntactic phenomena. The identity axiom (Ax) allows us to instantiate a fresh named *variable* x of some type τ . Finally, if given a context Γ and a variable x of type τ_1 we can derive a term s of type τ_2 , the implication introduction rule ($\multimap I$) allows us to build a function $\lambda x.s$ of type $\tau_1(\tau_2)$ from context Γ alone. With the addition of the above rules, we can now derive terms for more complicated sentences:

- (2) bessen die je me gaf
($die^{(np(s))(np(np))} \ (\lambda x. (gaf^{pron(np(pron(s)))} \ me^{pron} \ x^{np} \ je^{pron})^s)^{np(s)} \ bessen^{np}$)^{np}

The last item in our syntactic toolshed is a hint of type polymorphism, a telling example being the type $\forall \alpha : \alpha(\alpha)$ assigned to coordinators, where α is a *variable* ranging over types. The above recipe gives us the means to derive conjunctions of different syntactic categories in a uniform way:

- (3) a. ganzen en eenden
($en^{\alpha(\alpha)} \ ganzen^{np} \ eenden^{np}$)^{np}

- b. eenden eten rode en ganzen blauwe bessen

$$\left(\text{en}^{\alpha(\alpha)} \quad (\lambda xy. (x^{np(np(s))} (\text{rode}^{np(np)} y^{np}) \text{eenden}^{np})^s)^{(np(np(s))(np(s))} \right.$$

$$\left. (\lambda zw. (z^{np(np(s))} (\text{blauwe}^{np(np)} w^{np}) \text{ganzen}^{np})^s)^{(np(np(s))(np(s))} \right.$$

$$\left. \text{eten}^{np(np(s))} \text{bessen}^{np} \right)^s$$

where in Example (3-a) we set $\alpha := np$ for a simple noun phrase conjunction, whereas in (3-b) we set $\alpha := (np(np(s))(np(s)))$ for a conjunction of sentences sharing their phrasal head, but also the head of the object noun phrase. This last example is indicative of the treatment of elliptical constructions in a linear regime.

2.2 Semantics

The semantic logic we opt for is a higher-order simple predicate logic corresponding to the implication-only fragment of Intuitionistic Logic, or, in Curry-Howard terms, the simply typed λ -calculus. It is a close replica of our syntactic logic, modulo implication no longer being linear: complex types are now the type signatures of ordinary functions. In practical terms, the type forming operator is now an arrow \rightarrow rather than the lollipop \multimap , and a single rule of inference is added to our vocabulary:

$$\frac{\Gamma, \mathbf{x}_1^{\tau_1}, \mathbf{x}_2^{\tau_1} \vdash \mathbf{s}^{\tau_2}}{\Gamma, \mathbf{x}^{\tau_1} \vdash \mathbf{s}[\mathbf{x}/\mathbf{x}_1, \mathbf{x}/\mathbf{x}_2]^{\tau_2}} \textit{Contraction}$$

It suggests that if from some context Γ together with two distinct variables $\mathbf{x}_1, \mathbf{x}_2$ of the same type τ_1 we can derive a term \mathbf{s} of type τ_2 , then we can do the same with just a single instance of \mathbf{x} , provided we replace all occurrences of \mathbf{x}_1 and \mathbf{x}_2 in \mathbf{s} with \mathbf{x} . The effect of a non-linear semantic logic is that our semantic terms may now contain more than a single occurrence of terms appearing just once in the corresponding syntactic term, essentially permitting duplication of words when necessary (the utility of this will become evident in Example (5-e) later on).

Just like syntactic terms, semantic terms (otherwise called Lambda Logical Forms, or LLFs) are built up from variables and constant lexical items. LLFs are typed using a small set of atoms: np , n , pp , pr , and s_x , corresponding to noun phrase, common noun, prepositional phrase, particle, and sentence respectively.¹ A few sample LLFs are depicted in Example (4); to distinguish between syntactic and semantic expressions, we format lexical terms of the latter as boldface. For types we use the abbreviation $vp := np(s)$.

- (4) a. A woman who loves John is very happy
 $\mathbf{a}^{n(vp(s))} (\mathbf{who}^{vp(n(n))} (\mathbf{love}^{np(vp)} \mathbf{john}^{np}) \mathbf{woman}^n) (\mathbf{be}^{vp_{adj}(vp_{dcl})} (\mathbf{very}^{vp_{adj}(vp_{adj})} \mathbf{happy}^{vp_{adj}}))$
- b. Every man loves a woman
 $\mathbf{every}^{n(vp(s))} \mathbf{man}^n (\lambda x. \mathbf{a}^{n(vp(s))} \mathbf{woman}^n (\mathbf{love}^{np(vp)} x^{np}))$ Object narrow scope
 $\mathbf{a}^{n(vp(s))} \mathbf{woman}^n (\lambda y. \mathbf{every}^{n(vp(s))} \mathbf{man}^n (\lambda x. \mathbf{love}^{np(vp)} x^{np} y^{np}))$ Object wide scope

2.3 Reasoning with Natural Tableau

Natural Tableau is a signed tableau method specifically designed for a version of Natural Logic. Its core component is a set of inference rules, called tableau rules. During the reasoning process, these rules gradually break down the input logical forms, and different facets of the meaning are fleshed out. To avoid overloaded LLFs, we will omit types of lexical terms when appearing in tableau proofs.

1. In addition to the predicative adjectives feature (adj), a sentential clause category can be subcategorised as declarative (dcl), active past participial (pt), passive past participial (ps), present participial (ng), or a question (e.g. q or wh). These sentence category features are inspired by the Combinatory Categorical Grammar (CCG) treebank (Hockenmaier and Steedman 2007). Note that the category feature adj in CCG serves to prevent certain ungrammatical declarative sentences, e.g. “John happy” and “John very love a woman”.

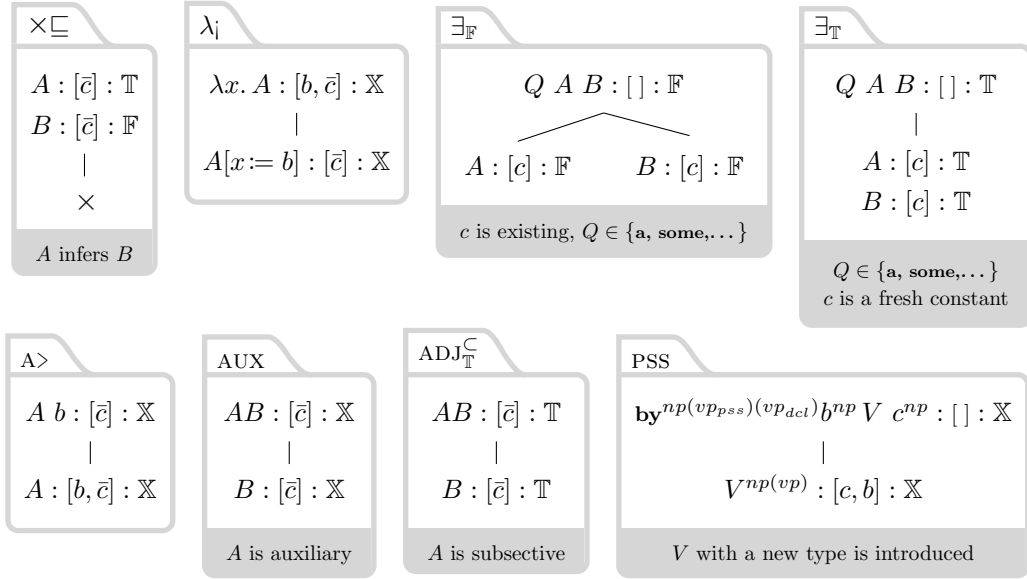


Figure 1: The (tableau) inference rules that are employed in the tableau proof of Figure 2. Each inference rule has its name and optional constraints that are explicitly stated below the rule. \bar{c} denotes a (possibly empty) list of terms. X is a variable over the truth signs T and F .

Figure 1 shows some of the tableau rules. Each rule has antecedent and consequent entries, where a tableau entry is a triplet of a λ -term, its (possibly empty) list of arguments, and a truth sign. For example, $A : [\bar{c}] : T$ means that when A is applied to its arguments (respecting the argument order), the resulting term is rendered as true. Consequent entries of a rule are usually shorter than the antecedent ones, which decomposes initial terms into smaller pieces. One special type of rule is a closure rule, e.g. $(\times \sqsubseteq)$. The rule spots inconsistencies (like an entity being A and not B , but at the same time A being more specific than B) and triggers the termination of the search. The best way to understand the rules and see how they work in tandem is to consider an actual tableau proof.

We illustrate a tableau proof in the style of Natural Tableau in Figure 2.² The tree-style proof is built to refute that the premise “a harmonica is played by a young boy” entails the hypothesis “a person sounds a musical instrument”. If the refutation fails, this serves as a proof for the entailment relation. The refutation is carried out by searching for a counterexample for the entailment relation, i.e. building a situation that makes the premise true and the hypothesis false. The tableau starts with this exact requirement: the LLFs of the premise and the hypothesis are set to be true and false in nodes ① and ②, respectively. The rest of the tableau is built by decomposing the semantics of the initial entries with the help of the inference rules. For example, (\exists_T) applies to ① and produces ③ and ④. The latter produces ⑤ via (λ_i) , and so forth. In the end, ① is decomposed as there are b “boy” (⑪) and h “harmonica” (③), and b “plays” h (⑩). This information is inconsistent with the semantics of ②, and it is expressed in terms of three closed branches closed due to the following inconsistencies: b not being “person” (⑫), h not being “musical instrument” (⑮), and b not “sounding” h (⑰). A tableau with its all branches closed means that the refutation failed, i.e. it was impossible to find a counterexample for the entailment relation. Therefore, the premise entails the hypothesis; Q.E.D.

2. The example is in English as Natural Tableau and its computational implementation, the LangPro theorem prover, were originally developed for reasoning with English sentences. Also following the LLF format in LangPro, lexical terms are represented as lemmas. Additional information about types and part-of-speech (pos) tags of the lexical terms are omitted for the sake of simplicity.

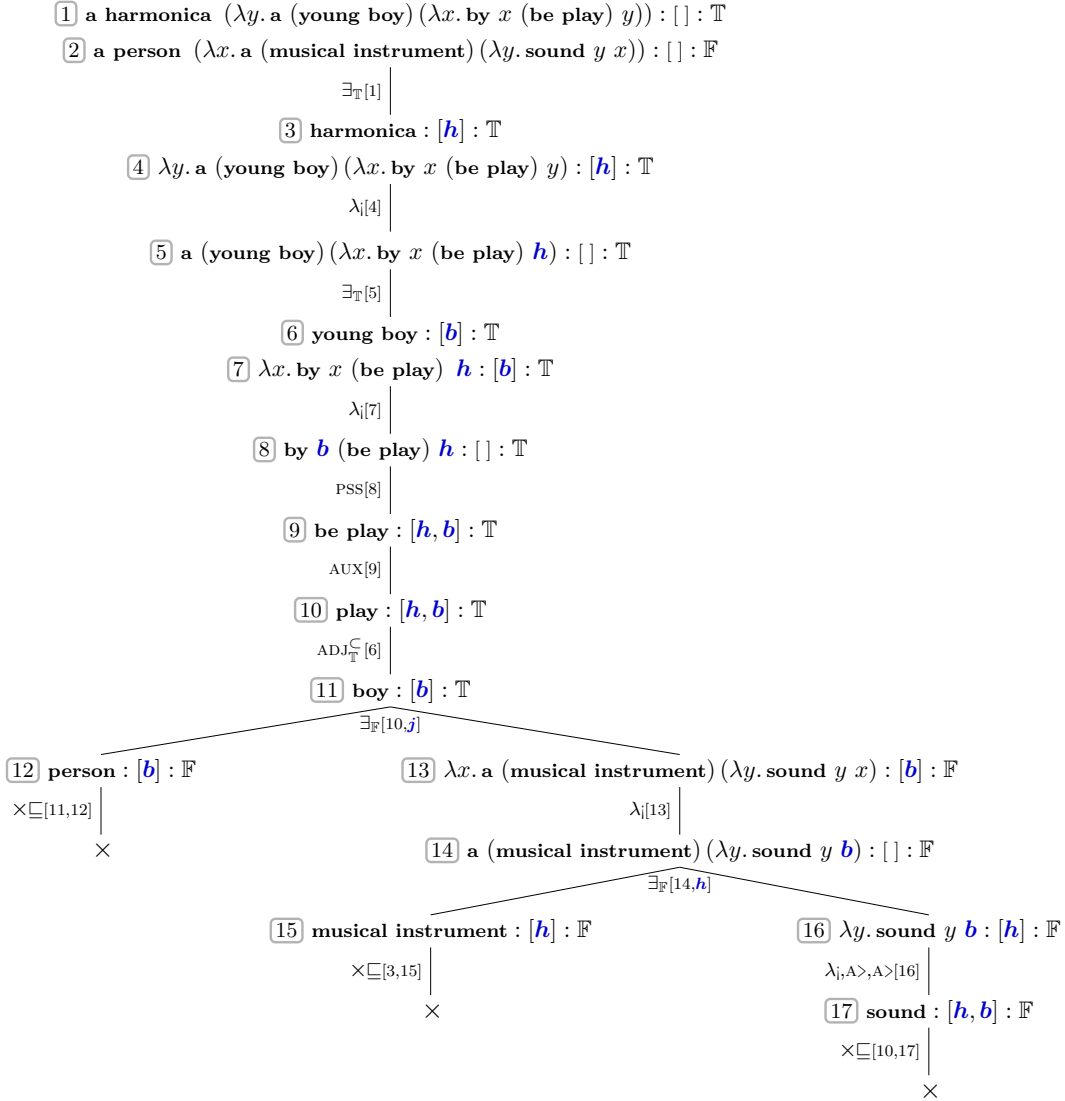


Figure 2: The closed tableau proves that the premise “a harmonica is played by a young boy” entails the hypothesis “a person sounds a musical instrument”.

The Natural Tableau method can also be used to learn from data via abductive reasoning – inference to the best explanation. For example, let’s assume that “musical instrument” is replaced with “French harp” in the hypothesis of the example in Figure 2. The proof for entailment would fail in case there is no knowledge available saying that a “harmonica” is a “French harp” (e.g. such knowledge is not available in WordNet). The idea behind abductive learning is that, given the correct/gold relation of such a problem (e.g. entailment), a tableau that attempts to prove the relation is constructed. If the proof is not found (i.e. the tableau is not closed), then the search starts for such knowledge that helps to close the tableau. In other words, abductive reasoning is used to infer the knowledge that supports the correct relation. In our example, such inferred knowledge would be “harmonica” being a sort of “French harp” ($\text{harmonica} \sqsubseteq \text{french harp}$), which suffices to find the proof for the entailment relation. This way, abductive learning helps infer new knowledge from labeled NLI problems, which can later be used for unseen problems.

3. Methodology

Having introduced the formal background, we now move on to describe our implementation of the automated theorem prover for Dutch and the computational machinery behind its components.

3.1 Parsing

We obtain syntactic λ -terms in the form described in §2.1 using two different parser pipelines: one based on Alpino, and another on Neural Proof Nets.³

Alpino is an attested wide-coverage parser for unrestricted written Dutch, creating parse structures in the form of dependency graphs (Bouma et al. 2001, van Noord 2006). Its grammar, in line with head-driven phrase structure grammars (Pollard and Sag 1994), consists of a set of manually specified phrase formation rules and a rich lexicon providing subcategorisation frames and dependency information. The format employed (graphs, rather than trees) provides the means to capture reentrancy; nodes correspond to words and phrases (labeled with syntactic category tags), and outgoing edges denote a dependency frame containing strictly one phrasal head, and zero or more complements and adjuncts. Due to the ambiguity inherent in frame assignment and rule application, Alpino may produce a multitude of dependency graphs for a single input sentence; these are evaluated and scored on the basis of a log-linear disambiguation model aided by a few hand-designed penalisation rules. We convert Alpino’s graphs to proofs of the syntactic type logic using the type extraction algorithm of Kogkalidis et al. (2020b); the algorithm traverses the parse graph, translating syntactic categories to atomic types and iteratively casting heads (resp. adjuncts) as linear functions that consume their complements (resp. phrasal parents) for each phrase.

Neural Proof Nets (NPN) is a formalism-specific neuro-symbolic parser composed of three parts (Kogkalidis et al. 2020a). First, a pretrained BERT model reads the tokenised input text and builds contextualised vectorial representations for each token (Delobelle et al. 2020). An autoregressive transformer stack then translates the encoded representations into a sequence of types aligned with each input work, handling lexical type assignment and disambiguation in context (Kogkalidis et al. 2019). Finally, a permutation module based on Sinkhorn networks (Mena et al. 2018) uses the format of multiplicative linear logic proof nets (Girard 1987) to tackle rule applications in parallel, transforming the type sequence into a proof proper.

3.2 Obtaining Lambda Logical Forms

Despite also being λ -expressions, LLFs differ from syntactic terms not only in their intended use, but also in their structure and types. Whereas syntactic terms capture the tectogrammatical structure underlying the sentence in a bottom-up fashion, LLFs are used to express the sentential meaning, and are processed top-down by the Natural Tableau inference rules. The conversion from the former to the latter is handled by a manually-designed pipeline, generally following (Abzianidze 2015b). The pipeline is depicted in Figure 3; it gradually simplifies syntactic expressions, homogenising parser inconsistencies and syntactic subtleties that would otherwise be attenuated in the output LLFs.

Simplifying lexical entries The syntactic terms come with fine-grained types built up from 31 atoms. On the other hand, Natural Tableau operates on LLFs typed using the following atoms: $\{n, np, s_x, pp, pr\}$ as detailed in §2.3. We translate syntactic terms into simple semantic terms using a many-to-one map from syntactic *atoms* to simplified *types*, as depicted in Table 2. This serves two functions: it first collates the (quite large) set of syntactic primitives to a more manageable size, but also casts syntactic atoms that hide their semantic frames into explicit functions.

3. Both pipelines generate proofs and terms enhanced with unary type- and term-level operators that specify dependency information on top of function-argument structures; we discard the dependency information while the structure of proofs and terms remains the same. We leave utilisation of the dependency information in the context of semantic reasoning as an open question for future work.




Figure 3: A procedure of obtaining LLFs from the syntactic terms of the parsers.

Syntactic Atoms (description & sign)	LLF Types	
Declarative sentence (verb at the 2 nd position)	s_{main}	s_{dcl}
Subordinate clause (verb final)	s_{sub}	s_{sub}
Pronoun	vnw	np
Preposition	vz	pr
Numeral	tw	np
aan het-infinitive group	ahi	$np(s_{ng})$
Verb	ww	$np(s_b)$
Passive/perfect participle	$part$	$np(s_{pt})$
te-infinitive group	ti	$np(s_{to})$
om te-infinitive-group	oti	
Adjectival Phrase	ap	$np(s_{adj})$
Adjective	adj	

Table 2: Mapping from typelological atoms to LLF types.

Syntactic constants for lexical entries are translated to semantic constants that respect the translation of their type signature. Example (5-a) shows the simple semantic term resulting from casting each of the syntactically flat atoms ti , $part$ and ww to the function type $vp := np(s)$.⁴

Fixing analyses Adjustment and correction of simple semantic terms is the most elaborate part of the conversion procedure. Since there is no clear cut between adjustments and corrections due to the differences in two styles of analysis motivated from top-down and bottom-up approaches, we will not distinguish them and call it fixing. During term fixing, pos tags of the lexical entries are used in addition to type information. Lexical terms are replaced with their lemmas (formatted in boldface). Below we illustrate representative instances of the term fixes for Dutch syntactic terms.⁵

Moving a determiner in a term structure above noun modifiers is the most applied rule. The instance of the rule application is shown in Example (5-b). Other fixing rules related to NPs are the rules that change the type of a verb or a preposition term that takes an argument of type n . Example (5-c) shows the changes in the types of “zijn” and “snijden” and the insertion of explicit quantifiers for existence and plurality for the bare NPs “hout” and “mannen” respectively.

Sometimes adjectives that act like nouns are analysed as predicative adjectives of type $np(s_{adj})$, like in (sim) of (5-d). The type of such nominal adjectives are set to n and the type of related lexical terms are changed accordingly. Example (5-d) also shows how the predicative PP “op een berg” is changed from an adjunct phrase to a complement of the copula “zijn”.

(5) a. (nld) om te vissen gebruikt

4. The simplification step also involves mapping the Alpino-style and Universal (Petrov et al. 2012) pos tags (coming from Alpino and spaCy, respectively, see §4.2) to the Penn Treebank-style. The latter is the tagset expected by an existing reasoning component (see §3.3).

5. Some of these fixes are the results of combinations of the already existing fixing rules from Abzianidze (2015b) specific to CCG derivation trees and the fixing rules specific to the structures of the Dutch syntactic terms.

- (syn) $\text{om}^{ti(part(part))} (\text{te}^{ww(ti)} \text{vissen}^{ww}) \text{gebruikt}^{part}$
(sim) $\text{om}^{vp(vp(vp))} (\text{te}^{vp(vp)} \text{vissen}^{vp}) \text{gebruikt}^{vp}$
(fix) $\text{om}^{vp(vp(vp))} (\text{te}^{vp(vp)} \text{vissen}^{vp}) \text{gebruiken}^{vp}$
- b. (nld) een grote bruine hond
(sim) $\text{grote}^{np(np)} (\text{bruine}^{np(np)} (\text{een}^{n(np)} \text{hond}^n))$
(fix) $\text{een}^{n(np)} (\text{groot}^{n(n)} (\text{bruin}^{n(n)} \text{hond}^n))$
- c. (nld) mannen zijn hout aan het snijden
(sim) $\text{zijn}^{vp(n(s))} (\text{aan_het}^{vp(vp)} (\text{snijden}^{n(vp)} \text{hout}^n)) \text{mannen}^n$
(fix) $\text{zijn}^{vp(np(s))} (\text{aan_het}^{vp(vp)} (\text{snijden}^{np(vp)} (\text{een}^{n(np)} \text{hout}^n))) (\text{s}^{n(np)} \text{man}^n)$
(llf) $\text{s}^{n(vp(s))} \text{man} (\text{zijn} (\text{aan_het} (\lambda x. \text{een}^{n(vp(s))} \text{hout} (\lambda y. \text{snijden} y^{np} x^{np}))))$
- d. (nld) een man in het blauw is op een berg
(sim) $\text{op}^{np(s(s))} (\text{een}^{n(np)} \text{berg}^n) (\text{is}^{vp} (\text{in}^{np(np(np))} (\text{het}^{vp(np)} \text{blauw}^{vp}) (\text{een}^{n(np)} \text{man}^n)))$
(fix) $\text{zijn}^{pp(vp)} (\text{op}^{np(pp)} (\text{een}^{n(np)} \text{berg}^n)) (\text{een}^{n(np)} ((\text{in}^{np(n(n))} (\text{het}^{n(np)} \text{blauw}^n)) \text{man}^n))$
(llf) $\text{het}^{n(vp(s))} \text{blauw} (\lambda x. \text{een}^{n(vp(s))} (\text{in} x^{np} \text{man}))$
 $(\lambda y. \text{een}^{n(vp(s))} \text{berg}^n (\lambda z. \text{zijn} (\text{op} z^{np} y^{np})))$
- e. (nld) een rode jas en kaki broek
(sim) $\text{en}^{\alpha(\alpha(\alpha))} (\lambda x. \text{rode}^{np(np)} (x^{n(np)} \text{jas}^n)) (\lambda y. \text{kaki}^{np(np)} (y^{n(np)} \text{broek}^n)) \text{een}^{n(np)}$
(fix) $\text{en}^{np(np(np))} (\text{een}^{n(np)} (\text{rood}^{n(n)} \text{jas}^n)) (\text{een}^{n(np)} (\text{kaki}^{n(n)} \text{broek}^n))$

Elliptical coordination constructions are modeled with syntactic terms containing λ -abstractions, as shown in (sim) of Example (5-e). We apply a non-linear rewriting rule to such constructions that distributes the argument over the coordinated function terms: in the example, “een” is distributed over “rode jas” and “kaki broek”. After the argument distribution, β -reductions are applied and the determiners are moved at the top level of NPs, as done in Example (5-b).

Type-raising NPs The final step in the conversion is to obtain LLFs from fixed terms. This is done by type-raising NPs with determiners/quantifiers. This procedure follows the algorithm described in Abzianidze (2016), which is already implemented in the LangPro theorem prover. Examples of LLFs with type-raised NPs are given for the sentences (5-c) and (5-d). All lexical terms retain their types except the determiners; their $n(np)$ type is replaced with $n(vp(s))$.

3.3 Natural Language Reasoning

Reasoning over the Dutch LLF is handled by LangPro (Abzianidze 2017), a Natural Tableau-based automated theorem prover. The prover, in its original English implementation, uses a CCG parser to parse and tag input sentences, and builds two tableaux (one for entailment, and one for contradiction⁶) while using the Princeton WordNet (Miller 1995) as a lexical knowledge base (KB). LangPro has been applied to a few NLI benchmarks, and its results rank high among logic-based NLI systems (Abzianidze 2016). In order to enable the acquisition of novel lexical knowledge from data, Abzianidze (2020) recently proposed a training methodology that models learning as abductive reasoning.

We extend the theorem prover to allow processing of Dutch sentences; switching between languages can be done easily by setting the corresponding flag. The adaptation process includes changes in two prover components: the inventory of tableau rules, and the knowledge base. We extend the scope of one tableau rule (ϵMOD) and add a new closure rule ($\times\perp$) to the rule inventory due to

6. To prove the contradiction relation between a premise and a hypothesis, a tableau starts with both premise and hypothesis marked with the true sign because a counterexample for the contradiction relation is when both of the sentences can be true.

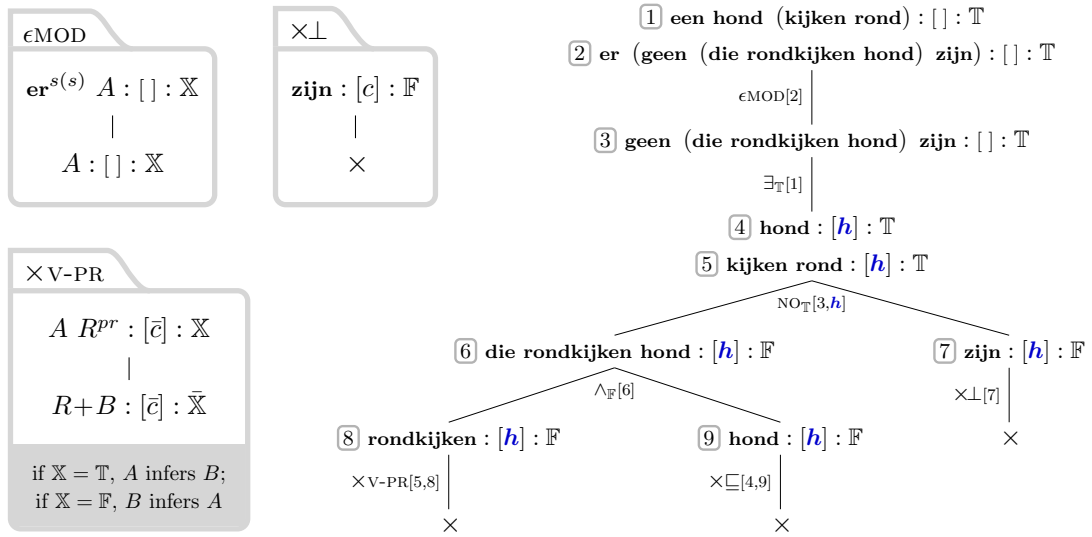


Figure 4: The tableau proves the SICK-NL problem 5220: “Een hond kijkt rond” contradicts “Er is geen hond die rondkijkt”. The proof uses new rules specially designed for Dutch. \bar{X} is a negated version of X , where T ’s negation is F and vice versa.

the analysis of Dutch expletive constructions differing from English ones. This contrast is shown in Examples (6) and (7). Dutch syntactic terms treat the expletive “er” as a clause modifier, while its English counterpart “there” is an argument of the main verb, following the CCG analysis. We also introduce a new rule ($\times V$ -PR) for Dutch phrasal verbs.

- (6) (eng) There is no dog looking around
 (11f) $\text{no}^{n(np)} (\text{wh}^{vp(n(n))} (\text{around}^{vp(vp)} \text{look}^{vp} \text{dog}^n) (\lambda x. \text{be}^{np(vp)} x \text{there}^{np})$
- (7) (nld) Er is geen hond die rondkijkt
 (11f) $\text{er}^{s(s)} (\text{geen}^{n(np)} (\text{die}^{vp(n(n))} \text{rondkijken}^{vp} \text{hond}^n) \text{zijn}^{vp})$

A tableau proof in Figure 4 illustrates the Dutch-specific additions to the tableau rule inventory. In order to prove that “Een hond kijkt rond” contradicts “Er is geen hond die rondkijkt”, the tableau method shows that there is no possible situation where both sentences are true; therefore, the tableau construction starts with entries 1 and 2 marked as true. In the proof, the (ϵ MOD) rule treats the expletive “er” as a semantically vacuous modifier (i.e. the identity function) when 3 is obtained from 2. The right-hand side branch is closed after ($\times\perp$) is applied to 7. Similarly to reasoning with English, the relative pronoun “die” is analysed as a logical conjunction. The middle branch is closed due to 4 and 9 contradicting each other. The left-hand side branch is closed with the help of the new ($\times V$ -PR) rule, which identified contradiction between 5 (it is true that “ h kijkt rond”) and 8 (it is false that “ h rondkijkt”).

Obviously, without a language-specific lexical database, the tableau prover would only be able to tackle dull logical relations. To further allow reasoning with lexical knowledge, we employ the Open Dutch WordNet (Postma et al. 2016). To make the database compatible with the theorem prover, we convert it to the Princeton-style prolog format.⁷ In addition to the antonymy and hyperonymy relations present in the resource, we also use near synonymy and cross-category near synonymy, which we cast as similarity and derivational morphosemantic relation, respectively.

7. <https://wordnet.princeton.edu/documentation/prologdb5wn>

We adopt a default approach of the theorem prover when extracting KB relations from the WordNet. In particular, two words are in a certain relation if there exist word senses of these words for which the corresponding WordNet relation holds. Put differently, it is an *all-sense* approach, where all corresponding word senses are considered when comparing two words. By opting for this approach, we avoid additionally complicating our pipeline by adding a word sense disambiguation system to it. Moreover, Abzianidze (2015a) showed that it works reasonably well for the English SICK dataset.

4. Experiments

In order to experimentally validate our methodology, we utilise SICK-NL, a recently created Dutch NLI dataset, described in §4.1. We perform a range of experiments involving various combinations of tools, settings and baselines, detailed in §4.2, and present our results in §4.3.

4.1 SICK-NL

The SICK dataset (Marelli et al. 2014b) is a collection of 6 076 sentences originating from image captions. These sentences are arranged in 9 840 problems of ordered pairs, made up of a premise and a hypothesis together with an inference label (neutral, entailment or contradiction) that signifies the one-directional logical inference from premise to hypothesis. The dataset is originally segmented in three parts: train, trial and test (4 500, 500 and 4 927 problems respectively) for the SemEval-14 shared task (Marelli et al. 2014a).

The Dutch version of SICK (Wijnholds and Moortgat 2021) is derived from the original through an automated translation process and a gold inference label transfer. The dataset consists of 6 060 unique sentences in total. The automatic translations were manually inspected, and erroneous ones were corrected on an individual sentence basis (i.e. not taking pairing contexts into account). Based on preliminary experiments with neural models, Wijnholds and Moortgat (2021) indicate that the Dutch counterpart is more difficult than the original; the main hypothesised reason is a reduction in lexical overlap between sentence pairs (despite a small reduction in average sentence length), owing to machine translation inflating the dataset’s vocabulary.

4.2 Experimental Setup

We experiment with syntactic expressions from both parsers. From Alpino, we request the global optimal parse for each sentence, imposing no time constraints. The graphs obtained are then converted into typological derivations with the aid of the extraction algorithm. In the process, a small portion of parses are discarded due to either underspecifying the sentence’s function/argument structure (e.g. resorting to vague discourse-level annotations) or failing rudimentary correctness checks. From Neural Proof Nets, we select the structurally correct analysis with the highest score that falls within a beam of width 6 (when at least one such exists). Table 3 reports sentence and dataset coverage from each parser pipeline.

Parser	Sentences Parsed	Problems Covered
Neural Proof Nets	5 812 (95,9%)	9 264 (94,1%)
Alpino	5 947 (98,1%)	9 611 (97,7%)

Table 3: Sentences parsed and problems covered with each parser.

We then transform syntactic analyses to logical forms as described in §3.2. In order to ease lexical lookup, we homogenise semantic constants by lemmatising them, using pos tags to disambiguate, and apply term conversions when necessary. Lemmas and tags are obtained from two sources:

Alpino, and spaCy’s large Dutch language model (Honnibal et al. 2020). Logical forms are then fed to the LangPro theorem prover – if a problem is missing a term for either premise or hypothesis, the prover’s prediction defaults to the neutral label.

We obtain an accuracy score on the test set (percentage of problems correctly classified) from each parser & tagger combination, as well as five ensemble models. Each ensemble aggregates the votes of equally-weighted models, prioritising non-neutral over neutral votes, and defaulting to neutral in case of conflict (e.g. entailment vs contradiction). We produce two ensembles over parsers, two over taggers, and one over all four taggers and parser combinations.

Next, we train each of the core models using abduction on the union of the training and trial portions of the dataset. We use the trial set in abductive learning since the theorem prover has no proper set of hyperparameters that can be tuned in the development phase. For the abductive learning we use the settings of Abzianidze (2020). Post-training, models are organised in ensemble pairs as before, without cross-model spilling of learned knowledge.

To quantitatively assess our models’ performance, we compare against established pretrained language models, fine-tuned as three-way sequence classifiers (a sequence being the concatenation of the premise and hypothesis sentences, as is standard practice). Following Wijnholds and Moortgat (2021), we use BERTje (de Vries et al. 2019), RobBERT (Delobelle et al. 2020) and mBERT (Devlin et al. 2019), but perform model selection on the basis of trial set accuracy, and average scores from five training instances.

4.3 Results

Table 4a presents the results for all parser & tagger combinations and ensembles, with and without abduction. Comparing individual components, we note that models perform better with (i) pos tags and lemmas coming from spaCy rather than Alpino, and (ii) parse structures coming from the Alpino pipeline rather than NPN. When it comes to abduction, trained models perform consistently better across the board, raising individual model performance by 1.38–1.83%. In line with previous work (Abzianidze 2015a, Martínez-Gómez et al. 2016), aggregating proofs from various model combinations substantially improves results. Our best performing model is the ensemble of four theorem provers using all cross combinations of parsers & taggers, where each of the provers has been trained using abduction; The ensemble model achieves a raw improvement of 1.22% over its best constituent (with the Alpino parser & spaCy tagger combination). We abbreviate this ensemble model as LangPro Σ^2 and use it for subsequent comparisons.

Unsurprisingly, and as Table 4b suggests, all BERT-based models outperform LangPro Σ^2 , with a maximum absolute difference of 3.2%. However, inspecting the confusion matrices of the systems in Table 5 reveals LangPro Σ^2 ’s merits, namely the high precision of its entailment and contradiction

Tagger	Parser			Model	Accuracy	Hybrid
	NPN	ALPINO	Σ			
Alpino	74.65 -1.50	75.87 -1.83	76.38 -1.75	LangPro Σ^2	78.8	–
spaCy	76.66 -1.38	77.61 -1.72	78.38 -1.58	BERTje	82.0	81.8
Σ	77.04 -1.40	77.98 -1.71	78.83 -1.62	RobBERT	81.7	82.6
				mBERT	79.9	80.6

(a) Accuracy of LangPro when using each parser & tagger combination, including ensembles (Σ) over parser, tagger or both. The right subcolumns report differences when no training with abduction is used.

(b) Performance of LangPro Σ^2 compared to fine-tuned neural baselines. The right column reports performance when LangPro Σ^2 proofs override neural predictions.

Table 4: Internal and external model comparisons on the test set of SICK-NL, with 56.4% of neutral-class baseline. The scores are a percentage of correctly classified problems.

%	LangPro Σ^2			BERTje			RobBERT			mBERT		
	E	C	N	E	C	N	E	C	N	E	C	N
Entailment	14.6	0.1	14.0	24.7	0.1	3.9	22.1	0.1	6.6	22.7	0.2	5.8
Contradiction	<0.1	9.8	4.8	0.7	12.7	1.3	0.6	12.5	1.5	0.6	12.0	2.0
Neutral	1.5	0.7	54.5	9.5	2.6	44.6	7.2	2.3	47.1	9.3	2.2	45.2

Table 5: Confusion matrices on the test set. The numbers represent a percentage of the test set.

predictions. Proofs generated by LangPro Σ^2 are reliable enough to safely override most neural models’ predictions, allowing the two types of systems to complement one another. Table 4b shows accuracy scores for all baselines, as well as hybrid models where LangPro Σ^2 ’s proofs (i.e. entailment and contradiction predictions) override predictions of the neural models. Evidently, LangPro Σ^2 can benefit RobBERT and mBERT but not BERTje. The reason behind the latter is that according to the gold labels, LangPro Σ^2 correctly proofs 39 proofs for entailment (31) and contradiction (8) problems which are wrongly classified by BERTje, but LangPro Σ^2 also provides false proofs for 48 neutral problems (the false proofs are discussed in §5.1), which outweigh the accuracy gain from the correct proofs.

The hybrid model that pairs LangPro Σ^2 and RobBERT outperforms all models, surpassing the previous benchmark of BERTje. Closer look at the predictions of LangPro Σ^2 and RobBERT reveal that RobBERT benefits most from LangPro Σ^2 ’s proofs (93) for entailment problems compared to mBERT (71) and BERTje (31). Table 5 also shows that RobBERT is the worst among the neural baselines in predicting entailment problems, but the adoption of LangPro Σ^2 ’s proofs results in the best performing hybrid model.

It is interesting to see the problems that all neural models failed at, but LangPro Σ^2 solved. Several of these problems are shown in Table 6. The problems seem easy, but for some reason all the neural baselines predict them as neutral. It is even more mysterious how all of them predict the comparable problem 169 correctly, which has the same premise as 175 paired with the hypothesis “Een familie kijkt naar een jongen die een honkbal slaat”. The reasoning capacity required to predict 169 correctly is sufficient for solving 175; we hypothesise that neural models give more weight to word sequence similarity when it comes to predicting entailment.

id/Label	Sentences
168	p Een kind slaat een honkbal
C	h Een kind mist een honkbal
175	p Een familie kijkt naar een kleine jongen die een honkbal raakt
E	h Een jongen slaat een honkbal
1556	p Een man draagt een boom
E	h Een man draagt een plant
897	p Mensen zitten op een strand vol zand bij de oceaan en genieten van een zonnige dag
C	h Er is niemand aan de wal
4470	p Een man schopt een voetbal
E	h Een man schopt een bal

Table 6: Based on the test set, a set of problems (cherry-picked from a total of 17 problems) that were misclassified by all neural models as neutral but solved by LangPro Σ^2 . The gold labels are abbreviated with the initial letters.




Figure 5: Error analysis for a sample of problems falsely classified as neutral.

5. Analysis

To gain a better insight into the model’s performance, we perform extensive qualitative analyses targeted at either specific components of the framework (namely, the syntactic parsers in §5.2 and the abductive learning in §5.3) or particular cases of interest (missing and imagined proofs by LangPro Σ^2 in §5.1). All the conducted analyses are based on the training part of SICK-NL to avoid eyeballing the problems from the test set.

5.1 Missing and Imagined Proofs

We begin by investigating the predictions of LangPro that differ from gold labels, treating the neutral label as the null hypothesis. To avoid confounding our analysis with the effects of abduction, we inspect predictions by untrained LangPro versions.

We first focus on *missing proofs* or type II errors, i.e. cases where *none* of the stand-alone models produce a proof (that is, a neutral prediction for a non-neutral gold label problem). We detect a total of 1038 such cases out of 4500, and randomly sample 60 of those, half between entailments and contradictions. Our findings are presented in Figure 5 and illustrating examples in Table 7. The majority of missing proofs (17) can be attributed to requiring commonsense reasoning or world knowledge, which goes beyond the capacities of the prover (e.g. problem 4701 requires knowing that “voetbalt” implies “speelt met een bal”). LangPro is also responsible for not delivering equally many proofs, the reason being the absence or malfunction of a structure altering rule (e.g. 8073 requires ignoring the auxiliary to equate the present continuous “staan te sparren” with the simple continuous “sparren”). Another 13 cases can be explained as requiring lexical relations not present

id/Label	Sentences
4701	p Een groep mannen voetbalt op het strand
E	h Een groep mannen speelt met een bal op het strand
8073	p Twee jongens in witte outfits en rode beschermende kleding staan te sparren op een mat
E	h Twee kinderen in witte outfits en rode beschermende kleding sparren op een mat
7375	p Een man staat op de top van de rotsen met wolken erachter
N	h Een persoon zit op een bergtop

Table 7: Problems that were misclassified by LangPro Σ^2 . The gold labels are abbreviated with the initial letters.




Figure 6: Error analysis for a sample of neutral problems falsely classified.

in the KB, and 7 more are due to noisy (erroneous or ambiguous) gold labels, partially caused by translation-induced meaning shifts (e.g. problem 1194 of Table 8). Finally, 6 are to be blamed on absurd parses.

Next, we turn our attention to *imagined proofs* or type I errors, i.e. cases where *all* of the stand-alone models produce a proof, whereas none was expected (that is, a non-neutral output with a neutral gold label). We detect and inspect a total of 33 such cases and present our findings in Figure 6; most of the errors (21) are due to inaccurate lexical relations with 11 of the remaining being debatable gold labels (Table 8).

Based on the above findings, we draw a number of conclusions pertaining to the framework as well as the dataset, and identify recurring patterns in the kinds of sentences we fail to properly analyze. First, we emphasise the 2 orders of magnitude difference between type I & II errors: there are almost 1 000 cases of missing proofs, but only about 30 cases of wrong proofs, which serves to show that LangPro Σ^2 is precise in its proofs but lacks high coverage, and thus makes for a good candidate first model in a hierarchical classification pipeline.

Witness to that, we remark that 33% of the false positives and 12% of the false negatives encountered are in fact plausible or outright correct, the issue lying with the label rather than the prediction! A portion of the mislabeled problems found are in agreement with prior analyses of

	id	Sentences	Label Change
1202	en	p A man is making a speech on a podium	N
	h	A man is speaking on a stage	↓
	nl	p Een man houdt een toespraak op een podium	?E
	h	Een man spreekt op een podium	
1818	en	p Some bells are ringing near a cook slicing peppers	N
	h	A cook is slicing some bell peppers	↓
	nl	p Er rinkelen wat belletjes bij een kok die paprika's in reepjes snijdt	?E
	h	Een kok snijdt wat paprika's in reepjes	
1194	en	p A man is speaking on a podium	C
	h	There is no man speaking on a podium	↓
	nl	p Een man spreekt op een lessenaar	?N
	h	Er spreekt geen man op een podium	

Table 8: Example problems with potential label discrepancies between English and Dutch.

the original SICK, common causes being a lack of an absolute reference frame, no clear distinction between alteration and contradiction, ungrammatical sentences, and annotation errors (Kalouli et al. 2017). Others, however, are unique to the Dutch translation, and can be pinpointed to the lexical choices of the machine translation system employed; Table 8 presents a few telling examples. In several cases (problems 1202 and 1818), two distinct source words are translated to the same target word, creating slight meaning shifts that affect the inference label. The issue is not exclusive to type I errors; in fact, shifts occur even more frequently the other way around, translating the same word differently depending on (sometimes irrelevant) context (problem 1194). Albeit not always catastrophic, translation-induced inconsistencies magnify the dataset’s difficulty: lexical inconsistencies increase the vocabulary size, and therefore the demand on the knowledge base, whereas grammatical inconsistencies necessitate a more exhaustive set of structure-altering rules.

Concerning Open Dutch WordNet, we note that the resource places a severe upper boundary on system performance, as 25% of the missing proofs are due to the absence of a needed lexical relation. This is not surprising, considering the scale of the database and the highly demanding nature of the task. What is, however, surprising is the frequency of relations that lead to unexpected proofs. Upon closer inspection, we distinguish two error cases. The first is due to relations that are just plain wrong, a striking example being `kat` \sqsubseteq `hond`, which contributes to a total of six imagined proofs.⁸ The other is more deeply rooted, and is associated with the all-sense approach we chose to adopt for simplicity (see §3.3). For instance, based on the standard (i.e. frequently used) senses, verbs like *liggen*, *lopen*, and *staan* are not hyponyms of *zitten*. But there is a sense of *zitten*, which means to occupy a certain position or area, and it is a hypernym of some senses of *liggen*, *lopen*, and *staan*. This makes the all-sense approach to adopt the relations like `liggen` \sqsubseteq `zitten`, `lopen` \sqsubseteq `zitten`, and `staan` \sqsubseteq `zitten` and to prove problems like 7375 in Table 7.⁹

Finally, the prevalence of separable verbs in Dutch can also be a source of stress for all components of our framework. Starting from the parsers, there is an apparent tension between optional adverbs and necessary but free-floating particles. LangPro itself then requires careful tuning on the treatment of each case (while accounting for possible errors) before deferring to the lexical database. The latter may often contain a relation between verbal cores but lack one for the full verbs, making derivations possible only if one selectively ignores particles; this, however, carries the danger of ignoring crucial parts of the sentential meaning.

5.2 Alpino vs. NPN

Our next analysis seeks to investigate the effect of parser choice on model performance. Perhaps strikingly, predictions that rely on the Alpino-based pipeline seem generally more reliable than predictions based on the NPN parser, despite the latter reportedly achieving a higher parsing accuracy (Kogkalidis et al. 2020a). We randomly sample and inspect 50 problems where exactly one of the parsers’ outputs leads to a proof. Our findings are presented in Figure 7. Both systems fail with almost equal frequency (26 failures from the Alpino-based pipeline and 24 from NPN), but the error sources are quite different between the two.

More than half (14) of the NPN failures arise from a lack of a parse. Of those, 5 problems contain comma-separated non-restrictive relative clauses (NPN has been trained with punctuation-free sentences), 3 contain very short and simple “aan het” constructions (incorrectly analysed in NPN’s training data) and 3 more contain simple transitive sentences where the subject is a conjunction (we hypothesise those to be training artifacts). Alpino, on the other hand, rarely fails at producing any parse (5 cases), but its output is more often wrong. There are 11 cases of PP attachment gone wrong, 8 cases of a modifier mislabeled as a predicate or vice-versa and 4 cases of severe issues in the

8. In the Open Dutch WordNet, while `hond.n.01` is a direct hyponym of `huisdier.n.01`, erroneously it is also a hypernym of it. This makes `kat.n.01` a hyponym of `hond.n.01` as it is a direct hyponym of `huisdier.n.01`.

9. It is possible to block such unwanted relations from the Open Dutch WordNet by specifically discarding certain problematic senses and hypernymy relations from it. But pushing the performance score as high as possible is not the main goal of the current paper.

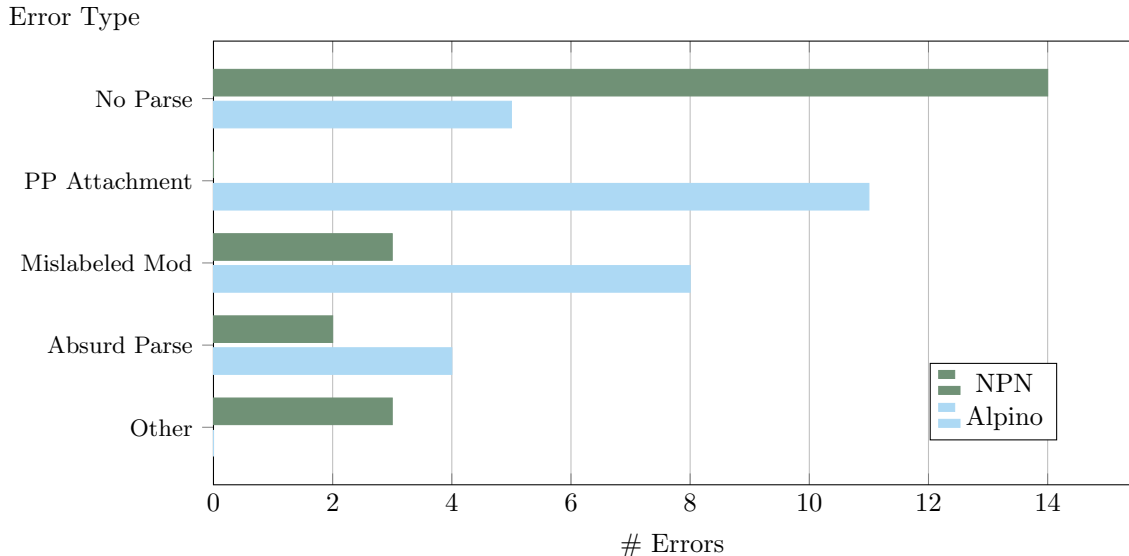


Figure 7: Error analysis for a sample of problems with parser disagreements.

predicted function/argument structure; NPN, in comparison, has 0, 3 and 2 of those respectively. Considering also Table 3, the verdict is that Alpino is more robust, boasting a higher coverage which in turn leads to more solved problems, but NPN is more accurate, boasting a higher proportion of correct parses; this serves to further justify our decision to ensemble the two, and is in line with the benefits observed when doing so.

5.3 Abduction

The abductive learning further boosts the performance of the theorem prover with the help of lexical knowledge induced during the training phase (see Table 4a). The learning bias of the abduction is to find the smallest set of relations over short phrases that explains (i.e. helps prove) the gold inference label for a premise-hypothesis pair. Here, we manually check and analyse the relations learned by the abduction from the train and trial parts of SICK-NL.

We consider the overlap between the relations learned by each version of the theorem prover differing in terms of the parser-tagger combinations. There are in total 161 such common relations learned. We classify each relation according to five categories. The *correct* and *wrong* categories are self-explanatory. *contextual* relations require a substantial amount of context to be considered justifiable. *Reversed* relations are reversed versions of correct subsumption relations, while *prepositional* are, as the name suggests, relations over prepositions. A distribution of the learned relations over the five categories with accompanying examples are shown in Table 9.

The majority of learned relations are correct, and either resemble WordNet-like entries, such as the antonym `leeg` | `vol` and the hyper/hyponym `pizza` | `voedsel`, or commonsense-like relations, e.g. `zwart` \sqsubseteq `donker`. The wrong relations cover examples caused by wrong lemmatisation (e.g. `wissen` \sqsubseteq `weten`), noisy gold labels (e.g. `suiker` \sqsubseteq `kruid`), and failure to correctly identify verb particle constrictions (e.g. `gieten` | `halen`). Worth discussing are also reversed relations. While some relations (e.g. `hond` \sqsubseteq `bulldog`) are induced from noisy gold labels, most of them are due to the learning bias of the abduction preferring relations with short phrases. This preference opts for learning `persoon` \sqsubseteq `fietser` rather than `persoon die op fietsen rijdt` \sqsubseteq `fietser`.

Type	Count (%)	Examples of learned relations
Correct	53 (32.9)	zwart \sqsubseteq donker : zwarte hond \sqsubseteq donkere hond leeg vol : schaatsbaan is leeg schaatsbaan is vol met mensen liggen in het gras rennen : hond ligt in het gras hond rent in het gras pizza \sqsubseteq voedsel : er is geen man die voedsel eet Een man eet een pizza lopen \sqsubseteq rennen : honden lopen snel samen \sqsubseteq honden rennen samen
Contextual	20 (12.4)	zwaard \sqsubseteq mes : snijdt een laars met een zwaard \sqsubseteq snijdt een laars met een mes boren sluiten : boort een gat sluit een gat klein \sqsubseteq jong : kleine jongen speelt \sqsubseteq jonge jongen speelt
Reversed	26 (16.1)	container \sqsubseteq plastic container : een container van plastic \sqsubseteq een plastic container hond \sqsubseteq bulldog : aap borstelt de hond aap borstelt geen bulldog persoon \sqsubseteq fietser : Iemand die op fietsen rijdt \sqsubseteq een fietser
Prepositional	27 (16.8)	op \sqsubseteq door : stel loopt het gangpad op \sqsubseteq stel loopt door het gangpad in \sqsubseteq aan : puppy knaagt aan een houten paal \sqsubseteq puppy bijt in een paal buiten \sqsubseteq rond : tijger loopt buiten een kooi \sqsubseteq tijger loopt rond een kooi
Wrong	35 (21.7)	wissen \sqsubseteq weten : man is het werk ... aan het wissen \sqsubseteq man wist het werk ... suiker \sqsubseteq kruid : voegt suiker toe aan het vlees \sqsubseteq voegt kruiden toe aan wat vlees spannen \sqsubseteq draaien : meisje ... spant een lint geen meisje ... dat een lintje draait gieten halen : giet olie in een koekenpan haalt de olie uit een koekenpan

Table 9: Types of relations learned with the abduction and their counts. Each relation instance comes with a context from the original problem as seen during abductive learning.

6. Related Work

There are a few logic-based systems that have been successfully applied to NLI. Good entry datasets for logic-based systems to NLI represent SICK (Marelli et al. 2014b) and FraCaS (Cooper et al. 1996) as the former contains relatively simple sentences mainly requiring reasoning with lexical and compositional knowledge, while the latter covers multi- and single-premised problems presupposing complex logic-based reasoning. Two logic-based NLI systems that stand out with their performance on these NLI datasets are *ccg2lambda* (Martínez-Gómez et al. 2016, Yanaka et al. 2018) and *LangPro* (Abzianidze 2017, Abzianidze 2020). While both systems are logic-based and use CCG parsers as a starting point, they differ in terms of the logical representations of sentences and the reasoning procedures. While *ccg2lambda* has already been applied to Japanese NLI (Mineshima et al. 2016), the current work represents the first cross-lingual application of *LangPro*. We believe that *ccg2lambda* could also be adapted to Dutch likewise *LangPro* to Japanese as both systems require comparable resources of syntactic parsing and lexical knowledge.

Categorial grammar-based syntactic trees represent a smooth starting point when it comes to obtaining logical forms of sentences since the grammars’ transparent syntax-semantics interface facilitates the meaning composition. That’s why most of the logic-based NLI systems have used CCG-based parsers, which are the best performing categorial grammar-based parsers for English. In this paper, we take advantage of a wide-coverage typological parser for Dutch, and shift our syntactic representations accordingly (Kogkalidis et al. 2020a).

An NLI dataset has been recently made available for Dutch (Wijnholds and Moortgat 2021), and during the writing this paper we are not aware of any Dutch NLI systems (other than the baseline neural models discussed in our experiments). We would like to stress here that a Dutch system

for semantic similarity (Marsi and Krahmer 2010) is not comparable to Dutch NLI systems as it is tackling a different NLP task. Additionally, it is measuring the similarity of syntactic trees without really reasoning about the meaning (the latter can also be attributed to the Dutch neural NLI models to some extent).¹⁰ Obviously one could adapt a semantic similarity-based system to NLI, but such an endeavour would offer little in terms of explainability compared to our reasoning-based framework.

7. Conclusion

Building on existing work, we have proposed a framework for the logical analysis of textual inference in Dutch. Our work has been motivated by the recent release of a Dutch translation of the SICK dataset (Wijnholds and Moortgat 2021), and is the first work specifically targeted to the dataset. As our entry point, we used the two available tools for acquiring type-theoretic analysis of written Dutch to parse the entirety of the dataset (Kogkalidis et al. 2020b, Kogkalidis et al. 2020a). Relying on the clean syntax-semantics interface offered by typological grammars and their close affinity to λ -calculi, we hand-designed a conversion scheme that first simplifies syntactic terms, before then casting them to semantic expressions. We then employed a high-horsepower Natural Tableau prover (Abzianidze 2017), and expanded upon it with new rules, aimed at addressing some of the quirks of Dutch. Supported with the lexical relations from the Open Dutch WordNet (Postma et al. 2016), the prover first learns domain-specific relations from the training data via an abductive learning component (Abzianidze 2020) and then predicts unseen problems based on formal proofs. We finally compared our system to strong neural baselines (Devlin et al. 2019, de Vries et al. 2019, Delobelle et al. 2020), and find them to be not only on equal standing, but also complementary to one another to some extent.

Contrary to neural alternatives, our proposed framework constitutes a “glass box” model, yielding answers not in the form of plain labels, but rather proofs which can be both human-inspected and machine-verified. This allows us to gain deeper insights into the problem in all of its aspects, including the peculiarities of the dataset (e.g. detecting noisy gold labels), the strengths and weaknesses of the components employed (e.g. identifying missing relations in the Open Dutch WordNet), and the methodological decisions we have followed. We have conducted an in-depth error analysis, which has shed light on the effects of automatic translation on the difficulty of the dataset, of word variation on lexical database stress, of parse choice on prominent error cases, and of the abduction on the quality of lexical knowledge extracted.

The extensive error analysis showed several directions for future work that can further improve our system’s performance. First, to surmount the issues related to wrong parses or inconsistent PP attachments, one could additionally consider logical forms generated from n-best parses from NPN and Alpino and/or employ yet another syntactic parser for Dutch (e.g. a Dutch CCG parser developed as part of the Parallel Meaning Bank (Abzianidze et al. 2017) is an obvious candidate for this). Second, to improve the abductive learning, a promising direction would be to incorporate embedding-based word similarity during training, e.g. give priority to relations with a higher cosine similarity between their arguments. Third, to better process Dutch particle verbs during term conversion and theorem proving, one could include specialised processing of verb lemmas that include a particle. Fourth, to procure more training data (Yanaka et al. 2019b) or phenomenon-specific evaluation sets (Yanaka et al. 2019a, Richardson et al. 2020, Yanaka et al. 2021) for Dutch NLI, one could automatically create labeled NLI pairs where our logic-based NLI system could serve as an integral component of the problem generation or as a sanity checker of the inference labels. Error analysis aside, we are finally curious to explore how dependency relations (provided by both

10. We include this comparison after one of the reviewers considered our work similar to the DAESO project in terms of scope and aim.

parsing frameworks, but not utilised here) can find use to term processing, lexical disambiguation and reasoning as a whole.

In hopes that our work will prove useful to future research on Dutch NLI and acknowledging the contributions of others that made it possible, we open-source our code and make it available at git.io/JzdGd.

Acknowledgements

We thank the organisers of the Natural Logic Meets Machine Learning (NaLoMA) workshop for hosting us, the reviewers for their suggestions on earlier drafts of this work, and the participants for attending. We also thank the members of the Utrecht NLP Reading Group for providing a friendly environment for discussions and rehearsals. Lasha is supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 742204). Konstantinos is supported by the Dutch Research Council (NWO) through the project “A composition calculus for vector-based semantic modelling with a localisation for Dutch” (360-89-070).

References

- Abzianidze, Lasha (2015a), A tableau prover for natural logic and language, *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Lisbon, Portugal, pp. 2492–2502. <https://www.aclweb.org/anthology/D15-1296>.
- Abzianidze, Lasha (2015b), Towards a wide-coverage tableau method for natural logic, in Murata, Tsuyoshi, Koji Mineshima, and Daisuke Bekki, editors, *New Frontiers in Artificial Intelligence*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 66–82.
- Abzianidze, Lasha (2016), *A natural proof system for natural language*, PhD thesis, Tilburg University.
- Abzianidze, Lasha (2017), LangPro: Natural language theorem prover, *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Association for Computational Linguistics, Copenhagen, Denmark, pp. 115–120. <https://www.aclweb.org/anthology/D17-2020>.
- Abzianidze, Lasha (2020), Learning as abduction: Trainable natural logic theorem prover for natural language inference, *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, Association for Computational Linguistics, Barcelona, Spain (Online), pp. 20–31. <https://www.aclweb.org/anthology/2020.starsem-1.3>.
- Abzianidze, Lasha, Johannes Bjerva, Kilian Evang, Hessel Haagsma, Rik van Noord, Pierre Ludmann, Duc-Duy Nguyen, and Johan Bos (2017), The Parallel Meaning Bank: Towards a multilingual corpus of translations annotated with compositional meaning representations, *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Association for Computational Linguistics, Valencia, Spain, pp. 242–247. <https://aclanthology.org/E17-2039>.
- Blackburn, Patrick and Johan Bos (2005), *Representation and Inference for Natural Language. A First Course in Computational Semantics*, CSLI.
- Blackburn, Patrick, Johan Bos, Michael Kohlhase, and Hans De Nivelle (2001), Inference and computational semantics, in Bunt, Harry, Reinhard Muskens, and Elias Thijsse, editors, *Computing Meaning: Volume 2*, Springer Netherlands, Dordrecht, pp. 11–28.

- Bouma, Gosse, Gertjan van Noord, and Robert Malouf (2001), Alpino: Wide-coverage computational analysis of Dutch, *Computational linguistics in the Netherlands 2000*, Brill Rodopi, pp. 45–59.
- Cooper, Robin, Dick Crouch, Jan Van Eijck, Chris Fox, Josef Van Genabith, Jan Jaspars, Hans Kamp, David Milward, Manfred Pinkal, Massimo Poesio, Steve Pulman, Ted Briscoe, Holger Maier, and Karsten Konrad (1996), *FraCaS: A Framework for Computational Semantics*, Deliverable D16.
- D’Agostino, Marcello, Dov M. Gabbay, Reiner Hähnle, and Joachim Posegga, editors (1999), *Handbook of Tableau Methods*, Springer Netherlands, Dordrecht.
- de Vries, Wietse, Andreas van Cranenburgh, Arianna Bisazza, Tommaso Caselli, Gertjan van Noord, and Malvina Nissim (2019), BERTje: A Dutch BERT model, *arXiv preprint arXiv:1912.09582*.
- Delobelle, Pieter, Thomas Winters, and Bettina Berendt (2020), RobBERT: a Dutch RoBERTa-based Language Model, *Findings of the Association for Computational Linguistics: EMNLP 2020*, Association for Computational Linguistics, Online, pp. 3255–3265. <https://www.aclweb.org/anthology/2020.findings-emnlp.292>.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2019), BERT: Pre-training of deep bidirectional transformers for language understanding, *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)*, pp. 4171–4186.
- Girard, Jean-Yves (1987), Linear logic, *Theoretical computer science* **50** (1), pp. 1–101, Elsevier.
- Glockner, Max, Vered Shwartz, and Yoav Goldberg (2018), Breaking NLI systems with sentences that require simple lexical inferences, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Association for Computational Linguistics, Melbourne, Australia, pp. 650–655. <https://aclanthology.org/P18-2103>.
- Gururangan, Suchin, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith (2018), Annotation artifacts in natural language inference data, *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, Association for Computational Linguistics, New Orleans, Louisiana, pp. 107–112. <https://www.aclweb.org/anthology/N18-2017>.
- Hockenmaier, Julia and Mark Steedman (2007), CCGbank: A corpus of CCG derivations and dependency structures extracted from the Penn Treebank, *Comput. Linguist.* **33** (3), pp. 355–396, MIT Press, Cambridge, MA, USA.
- Honnibal, Matthew, Ines Montani, Sofie Van Landeghem, and Adriane Boyd (2020), spaCy: Industrial-strength Natural Language Processing in Python. <https://doi.org/10.5281/zenodo.1212303>.
- Kalouli, Aikaterini-Lida, Livy Real, and Valeria de Paiva (2017), Textual inference: getting logic from humans, *IWCS 2017 — 12th International Conference on Computational Semantics — Short papers*. <https://aclanthology.org/W17-6915>.
- Kogkalidis, Konstantinos, Michael Moortgat, and Richard Moot (2020a), Neural proof nets, *Proceedings of the 24th Conference on Computational Natural Language Learning*, Association for Computational Linguistics, Online, pp. 26–40. <https://www.aclweb.org/anthology/2020.conll-1.3>.

- Kogkalidis, Konstantinos, Michael Moortgat, and Richard Moot (2020b), *Æthel: Automatically extracted typological derivations for Dutch*, *Proceedings of The 12th Language Resources and Evaluation Conference*, European Language Resources Association, Marseille, France, pp. 5259–5268. <https://www.aclweb.org/anthology/2020.lrec-1.647>.
- Kogkalidis, Konstantinos, Michael Moortgat, and Tejaswini Deoskar (2019), Constructive typological supertagging with self-attention networks, *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pp. 113–123.
- Marelli, Marco, Luisa Bentivogli, Marco Baroni, Raffaella Bernardi, Stefano Menini, and Roberto Zamparelli (2014a), SemEval-2014 task 1: Evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment, *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Association for Computational Linguistics, Dublin, Ireland, pp. 1–8. <https://aclanthology.org/S14-2001>.
- Marelli, Marco, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli (2014b), A SICK cure for the evaluation of compositional distributional semantic models, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, European Language Resources Association (ELRA), Reykjavik, Iceland, pp. 216–223. <http://www.lrec-conf.org/proceedings/lrec2014/pdf/363.Paper.pdf>.
- Marsi, Erwin and Emiel Krahmer (2010), Automatic analysis of semantic similarity in comparable text through syntactic tree matching, *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, Coling 2010 Organizing Committee, Beijing, China, pp. 752–760. <https://aclanthology.org/C10-1085>.
- Martínez-Gómez, Pascual, Koji Mineshima, Yusuke Miyao, and Daisuke Bekki (2016), ccg2lambda: A compositional semantics system, *Proceedings of ACL-2016 System Demonstrations*, Association for Computational Linguistics, Berlin, Germany, pp. 85–90. <https://aclanthology.org/P16-4015>.
- McCoy, Tom, Ellie Pavlick, and Tal Linzen (2019), Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, pp. 3428–3448. <https://aclanthology.org/P19-1334>.
- Mena, Gonzalo, David Belanger, Scott Linderman, and Jasper Snoek (2018), Learning latent permutations with Gumbel-Sinkhorn networks, *International Conference on Learning Representations*.
- Miller, George A. (1995), Wordnet: A lexical database for English, *Communications of the ACM* **38** (11), pp. 39–41, ACM, New York, NY, USA.
- Mineshima, Koji, Ribeka Tanaka, Pascual Martínez-Gómez, Yusuke Miyao, and Daisuke Bekki (2016), Building compositional semantics and higher-order inference system for a wide-coverage Japanese CCG parser, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Austin, Texas, pp. 2236–2242. <https://aclanthology.org/D16-1242>.
- Moot, Richard and Christian Retoré (2012), *The logic of categorial grammars: a deductive account of natural language syntax and semantics*, Vol. 6850, Springer.
- Morrill, Glyn V (2012), *Type logical grammar: Categorical logic of signs*, Springer Science & Business Media.

- Moss, Lawrence S. (2010), Natural logic and semantics, in Aloni, Maria, Harald Bastiaanse, Tikitou de Jager, and Katrin Schulz, editors, *Logic, Language and Meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, Revised Selected Papers*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 84–93.
- Muskens, Reinhard (2010), An analytic tableau system for natural logic, in Aloni, Maria, Harald Bastiaanse, Tikitou de Jager, and Katrin Schulz, editors, *Logic, Language and Meaning*, Vol. 6042 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 104–113.
- Petrov, Slav, Dipanjan Das, and Ryan McDonald (2012), A universal part-of-speech tagset, *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, European Language Resources Association (ELRA), Istanbul, Turkey, pp. 2089–2096.
- Poliak, Adam, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme (2018), Hypothesis only baselines in natural language inference, *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, Association for Computational Linguistics, New Orleans, Louisiana, pp. 180–191. <https://www.aclweb.org/anthology/S18-2023>.
- Pollard, Carl and Ivan A Sag (1994), *Head-driven phrase structure grammar*, University of Chicago Press.
- Postma, Marten, Emiel van Miltenburg, Roxane Segers, Anneleen Schoen, and Piek Vossen (2016), Open Dutch WordNet, *Proceedings of the Eight Global Wordnet Conference*, Bucharest, Romania.
- Richardson, Kyle, Hai Hu, Lawrence S. Moss, and Ashish Sabharwal (2020), Probing natural language inference models through semantic fragments, *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, AAAI Press, pp. 8713–8721. <https://aaai.org/ojs/index.php/AAAI/article/view/6397>.
- Sanchez, Ivan, Jeff Mitchell, and Sebastian Riedel (2018), Behavior analysis of NLI models: Uncovering the influence of three factors on robustness, *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Association for Computational Linguistics, New Orleans, Louisiana, pp. 1975–1985. <https://aclanthology.org/N18-1179>.
- Sørensen, Morten Heine and Pawel Urzyczyn (2006), *Lectures on the Curry-Howard isomorphism*, Elsevier.
- Tsuchiya, Masatoshi (2018), Performance impact caused by hidden bias of training data for recognizing textual entailment, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, European Language Resources Association (ELRA), Miyazaki, Japan. <https://aclanthology.org/L18-1239>.
- van Benthem, Johan (2008), Natural logic: A view from the 1980s, in M. K. Chakraborty, B. Lowe, M. N. Mitra and S. Sarukkai, editor, *Logic, Navya-Nayaya & Applications. Homage to Bimal Krishna Matilal*, Vol. 15 of *Studies in Logic*, London College Publications.
- van Noord, Gertjan (2006), At last parsing is now operational, *Actes de la 13ème conférence sur le Traitement Automatique des Langues Naturelles. Conférences invitées*, ATALA, Leuven, Belgique, pp. 20–42. <https://www.aclweb.org/anthology/2006.jeptalnrecital-invite.2>.
- Wadler, Philip (1993), A taste of linear logic, *International Symposium on Mathematical Foundations of Computer Science*, Springer, pp. 185–210.

- Wijnholds, Gijs and Michael Moortgat (2021), SICK-NL: A dataset for Dutch natural language inference, *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, Association for Computational Linguistics, Online, pp. 1474–1479. <https://aclanthology.org/2021.eacl-main.126>.
- Yanaka, Hitomi, Koji Mineshima, and Kentaro Inui (2021), Exploring transitivity in neural NLI models through veridicality, *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, Association for Computational Linguistics, Online, pp. 920–934. <https://aclanthology.org/2021.eacl-main.78>.
- Yanaka, Hitomi, Koji Mineshima, Daisuke Bekki, Kentaro Inui, Satoshi Sekine, Lasha Abzianidze, and Johan Bos (2019a), Can neural networks understand monotonicity reasoning?, *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, Association for Computational Linguistics, Florence, Italy, pp. 31–40. <https://aclanthology.org/W19-4804>.
- Yanaka, Hitomi, Koji Mineshima, Daisuke Bekki, Kentaro Inui, Satoshi Sekine, Lasha Abzianidze, and Johan Bos (2019b), HELP: A dataset for identifying shortcomings of neural models in monotonicity reasoning, *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, Association for Computational Linguistics, Minneapolis, Minnesota, pp. 250–255. <https://aclanthology.org/S19-1027>.
- Yanaka, Hitomi, Koji Mineshima, Pascual Martínez-Gómez, and Daisuke Bekki (2018), Acquisition of phrase correspondences using natural deduction proofs, *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Association for Computational Linguistics, New Orleans, Louisiana, pp. 756–766. <https://aclanthology.org/N18-1069>.