

Veridicality and Neg-Raising

Will Gantt, Ben Kane, Ruoyang Hu

February 17, 2020

Keywords: natural language inference, veridicality, negation-raising, factivity

1 Research Problem & Background

The ability of humans to generate meaningful inferences from text or speech is essential to our use of language. Within the field of natural language processing, natural language inference (NLI) refers to the task of automatically determining entailments from a given sentence, i.e. facts that can be said to be true assuming that the given sentence was true. One important subset of inferences are *lexiosyntactic inferences*: inferences which are related to the interaction between specific lexical items and their syntactic context [5]. For example, consider the following four entailments:

1. Joe believes that Bo left. \nrightarrow Bo left.
2. Joe knows that Bo left. \leadsto Bo left.
3. Joe doesn't believe that Bo left. \leadsto Joe believes that Bo didn't leave.
4. Joe doesn't know that Bo left. \leadsto Bo left.

It should be evident from the above that the conditions under which certain inferences can be drawn are, in general, not obvious, and depend on interactions between lexical information and syntactic context.

Many lexical semantic distinctions have been discussed in the linguistics literature related to lexiosyntactic inferences. In this project, we are particularly interested in two potentially related phenomenon: *veridicality* and *negation-raising*.

The former refers to an assertion of the truth of some proposition. The sentence “Romeo loves Juliet” asserts Romeo’s affection for Juliet. Generally, a verb ‘V’ is veridical if ‘DP V S’ entails ‘S’ [4]. Assertions of falsehood are also considered “veridical”, but are of opposite polarity: “Romeo does not love Juliet” still makes a claim about Romeo’s affection for Juliet - namely, that it’s false. In the above examples, the successful inference in (2) shows that ‘know’ is veridical, whereas the failure of the inference in (1) shows that ‘believe’ is not veridical.

Negation-raising (“neg-raising”) is a phenomenon in which the negation of a predicate with a subordinate clause can be interpreted as though the negation were *in* the subordinate clause itself [1]. This is demonstrated by inference (3) in the examples above. By contrast, ‘know’ does not have neg-raising properties, since the inference licensed in (4) is “Bo left”, rather than “Joe knows that Bo didn’t leave”.

Clearly, it is crucial for any NLI system to be able to deal with these types of inferences in a way that is accurate and robust. This necessitates understanding the relationship between phenomenon such as neg-raising and veridicality on a finer level.

2 Research Question

Previous studies have attempted to study the relationship between neg-raising and properties of predicates or subordinate clause structure [1]. However, the specific relationships between veridicality and neg-raising have not yet been explored.

In this project, we aim to investigate the relationships between veridicality, factivity (a strengthening of the veridicality condition which requires that a verb presupposes its declarative complement), and neg-raising properties of classes of predicates. Specifically, we want to measure the distribution of neg-raising properties across predicates (focusing primarily on desire, belief, and emotive predicates, such as “want”, “think”, and “regret”, respectively), and create a model to predict veridicality inferences based on neg-raising properties.

Potential further directions include trying to analyze finer-grain distinctions within lexical classes. For instance, linguists have speculated “semi-factive” and “semi-neg-raising” verbs which exhibit various behavior in different linguistic contexts; for instance, “know” can act as a factive in some

contexts, but as a non-factive in other contexts.

3 Research Approach

3.1 Dataset

Aaron White (in the UR linguistics department) and others have compiled two corpora of English sentences - one annotated for veridicality and one annotated for neg-raising.^{1,2} These datasets contain 3938 and 7936 sentences respectively, and 773 and 925 unique predicates respectively.

Some additional data collection may be required, particularly in order to obtain data on “semi-neg-raising” verbs mentioned previously. Should this be necessary, the code used in [1] to collect neg-raising data is available, and only minor modifications should be necessary to collect the requisite data.

3.2 Model

Some previous studies on lexiosyntactic inferences have relied on a model of semantic selection (S-selection) originally put forward in [3]. The goal of this model is to induce the semantic type signatures of verbs (i.e. the types of arguments that the relation denoted by the verb can take) from the syntactic distribution of verbs (i.e. the syntactic frames in which the verb is judged to be acceptable), through a matrix factorization approach. The work in [1] extends this model with another module associating the semantic type of the verb with the inference patterns it licenses.

One other approach which may be promising is the use of a Combinatorial Categorical Grammar (CCG) based model. CCGs have been frequently used in linguistics to perform syntactic and semantic analysis. A CCG can be induced using tree-LSTMs [2], allowing for a highly structured model. Another alternative would be using a BERT-based language model. An initial obstacle for this project will be determining the best model to use for analysis based on structural considerations. Additionally, model implementation is likely to be time-consuming, though this may be somewhat mitigated by open-source libraries like PyTorch for neural models (if this is the direction we decide to take) and existing code from Aaron or his collaborators.

¹<http://megaattitude.io/projects/mega-negraising/>

²<http://megaattitude.io/projects/mega-veridicality>

References

- [1] Hannah Youngeun An and Aaron Steven White. The lexical and grammatical sources of neg-raising inferences, 2019.
- [2] Mike Lewis, Kenton Lee, and Luke Zettlemoyer. Lstm CCG parsing. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 221–231, San Diego, California, June 2016. Association for Computational Linguistics.
- [3] Aaron White and Kyle Rawlins. A computational model of s-selection. *Semantics and Linguistic Theory*, 26:641, 10 2016.
- [4] Aaron Steven White and Kyle Rawlins. The role of veridicality and factivity in clause selection. In *Proceedings of the 48th Annual Meeting of the North East Linguistic Society*, 2018.
- [5] Aaron Steven White, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. Lexicosyntactic inference in neural models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4717–4724, Brussels, Belgium, October-November 2018. Association for Computational Linguistics.