

Annotating Factive Verbs

Alvin Grissom II, Yusuke Miyao

National Institute of Informatics, Japan
Hitotsubashi 2-1-2, Chiyoda-ku, Tokyo 101-8430 JAPAN
acgrissom@gmail.com, yusuke@nii.ac.jp

Abstract

We have created a scheme for annotating corpora designed to capture relevant aspects of factivity in verb-complement constructions. Factivity constructions are a well-known linguistic phenomenon that embed presuppositions about the state of the world into a clause. These embedded presuppositions provide implicit information about facts assumed to be true in the world, and are thus potentially valuable in areas of research such as textual entailment. We attempt to address both clear-cut cases of factivity and non-factivity, as well as account for the fluidity and ambiguous nature of some realizations of this construction. Our extensible scheme is designed to account for distinctions between claims, performatives, atypical uses of factivity, and the authority of the one making the utterance. We introduce a simple XML-based syntax for the annotation of factive verbs and clauses, in order to capture this information. We also provide an analysis of the issues which led to these annotative decisions, in the hope that these analyses will be beneficial to those dealing with factivity in a practical context.

Keywords: factivity, annotation, semantics

1. Introduction

This paper introduces our work on building an annotated corpus which captures the linguistic notion of *factivity* and addresses some of the challenges inherent to the task.

Presuppositions are facts assumed to be true at the time of an utterance, and factivity is a kind of presupposition that a part of an utterance is presupposed to be true. There are many ways in which factivity can be realized; here, we focus specifically on *factive verbs*, which are verbs that take a clause as a complement and introduce a presupposition that the complement clause is true.

Inherently, factivity is an indispensable factor for the recognition of textual entailment, factuality, etc. However, to our knowledge, there are no corpora which annotate factivity in quite this way, while there are other kinds of semantically annotated corpora, such as PropBank and FactBank. We devise an annotation scheme which captures the relationship between a factive verb, its complement clause, and the degree to which it is factive. This work constitutes a first step towards the recognition of various kinds of factivity in real-world texts.

2. Problem

In the field of pragmatics, the notion of presupposition (Kiparsky, 1970; Beaver, 2001; Gazdar, 1979) has been studied extensively in linguistics literature. Presupposition is defined as a fact which is assumed to be true at the time of an utterance, but not explicitly mentioned. In natural language text, this information is implicit in the text, but not explicitly enumerated. Extracting presuppositions from text allows for the extraction of implicit information, potentially valuable for textual entailment systems.

- (1) He reached the third door on the left.
Presupposition: There is a first door and a second door on the left.
- (2) Cain does not work at a bank anymore.
Presupposition: Cain once worked at a bank.

Presuppositions differ from other kinds of entailments in some interesting ways; for instance, presuppositions are believed to be blocked by certain linguistic phenomena, such as negation, though there are known complications (Heim, 1990).

This work is focused on factive verbs and their complements, a specific type of presupposition. When verbs function factively, they presuppose the truth of the complement of the verb.

- (3)
 - a. I regret [that] I came to school late.
 - b. I came to school late, I regret.
⇒ I came to school late.
- (4) I propose [that] I came to school late.

In example (3) a., the complement *I came to school late* is assumed to be true. Regardless of whether the person speaking is truthful, the presupposition is projected lexically by *regret*. This is also true of the syntactic variant, (3) b. The same cannot be said of (4), however, since something proposed is not necessarily something done.

In the computational linguistics literature, there has not, to our knowledge, been much concentrated research into the notion of factive verbs and factivity annotation or classification. Factivity is largely determined lexically. For example, some verbs, such as emotive verbs (e.g., love, hate, regret) are nearly always factive (Giannakidou, 2006). However, a complication in this task is that factivity is sometimes not only lexically determined but also dependent on its syntactic realization, context, etc. Indeed, we have observed that, while there appears to be a factivity continuum – that is, some verbs are almost always used factively, while others never are – verbs may be coerced in certain situations to behave in counterintuitive ways, as is noted in the literature (Nairn et al., 2006).

Complicating annotation efforts somewhat are ambiguities.

- (5) I'm afraid I don't know.

(6) I'm afraid that I might be wrong.

A moment's reflection will reveal that (5) is at least semantically two-ways ambiguous: either (a) I am afraid *of the fact that* I don't know, or (b) I am afraid of the *possibility* that I don't know. In the former case, *I don't know* is factive; in the latter case, it is not. One might postulate the same ambiguity for (6), though pragmatic considerations make the former interpretation more strained without context. In one case, *afraid* functions as sort of surrogate for *regret*, making the following clause factive; in the other, *afraid* functions also as a verb of *uncertainty*, and thus is non-factive. One can also imagine explanations both lexical and syntactic for such ambiguity, but we shall deem the verb factive if the clause following it is factive in a given situation. In some instances, this cannot be determined without context. We allot a tag to mark this. Now, consider the examples of *coercion*.

(7) I remember [that] John came home at 12:00.

(8) I remember John coming home at 12:00, but he actually came home at 1:00.

In (8), *remember* is coerced into being non-factive by the contrastive conjunction which follows. In most circumstances, as in (7), *remember* is factive, but we cannot assume that every instance of *remember* is indicative of factivity. We allocate an optional tag for coercion, for those (probably rare) cases in which it is glaringly obvious that a verb which would typically not be factive is behaving in such a way (or vice versa), due to some unusual phenomenon.

There is, in addition, a special class of such verbs, the factuality of which depend upon the assumptions of the one making the determination. This class of verbs, which we shall simply call *claim* verbs, make explicit claims. The factivity of them is not obvious. Consider the following.

(9) The government announced that it would support the treaty.

(10) The government announced that it would achieve world peace.

(11) I'm afraid [that] I didn't go.

(12) John admitted that he cheated.

(13) The teacher informed me of John's misbehavior.

Some of these, especially *admit* and the specialized use of *afraid*, are difficult to categorize. Working with our original definition of factivity, independent of the factuality of the statement, *admit* seems to be factive, though there are situations in which it may not be. In such a case, it should be annotated both as factive and as a claim. Others, such as *announce*, are similarly difficult, since *announce* has the sense of *revaling* and the sense of simply stating something which may or may not be true. They should be judged based on their context and marked as ambiguous if they remain ambiguous even with the context. We expect that enough such annotations can provide the data for determining the factivity based on contextual information.

There are some verbs which change the state of the world by virtue of their very utterance (Searle, 1969). These are called performatives or speech acts, and, according to (Austin, 1975), they are neither true or false, but rather felicitous or infelicitous. Before the utterance, the complement clause is untrue; after the utterance, it is true.

(14) John [hereby] requires that all people must pay taxes.

(15) We, the United States, [hereby] declare the treaty null.

(16) The United States declared the treaty null.

(17) # Bob declared the treaty null.

(18) I emphasize that John has been here for three months.

(19) I [hereby] pronounce you man and wife.

(20) I [hereby] apologize for that break of contract.

There is a relevant difference between the performatives in 18 and 20, versus those in 15 and 19. The former two are nearly always factive, because they presuppose the truth of the complement clause at the time of utterance; the latter are, strictly speaking, non-factive, since the complement clause is only true *after* the utterance by an entity with the requisite authority. Note that while (16) is not performative, it differs from a performative sentence only in the tense of the verb and the authority of the agent. We shall deem past tense descriptions of prior performatives, such as 16, to be factive. The complement clauses in these cases are true (or were true at the point of reference, as with all past tense factives). Utterance 17 is not a performative because Bob does not have the authority to declare a treaty null. It is non-factive. In the case of non-factive performatives, such as 15, before the utterance, the complement clause is false; after the utterance, the complement clause is true. While they contain interesting information in and of themselves, they are, strictly speaking, not factive, though they are factual statements. When annotating, focusing on the verb itself, or on the entire sentence, can lead to inconsistent and misleading determinations. We therefore emphasize that an annotator, especially in the case of performatives, should focus on the complement clause and, in particular, whether this clause is true at the time of the utterance, to avoid confusion. This strategy should be effective in avoiding most of the philosophical complexities of utterances which make themselves true by virtue of being uttered.

3. Corpus Annotation

We adopt a relatively simple XML-based annotation scheme which marks not only the factivity of a verb and its clause, but also the lemma, whether the verb has been nominalized, and whether the factivity is ambiguous. The factivity is ambiguous when it is not obvious whether or not it is factive. This will often be due to an unclear sense of the term, even in context.

For our initial work, we used the PASCAL Recognizing Textual Entailment 2 (RTE2) textual entailment corpus (Bar-Haim et al., 2006), and augmented it with our annotation scheme for factivity. We believe that this may aid in

Table 1: XML Tags for annotating factive verbs

<fcue>	marks factive verb
<f>	marks factive clause
<nfcue>	marks non-factive verb
<nf>	marks a non-factive clause
<ambig> (opt.)	sense slightly ambiguous
<scue>	marks claim verb
<c>	marks claim
<break>	indicates break in sentence
<coercion> (opt.)	unusual (non)-factivity
<agent> (opt.)	identifies the agent

research into the relationship between factivity and textual entailment.

In our current work, we are only concerned with verb-complement pairs. We are not annotating, at this point, adjectives, adverbs, or other kinds of presupposition. We provide an identical scheme for the annotation of non-factive verbs, as well, since they may be embedded in a factive clause or vice versa.

The tags employed are shown in Table 1. The <fcue> tag is used to explicitly mark verbs which are functioning factively. Consider the following example from the RTE2 data set.

```
<t><agent type="authority">
East Timor's Prime Minister
</agent>
<ambig>
<scue id="8">says</scue>
</ambig>
<c ref="8">he<ncue id="9">expects</ncue>
<nf ref="9">within two months to
announce a deal with Australia on
developing a disputed oil and
gas field</nf></c>.</t>
<h>East Timor expects a deal with Australia.</h>
```

(Note that the <h> tag is from RTE and refers to the hypothesis of the sentence.) In this example, we have the minimum number of attributes, *id* and *ref*, to link clauses to verbs. A more complete version would have lemmas. Since *say* is obviously specifying a claim, it is marked as such. Some might argue that, since a Prime Minister is an authority figure, this should be treated as factive, but we have chosen to distinguish between the two, allowing for both tags. We mark *says* as ambiguous, because there seems to be authority to the claim, functioning as the verb *announce*, which is one of the more consistently ambiguous verbs. Indeed, in the RTE2 corpus, the given hypothesis assumes that what the Prime Minister says is factual. This is a difficult issue. In this example, we did not annotate *announce*, because it is not functioning factively in context.

```
<agent type="authority">
Mr. Fitzgerald
```

```
</agent>
<fcue lemma="reveal" id="1" form="v">revealed
</fcue>
<f ref="1">he was one of several
top officials who
<scue lemma="tell" id="2">told
</scue> Mr. Libby in June 2003 that
<c ref="2">Valerie Plame, wife of the
former ambassador Joseph Wilson,
worked for the CIA.</c></f>
```

Here, the <fcue> tag marks the factive verb, *reveal*. This verb has a unique identifier in the document – in this case, 1 – and it appears in the form of a verb; hence, *form*="v". The following tag is the <f> tag, which marks a factive clause. In our annotation, it should always have a corresponding <fcue>. The *ref*="1" indicates that there is a link between the *fcue* with *id* 1 and the factive clause. In this case, we also have an embedded clause involving an explicit claim, making use of the *told*. In this XML-style syntax, this is captured, and we can separate the clauses from each other.

At times, the determination may be difficult, as in the following example from RTE2.

```
<agent type="authority">
Ukraine's election commission
</agent>
<fcue lemma="declare" form="v" id="1">
declared
</fcue>
<f ref="1"> the Kremlin-backed prime minister,
Viktor Yanukovich, the winner of the country's
bitterly disputed presidential election</f>,
sharpening a crisis sparked by the opposition
candidate's allegations that the vote was
fraudulent.
```

In this case, *declare* is being used in a speech act. Recall that speech acts cause the state of the world to change by virtue their utterance (Austin, 1975). For our purposes, speech acts will be considered factive, but there is nothing preventing one from tagging such a verb as both a claim and as factive. Note that we have not annotated *disputed*, since, for the purposes of this paper, we limit ourselves to verb-complement structures.

In some cases, there may be a disjoint clause as in the following sentence:

- (21) The problem, he revealed, would not be resolved simply.

In this example, the clause is separated by *he revealed*. This common phenomenon is annotated with the <break> tag.

```
<f ref="4">The problem,
<break>he<fcue lemma="reveal" id="4">
revealed
</fcue>
</break>, would not be resolved simply.</f>
```

In the RTE2 test set, we annotated 127 instances of verbs that specify claims (e.g., *say*, *announce*, etc.), 81 instances

of factivity, and 17 instances of non-factivity (e.g., *hope*, *think*), out of the 800 test-hypothesis pairs.

To determine ambiguous cases, such as coercion, we are currently relying on intuition, but we believe that there may be ways of quantifying this, described in the section on future work. In addition, there are syntactic methods for detecting certain kinds of ambiguity, but they are beyond the scope of this paper. At present, tags marking ambiguity are essentially warnings, without much specific information. Since our scheme is XML-based, it is extensible without compromising backward compatibility, and some of the details of the attributes may vary based on what is deemed useful for a particular corpus, such as the broad category of “claim.” Based on our experience with the RTE-2 corpus, in general, the kinds of claims that warrant special attention with regard to factivity are, at least, *speech acts* and *authoritative claims*.

4. Related Work

There is a natural relationship between factivity and event extraction, as often, but not always, facts presupposed explicitly refer to events. Pustejovsky et al. (Pustejovsky et al., 2003) discuss *TimeBank*, which annotates events, times, and temporal relations, and events may be labeled for factivity. *FactBank* (Saurí and Pustejovsky, 2009), an additional layer atop *TimeBank*, annotates “factuality” of the clauses. The key difference between factuality and factivity is this: factivity refers to a fact that is presupposed in an utterance, independent of the veracity of the utterance in the real world; factuality refers such veracity. While factive usages do generally betray a certain level of commitment to the factuality of a proposition, we are not directly concerned with factuality. Prasad et al. (Prasad et al., 2006) describe an annotation scheme for “abstract objects” (e.g., propositions) in the Penn Discourse Treebank, designed to capture, among other things, the degree of factuality of the objects.

Factivity has been considered an important factor in textual entailment recognition, although this problem was not studied intensively in natural language processing literature. Nairn et al. (Nairn et al., 2006) presented a detailed analysis of the notion of “relative polarity,” which concerns the relationship between the factuality of a matrix clause and its complement clause. This literature discusses factive verbs as a part of relative polarity, although it is primarily concerned with theoretical considerations and does not provide resources or tools for factive verbs. Clausen and Manning (Clausen and Manning, 2009) used lexical factivity triggers in a tree structure for natural language inference. They note that lexically triggered presuppositions can expand the information available for solving a particular problem without deep semantic interpretation. Their goal was neither annotation nor recognition of factivity, but inference based on prior knowledge of the factivity or antifactivity of a given verb. We believe that proper annotation of factivity, and presupposition in general, will benefit these areas of research. Note that (Garoufi, 2007), in a study of textual entailment datasets, found that counter-factivity occurred very infrequently. We do not address it in our current annotation scheme; it is subsumed under the

notion of non-factivity.

5. Conclusion

We have devised a scheme for annotating factive verbs in real-world texts. Sufficiently developed corpora should allow natural language processing researchers to develop new means of factivity detection and allow linguists to research these phenomena quantitatively. Following this scheme, the next step will be to account for other kinds of factivity, beyond verbs and their complements.

6. Future Work

We would like to use this for textual entailment, as the behavior of the entailments is predictable if one can determine the (non)factivity of an expression. In addition, we have done some preliminary work on automatically identifying factive verbs from a large corpus. We used UKWAC (Ferraresi et al., 2008), a 2 billion word English corpus from the web, to search for certain syntactic patterns that we believed were likely to yield a (non)factive verb. To do this, we created a program to search for syntactic patterns in a large corpus, and computed the odds that each verb would occur in the “factive” or “non-factive” category. The criteria used for determining these syntactic patterns are beyond the scope of this paper, but the following shows some fundamental idea. By searching, for example, for a past participle, followed by any verb, and then the word *that*, we found that the verb in question was often factive.

- VHZ/have VVN/**report** IN/that
- VHZ/have VVN/**notice** IN/that
- VHZ/have VVN/**suggest** DT/that/
- VHZ/have VVN/**conclude** IN/that

Likewise, the following are also from our “factive” syntactic category:

- NN/industry MD/should VV/**recognise** IN/that/
- NP/Microsoft MD/should/ VV/**acknowledge** IN/that/
- NN/paragraph/ MD/should/ VB/**be** IN/that/
- NN/evidence MD/should/ VV/**show** IN/that/

The verbs yielded are not perfect – nor should be expect them to be – but they do tend to skew toward factivity or non-factivity. We have begun experimenting with measuring these occurrences of verbs in these syntactic frames statistically by making assumptions concerning which categories are factive and which are not, as we believe that interesting and potentially useful linguistic insights may be gleaned from such an analysis.

7. References

J.L. Austin. 1975. *How to do things with words*. Harvard Univ Pr.

- R. Bar-Haim, I. Dagan, B. Dolan, L. Ferro, D. Giampiccolo, B. Magnini, and I. Szpektor. 2006. The second PASCAL recognising textual entailment challenge. In *Proceedings of the second PASCAL challenges workshop on recognising textual entailment*, pages 1–9. Citeseer.
- D.I. Beaver. 2001. *Presupposition and assertion in dynamic semantics*. CSLI publications.
- D. Clausen and C.D. Manning. 2009. Presupposed content and entailments in natural language inference. In *Proceedings of the 2009 Workshop on Applied Textual Inference*, pages 70–73. Association for Computational Linguistics.
- A. Ferraresi, E. Zanchetta, M. Baroni, and S. Bernardini. 2008. Introducing and evaluating ukwac, a very large web-derived corpus of english. In *Proceedings of the 4th Web as Corpus Workshop (WAC-4) Can we beat Google*, pages 47–54.
- K. Garoufi. 2007. Towards a better understanding of applied textual entailment.
- G. Gazdar. 1979. *Pragmatics: Implicature, presupposition, and logical form*. Academic Press New York.
- A. Giannakidou. 2006. Only, emotive factive verbs, and the dual nature of polarity dependency. *Language*, 82(3):575–603.
- I. Heim. 1990. Presupposition projection. In *Reader for the Nijmegen Workshop on Presupposition, Lexical Meaning, and Discourse Processes, Univ. of Nijmegen*.
- P. Kiparsky. 1970. Semantic rules in grammar. *The Nordic languages and modern linguistics*, pages 262–285.
- R. Nairn, C. Condoravdi, and L. Karttunen. 2006. Computing relative polarity for textual inference. In *Proceedings of ICoS-5*. Citeseer.
- R. Prasad, N. Dinesh, A. Lee, A. Joshi, and B. Webber. 2006. Attribution and its annotation in the Penn Discourse Treebank. *Traitement Automatique des Langues*, 47(2):43–63.
- J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizauskas, A. Setzer, D. Radev, B. Sundheim, D. Day, L. Ferro, et al. 2003. The TimeBank corpus. In *Corpus Linguistics*, volume 2003, page 40.
- R. Sauri and J. Pustejovsky. 2009. FactBank: a corpus annotated with event factuality. *Language Resources and Evaluation*, 43(3):227–268.
- J.R. Searle. 1969. *Speech acts: An essay in the philosophy of language*. Cambridge university press.