

Polling the Blogosphere: a Rule-Based Approach to Belief Classification

Jason S. Kessler

Department of Computer Science
Indiana University
Bloomington, IN 47405
jaskessl@cs.indiana.edu

Abstract

The research described here is part of a larger project with the objective of determining if a writer believes a proposition to be true or false. This task requires a deep understanding of a proposition's semantic context, which is far beyond NLP's state of the art. In light of this difficulty, this paper presents a shallow semantic framework that addresses the sub-problem of finding a proposition's truth-value at the sentence level. The framework consists of several classes of linguistic elements that, when linked to a proposition through specific lexico-syntactic connectors, change its truth-value. A pilot evaluation of a system implementing this framework yields promising results.

Introduction

Among the many difficulties faced in constructing a search engine, determining how to display results is one of the most challenging. For unambiguous queries, however, the classical approach of returning a list of links and their summaries which are ordered by query-similarity and authority is an efficient way of directing users to relevant sources of information. For instance, searching Google for "ICWSM" returns the conference's web page as the first result. On the other hand, an ambiguous query like "apple," which could refer to a company or a fruit, may require a disambiguation step, enabling the users to select a desired word sense. The problem of disambiguating queries sense-wise has been approached with a reasonable degree of success using statistical methods (Liu, Yu, & Meng 2005). Search engines such as Clusty¹ have built user interfaces to address this sort of ambiguity.

However, query results might cluster into categories that do not correspond to semantic ambiguities. For instance, a Web search for a controversial issue such as the ongoing war in Iraq returns pages that are written with an anti-war perspective, a pro-war perspective, or a neutral perspective. When searching for pages about controversial topics, users may be searching for information written exclusively from a particular perspective or may simply benefit from seeing the page's perspective identified along with its summary.

The problem of identifying a document's perspective, however, is not exclusively a lexical semantic problem². Unlike the "apple" example, pages written from different perspectives concerning the Second Iraq War all use the same sense of the term (i.e., they all talk about the same war) and thus require a different method of classification or clustering. Some preliminary work has been done on this problem. For example, Lin & Hauptmann (2006) used differences in term-distributions to classify documents according to their author's political orientation.

While bag-of-words methods can be used to categorize results from many types of queries, including those discussed above, there are still some types of queries that yield result-sets which require a more complex means to be categorized meaningfully. An example is a query for which the user wishes to find documents that indicate their authors' acceptance or rejection of a specific proposition. Although light, statistical approaches like the ones discussed above can help uncover a document's general slant, they cannot be used to reason about a document's specific semantic content. Instead, a semantic representation is required to make this categorization.

For example, suppose a user wants to find pages that contain agreement or disagreement with the proposition that America is winning the war in Iraq. A naive method would be to search for pages containing the literal proposition, considering that results set to be the "agree" cluster and search on the negation of the statement for the "disagree" cluster. Therefore, the "agree" cluster would consist of pages containing the strings "the U.S. is winning the war in Iraq," while the disagree cluster would be formed via searching "the U.S. is not winning the war in Iraq."

This approach suffers from both recall and precision problems. Recall problems occur because writers might not phrase the proposition in the same way as the search string. Moreover, the proposition may not be stated, but merely alluded to (e.g., "The U.S. is winning both of its wars.") Solutions to the problem of finding multiple linguistic realization of the same proposition lie in fields such as paraphrase generation (e.g., Barzilay & Lee (2003)).

The focus of this paper is problems with precision,

namely, that simply mentioning a proposition does not mean the author agrees with it. Natural language provides a plethora of ways to contextually negate propositions. To illustrate, consider the following contexts a Google search for “the U.S. is winning the war in Iraq” yielded,

- (1) a. Buy this book if you believe its premise that **the U.S. is winning the war in Iraq** against the insurgency.
- b. He disagreed with President Bush’s assessment earlier in the day that **the U.S. is winning the war in Iraq**.
- c. And then later, he was asked if the **U.S. is winning the war in Iraq**.

The naive approach would incorrectly assign each context in (1) to the “agree” cluster. A better system would place (1-b) in the “disagree” cluster and recognize (1-a) and (1-c) are in a neutral context. Moreover, the negative assertion in (1-b) was presumably attributed to a source other than the author. This should be recognized when the results are displayed to the user.

This paper addresses this precision problem, presenting a system that, given a proposition and its sentential context, determines if its writer agrees with, denies, or takes no stance towards the proposition’s truth-value.

In addition to information retrieval, this system addresses problems that are faced in multi-document summarization and question-answering. In multi-document summarization, propositions that are negated in some documents and not in others are candidates to be included multiple times within the summary in both their negated and non-negated contexts. If a question-answering system is posed a question such as “Who thinks [PROPOSITION]?” the system can identify passages containing positive assertions of the proposition and, with the help of a system such as Stoyanov & Cardie (2006), return a list of their sources.

The blogosphere is a particularly apt domain to be searched using this approach. Because blogs tend to be heavily opinionated, to reflect world-views outside the mainstream, and to contain large quantities of information, the blogosphere will likely contain a diversity of beliefs regarding many propositions. The program proposed in this paper, which identifies these conflicts and provide users a way of finding information about all sides, has the potential to perform a valuable service.

A Linguistic Perspective

A formalism for describing this phenomena minimally must cover two properties: the set and behavior of possible propositional truth values and the means writers use to assign propositions truth values. If we leave aside pragmatic phenomena such as sarcasm and only focus on intrasentential cases, several problems still arise.

The simplifying assumption is made that writers categorize each proposition they use into one of three categories. Writers can believe, deny, or take no stance toward a proposition. Thompson & Yiyun (1991) use the term “stance” to describe “the attitude.. ..towards the validity of infor-

mation reported or opinion,” a definition that is adopted in this paper. While making a distinction between the document author’s stance and a reported speaker’s stance, they identify three types of stance which can be summarized as the assertion of information’s validity, its invalidity, or giving no indication of its validity or invalidity. In this paper, we will refer to the stance of the writer as **positive**, **negative**, or **neutral**. This ternary also corresponds Helbig & Glckner’s 2007 “real,” “unreal,” and “hypothetical” facticity categories. Nairn, Condoravdi, & Karttunen (2006) describes complement clauses as having a “relative polarity.” The term “polarity” is avoided in this paper because it is also used to discuss positive and negative sentiment (e.g., Wilson, Wiebe, & Hoffmann (2005).) Sentiment analysis deals with writers emotions and not necessarily their beliefs. For example, in “Bill is upset because the bar closed,” Bill expresses negative sentiment about the bar closing, but also expresses positive belief that it has closed. In this paper, the problem of identifying stance of the writer towards any proposition is pursued.

Note that, in this paper, **stance** will be used in the same way as “relative polarity.” Thus, stance will be used as a property of a proposition as well as of the writer.

The second property that this formalism needs to capture is how writers can communicate their stance. Writers have a number of tools at their disposal for declaring their stance. For example, researchers have looked into how matrix verbs can negate or presuppose their clausal complements (Karttunen 1971). A broad view is taken on how stance can be communicated, not only including matrix verbs but co-referential noun phrases, discourse connectives, and conditionals. These term classes, or **Veridicality Elements** (VEs) can all be used to change the stance toward their argument proposition.

When used in this paper, the term “veridicality” refers to how a VE can be used to change the writer’s stance towards a proposition. Giannakidou (2002), uses the term to describe propositional operators that entail the truth of its argument. She contrasts it to nonveridicality, where no truth value is entailed, and antiveridicality, where the negation is entailed.

Different VEs affect their argument propositions in different ways. For instance, the VE “asserts” indicates the source of “asserts” believes the argument proposition to be true. On the other hand, characterizing a proposition using the VE “falsehood” would indicate the speaker believes the proposition to be false. **Veridicality Orientations** (VOs) are general categories describing how these VEs affect their argument proposition. For instance, “falsehood” carries a negative VO. Syntactic constructions that pair VEs to their argument propositions are referred to as **Veridicality Transformations** or VTs. VTs merge their VEs to their argument propositions. Veridicality transformations are ordered so that certain transformations happen before others. The merging process helps to preserve this ordering.

Consider a rough illustration of how a simplified version of the framework functions. Consider the sentence (2-a) (repeated from (1-b)). Assume it has two VEs, “disagreed” and “assessment,” each affecting the proposition “the U.S. is

winning...”³ “Assessment” acts as a characterization of the proposition while “disagreed” describes the speaker’s stance towards the proposition.

Intuitively, we can see that “assessment” is co-referential to the proposition. This step leads to the reification of the arguments of “disagreed”, allowing the verb to be interpreted. Because it’s impossible to interpret a verb without reified arguments, we need to ensure that VTs which apply to siblings of propositions to take place ahead of VTs which apply head words to the proposition.

In (2-b), the first VT has taken place, connecting “assessment” to its sibling proposition, the clausal complement of its head verb. Assessment has a positive VO, thus maintaining the writer’s positive stance towards its proposition. The adjective phrase “earlier in the day” is not part of the VT and thus not included in the new proposition. While the head verb “disagreed” is part of the VT, it is not included in the proposition because its position could hold a VE. The final VT, shown in (2-c), matches the VE “disagreed” to its clausal complement, the proposition. The negative VO of “disagreed” is transferred to the proposition, and its stance becomes negative.

- (2) a. He [disagreed]_{VE} with President Bush’s [assessment]_{VE} earlier in the day that [the U.S. is winning the war in Iraq]_{POS}. STANCE·
- b. He [disagreed]_{VE} with President Bush’s [[assessment]_{VE} (earlier in the day) that the U.S. is winning the war in Iraq]_{POS}. STANCE·
- c. He [[disagreed]_{VE} (with President Bush’s) [assessment]_{VE} (earlier in the day) that the U.S. is winning the war in Iraq]_{NEG}. STANCE·

While two VTs and two types of VEs are discussed in this example, there are many more. The next section presents a system for implementing this framework in conjunction with the VTs and VE categories defined in this paper.

A Rule-Based Approach

A rule-based approach is used to detecting intrasentential propositional stance. It begins with a dependency parse of a sentence, treating a specified proposition to be an atomic node, *P*. Next, the system labels expressions that belong to VE classes specified in Section (2). The labeling only takes into account lexical information. Classes are derived from the VO and part-of-speech.

The system uses a sequence of VTs, ordered to prioritize the reification of arguments before the head words are applied. VTs are based on patterns over a dependency parse graph incorporating *P* and VE classes. Given an annotated dependency parse, the system finds the first matching VT, applies it, and, starting from the front of the sequence, looks

³In the fully described system, “Bush” would be treated as a SOURCE VE. For simplicity, this aspect is ignored.

for another matching VT over the transformed parse graph. This process continues until no matching patterns are found.

Preprocessing

Toutanova *et al.* (2003) is used to perform part-of-speech tagging, while de Marneffe, MacCartney, & Manning (2006) is used to perform a dependency parse. de Marneffe, MacCartney, & Manning (2006) provides a richer set of grammatical relations than more widely-used CONNL based parsers. Next, the system merges the nodes that form the propositional query, calling this new node *P*.

Because merging nodes in a tree can produce cycles and nodes with multiple parents, the modified dependency tree is no longer a tree and should be referred to as a dependency graph.

Veridicality Elements

Figure 2 presents the VOs the system employs. Note these correspond to Nairn, Condoravdi, & Karttunen’s 2006 “two-way implicative,” “factive,” “counterfactive,” and “neutral,” which are discussed further in Related Work.

12 VE classes based on VO and part-of-speech have been identified. To populate VE classes, around 40 hand-picked example sentences were collected, their VEs and stances were manually identified, the on-line thesaurus⁴ was used to find expressions to add to each category. Additionally, an approach inspired by Brin (1998) was used to identify other common VEs.

The first step in this approach was to find an exact-phrase query in which two VEs would reliably be present. For instance, “I agree with the assertion that” would be a commonly found string that contains two VEs. Next, one VE would be made into a wild card character, resulting in the search string “I * with the assertion that.” The wild card is generally filled by other VEs, which can then be manually assigned to a class. Next, the most commonly found VEs in that search were substituted for the wild card while the other VE was made a wild card character, leading to queries “I take issue with the * that” and “I quibble with the * that”. Each of the queries leads to nouns that are likely VEs and can be manually categorized. The newly discovered VEs can then be used to repeat the process.

While the set of VE classes employed in the system (see Figure 1) is far from complete, it is serviceable, and provides a wider range of classes than an earlier system (Nairn, Condoravdi, & Karttunen 2006) which only uses verbs. Here, verbs make up 4/11 classes.

Positive verbs are verbs that can be negated and have their semantic orientation flipped. For instance, “not agree” inverts the orientation of agree. While “not say” does not strictly flip the orientation of “say” (“say” becomes neutral,) its negation creates a considerable pragmatic force towards inversion. Negative verbs have symmetrical issues.

Factive verbs and rare counter-factive verbs lock in their positive and negative veridicality.

Sources of propositions (e.g., the speaker (Wiebe 2002)) can become negative VEs if they are used in a quantifier-

⁴<http://thesaurus.reference.com/>

Class Name	Example Expressions	Veridicality Orientation	Expression List Size
Positive Verbs	agree, say, believe	Positive	72
Negative Verbs	challenge, deny, dismiss	Negative	42
Factive Verbs	sift, extrapolate, admit, go beyond	Factive	59
Counter-Factive Verbs	pretend, make believe, wish	Counter-Factive	4
Neutral Verbs	address, contest, explain	Neutral	20
Negative Sources	nobody, fools, idiots	Negative	42 ⁵
Conditionals	if, when, on the condition that	Neutral	17 ⁶
Causal	although, because, whereas	Factive	6
Positive Nouns	argument, contention, idea	Positive	99
Negative Nouns	apostasy, delusion, dream	Negative	122
Positive Adjectives	true, undeniable, trustworthy	Positive	58
Negative Adjectives	false, counter-factual, amiss	Negative	94

Figure 1: These are the VE classes that have been identified. Expression List Size refers to the number of single or multi-word expressions collected for each class.

Veridicality Orientation	Explanation	Can it be changed by a negation or VE?
Positive	Preserves propositional stance.	Yes.
Negative	Inverts propositional stance.	Yes.
Neutral	Indicates writer takes no stance.	No.
Factive	Indicates presupposition of positive stance.	No.
Counter-Factive	Indicates presupposition of negative stance.	No.

Figure 2: Explanation of Veridicality Orientations

like way. For instance, in the sentence “Nobody believes that the sky is green,” “nobody” negates “believes.” The orientation of sources becomes more complicated when a pejorative term is used as the source. If a pejorative source is used unmodified (e.g., “Idiots say that...”) it has a clear negative orientation. However, if it is used with the modifier “even,” it becomes positive (e.g., “Even an idiot...”) What words are pejorative can be dependent on the author. For instance the word “liberal” might be pejorative if written by a conservative but might not when written by a liberal. The author created a small list of pejorative sources, and has supplemented it from the negative polar class described in Wilson, Wiebe, & Hoffmann (2005).

Conditionals are considered to be VEs. For the purposes of the system, arguments of conditionals are considered to be neutral as detecting when an antecedent of a condition is presupposed would require robust semantic inference. Causals commit the writer to a positive stance and are thus factive.

Nouns and adjectives share similar VO properties to their verb counterparts. Identifying factive and counter-factive nouns and adjectives is clearly future work.

Veridicality Transformations

⁵This also includes expressions marked “MPQA-negative” by OpinionFinder Wilson, Wiebe, & Hoffmann (2005).

⁶These are based on discourse connectives used in Hutchinson (2004). The connectives chosen all imply that one argument is contingent on another, allowing the author not to take a veridical stance on either argument. Section 3.1.2 of Polanyi & Zaenen (2005) describes antecedents of conditionals as not being entailed. The au-

thors serve as the means for VEs to interact with propositions and modify their stance. Furthermore, VTs register that modification in a way that allows other VEs in a sentence to modify the proposition.

The system uses the set of hand-written VTs described in Figure 3. The system iterates through the list in the order shown, attempting to match a syntactic pattern (listed in “Syntactic Form”) over the sentence’s dependency parse tree. If the VE class listed isn’t found but the syntactic form matches the sentence, the word in the VT slot is treated as a VE of positive orientation. The VT is then applied, and the process starts over from the beginning of the list. If no matching patterns are found, the proposition’s current stance is returned as the final stance.

Applying a VT to a VE and proposition involves four steps. First, the VE’s contextual orientation is identified. The contextual identification involves determining if the VE is negated, using a modified version of Wilson, Wiebe, & Hoffmann (2005)’s valence shifters. Factives and counter-factives are exempt from this process. Otherwise, negative VEs are made positive and vice versa. In addition to specifying its VE, the VT also defines which nodes will be merged into the query. This is the second step in VT application. The VE is always merged and others may be as well. See the “Nodes to Merge” column for a list of nodes a particular VT merges. Additional nodes to merge are found when identifying contextual orientation, and in the case of noun VEs, their full NPs.

The third step involves finding the stance of the expanded thors go on to give an example where a modal in the consequent leads the reader to take a negative stance toward the antecedent.

proposition. If the existing proposition's stance was given through a factive, counter-factive, or neutral VE, the new proposition inherits it. Furthermore, if the VE's contextual orientation is positive, the existing proposition's stance is maintained. If the VE is contextually negative, the new proposition's stance is inverted.

The fourth step is to merge the assembled nodes into a new proposition.

The process of identifying VEs involved examining dependency parses of the same sample sentences used to form the core VE set.

Some Examples

For each of the examples listed, the semantic diagram produced by the system is displayed. The diagram of the dependency parse tree is read radially outward from the proposition, which is a circumscribed node. The rectangles represent each intermediate proposition and their stance is shown. The final stance is that of the outermost rectangle.

Example 1

This example is entirely positive. As seen in Figure (2), the innermost VT is a Single PP, leaving the proposition positive. The Positive Noun "statement" is the VT's operating VE. The determiner "the" is included in the nominalization of statement, and, as specified in the VT, the term is also included in the new proposition.

Next, the Subject Source VT occurs, based on the source "Bill," which is positive by default. The Verb Complement VT occurs last, so as to not remove the parent node (its VE) that the other VTs relied on.

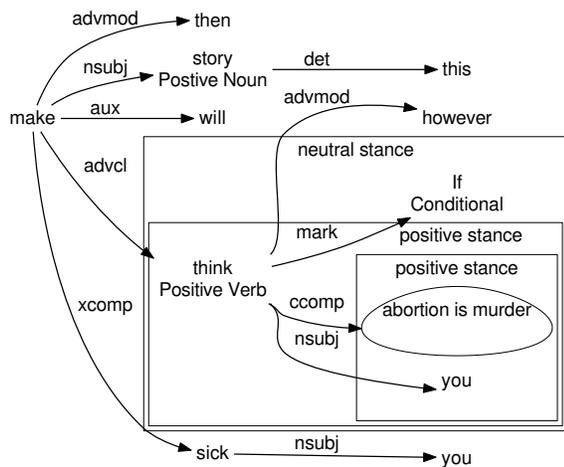


Figure 5: If however you think abortion is murder, then this story will make you sick.

Example 2 The example seen in Figure 5 is of neutral stance. The structure is similar to the one mentioned above, with a nominal source and a complement taking verb. The conditional "If" serves to show that the proposition is neutral

because there are no factive or counter-factive VEs within its scope.

Evaluation

A simple evaluation was performed on this proof-of-concept system.

Search results from Google Blog Search⁷ were used as the data source for the evaluation. Sentences from the first 100 search results using the controversial, propositional query of "abortion is murder" were extracted. Sentences 150 words long or more were thrown out, along with the instances where the sentence contained only the search string (e.g., "Abortion is murder!") Furthermore, unparseable sentence fragments were excluded from the data-set, as well as six sentences that caused parser errors. The sentences excluded were either trivial cases, hopelessly ungrammatical, part of long, automatically generated sentences designed to deceive spam-filters or artificially inflate post ranking, or, in three cases, contained quotation-mark characters that could not be processed. Seven sentences in the corpus were questions, which are not currently addressed. One sentence contained two uses of the proposition; each use was evaluated separately. These issues left the corpus with 48 sentences. The distribution of VOs was 27 positive, 3 negative and 18 neutral.

The system evaluated 39 sentences correctly, giving it a precision of 0.81. Counting rhetorical questions, the system had recall of 0.71. The baseline precision, considering all sentences positive, would be 0.56, giving the system a 25 point improvement over the baseline. The system correctly labeled all VEs and VTs for 31 sentences. While the sample size is very small, this level of performance is promising.

Related Work

The VE-VO relationship that is used corresponds to the relationships described in Sauri, Verhagen, & Pustejovsky (2006). Their relationship classifier was not employed because it is based on a relatively small corpus designed to assess temporal relationships between events and not exclusively involving propositions. Their system misses many simple relations (e.g. direct object ⇒ sibling proposition) that are present in the system presented here.

Event recognition is a rapidly growing field. Systems for general event categorization (Sauri *et al.* 2005) have been developed, as well as for opinion-oriented events (Choi, Breck, & Cardie 2006). Also, event recognition systems related to reported speech events have been developed (Wiebe, Bruce, & Duan 1997).

Nairn, Condoravdi, & Karttunen (2006), described earlier, has a more expansive VO set. The authors define five "implication signatures" of complement taking verbs, indicating their entailment or presupposition under positive and negative relative polarity. The first signature they define "two-way implicative," which is used for verbs whose entailment can be flipped when under negative relative polarity. The second two categories are "one-way implicatives."

⁷<http://blogsearch.google.com>

Pattern Name	Syntactic Form	Nodes to Merge	Example
Double Source	$X \xrightarrow{\text{prep}} _ \xrightarrow{\text{pobj}} _ \xrightarrow{\text{prep}} _ \xrightarrow{\text{pobj}} _ \xrightarrow{\text{poss}} _ (VE: - \text{Source})$ $X \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	P, VE	Sam argues in defense of the idiot 's assertion that P
Double PP	$X \xrightarrow{\text{prep}} _ \xrightarrow{\text{pobj}} _ \xrightarrow{\text{prep}} PREP \xrightarrow{\text{pobj}} (VE: \pm \text{Noun})$ $X \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	$P, VE, PREP$	Sam argues in defense of the assertion that P
Single Source	$X \xrightarrow{\text{prep}} _ \xrightarrow{\text{pobj}} _ \xrightarrow{\text{poss}} _ (VE: - \text{Source})$ $X \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	P, VE	Sam agrees with the idiot 's assertion that P
Single Passive Source	$X \xrightarrow{\text{prep}} \text{by} \xrightarrow{\text{pobj}} _ (VE: \pm \text{Source})$ $X \xrightarrow{\text{auxpass}} _$ $X \xrightarrow{\text{nsubjpass}} _ P$	P, VE, by	It was argued by the idiot that P
Passive	$(VE: \text{Verb}) \xrightarrow{\text{auxpass}} AUX$ $VE \xrightarrow{\text{nsubjpass}} P$	P, VE, AUX	It was argued by the idiot that P
Single PP	$X \xrightarrow{\text{prep}} PREP \xrightarrow{\text{pobj}} _ (VE: \pm \text{Noun})$ $X \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	$P, VE, PREP$	She agreed with the assertion that P
DO Characterization	$X \xrightarrow{\text{dobj}} _ (VE: \pm \text{Noun})$ $X \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	VE, P	Bill uttered the falsehood that P
Subject Source	$P \xrightarrow{\text{nsubj}} _ (VE: \pm \text{Source})$ $X \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	VE, P	The idiot [said] that P
Verb Complement	$(VE: \text{Verb}) \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} P$	VE, P	[He] lied that P
Relative Clause	$(VE: \text{Noun}) \xrightarrow{\text{rmod}} P$	VE, P	The idiot [who] P
Adjective Modification	$(VE: \text{Adjective}) \xrightarrow{\text{cop}} BE$ $VE \xrightarrow{\text{nsubj}} P$	VE, P, BE	It is true that P
Conditional Consequent	$P \xrightarrow{\text{ccomp}} _ \xrightarrow{\text{xcomp}} _ \xrightarrow{\text{advmod}} _ \xrightarrow{\text{advcl}} _ \text{mark} \xrightarrow{\text{advmod}} _ \xrightarrow{\text{advcl}} _ (VE: \text{conditional})$	VE, P	If he is on time, P .
Conditional Consequent	$P \xrightarrow{\text{mark}} _ \xrightarrow{\text{advmod}} _ \xrightarrow{\text{advcl}} _ (VE: \text{conditional})$	VE, P	P as soon as next week.
Conditional Consequent	$X \xrightarrow{\text{mark}} _ \xrightarrow{\text{advmod}} _ \xrightarrow{\text{advcl}} _ P$ $X \xrightarrow{\text{dobj}} _ \xrightarrow{\text{nsubj}} _ \xrightarrow{\text{nsubjpass}} _ (VE: \text{conditional})$	VE, P	P as soon as next week.
Conditional Antecedent	$P \xrightarrow{\text{complm}} _ \xrightarrow{\text{dep}} _ \xrightarrow{\text{mark}} _ (VE: \text{conditional})$	P, VE	If P
Causal	$P \xrightarrow{\text{complm}} _ \xrightarrow{\text{dep}} _ (VE: \text{causal})$	P, VE	While P

Figure 3: Veridicality Transformations. Syntactic forms are expressed over dependency relations. P represents the proposition. $(VE: VE \text{ class})$ is a node of a particular VE class labeled VE . $_$ is used as a placeholder for nodes that are not integrated into the proposition. The patterns appear in sequence.

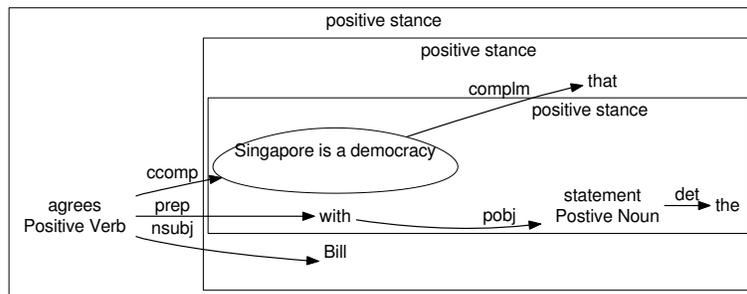


Figure 4: Bill agrees with the statement that Singapore is a democracy.

These are verbs that only have an entailment when under positive relative polarity (“one-way +implicatives”) or that only entail under negative relative polarity (“one-way +implicatives.”) A fourth category includes factives and counterfactuals, that act like “two-way implicatives” except they presuppose instead of entailing. The fifth category, “neutral” covers verbs with no entailment or presupposition. de Paiva *et al.* (2007) describes an extension of this research, reporting that certain nominal deverbals have been given implication signatures based on the transformation process described in Gurevich *et al.* (2006). It should be noted that no direct evaluation of either system was provided.

In Somasundaran *et al.* (2007), the authors used statistical techniques borrowed from subjectivity recognition to detect the presence of arguing in a sentence. It appears their definition of arguing corresponds to the presence of non-neutral veridicality. In their related work section, they make a clear distinction between the task of arguing detection and subjectivity/sentiment detection. This paper is a logical extension their research, unraveling the structure of arguing as opposed to detecting its presence.

Bethard *et al.* (2004) presents a system to identify opinion-bearing propositions and connect them to their sources. As stated in Somasundaran *et al.* (2007), this paper’s task is different from subjectivity recognition.

Wilson, Wiebe, & Hoffmann (2005) gives a system that computes the contextual polarity of subjective expressions. This polarity differs from my definition of veridicality orientation. For instance “sad” is of negative polarity but is of factive orientation because it preserves the truth value of what it modifies. However, I do employ the concept of negative polarity to detect when source position is used as a negative VE.

Polanyi & Zaenen (2005) proposes a method for identifying writers’ and opinion holders’ attitudes regarding discourse entities in a similar manner to the one presented in this paper, describing a framework for understanding how evaluative terms, modals, negatives and intensifiers can interact with each other to label positive and negative attitudes and assign them an integral value.

Future Work

Integrating information from the TimeBank corpus (Pustejovsky *et al.* 2003) would enable the system to take into account temporal information as well as make building a statistical model easier. TimeBank also includes epistemic modality annotations, which could enable the system to represent degrees of possibility.

Fauconnier (1985) motivates a number of new ways of further exploring the problem. At the most basic, taking into account temporal and locative event scope can help get past the agree/disagree/neutral ternary by establishing a framework for giving the veridicality conditions. Also, by having a framework for understanding references to mental spaces, we can identify expressions which refer to other’s stances on a proposition and find the veridicality of those.

The system could benefit from an expansion of its semantic orientations. Nairn, Condoravdi, & Karttunen (2006) describes the “implication signatures” of complement taking verbs that are half-negative and half-neutral or half-positive, half-neutral. For instance, a verb such as “attempt to” is uncertain in its non-negated form. However, when it is negated, it becomes negative.

Questions are not addressed, which make up a sizable portion of sentences containing controversial propositions. Most importantly, more testing needs to be carried out on larger corpora.

Conclusion

I have defined a framework for reasoning about how writers express their beliefs about propositions. The framework involves using hand-written patterns over a dependency parse relating veridicality-bearing elements to propositions. The system determines how these elements nest and, based on their nesting, returns the writer’s stance. On a pilot study over a small set of declarative sentences extracted from blogs, a precision of 0.81 and recall of were achieved.

Acknowledgements

I would like to thank Mike Gasser, Jay Powell, Shakila Shayan, Kam Woods, and the anonymous reviewers for their helpful feedback.

References

- Barzilay, R., and Lee, L. 2003. Learning to paraphrase: an unsupervised approach using multiple-sequence alignment. In *NAACL '03*, 16–23. Morristown, NJ, USA: Association for Computational Linguistics.
- Bethard, S.; Yu, H.; Thornton, A.; Hatzivassiloglou, V.; and Jurafsky, D. 2004. Automatic extraction of opinion propositions and their holders. In *Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*.
- Brin, S. 1998. Extracting patterns and relations from the world wide web. In *WebDB Workshop at 6th International Conference on Extending Database Technology, EDBT'98*.
- Choi, Y.; Breck, E.; and Cardie, C. 2006. Joint extraction of entities and relations for opinion recognition. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, 431–439. Sydney, Australia: Association for Computational Linguistics.
- de Marneffe, M.-C.; MacCartney, B.; and Manning, C. D. 2006. Generating typed dependency parses from phrase structure parses. In *LREC 2006*.
- de Paiva, V.; Bobrow, D. G.; Condoravdi, C.; Crouch, R.; Karttunen, L.; King, T. H.; Nairn, R.; and Zaenen, A. 2007. Textual inference logic: Take two. In *Workshop on Contexts and Ontologies, Representation and Reasoning, CONTEXT 2007*.
- Fauconnier, G. 1985. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. MIT Press.
- Giannakidou, A. 2002. Licensing and sensitivity in polarity items: from downward entailment to nonveridicality. In Andronis, M.; Pycha, A.; and Yoshimura, K., eds., *CLS 38: Papers from the 38th Annual Meeting of the Chicago Linguistic Society, Parasession on Polarity and Negation*.
- Gurevich, O.; Crouch, R.; King, T. H.; and de Paiva, V. 2006. Deverbals in knowledge representation. In *FLAIRS '06*.
- Helbig, H., and Glckner, I. 2007. The role of intensional and extensional interpretation in semantic representations - the intensional and preextensional layers in multinet. In herausgegeben von Sharp, B., and Zock, M., eds., *NLPCS '07*, 92–105. INSTICC Press.
- Hutchinson, B. 2004. Acquiring the meaning of discourse markers. In *ACL*, 684–691.
- Karttunen, L. 1971. The logic of english predicate complement constructions. *Indiana University Linguistics Club*.
- Lin, W. H., and Hauptmann, A. 2006. Are these documents written from different perspectives?: a test of different perspectives based on statistical distribution divergence. In *ACL '06*, 1057–1064. Morristown, NJ, USA: Association for Computational Linguistics.
- Liu, S.; Yu, C.; and Meng, W. 2005. Word sense disambiguation in queries. In *CIKM '05: Proceedings of the 14th ACM international conference on Information and knowledge management*, 525–532. New York, NY, USA: ACM.
- Nairn, R.; Condoravdi, C.; and Karttunen, L. 2006. Computing relative polarity for textual inference. In *ICoS-5*, 67–76.
- Polanyi, L., and Zaenen, A. 2005. Contextual valence shifters. In Shanahan, J. G.; Qu, Y.; and Wiebe, J., eds., *Computing Attitude and Affect in Text*. Dordrecht, The Netherlands: Springer Verlag.
- Pustejovsky, J.; Hanks, P.; Sauri, R.; See, A.; Gaizauskas, R.; Setzer, A.; Radev, D.; Sundheim, B.; Day, D.; Ferro, L.; and Lazo, M. 2003. The TIMEBANK corpus. In *Proceedings of Corpus Linguistics 2003*, 647–656.
- Sauri, R.; Knippen, R.; Verhagen, M.; and Pustejovsky, J. 2005. Evita: a robust event recognizer for qa systems. In *HLT '05: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 700–707. Morristown, NJ, USA: Association for Computational Linguistics.
- Sauri, R.; Verhagen, M.; and Pustejovsky, J. 2006. Slinket: A partial modal parser for events. In *Language Resources and Evaluation Conference, LREC*.
- Somasundaran, S.; Wilson, T.; Wiebe, J.; and Stoyanov, V. 2007. Qa with attitude: Exploiting opinion type analysis for improving question answering in on-line discussions and the news. In *International Conference on Weblogs and Social Media*.
- Stoyanov, V., and Cardie, C. 2006. Toward opinion summarization: Linking the sources. In *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, 9–14. Sydney, Australia: Association for Computational Linguistics.
- Thompson, G., and Yiyun, Y. 1991. Evaluation in the reporting verbs used in academic papers. In *Applied Linguistics*, volume 12(4), 365–382.
- Toutanova, K.; Klein, D.; Manning, C. D.; and Singer, Y. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *NAACL '03*, 173–180. Morristown, NJ, USA: Association for Computational Linguistics.
- Wiebe, J., and Mihalcea, R. 2006. Word sense and subjectivity. In *ACL '06*, 1065–1072. Morristown, NJ, USA: Association for Computational Linguistics.
- Wiebe, J.; Bruce, R.; and Duan, L. 1997. Probabilistic Event Categorization. In *RANLP*.
- Wiebe, J. 2002. Instructions for annotating opinions in newspaper articles. Technical Report TR-02-101, Department of Computer Science, University of Pittsburgh, Pittsburgh, PA.
- Wilson, T.; Wiebe, J.; and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP*.