

Determining Modality and Factuality for Text Entailment

Roser Saurí
Lab for Linguistics and Computation
Computer Science Department
Brandeis University
roserr@cs.brandeis.edu

James Pustejovsky
Lab for Linguistics and Computation
Computer Science Department
Brandeis University
jamesp@cs.brandeis.edu

Abstract

Recognizing textual entailment (TE) is a complex task involving knowledge from many different sources. One major source of information in this task is event factuality, since the inferences derivable from factual eventualities are different from those judged as possible or as non-existent. Some TE systems already factor in factuality features at the local level, but determining the factuality of events more generally involves dealing with information that is non-local to a particular textual event. In this paper, we present a tool providing events with their factuality values, characterized as pairs of modality and polarity features. In previous work, we identified polarity and modality at the local context with a performance of 92% precision and 56% recall. The research presented here extends and enhances our algorithm to incorporate the influence of non-local context as well as the identification of sources.

1 Introduction

Recognizing textual entailment is a complex task involving knowledge from many different sources. Beyond such basic steps such as predicate argument identification or anaphora resolution, it requires the use of lexicons populated with typing and subtyping information; monotonic behaviour; as well as access to temporal, spatial, and causal relations and their respective reasoning components.

An additional basic component for this task is determining the factuality of events. In many cases, the availability of rich resources and reasoning systems such as those mentioned above will not suffice, if the factual status of events in the text is not known. Consider (1):¹

¹Example from the 3rd Recognising Textual Entailment PASCAL challenge, RTE 2007. The task consists in determining whether a segment provided as text (T) entails a second segment H, the hypothesis. Segment T' in (2) is a modification of the original T.

(1) **T**: A senior Russian politician has hailed a decision by Uzbekistan to shut down a United States military base.

H: Uzbekistan closes United States military base. [Id=180]

(2) **T'**: A senior Russian politician hailed the refusal by Uzbekistan to shut down a United States military base.

A graph matching-based system informed with an adequate lexicon seems enough to identify the equivalence of *shut down* (in T) and *close* (in H). But the same approach would predict a wrong entailment given the alternative text in (2), where the shut down event is presented as a “counterfact”. That is, inferences derivable from factual events are different from those judged as possible or non-existent.

Event *factuality* is understood here as the category conveying whether event-denoting expressions are presented as corresponding to real situations in the world, situations that have not happened, or situations of uncertain status. As such, it is impacted by values from the systems of *polarity* (positive vs. negative) and *modality* (degree of certainty of what is asserted; e.g., possible vs. certain).

Recent work from different areas in the field points out the need for systems to be sensitive to this additional level of information: from highly domain-oriented disciplines such as bioinformatics [12], to more genre-oriented tasks. For example, [10] discusses the relevance of veridicity for IE.

In the scope of textual entailment, factuality has been taken as a basic feature in some of the systems participating in (or using the data from) previous PASCAL RTE challenges. For example, [18] copes with intensional contexts, [5] sets features accounting for the presence of polarity, modality, and factivity markers in the textual fragments, and [17] checks for polarity and modality scoping over matching nodes in a graph.

The use of factuality features does not seem essential, according to the ablation studies in some of the work mentioned above. As it happens, however, locally-derived polarity and modality features are not rich enough to adequately determine the factual status of events. Event factuality involves local but also non-local information. For instance, in (3), the factuality of *return* is conditioned by a

modal auxiliary (*may*) which is syntactically positioned two levels of embedding higher in the tree.

- (3) **T:** The Italian parliament may approve a draft law allowing descendants of the exiled royal family to return home.
H: Italian royal family returns home. [RTE3, Id=357]

Different discourse participants may present divergent perspectives about the factuality nature of the very same event. Hence, correctly identifying these sources is equally crucial for text entailment tasks. In (4), for example, the author (H1), John (H2), and Mary (H3) have a different perspective on the event of John being sick:

- (4) **T:** Mary regrets that John does not know he is sick.
H1: John is sick. **True**
H2: John is aware of it. **False**
H3: Mary knows John is sick. **True**

We present here a *factuality profiler*, a tool providing events with the factuality values that the relevant sources assign to them. Factuality values are characterized as pairs of modality and polarity features. The profiler is based on a top-down feature setting algorithm, which walks over dependency trees generated by a standard dependency parser along the lines of [4].

The next section gives an overview of previous work at the basis of our current research. Section 3 introduces the theoretical foundation of our system, while section 4 reviews the main linguistic means for expressing factuality. The system is presented in section 5.

2 Previous Work

In [16], we presented a suite of tools for identifying both polarity and modality using grammatical items as well as subordination contexts. Grammatical markers of polarity and modality were identified by EvITA, a non-domain specific event recognizer. We reported a performance of 74.55% precision and 78.61% recall, for a resulting F1-measure of 76.53%. The accuracy ratio (i.e., the percentage of values EvITA marked according to the gold standard) is 98.03% for grammatical polarity, and 97.04% for modality.

On the other hand, modality and factuality information expressed through subordination contexts are identified by SlinkET, with a performance of 92% precision and 56% recall (70% F1-measure).

These tools are however limited in that (a) they do not account for the effect that multiple embeddings have on the factuality of events, and (b) they do not distinguish between different sources but offer only the author perspective, and thus entailments H2 and H3 in (4) cannot be accounted for.

Given these limitations, we have extended and enhanced our algorithm to incorporate the influence of non-local context as well as the identification of sources and the subsequent determination of the factuality of events. We have

written the specification of the algorithm to accomplish these goals and we are currently annotating the test corpus against which we will evaluate the system.

Our work relates to previous research in the field aiming at similar goals. [13] is a first attempt at modeling the percolation of the polarity feature down the tree. This work, however, does not cope with modality. On the other hand, our notion of source is very close to that proposed in [3] and [2] for sources of opinions.

3 Event Factuality in Text

3.1 What is Event Factuality

Eventualities in discourse can be couched in terms of a veridicality axis that ranges from truly factual to counterfactual, passing through a whole spectrum of degrees of modality. In some contexts, the factual status of events is presented with absolute certainty. Events are then characterized as *facts* (5) or *counterfacts* (8). Other contexts introduce different shades of uncertainty. Depending on the polarity, events are then qualified as *possibly factual* (6) or *possibly counterfactual* (7).

- (5) Five U.N. inspection teams visited a total of nine other sites.
(6) United States may extend its naval quarantine to Jordan’s Red Sea port of Aqaba.
(7) They may not have enthused him for their particular brand of political idealism.
(8) The size of the contingent was not disclosed.

We therefore characterize **factuality** as involving **polarity** (positive vs. negative) and **epistemic modality** (speaker’s degree of commitment to the truth of what is asserted [14]). We define the factuality space based on Horn’s analysis of modality and its interaction with polarity [7], which is grounded on both linguistic tests and logical relations at the basis of the Aristotelian Square of Opposition. The system is shown in Table 1:

Table 1. Factuality values

	Positive (+)	Negative (-)	Unknown
Certain	Fact: <CT,+>	Counterfact: <CT,->	Certain but unknown output: <CT,UN>
Probable	Probable: <PR,+>	Not probable: <PR,->	(NA)
Possible	Possible: <PS,+>	Not certain: <PS,->	(NA)
Unknown	(NA)	(NA)	Unknown or uncommitted: <UN,UN>

The factual value of events is then presented as a tuple $\langle mod, pol \rangle$, containing a modality and a polarity value.²

²Semantically, this can be interpreted as: $Val(mod)(Val(pol)(e))$ -i.e., the modal value scopes over the polarity value.

The polarity axis divides into positive, negative, and unknown, while the modality axis distinguishes among certain (CT), probable (PR), possible (PS), and unknown (UN). The *unknown* values are added to account for cases of uncommitment. A source can be fully uncommitted (<UN,UN>) to the factuality of a given event (e.g., *he* with regard to *came* in: **He** *didn't say that she came*), or only partially uncommitted (<CT,UN>; as in: **He** *knows whether she came*).

3.2 The Factuality Profile of Events

Whenever there is a mention of an event, there is a **commitment act** towards the factuality of that event, performed by a particular source. Thus, **sources** are understood here as the cognitive individuals assigning factuality values to events mentioned in text.

By default, events mentioned in discourse always have an implicit source, viz., the author of the text. Additional sources are introduced in discourse by means of what we call epistemic predicates: predicates of reporting (*say, tell*), knowledge and opinion (e.g., *believe, know*), psychological reaction (*regret*), etc. They are so called because they express the epistemic commitment of one of their arguments (generally realized as the grammatical subject) towards the event expressed in their syntactic complement.

We identify the source referred to by the subject as the **cognizer** of the complement event. On the other hand, the author of the sentence will be referred to as its **anchor**, since it is the source that presents the commitment act of the cognizer towards the complement event.³

The status of the additional sources is, however, different than that of the text author. For instance, in (4) the reader only knows that Mary knows about event *e* (John being sick) according to what the author asserts –in other words, the reader does not have direct access to the factual assessment of Mary about *e*. Thus, we need to appeal to the notion of **nested source** as presented in [19]. *Mary* is not a licit source of the factuality of event *e*, but *Mary* according to the author instead, represented here as *mary_anchor*.⁴

The factuality value assigned to events in text must be relative to the relevant sources at play. Only under this assumption it is possible to account for the potential contradictions between factual values assigned to the same event, as is common in news reports. Consequently, events can be characterized by multiple factual values (one for each relevant source). We call the set of factual values that different sources assign to event *e* as its **factuality profile**:

$$p_e = \{ \langle f, s \rangle \mid f \text{ is a factuality value of the form } \langle \text{mod}, \text{pol} \rangle \text{ \& } s \text{ is a relevant source for that event} \}$$

Our tool for identifying the polarity and modality values of events is therefore referred to as a **factuality profiler** (FP).

³A more precise definition will be provided in section 5.4.

⁴Equivalent to the notation <*author,mary*> in Wiebe's work.

4 Factuality Markers

Event factuality in natural language is marked by both lexical items and syntactic constructions.

4.1 Lexical Markers

Situation Selecting Predicates (SSPs). These are predicates (verbs, nouns, or adjectives) that select for an argument denoting an event (or situation) of any sort. Syntactically, they subcategorize for *that*-, gerundive-, and *to*-clauses, or NPs headed by event-denoting nouns. The SSPs in (9) are in bold face; their embedded events, underlined.

- (9) a. Uri Lubrani also **suggested** Israel was willing to withdraw from southern Lebanon.
 b. Kidnappers **kept** their promise to kill a store owner they took hostage.

SSPs contribute to characterizing the factuality of the event denoted by its complement. For example, complements to weak assertive predicates [6] (*think, suppose*) are depicted as not totally certain; complements of reporting predicates [1] are presented as certain according to a particular source; factive (*regret, know*) and implicative predicates (*manage, prevent*) characterize their embedded complements as either factual or counterfactual [11, 8, 9]; and arguments of volition and commitment predicates (*wish; offer*) are presented as possible in a future temporal reference.

Modal Particles. They include modal auxiliaries (*could, may, must*); also, clausal and sentential adverbial modifiers (*maybe, likely, possibly*).

Polarity Particles. Including elements of a varied nature: adverbs (*not, until*), quantifiers (*no, none*), etc. They switch the polarity of its context. When scoping over a modal particle, they also affect the interpretation of the modal.

4.2 Syntactic Contexts

Syntactic structures introducing modality involve the presence of two clauses, generally one embedded under the other. In some cases, the embedded event is presupposed as holding; e.g., relative clauses (10), cleft sentences (11), and subordinated temporal clauses.

- (10) *Rice, who became secretary of state two months ago today, took stock of a period of tumultuous change.*
 (11) *It was Mr. Bryant who, on July 19, 2001, asked Rep. Bartlett to pen and deliver a letter to him.*

In others, the event denoted by the embedded clause is intensional in nature; e.g., purpose clauses (12) and conditional constructions (13).

- (12) *The environmental commission must adopt regulations to ensure people are not exposed to radioactive waste.*
- (13) *EZLN will return to the negotiating table if the conflict zone is **demilitarized**.*

Currently, our work focuses exclusively on factuality expressed by means of lexical markers.

5 Building a Factuality Profiler

5.1 Computational Approach

We compute the factuality value of each event by means of a top-down algorithm traversing a dependency tree. Two reasons motivate a top-down approach. The first one is of an empirical nature. As seen in section 4, subordination is directly involved in the factual characterization of events, and due to natural language recursivity, the factuality of a given event may depend on non-local information, located several levels higher in the tree (cf. example (3)).

The second reason for a top-down approach is methodological. We conceive the factuality profiler as a naïve and neutral decoder; naïve in that it assumes that sources are trustworthy, based on the Gricean maxim of quality; and neutral, because it considers all sources as equally reliable. That is, a factuality profiler will represent information presented in the text as true, without questioning anyone’s view or adopting a particular side.⁵

Since we view the factuality profiler as a naïve decoder, each sentence will have its top context level initiated with a default factuality value of $\langle CT, + \rangle$, which will be potentially modified by the factuality markers available at different context levels. The factuality of every event, e , is directly derived from the factuality information available at its context.

5.2 Lexical Information

We detail here the information needed for each type of marker. The data has been obtained from exploring a fragment of the American National Corpus (ANC).

Polarity Particles. These shift the polarity value of the current context: from positive to negative and vice versa. In (14a) for example, the factual value assigned by Mary (m_{s_0}) to event e (John being sick) is $\langle CT, + \rangle$, whereas in (14b) it is switched to $\langle CT, - \rangle$, due to the presence of *not*:

- (14) a. Mary said John is sick $_e$.
 b. Mary said John is not sick $_e$.
 c. Mary did **not** say John is not sick $_e$.

⁵A later postprocessing can use different weights in order to favor one source as more reliable than another.

In some cases, however, the clause where a polarity marker is used is in the scope of a higher factuality marker that sets the contextual polarity as underspecified. In (14c) the predication *did not say* has opened an underspecified (or uncommitted) context, and hence event e is now assessed as $\langle UN, UN \rangle$.

We therefore need to account for the effects of polarity considering this third contextual polarity value as well. Table 2 shows the interaction of contextual polarity (columns) and the polarity value contributed by a new marker (rows).

Table 2. Polarity value given context polarity.

Marker value	Contextual polarity		
	+	-	UN
+	+	-	UN
-	-	+	UN

Modality Markers. Each marker expresses a particular modal degree. For example, *may* denotes possibility, *likely* probability, and *must* certainty. However, the use of a given modal particle in a clause does not necessarily color the predication with the inherent value of that modal. Similar to polarity markers, the resulting value depends on the interaction between the modal element and the current modality and polarity of the context. Consider:

- (15) a. Mary **denies** [John may have been sick $_e$].
 b. Mary **says** it’s **possible** [John may have been sick $_e$].

In (15a), *may* is used in a context of negative polarity and absolute certainty set by the SSP *deny*, whereas in (15b), it is used in a context of positive polarity and modality set to possible. Because of this, in the first example, event e (underlined) is presented as a counterfact according to Mary ($\langle CT, - \rangle$), but in the second, as a possibility ($\langle PS, + \rangle$).

Table 3 illustrates the interaction between the contextual modality (columns) and the modal value contributed by the marker (rows).⁶ Note that in contexts of positive polarity, the lowest value prevails, whereas in negative contexts, the dominant modal value is the one provided by the context.

Table 3. Modality value given context factuality

Marker	Contextual factuality							
	Polarity = +				Polarity = -			
	CT	PR	PS	UN	CT	PR	PS	UN
CT	CT	PR	PS	UN	CT	PR	PS	UN
PR	PR	PR	PS	UN	CT	PR	PS	UN
PS	PS	PS	PS	UN	CT	PR	PS	UN

Situation Selecting Predicates (SSPs). Two types are distinguished:

⁶It has been compiled by exploring corpus data as well as made up examples. Combinations with mid values (probability) are highly unusual; the resulting values are only estimated.

i. Epistemic SSPs. These express the epistemic commitment of one of their arguments (generally the subject) towards the event expressed in the complement. In addition, they convey the epistemic attitude of the sentence’s author.

In section 3.2, we introduced the roles of *cognizer* and *anchor* to refer to the grammatical subject and text author sources, respectively. Depending on the contextual polarity and modality values, these two sources will judge the factuality of the complement event in different ways. Consider the factual assessment of event e by both the anchor (a) and the cognizer (c) in the sentences below –where $f(e, a)$ expresses the factual value that the source satisfying the anchor role assigns to event e , and $f(e, c)$ the factual value assigned to that same event by the source playing the cognizer role:

- (16) a. Mary **knows** that John is sick _{e} $f(e, a) = \langle CT, + \rangle$
 $f(e, c) = \langle CT, + \rangle$
 b. Mary **does not know** that John is sick _{e} $f(e, a) = \langle CT, + \rangle$
 $f(e, c) = \langle UN, UN \rangle$
 c. Mary **may have known** John is sick _{e} $f(e, a) = \langle CT, + \rangle$
 $f(e, c) = \langle UN, UN \rangle$

In (16a) both the anchor (text author) and the cognizer (Mary) judge the event of John being sick (e) as a fact: $\langle CT, + \rangle$. Contrast this with sentences (16b-c), where the main predicate is either under negation or within a modal scope. Now the anchor source judges e as a fact as in the previous case (due to the presuppositional nature of the predicate *know*), but the cognizer is presented as not knowing e (or *possibly* not knowing it). Hence, the judgement of e relative to c must be uncommitted: $\langle UN, UN \rangle$.

Consequently, epistemic SSPs in the lexicon will need to have their factuality assignments represented not only depending on the polarity and modality of the context, but also according to the different source roles. *Know* and *say*, for example, behave as illustrated in Table 4.⁷

Table 4. Epistemic predicates

		Contextual factuality			
		mod=CT		mod<CT	
		pol=+	pol=-	pol=+	pol=-
<i>know</i>	(a)	CT,+	CT,+	CT,+	CT,+
	(c)	CT,+	U	U	U
<i>say</i>	(a)	U	U	U	U
	(c)	CT,+	U	U	U

ii. Non-epistemic SSPs. These do not contribute any additional source –the only source at play here is the text generator. They include (semi-)implicative predicates (e.g., *fail*, *manage*), but also other types of predicates introducing a future event as their complement; e.g., volition (*want*),

⁷For visual clarity, factuality values $\langle U, U \rangle$ are simplified as $\langle U \rangle$. Similarly, we also omit the column for underspecified contexts.

commissive (*offer*), and command (*require*) predicates. In the case of predicates introducing future events, the complement event is always of an intensional nature. Hence, the source is uncommitted with respect to its factual value ($\langle U, U \rangle$) regardless of the contextual polarity and modality (last two rows in Table 5).

In the case of implicative predicates [8], the factuality status of the embedded event is assessed as unknown when the implicative predicate is in a context where modality is lower than certain ($mod < CT$). But in contexts of certainty, the embedded event is assessed as factual ($\langle CT, + \rangle$) or counterfactual ($\langle CT, - \rangle$), depending on the SSP type and the contextual polarity (six first rows in Table 5).⁸

Table 5. Non-epistemic predicates

	Contextual factuality			
	mod=CT		mod<CT	
	pol=+	pol=-	pol=+	pol=-
<i>manage</i>	CT,+	CT,-	U	U
<i>cause</i>	CT,+	U	U	U
<i>attempt</i>	U	CT,-	U	U
<i>fail</i>	CT,-	CT,+	U	U
<i>refuse</i>	CT,-	U	U	U
<i>hesitate</i>	U	CT,+	U	U
<i>want</i>	U	U	U	U
<i>offer</i>	U	U	U	U

5.3 Terminological Framework

5.3.1 Embedding levels

Consider each sentence, S , as consisting of one or more **levels, l , of dependency embedding**. Every time an epistemic SSP is consumed, a new level is introduced. Thus, a simple sentence (17a) has one level, while a sentence with two epistemic SSPs, like (17b), has three. We identify each level by its depth of embedding. The sentence root level has depth 0, corresponding to l_0 .

- (17) a. [_{l_0} Mary slept soundly].
 b. [_{l_0} Mary regrets [_{l_1} John does not know [_{l_2} he is sick]].

Each level l_n contains **one (or more) events, e_n** , the factuality of which is evaluated relative to the **set of relevant sources at that level, RS_n** . Each level may also introduce a new source, s_n . If $n=0$, the new source is s_0 (used by default to refer to the text author). If $n>0$, the introduction of a new source is optional.

⁸Note that this type of predicates only affect polarity. This is the data used in [13] to compute local polarity.

5.3.2 Source Roles

As seen, the factuality value assessed by each relevant source $s \in \text{RS}_n$ depends on the role that source satisfies: anchor or cognizer (cf. examples (16)). We therefore need a way to identify such roles.

Identifying the anchor and the cognizer of e_2 in sentence (18) is trivial: the former is the text author (s_0) while the later is the subject of e_1 relative to that author (s_j-s_0). But who is the anchor of e_2 when the sentence is embedded under another SSP (19-21)?

(18) John $_{s_j}$ **knows** $_{e_1}$ he is **sick** $_{e_2}$.

(19) Mary $_{s_m}$ **says** $_{e_0}$ John $_{s_j}$ **knows** $_{e_1}$ he is **sick** $_{e_2}$.

(20) Mary $_{s_m}$ is **aware** $_{e_0}$ John $_{s_j}$ **knows** $_{e_1}$ he is **sick** $_{e_2}$.

(21) Mary $_{s_m}$ is not **aware** $_{e_0}$ John $_{s_j}$ **knows** $_{e_1}$ he is **sick** $_{e_2}$.

In (19) *Mary* is the author of the embedded sentence, as reported by s_0 ; thus, the anchor of e_2 is s_m-s_0 . This same source is also the anchor of e_2 in (20). Note however that *Mary* is not reported as having uttered sentence (18). She did not “present the commitment act of the cognizer towards” e_2 (as we defined anchor in section 3.2). Yet, she is in the adequate mental state for reporting that commitment act. We therefore broaden the notion of anchor to include not only those sources directly expressing a particular commitment act (*Mary says*), but also those that are reported as being in the mental state for expressing one (*Mary is aware, knows*).

Finally, what is the anchor of e_2 in (21)? Source s_m-s_0 is neither expressing a commitment act nor in the mental state for doing it. The source that is effectively expressing the commitment act of John knowing about e_2 is the text author, s_0 . This is thus the anchor of e_2 . Stated generally then, the anchor of an embedded event e is the source that is not uncommitted to event e' , the SSP which embeds e . Applied to our examples (19-21), the anchor of embedded event e_2 (John being sick) is the source that is not uncommitted to event e_1 (John knowing he is sick).

Once the anchor is identified, we can obtain the cognizer. We know it involves the grammatical subject of the SSP embedding our event e_2 ; that is, s_j . We also know that it must necessarily be a nested source; more precisely, that it must be nested by the anchor, because this is the source (being in the mental state for) eliciting the cognizer assessment of e_2 . Hence, in examples (19-20) the cognizer of e_2 is $s_j-s_m-s_0$, while in (21) it is s_j-s_0 . Put again in general terms, the cognizer of an event, e , is the source expressed by the grammatical subject of the SSP embedding e , nested relative to the anchor source.

The two roles can be defined as follows:

Anchor: At level l_n the *anchor* of event e_n is satisfied by those sources $s \in \text{RS}_n$ that are not uncommitted towards the factuality of e_{n-1} , the event embedding e_n .

Source: At level l_n the *source* of event e_n is satisfied by the grammatical subject of the SSP expressing event e_{n-1} , nested relative to the anchor source.

5.4 Algorithm

The core procedure of the FP applies top-down traversing a dependency tree. At the beginning of each new level, l_n , factuality values are set for that level. We can refer to them as *contextual factuality values*, CF_n . These must be set relative to the relevant sources in RS_n , since different sources may assess the same event in different ways within the same clause. The set of contextual factuality values at level l_n can be defined as: $\text{CF}_n = \{ \langle f, s \rangle \mid f \text{ is a factuality value \& } s \text{ is a relevant source in } l_n \}$.

By default, at level l_0 the set CF_n contains only the value $\langle \text{CT}, + \rangle$ relative to the text author (i.e., $\langle \langle \text{CT}, + \rangle, s_0 \rangle$). This applies the naïve decoder assumption.

Additional markers within each level l_n (polarity and modality particles, non-epistemic SSPs) may modify the values in CF_n , according to the information in tables 2, 3 and 5. When an event, e_n , is found, its factuality profile, p_{e_n} , is the set of contextual factuality values CF_n available at that point.

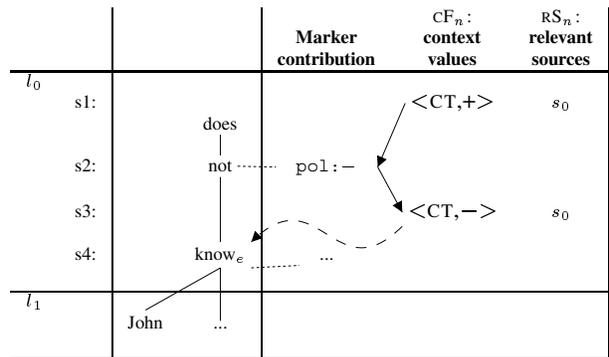


Figure 1. *John does not know he is sick.*

Figure 1 illustrates the initial steps of this process: in step 1 of level l_0 (s1), the profiler sets the factuality values for the relevant sources in that level –only the text author, s_0 . We are at the top of the tree, and hence this a default step satisfying the naïve decoder assumption. Then, the profiler travels down the tree until it finds the polarity marker *not* (s2), which switches the polarity initially assigned to the context, from positive to negative. The resulting contextual factuality value is that in (s3). Next, the verb *know* is found, which denotes an event. Its factuality profile, p_{know_e} , is the set of factuality values relative to the relevant sources in that level: $p_{know_e} = \{ \langle \langle \text{CT}, + \rangle, s_0 \rangle \}$.

The FP algorithm is provided in Algorithm 1. In what follows, we discuss its three basic components.

Algorithm 1 Factuality Profiler

```
1:  $n \leftarrow 0$ 
2: set level  $l_n$ 
3: for all  $i$  in TREE do
4:   if  $i$  is an event then
5:     obtain the factuality profile of  $i$ ,  $p_i$ 
6:   end if
7:   if  $i$  is an epistemic ssp then
8:      $n \leftarrow n + 1$ 
9:     set level  $l_n$ 
10:  else if  $i$  is another type of marker then
11:    update contextual factuality,  $CF_n$ 
12:  end if
13: end for
```

1. Set Level l_n (lines 2 and 9). This function is called every time a new level is triggered, be it at the top of the tree (lines 1-2) or when an epistemic SSP is found (lines 7-9). It executes the following steps:

i. Identify the set of relevant sources at the current level, RS_n . The procedure is defined inductively:

- i.* The set of relevant sources at level l_0 contains only a non-nested source which corresponds to the text author: $RS_0 = \{s_0\}$.
- ii.* The set of relevant sources at level l_n , where $n > 0$, is:
 $RS_n = RS_{n-1} \cup \{s_n \text{--} s \mid s_n \text{ is the new source introduced at level } l_n \text{ \& } s \in RS_{n-1}\}$

Clause *ii.* states that the set, RS_n , at level, l_n , contains (a) the relevant sources at the previous level, l_{n-1} (first part of the union); and (b) nested sources constituted of the new source introduced in the level, s_n , and a relevant source of the preceding level. We cannot fully justify this due to space limitations, but the reason is implicitly illustrated in Fig. 2. Complete details are given in [15].

ii. For each $s \in RS_n$, identify its role: anchor and source. This is a vacuous step at the top level, but it becomes fundamental as soon as new sources are introduced. Recall that the factuality of events embedded under SSPs is obtained from the lexical information provided by the SSP (cf. tables 4-5). In the case of epistemic predicates, anchor and cognizer can assign different factuality values. We thus need to identify what source has which role, in order to correctly compute the factuality that each of them assigns to the event.

The definitions of anchor and cognizer roles were provided in section 5.3.2. More than one source can satisfy these roles, and therefore events can have multiple anchors and cognizers. The sets of anchors and cognizers of event e are identified as $a(e)$ and $c(e)$, respectively, and computed as follows:

i. At level l_0 : $a(e_0) = \{s_0\}$ and $c(e_0) = \{s_0\}$.

ii. At level l_n , for $n > 0$:

$a(e_n) = \{s \mid s \in RS_{n-1} \text{ \& } f(e_{n-1}, s) \neq \langle U, U \rangle\}$ and
 $c(e_n) = \{s_n \text{--} s_a \mid s_n \text{ is the new source introduced in level } l_n \text{ \& } s_a \in a(e_n)\}$.

iii. Set the contextual factuality values, CF_n . Each level can have more than one relevant source. Thus, contextual factuality is not absolute but needs to be relative to each source relevant in the level. For example, in Fig. 2, the use of *be aware* under negation sets the contextual factuality of level l_1 as $\langle CT, + \rangle$ according to source s_0 (the text author), but $\langle U, U \rangle$ according to source $s_1 \text{--} s_0$ (*Mary*). Conceiving contextual factuality as relative to the relevant sources results in an assignment of appropriate values down the tree, which ensures that at any point the factuality of a given event can be adequately assessed according to the sources at play. Context factuality is obtained as follows:

i. At level l_0 : $CF_n = \{\langle \langle CT, + \rangle, s_0 \rangle\}$

ii. At level l_n , for $n > 0$: $CF_n = \{\langle f, s \rangle \mid s \in RS_n \text{ \& } f = Lex(e_{n-1}, c_{e_{n-1}}, s_r)\}$

Clause *i.* applies the assumption of the naïve decoder. In clause *ii.*, function *Lex* performs a search in the lexical base (here, tables 4-5). Its parameters are: (a) e_{n-1} , referring to the epistemic SSP in the previous level l_{n-1} that triggered the current level; (b) $c_{e_{n-1}}$, the contextual factuality under which e_{n-1} is used;⁹ and (c) s_r , the role of s (anchor, cognizer, or none).

2. Update contextual factuality, CF_n (line 11). When a polarity or modality particle, or a non-epistemic SSP is found in l_n , FP updates the contextual factuality values $v \in CF_n$ according to the information in tables 2, 3, and 5.

3. Obtain the factuality profile of e , P_e (line 5). Applied when an event is found. Due to the on-the-fly updating of the contextual factuality values in CF_n whenever a new level is set (stage 1 above) or a new marker is found (stage 2), the event profile is in fact already computed. The factuality profile for event, e_n , p_{e_n} , corresponds to the set of contextual factuality values CF_n available at that point.

6 Results and Concluding Remarks

Figures 2 and 3 illustrate the factuality profiles computed for each event in sentences (22) and (23), respectively.

(22) $Mary_{s_1}$ is not **aware** $_{e_0}$ $John_{s_1}$ **knows** $_{e_1}$ he is **sick** $_{e_2}$.

(23) $Mary_{s_1}$ **regrets** $_{e_0}$ $John_{s_1}$ does not **know** $_{e_1}$ he is **sick** $_{e_2}$.

⁹Since there are potentially more than one contextual factuality values, obtaining the proper one requires an additional mechanism that we ignore here for reasons of space.

		FACTUALITY PROFILES:		
TOP		Relevant Sources	Source Roles	Factual Assignm.
l_0	<pre> graph TD TOP --> s0 TOP --> is is --> not not --> aware_e0 </pre>			
l_1	<pre> graph TD TOP --> Mary_s1 TOP --> knows_e1 </pre>	s_0 $s_1 \rightarrow s_0$	(a, c) (a) (c)	$f:\langle CT,+ \rangle$ $f:\langle U,U \rangle$
l_2	<pre> graph TD TOP --> John_s2 TOP --> sick_e2 </pre>	s_0 $s_1 \rightarrow s_0$ $s_2 \rightarrow s_0$ $s_2 \rightarrow s_1 \rightarrow s_0$	(a) (a) (c) (c)	$f:\langle CT,+ \rangle$ $f:\langle U,U \rangle$ $f:\langle CT,+ \rangle$ $f:\langle U,U \rangle$

Figure 2. Factuality profiles for events in (22).

The different factuality markers employed in each sentence (e.g., *aware* vs. *regret*) as well as their structural configuration (e.g., differences in negation scope), have consequences in the assessment of the factuality of e_2 by each relevant source.

Sentence (23) corresponds to that in (4), repeated here as (24). Note that entailments (H1-H3) are correctly predicted by the factuality profiler.

- (24) **T:** $Mary_{s_1}$ **regrets** $_{e_0}$ $John_{s_1}$ **does not know** $_{e_1}$ **he is sick** $_{e_2}$.
H1: John is sick. **True** $f(s_0, e_2) = \langle CT,+ \rangle$
H2: John is aware of it. **False** $f(s_2 \rightarrow s_0, e_2) = \langle U,U \rangle$
H3: Mary knows John is sick. **True** $f(s_1 \rightarrow s_0, e_2) = \langle CT,+ \rangle$

The factuality profiler presented here is still under development. We are enlarging the lexicon and plan to incorporate syntactically-triggered factuality in the near future. Currently, we are annotating the corpus against which we will evaluate our system, TimeBank1.2.¹⁰ TimeBank contains 7935 manually identified events, 1457 of which (18.4%) have been classified as epistemic SSPs in the scope of the current project, with an interannotator agreement of $k=0.91$ over 80% of the corpus.

References

- [1] S. Bergler. *Evidential Analysis or Reported Speech*. PhD thesis, Brandeis University, 1992.
- [2] S. Bethard, H. Yu, A. Thornton, V. Hatzivassiloglou, and D. Jurafsky. Automatic extraction of opinion propositions and their holders. In *Proceedings of AAAI Spring Symposium on Exploring Attitude and Affect in Text*, 2004.
- [3] Y. Choi, C. Cardie, E. Riloff, and S. Patwardhan. Identifying sources of opinions with conditional random fields and extraction patterns. In *HLT/EMNLP 2005*, 2005.

¹⁰<http://www.timeml.org/site/timebank/timebank.html>

		FACTUALITY PROFILES:		
TOP		Relevant Sources	Source Roles	Factual Assignm.
l_0	<pre> graph TD TOP --> s0 TOP --> regrets_e0 </pre>	s_0	(a, c)	$f:\langle CT,+ \rangle$
l_1	<pre> graph TD TOP --> Mary_s1 TOP --> does does --> not not --> know_e1 </pre>	s_0 $s_1 \rightarrow s_0$	(a) (c)	$f:\langle CT,+ \rangle$ $f:\langle CT,- \rangle$ $f:\langle CT,- \rangle$
l_2	<pre> graph TD TOP --> John_s2 TOP --> sick_e2 </pre>	s_0 $s_1 \rightarrow s_0$ $s_2 \rightarrow s_0$ $s_2 \rightarrow s_1 \rightarrow s_0$	(a) (a) (c) (c)	$f:\langle CT,+ \rangle$ $f:\langle CT,+ \rangle$ $f:\langle U,U \rangle$ $f:\langle U,U \rangle$

Figure 3. Factuality profiles for events in (23).

- [4] M.-C. de Marneffe, B. MacCartney, , and C. D. Manning. Generating typed dependency parses from phrase structure parses. In *LREC 2006*, 2006.
- [5] M.-C. de Marneffe, B. MacCartney, T. Grenager, D. Cer, A. Rafferty, and C. D. Manning. Learning to distinguish valid textual entailments. In *RTE-2*, 2006.
- [6] J. B. Hooper. On assertive predicates. In J. Kimball, ed., *Syntax and semantics, IV*, 91–124. Academic Press, 1975.
- [7] L. R. Horn. *A Natural History of Negation*. University of Chicago Press, Chicago, 1989.
- [8] L. Karttunen. Implicative verbs. *Language*, 47, 1970.
- [9] L. Karttunen. Some observations on factivity. *Papers in Linguistics*, 4:55–69, 1971.
- [10] L. Karttunen and A. Zaenen. Veridicity. In G. Katz, J. Pustejovsky, and F. Schilder, editors, *Dagstuhl Seminar Proceedings*. IBFI, 2005.
- [11] P. Kiparsky and C. Kiparsky. Fact. In M. Bierwisch and K. E. Heidolph, editors, *Progress in Linguistics. A Collection of Papers*, pp. 143–173. Mouton, The Hague, 1970.
- [12] M. Light, X. Y. Qiu, and P. Srinivasan. The language of Bioscience: Facts, speculations, and statements in between. In *BioLINK 2004*, pp. 17–24, 2004.
- [13] R. Nairn, C. Condoravdi, and L. Karttunen. Computing relative polarity for textual inference. In *ICoS-5*, 2006.
- [14] F. R. Palmer. *Mood and Modality*. Cambridge University Press, Cambridge, England, 1986.
- [15] R. Saurí. *A Factuality Profiler for Eventualities in Text*. PhD thesis, Brandeis University, forthcoming.
- [16] R. Saurí, M. Verhagen, and J. Pustejovsky. Annotating and recognizing event modality in text. In *FLAIRS 2006*, 2006.
- [17] R. Snow and L. Vanderwende. Effectively using syntax for recognizing false entailment. In *HLT-NAACL 2006*, 2006.
- [18] M. Tatu and D. Moldovan. A semantic approach to recognizing textual entailment. In *HLT/EMNLP*, 2005.
- [19] J. Wiebe, T. Wilson, and C. Cardie. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2), 2005.