

Computing relative polarity for textual inference

Rowan Nairn, Cleo Condoravdi, Lauri Karttunen

Palo Alto Research Center

rnairn@gmail.com, condorav@parc.com, Lauri.Karttunen@parc.com

Abstract

Semantic relations between main and complement sentences are of great significance in any system of automatic data processing that depends on natural language. In this paper we present a strategy for detecting author commitment to the truth/falsity of complement clauses based on their syntactic type and on the meaning of their embedding predicate. We show that the implications of a predicate at an arbitrary depth of embedding about its complement clause depend on a globally determined notion of relative polarity. We, moreover, observe that different classes of complement-taking verbs have a different effect on the polarity of their complement clauses and that this effect depends recursively on their own embedding. A polarity propagation algorithm is presented as part of a general strategy of canonicalization of linguistically-based representations, with a view to minimizing the demands on the entailment and contradiction detection process.

1 Introduction

In a 1971 article titled “The Logic of English Predicate Complement Constructions” [9] Lauri Karttunen, 29, wrote:

It is evident that logical relations between main sentences and their complements are of great significance in any system of automatic data processing that depends on natural language. For this reason, the systematic study of such relations, of which this paper is an example, will certainly have a great practical value, in addition to what it may contribute to the theory of the semantics of natural languages.

It is only now that this 35-year old prediction is becoming a reality in the context of automated question answering and reasoning initiatives such as the PASCAL Textual Entailment Challenge (see [7]) and the ARDA-sponsored AQUAINT project (see [10], [12], [4]).

Recognizing whether a given piece of text can be strictly or plausibly inferred from, or is contradicted by, another piece of text is, arguably, a minimal

criterion for Natural Language Understanding (see [2]). We call this task LOCAL TEXTUAL INFERENCE. Textual inferences may be based on purely linguistic knowledge, assumptions about language use by collaborative rational agents, knowledge about the world, or any combination thereof. The semantics of complement constructions is an important part of local textual inference. It has the added advantage of carving out a well-circumscribed domain of inferences based primarily on linguistic knowledge.

A system that computes textual inferences should be able to deduce, for example, that (1b) and (1c) follow from (1a).

- (1) a. Ed forgot to close the door.
- b. Ed intended to close the door.
- c. Ed did not close the door.

There is a clear difference between the two embedding predicates *forget to* and *intend to*. (1c) does not follow from (1b). A speaker or author of (1b) may well BELIEVE in the truth of (1c) but he is not COMMITTED to it by virtue of having said (1b). In the following we focus on cases where the author's commitment to the truth of a complement clause arises solely from the larger sentence it belongs to, leaving aside other sources of information about the beliefs of the author. The author of (1a) is committed to both (1b) and (1c) but due to different aspects of the meaning of *forget to*, as we will show shortly.

The fact that forgetting to do something entails not doing it does not arise solely from the meaning of the verb *forget* but depends also on the type of its complement. Consider the difference between *forget to* and *forget that*.

- (2) a. Ed forgot that the door was closed.
- b. The door was closed.

(2a) commits the author to the view that the complement (2b) is true rather than false. Furthermore, with *forget that* this commitment is preserved under negation and in questions.

- (3) a. Ed did not forget that the door was closed.
- b. Did Ed forget that the door was closed?

(2a), (3a) and (3b) are alike in committing the speaker to (2b). The difference between *forget that* and *forget to* is striking.

- (4) a. Ed did not forget to close the door.
- b. Did Ed forget to close the door?

In contrast to (1a), in a narrative text (4a) commits the author to the view that Ed closed the door, the opposite of (1b).¹ (4b) is noncommittal either way.

The different semantic behaviors of *forget that* and *forget to* have been known for a long time. There is a large body of linguistic literature, start-

¹ In a spoken dialogue it is of course possible, typically with a special intonation, to use (4a) to contradict (1a): *Ed didn't "forget" to close the door. He never intended to do it.*

ing with Kiparsky & Kiparsky 1971 [11] and Karttunen 1971 [8], about FAC-TIVE constructions such as *forget/remember/know/... that* and IMPLICATIVE constructions such as *forget/remember/manage/bother/... to*. A common view is that factive constructions PRESUPPOSE rather than ENTAIL that the complement sentence is true.² Implicative constructions have entailments and some of them also carry presuppositions. For example, (1a) entails (1c) and presupposes (1b). (4a) carries the same presupposition as (1a) but the opposite entailment. While the entailments of implicative constructions are generally quite clear, it is often difficult to pin down exactly what is being presupposed. It may be argued, for example, that (1b) is too specific. Maybe the presupposition is more vague: *Ed ought to have closed the door or Ed was expected to close the door*. All the examples in (5) entail that Ed did not open the door but presuppose a different reason for this fact.

(5) Ed didn't manage/bother/dare/happen to open the door.

In this paper we focus on building a partial computational semantics for implicative constructions ignoring for the time being the presuppositional aspects of their meaning. However, we handle simple factive constructions and the interaction between implicative and factive verbs. The work was carried out in the context of the AQUAINT project using the XLE engine for parsing and semantic analysis.³ The AQUAINT project conducted a PASCAL-like experiment on local textual inferences based on a more nuanced task. Given a sentence *A*, we may conclude either that *B* is TRUE or that *B* is FALSE or that the answer is UNKNOWN, that is, *B* or its negation cannot be inferred from *A* alone. In contrast, the PASCAL test collapses FALSE and UNKNOWN into FALSE.⁴

We faced two initial challenges. The first is that there are several types of implicative verbs. Some yield an entailment in both affirmative and negative environments but there are others, “one-way implicatives”, that yield entailments only in one or the other environment. Furthermore, the entailment may be either positive or negative depending on the polarity of the environment. For example, *forget to* yields a negative entailment in a positive environment, (1a), and a positive entailment in a negative environment, (4a). But *manage to* works in the opposite way. This type of semantic information is not available in or deducible from any public lexical database such as WordNet, VerbNet or FrameNet. We had to compile ourselves a table of “implication signatures” for a large class of complement-taking constructions.

The second challenge is that implicative and factive constructions may be stacked together. The polarity of the environment of an embedding predicate is determined relatively to the chain of predicates or sentential operators it is in the scope of. Although it may not be obvious at the first glance, (6)

² This is not to say that there is a common view on how the notion of presupposition should be construed theoretically.

³ <http://www2.parc.com/istl/groups/nltt/xle/>

⁴ For a critical look at the PASCAL task, see Zaenen, Karttunen and Crouch [12].

commits the author to the view that Ed did not open the door.

(6) Ed didn't manage to remember to open the door.

In 6 *remember* is in a positive clause but the relative polarity of that clause is negative. The computation of relative polarity must be a recursive process.

2 Implication signatures

We focused on complement-taking verbs, especially those that take infinitival or *that* complements. Taking the verbs in order of decreasing frequency in the British National Corpus (BNC),⁵ we determined their natural implications (if any). Judgments were based on agreement by multiple annotators using resources such as Google search and the Linguist's Search Engine to sample the relevant constructions in the wild. In particular cases it can be difficult to decide between ENTAILMENTS, that is, what the author is actually committed to, and CONVERSATIONAL IMPLICATURES, that is, what a reader/hearer may feel entitled to infer. For example, *Ed did not refuse to participate* might lead the hearer to conclude that Ed participated. But the speaker could continue with *He was not even eligible* indicating the opposite. For this reason we classify *refuse to* as a one-way implicative. Of the 1250 relevant verbs in our lexicon we classified 400 on a first pass. Roughly a third of those carried some kind of implication: a positive or negative entailment, a factive or a counterfactive presupposition. Conversational implicatures were flagged for later attention. Figure 1 shows the classifications of the resulting lookup table.

	Word in subcat frame	Relative Polarity	
		(+) positive	(-) negative
		Entailment	
Two-way implicatives	<i>manage to</i> <i>forget to</i>	(+) positive (-) negative	(-) negative (+) positive
One-way +implicatives	<i>force to</i> <i>refuse to</i>	(+) positive (-) negative	none none
One-way -implicatives	<i>attempt to</i> <i>hesitate to</i>	none none	(-) negative (+) positive
		Presupposition	
Factives Counterfactives	<i>forget that</i> <i>pretend that</i>	(+) positive (-) negative	(+) positive (-) negative
Neutral	<i>want to</i>	Entailment/Presupposition	
		none	none

Fig. 1. Some examples from our verb markup table

⁵ <http://www.natcorp.ox.ac.uk/>

3 Theoretical and technical prerequisites

Our approach to textual inference relies on parsed text that is further transformed by a process of `CANONICALIZATION`. The mechanism for entailment and contradiction detection (`ECD`) combines structural matching and inference-based techniques. It operates on packed representations, encoding ambiguities, without the need for disambiguation. We will not discuss `ECD` any further here. Instead we will focus on describing in more detail some of the relevant features of the representations on which it operates.

Input text is syntactically analyzed by the `XLE` parser, based on a broad coverage, hand-coded grammar of English. Linguistic semantic representations are constructed from the parse output, using `SKOLEMIZATION` and flattening embedded structures to clausal form. These logical forms are in turn canonicalized to more uniform representations via packed term rewriting as described in Crouch [3]. The implication projection algorithm to be described in the next section forms part of this component of canonicalization and is implemented as a set of recursive rewrite rules that operate on packed representations.⁶

The canonicalized representations that are input to `ECD` are essentially a kind of description logic with contexts.⁷ Roughly, each verbal predication corresponds to a constructed concept, an event type with role restrictions. The main concept is provided by a mapping of the verbal predicate to a concept in some background ontology. The role restrictions come from various arguments and modifiers. The constructed concept is named by the skolem introduced by the verbal predicate. Flattening replaces embedded expressions with complex internal structure, such as clausal complements, with atomic first order terms, contexts. The information about the level of embedding of an expression is preserved by associating its content with the corresponding context. Negation and intensional operators also trigger the introduction of new contexts. Contexts thus serve as scope markers since their use enables globally represented information, such as the scope of operators, to be made locally accessible.

The content of the top level context, designated as `t`, represents what the author of the sentence is taken to be committed to. In general, we tie truth of a sentence to the `INSTANTIABILITY` of the skolem corresponding to its head predicate. This, in effect, amounts to the familiar existential closure over events: if the skolem corresponding to a clause’s head predicate denotes an event description, an instantiability declaration for that skolem means that the event description is instantiated. Therefore, an implication that a complement clause is true/false can be construed as an existential/negative existential implication, which in our terms is an implication about the instantiation/non-instantiation of the event type described by the embedded clause.

⁶ Packing is `XLE`’s mechanism for ambiguity management and operates independently of canonicalization and inference.

⁷ For more details see Bobrow *et al.* [1], Crouch [3] and Condoravdi *et al.* [2].

Instantiability is always relative to a context, in the simplest case the context of origin of the skolem. In order to become author commitment, an instantiability declaration has to be associated with the top level context τ . When two contexts stand in certain relations to one another, in particular the relations of veridicality and antiveridicality, information can be inherited from one to another. Lifting rules lift assertions from a lower context to a higher context, either as they are, when the two contexts are veridical to one another, or by switching the polarity of instantiability assertions, when the two contexts stand in an antiveridical relation. Negation introduces a context that is antiveridical with respect to the immediately higher context. To illustrate, (7) gives the contextual structure for a negative sentence like *Ed didn't leave Paris* and (8) the corresponding instantiability assertions (*leave_ev57* is the name for the constructed event type of Ed leaving Paris). One important thing to note is that the assertion `instantiable(leave_ev57)` in `not58` is lifted as `uninstantiable(leave_ev57)` to the top level context τ , thus capturing the intuitive meaning that the event type of Ed leaving Paris was not instantiated.

- (7) `context(τ)`
 `context(not58)` *new context triggered by negation*
 `context_relation(not τ not58)`
 `antiveridical(not58 τ)` *interpretation of negation*
- (8) `not58: instantiable(leave_ev57)`
 `τ : uninstantiable (leave_ev57)` *entailment of negation*

Lexical entailments and presuppositions are similarly overtly spelled out in the representations operated on by ECD. This way the process of canonicalization prepackages some of the local textual inferences. The challenge of course is to figure out which context the relevant instantiability assertions ought to be lifted to, which is what the implication projection algorithm determines.

4 The implication projection algorithm

Aside from the onerous task of classifying hundreds of verbs, the complications of this problem stem from the interaction of multiple embedded clauses. As mentioned previously, the entailment yielded by a complement-taking construction is dependent on the polarity of the context it appears in. This polarity in turn is not locally determined but dependent on the embedding structure of contexts. Therefore, a verb in a negative clause is not necessarily in a negative environment since the negativity of a *not* may be neutralized by another negative, as for example in (9).

- (9) Ed refused not to attempt to leave.

Here the normal negative entailment licensed by *not attempt* is neutralized by the negative polarity setting due to the higher-level predicate *refuse*. Notice

that *refuse* does not simply negate the entailment. It cancels it entirely. Embedding within a verb such as *refuse* can also license entailments that were not available previously. Consider (10a), which is compatible with either (10b) or (10c).

- (10) a. Ed attempted to leave.
 b. Ed left.
 c. Ed didn't leave.

(11), on the other hand, implies (10c).

- (11) Ed refused to attempt to leave.

Evidently, it is not enough to look at the immediate outer context of a complement construction. The polarity of any context depends on the sequence of potential polarity switches stretching back to the top context. Each complement-taking verb, operating on its parent context's polarity, either switches, preserves or simply sets the polarity for its embedded context, as specified by an entry in the lookup table.

Furthermore, this means that polarity is a relative notion. If the sequence of polarity switches was started at a level below the top context then the final polarity value might turn out different. Thus when we talk about the polarity of a context we mean polarity relative to some ancestor context. Normally, it is the top context which interests us the most, but it may be useful to infer the implications of a clause for other contexts. For example, it is probably useful to infer (12b) from (12a). The algorithm provides for this generality.

- (12) a. John believes that Ed managed to leave.
 b. John believes that Ed left.

Every context C then has associated with it a set of ancestor contexts relative to which its polarity is positive (denoted \oplus_C) and a set of contexts relative to which its polarity is negative (denoted \ominus_C). Every context, including the top one, is positive relative to itself. The polarity sets of a context are computed in terms of its parent's sets ($\oplus_{p(C)}$ and $\ominus_{p(C)}$) with reference to the verb ($V_{p(C),C}$) which links the two contexts and its signature in the lookup table ($sig^e(V_{p(C),C})$) where the environment superscript e is either positive $+$ or negative $-$.

$$\oplus_C =_{def} \{C\} \cup \begin{cases} \oplus_{p(C)} & \text{if } sig^+(V_{p(C),C}) = + \\ \ominus_{p(C)} & \text{if } sig^-(V_{p(C),C}) = + \\ \emptyset & \text{otherwise} \end{cases}$$

$$\ominus_C =_{def} \begin{cases} \oplus_{p(C)} & \text{if } sig^+(V_{p(C),C}) = - \\ \ominus_{p(C)} & \text{if } sig^-(V_{p(C),C}) = - \\ \emptyset & \text{otherwise} \end{cases}$$

Figure 2 shows the example sentence *Ed did not forget to force Dave to leave* parsed and with relative polarities assigned to each context. To get to this

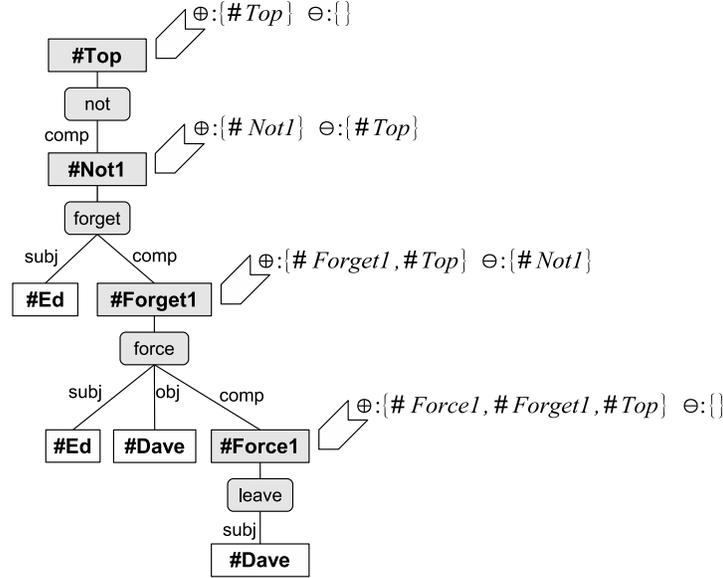


Fig. 2. After the polarity propagation pass

situation the algorithm first assigns the top context the polarity sets $\{\#Top\}$ and \emptyset . It then recursively computes the polarity sets for each embedded context using the context-linking verb as an index to the lookup table. *Not* is treated in the same way as *forget to* – they both invert the polarity sets. *Force* is a one-way implicative that disregards the negative polarity set of its parent.

Recall that we needed to work out which concepts should be instantiated in which contexts and, now that we have marked the contexts appropriately with relative polarities, we can extract that information. The head event skolem of a context, and presumably all its role fillers, should be made instantiable not only in the context it arises in but also in all contexts relative to which its originating context has positive polarity. Similarly, an event should be made uninstan- tiable in all contexts relative to which its originating context has negative polarity.

$$\begin{aligned} \text{instantiables}(C) &=_{\text{def}} \{\text{head}(C') \mid C \in \oplus_{C'}\} \\ \text{uninstantiables}(C) &=_{\text{def}} \{\text{head}(C') \mid C \in \ominus_{C'}\} \end{aligned}$$

From the polarity marking in Figure 2 we can conclude that the event concept corresponding to the sentence *Dave left* is in fact instantiable at the top level (as well as in the $\#Force$ and $\#Forget$ contexts) and thus we can attribute it as a commitment of the speaker.

5 Conclusion and Further Work

The present study is, as far as we know, the first systematic implementation of textual inferences arising from the six types of implicative verbs presented in Figure 1 and their interaction with factive verbs.

In this work we have focused on cases where the judgement of whether the author is committed to the truth or the falsity of a complement clause can be made reliably from the sentence in question. Further work is needed at least in the following three areas.

Lexicographic gaps. In our classification we only considered simple verbal and adjectival complements. We have yet to study and determine the semantics of complement constructions associated with nominals in collocations such as *take the trouble to*, *have the foresight to*, *take time to*, for which there is virtually no literature.

Conversational implicatures. It is well known that constructions such as *be able to* yield a negative entailment in a negative environment. *Ed was not able to open the door* entails *Ed did not open the door*. There is no entailment in the corresponding affirmative sentence. Yet, if the author writes *Ed was able to open the door* and says nothing to indicate that the door was not opened, the reader is likely to infer, and justifiably so, that Ed opened the door. This kind of CONVERSATIONAL IMPLICATURE is cancelable (Grice [6]). It is not a contradiction to say *Ed was able to open the door but he kept it closed*. If a student asks his professor *Did you have the time to read my paper?* and the professor answers *Yes* but has not read the paper, the answer can be literally true and very misleading at the same time.⁸

Degrees of “factivity”. Factive verbs and constructions do not constitute a uniform class. Looking at the pattern of usage of verbs such as *mention that*, *report that*, *say that*, etc. on Google, we observed that in cases such as *He did not mention that Coalition allies now plan to leave* it was virtually always clear from the context that the author believed the complement to be true. The verb *report* is similar to *mention* but there are also cases where *...did not report that X* was meant to suggest that *X* is false. On the other hand, *...did not deny that X* suggests that *X* is true, whereas *...denied that X* is noncommittal with respect to *X*.

Acknowledgements

This material is based in part on work funded by the U.S. Government, and any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Government.

⁸ For a seminal paper on INVITED INFERENCES, see [5].

References

- [1] Bobrow, D., C. Condoravdi, R. Crouch, R. Kaplan, L. Karttunen, T. King, V. de Paiva and A. Zaenen, *A basic logic for textual inference*, in: *Proceedings of the AAAI Workshop on Inference for Textual Question Answering*, Pittsburgh, PA, 2005, <http://www2.parc.com/istl/groups/nlitt/papers/textual-inference.pdf>.
- [2] Condoravdi, C., R. Crouch, R. Stolle, V. de Paiva and D. Bobrow, *Entailment, intensionality and text understanding*, in: *Proceedings of the Workshop on Text Meaning, Human Language Technology Conference (HLT-NAACL-2003)*, Edmonton, Canada, 2003, <http://www2.parc.com/spl/members/stolle/Papers/condoravdi-textmeaning.pdf>.
- [3] Crouch, R., *Packed rewriting for mapping semantics to KR*, in: *Proceedings of the Sixth International Workshop on Computational Semantics*, Tilburg, the Netherlands, 2005, http://www2.parc.com/istl/groups/nlitt/papers/iwcs05_crouch.pdf.
- [4] Crouch, R., R. Sauri and A. Fowler, *AQUAINT pilot knowledge-based evaluation: Annotation guidelines* (2005), http://www2.parc.com/istl/groups/nlitt/papers/aquaint_kb_pilot_evaluation_guide.pdf.
- [5] Geis, M. and A. Zwicky, *On invited inferences*, *Linguistic Inquiry* **2** (1971), pp. 561–565.
- [6] Grice, H. P., *Logic and conversation*, in: P. Cole and J. L. Morgan, editors, *Speech Acts*, Academic Press, New York, NY, 1989 pp. 41–58.
- [7] Ido Dagan, O. G. and B. Magnini, *The PASCAL recognising textual entailment challenge*, in: *Proceedings of the PASCAL Challenges Workshop on Recognising Textual Entailment*, Southampton, U.K., 2005, http://www.cs.biu.ac.il/~glikmao/rte05/dagan_et_al.pdf.
- [8] Karttunen, L., *Implicative verbs*, *Language* **47** (1971), pp. 340–358.
- [9] Karttunen, L., *The logic of English predicate complement constructions* (1971), distributed by the Indiana University Linguistics Club. http://www2.parc.com/istl/members/karttune/publications/english_predicate.pdf.
- [10] Karttunen, L. and A. Zaenen, *Veridicity*, in: G. Katz, J. Pustejovsky and F. Schilder, editors, *Annotating, Extracting and Reasoning about Time and Events*, number 05151 in Dagstuhl Seminar Proceedings (2005), <http://drops.dagstuhl.de/opus/volltexte/2005/314>.
- [11] Kiparsky, P. and C. Kiparsky, *Fact*, in: D. Steinberg and L. Jakobovits, editors, *Semantics. An Interdisciplinary Reader*, Cambridge University Press, Cambridge, England, 1971 .
- [12] Zaenen, A., L. Karttunen and R. Crouch, *Local textual inference: can it be defined or circumscribed?*, in: *Workshop on the Empirical Modeling of Semantic Equivalence and Entailment*, Ann Arbor, MI, 2005, <http://www2.parc.com/istl/members/karttune/publications/acl2005workshop.pdf>.