A Uniform Treatment of Pragmatic Inferences in Simple and Complex Utterances and Sequences of Utterances

Daniel Marcu and Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 1A4
{marcu,gh}@cs.toronto.edu

Abstract

Drawing appropriate defeasible inferences has been proven to be one of the most pervasive puzzles of natural language processing and a recurrent problem in pragmatics. This paper provides a theoretical framework, called *stratified logic*, that can accommodate defeasible pragmatic inferences. The framework yields an algorithm that computes the conversational, conventional, scalar, clausal, and normal state implicatures; and the presuppositions that are associated with utterances. The algorithm applies equally to simple and complex utterances and sequences of utterances.

1 Pragmatics and Defeasibility

It is widely acknowledged that a full account of natural language utterances cannot be given in terms of only syntactic or semantic phenomena. For example, Hirschberg (1985) has shown that in order to understand a scalar implicature, one must analyze the conversants' beliefs and intentions. To recognize normal state implicatures one must consider mutual beliefs and plans (Green, 1990). To understand conversational implicatures associated with indirect replies one must consider discourse expectations, discourse plans, and discourse relations (Green, 1992; Green and Carberry, 1994). Some presuppositions are inferrable when certain lexical constructs (factives, aspectuals, etc) or syntactic constructs (cleft and pseudo-cleft sentences) are used. Despite all the complexities that individualize the recognition stage for each of these inferences, *all* of them can be defeated by context, by knowledge, beliefs, or plans of the agents that constitute part of the context, or by other pragmatic rules.

Defeasibility is a notion that is tricky to deal with, and scholars in logics and pragmatics have learned to circumvent it or live with it. The first observers of the phenomenon preferred to keep defeasibility outside the mathematical world. For Frege (1892), Russell (1905), and Quine (1949) "everything exists"; therefore, in their logical systems, it is impossible to formalize the cancellation of the presupposition that definite referents exist (Hirst, 1991; Marcu and Hirst, 1994). We can taxonomize previous approaches to defeasible pragmatic inferences into three categories (we omit here work on defeasibility related to linguistic phenomena such as discourse, anaphora, or speech acts).

1. Most linguistic approaches account for the defeasibility of pragmatic inferences by analyzing them in a context that consists of all or some of the previous utterances, including the current one. Context (Karttunen, 1974; Kay, 1992), procedural rules (Gazdar, 1979; Karttunen and Peters, 1979), lexical and syntactic structure (Weischedel, 1979), intentions (Hirschberg, 1985), or anaphoric constraints (Sandt, 1992; Zeevat, 1992) decide what presuppositions or implicatures are projected as pragmatic inferences for the utterance that is analyzed. The problem with these approaches is that they assign a dual life to pragmatic inferences: in the initial stage, as members of a simple or complex utterance, they are defeasible. However, after that utterance is analyzed, there is no possibility left of cancelling that inference. But it is natural to have implicatures and presuppositions that are inferred and cancelled as a sequence of utterances proceeds: research in conversation repairs (Hirst et al., 1994) abounds in such examples. We address this issue in more detail in section 3.3.

2. One way of accounting for cancellations that occur later in the analyzed text is simply to extend the boundaries within which pragmatic inferences are evaluated, i.e., to look ahead a few utterances. Green (1992) assumes that implicatures are connected to discourse entities and not to utterances, but her approach still does not allow cancellations across discourse units.

3. Another way of allowing pragmatic inferences to be cancelled is to assign them the status of defeasible information. Mercer (1987) formalizes presuppositions in a logical framework that handles

defaults (Reiter, 1980), but this approach is not tractable and it treats natural disjunction as an exclusive-or and implication as logical equivalence.

Computational approaches fail to account for the cancellation of pragmatic inferences: once presuppositions (Weischedel, 1979) or implicatures (Hirschberg, 1985; Green, 1992) are generated, they can never be cancelled. We are not aware of any formalism or computational approach that offers a unified explanation for the cancellability of pragmatic inferences in general, and of no approach that handles cancellations that occur in sequences of utterances.

It is our aim to provide such an approach here. In doing this, we assume the existence, for each type of pragmatic inference, of a set of necessary conditions that must be true in order for that inference to be triggered. Once such a set of conditions is met, the corresponding inference is drawn, but it is assigned a defeasible status. It is the role of context and knowledge of the conversants to "decide" whether that inference will survive or not as a pragmatic inference of the structure. We put no boundaries upon the time when such a cancellation can occur, and we offer a unified explanation for pragmatic inferences that are inferable when simple utterances, complex utterances, or sequences of utterances are considered.

We propose a new formalism, called "stratified logic", that correctly handles the pragmatic inferences, and we start by giving a very brief introduction to the main ideas that underlie it. We give the main steps of the algorithm that is defined on the backbone of stratified logic. We then show how different classes of pragmatic inferences can be captured using this formalism, and how our algorithm computes the expected results for a representative class of pragmatic inferences. The results we report here are obtained using an implementation written in Common Lisp that uses Screamer (Siskind and McAllester, 1993), a macro package that provides nondeterministic constructs.

2 Stratified logic

2.1 Theoretical foundations

We can offer here only a brief overview of stratified logic. The reader is referred to Marcu (1994) for a comprehensive study. Stratified logic supports one type of indefeasible information and two types of defeasible information, namely, infelicitously defeasible and felicitously defeasible. The notion of infelicitously defeasible information is meant to capture inferences that are anomalous to cancel, as in:

(1) * John regrets that Mary came to the party but she did not come.

The notion of felicitously defeasible information is meant to capture the inferences that can be cancelled without any abnormality, as in:

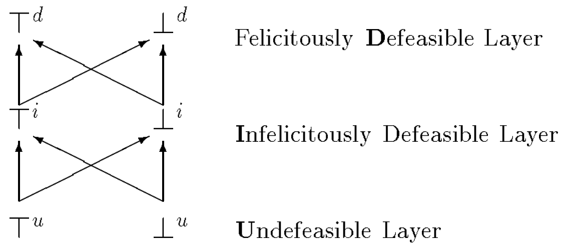

Figure 1: The lattice that underlies stratified logic

(2) John does not regret that Mary came to the party because she did not come.

The lattice in figure 1 underlies the semantics of stratified logic. The lattice depicts the three levels of strength that seem to account for the inferences that pertain to natural language semantics and pragmatics: indefeasible information belongs to the u layer, infelicitously defeasible information belongs to the i layer, and felicitously defeasible information belongs to the d layer. Each layer is partitioned according to its polarity in truth, $\top^u, \top^i, \top^d$, and falsity, $\bot^u, \bot^i, \bot^d$. The lattice shows a partial order that is defined over the different levels of truth. For example, something that is indefeasibly false, $\bot^u$, is stronger (in a sense to be defined below) than something that is infelicitously defeasibly true, $\top^i$, or felicitously defeasibly false, $\bot^d$. Formally, we say that the u level is stronger than the i level, which is stronger than the d level: $u < i < d$. At the syntactic level, we allow atomic formulas to be labelled according to the same underlying lattice. Compound formulas are obtained in the usual way. This will give us formulas such as $regrets^u(John, come(Mary, party)) \rightarrow come^i(Mary, party))$, or $(\forall x)(\neg bachelor^u(x) \rightarrow (male^d(x) \wedge adult^d(x)))$. The satisfaction relation is split according to the three levels of truth into u-satisfaction, i-satisfaction, and d-satisfaction:

Definition 2.1 *Assume σ is an $\mathcal{SL}$ valuation such that $t_i^\sigma = d_i \in \mathcal{D}$ and assume that $\mathcal{SL}$ maps n-ary predicates p to relations $R \subset \mathcal{D} \times \ldots \times \mathcal{D}$. For any atomic formula $p^x(t_1, t_2, \ldots, t_n)$, and any stratified valuation σ, where $x \in \{u, i, d\}$ and t_i are terms, the x-satisfiability relations are defined as follows:*

- $\sigma \models^u p^u(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^u$
- $\sigma \models^u p^i(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^u \cup \overline{R^u} \cup R^i$
- $\sigma \models^u p^d(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^u \cup \overline{R^u} \cup R^i \cup \overline{R^i} \cup R^d$
- $\sigma \models^i p^u(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^i$
- $\sigma \models^i p^i(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^i$
- $\sigma \models^i p^d(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^i \cup \overline{R^i} \cup R^d$
- $\sigma \models^d p^u(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^d$

- $\sigma \models^d p^i(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^d$
- $\sigma \models^d p^d(t_1, \ldots, t_n)$ *iff* $\langle d_1, \ldots, d_n \rangle \in R^d$

Definition 2.1 extends in a natural way to negated and compound formulas. Having a satisfaction definition associated with each level of strength provides a high degree of flexibility. The same theory can be interpreted from a perspective that allows more freedom (u-satisfaction), or from a perspective that is tighter and that signals when some defeasible information has been cancelled (i- and d-satisfaction).

Possible interpretations of a given set of utterances with respect to a knowledge base are computed using an extension of the semantic tableau method. This extension has been proved to be both sound and complete (Marcu, 1994). A partial ordering, $\leq$, determines the set of *optimistic* interpretations for a theory. An interpretation m_0 is preferred to, or is more *optimistic* than, an interpretation m_1 ($m_0 \leq m_1$) if it contains more information and that information can be more easily updated in the future. That means that if an interpretation m_0 makes an utterance true by assigning to a relation R a defeasible status, while another interpretation m_1 makes the same utterance true by assigning the same relation R a stronger status, m_0 will be the preferred or *optimistic* one, because it is as informative as m_1 and it allows more options in the future (R can be defeated).

Pragmatic inferences are triggered by utterances. To differentiate between them and semantic inferences, we introduce a new quantifier, $\forall^{Ut}$, whose semantics is defined such that a pragmatic inference of the form $(\forall^{Ut}\vec{v})(\alpha_1(\vec{v}) \rightarrow \alpha_2(\vec{v}))$ is instantiated only for those objects $\vec{t}$ from the universe of discourse that pertain to an utterance having the form $\alpha_1(\vec{t})$. Hence, only if the antecedent of a pragmatic rule has been uttered can that rule be applied. A meta-logical construct *uttered* applies to the logical translation of utterances. This theory yields the following definition:

Definition 2.2 *Let Φ be a theory described in terms of stratified first-order logic that appropriately formalizes the semantics of lexical items and the necessary conditions that trigger pragmatic inferences. The semantics of lexical terms is formalized using the quantifier $\forall$, while the necessary conditions that pertain to pragmatic inferences are captured using $\forall^{Ut}$. Let $uttered(u)$ be the logical translation of a given utterance or set of utterances. We say that utterance u* <u>pragmatically implicates</u> *p if and only if p^d or p^i is derived using pragmatic inferences in at least one optimistic model of the theory $\Phi \cup uttered(u)$, and if p is not cancelled by any stronger information ($\neg p^u, \neg p^i, \neg p^d$) in any optimistic model schema of the theory. Symmetrically, one can define what a negative pragmatic inference is. In both cases, $\Phi \cup uttered(u)$ is u-consistent.*

2.2 The algorithm

Our algorithm, described in detail by Marcu (1994), takes as input a set of first-order stratified formulas Φ that represents an adequate knowledge base that expresses semantic knowledge and the necessary conditions for triggering pragmatic inferences, and the translation of an utterance or set of utterances $uttered(u)$. The algorithm builds the set of all possible interpretations for a given utterance, using a generalization of the semantic tableau technique. The model-ordering relation filters the *optimistic* interpretations. Among them, the defeasible inferences that have been triggered on pragmatic grounds are checked to see whether or not they are cancelled in any optimistic interpretation. Those that are not cancelled are labelled as pragmatic inferences for the given utterance or set of utterances.

3 A set of examples

We present a set of examples that covers a representative group of pragmatic inferences. In contrast with most other approaches, we provide a consistent methodology for computing these inferences and for determining whether they are cancelled or not for all possible configurations: simple and complex utterances and sequences of utterances.

3.1 Simple pragmatic inferences

3.1.1 Lexical pragmatic inferences

A factive such as the verb *regret* presupposes its complement, but as we have seen, in positive environments, the presupposition is stronger: it is acceptable to defeat a presupposition triggered in a negative environment (2), but is infelicitous to defeat one that belongs to a positive environment (1). Therefore, an appropriate formalization of utterance (3) and the requisite pragmatic knowledge will be as shown in (4).

(3) John does not regret that Mary came to the party.

(4)

$$\left\{\begin{array}{l} uttered(\neg regrets^u(john, \\ \qquad\qquad come(mary, party))) \\ (\forall^{Ut}x, y, z)(regrets^u(x, come(y, z)) \rightarrow \\ \qquad come^i(y, z)) \\ (\forall^{Ut}x, y, z)(\neg regrets^u(x, come(y, z)) \rightarrow \\ \qquad come^d(y, z)) \end{array}\right.$$

The stratified semantic tableau that corresponds to theory (4) is given in figure 2. The tableau yields two model schemata (see figure 3); in both of them, it is defeasibly inferred that *Mary came to the party.* The model-ordering relation $\leq$ establishes m_0 as the optimistic model for the theory because it contains as much information as m_1 and is easier to defeat. Model m_0 explains why *Mary came to the party* is a presupposition for utterance (3).

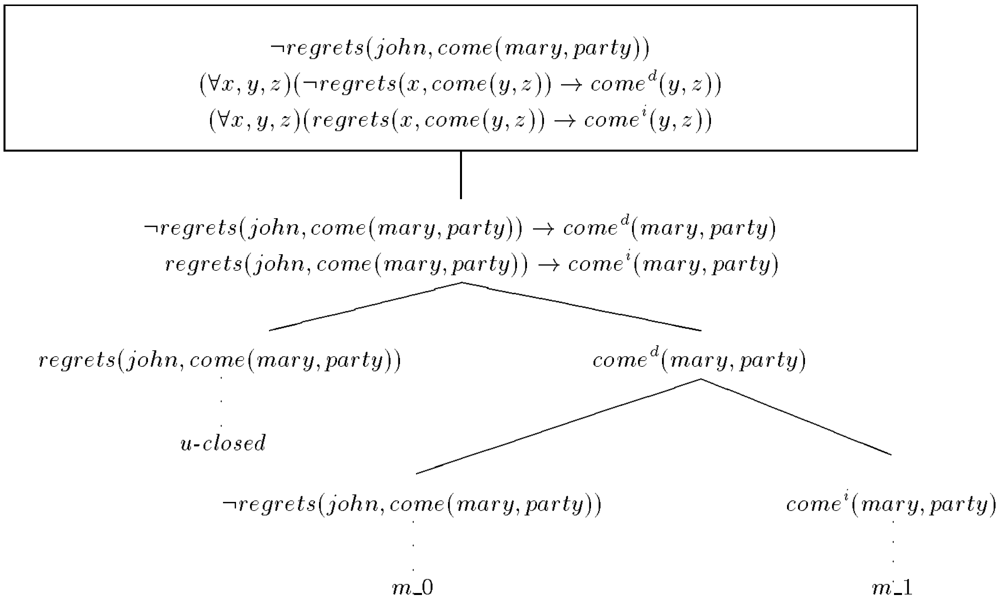

Figure 2: Stratified tableau for *John does not regret that Mary came to the party.*

Schema #	*Indefeasible*	*Infelicitously defeasible*	*Felicitously defeasible*
m_0	$\neg regrets^u(john, come(mary, party)$		$come^d(mary, party)$
m_1	$\neg regrets^u(john, come(mary, party)$	$come^i(mary, party)$	$come^d(mary, party)$

Figure 3: Model schemata for *John does not regret that Mary came to the party.*

Schema #	*Indefeasible*	*Infelicitously defeasible*	*Felicitously defeasible*
m_0	$went^u(some(boys), theatre)$ $\neg went^u(all(boys), theatre)$		 $\neg went^d(most(boys), theatre)$ $\neg went^d(many(boys), theatre)$ $\neg went^d(all(boys), theatre)$

Figure 4: Model schema for *John says that some of the boys went to the theatre.*

Schema #	*Indefeasible*	*Infelicitously defeasible*	*Felicitously defeasible*
m_0	$went^u(some(boys), theatre)$ $went^u(most(boys), theatre)$ $went^u(many(boys), theatre)$ $went^u(all(boys), theatre)$		 $\neg went^d(most(boys), theatre)$ $\neg went^d(many(boys), theatre)$ $\neg went^d(all(boys), theatre)$

Figure 5: Model schema for *John says that some of the boys went to the theatre. In fact all of them went to the theatre.*

3.1.2 Scalar implicatures

Consider utterance (5), and its implicatures (6).

(5) John says that some of the boys went to the theatre.

(6) Not {many/most/all} of the boys went to the theatre.

An appropriate formalization is given in (7), where the second formula captures the defeasible scalar implicatures and the third formula reflects the relevant semantic information for *all*.

(7)

$$\begin{cases} uttered(went(some(boys), theatre)) \\ went^u(some(boys), theatre) \rightarrow \\ \qquad (\neg went^d(many(boys), theatre) \wedge \\ \qquad \neg went^d(most(boys), theatre) \wedge \\ \qquad \neg went^d(all(boys), theatre)) \\ went^u(all(boys), theatre) \rightarrow \\ \qquad (went^u(most(boys), theatre) \wedge \\ \qquad went^u(many(boys), theatre) \wedge \\ \qquad went^u(some(boys), theatre)) \end{cases}$$

The theory provides one optimistic model schema (figure 4) that reflects the expected pragmatic inferences, i.e., *(Not most/Not many/Not all) of the boys went to the theatre.*

3.1.3 Simple cancellation

Assume now, that after a moment of thought, the same person utters:

(8) John says that some of the boys went to the theatre. In fact all of them went to the theatre.

By adding the extra utterance to the initial theory (7), $uttered(went(all(boys), theatre))$, one would obtain one optimistic model schema in which the conventional implicatures have been cancelled (see figure 5).

3.2 Complex utterances

The Achilles heel for most theories of presupposition has been their vulnerability to the projection problem. Our solution for the projection problem does not differ from a solution for individual utterances. Consider the following utterances and some of their associated presuppositions (11) (the symbol $\triangleright$ precedes an inference drawn on pragmatic grounds):

(9) Either Chris is not a bachelor or he regrets that Mary came to the party.

(10) Chris is a bachelor or a spinster.

(11) $\triangleright$ Chris is a (male) adult.

Chris is not a bachelor presupposes that *Chris is a male adult*; *Chris regrets that Mary came to the party* presupposes that *Mary came to the party.* There is no contradiction between these two presuppositions, so one would expect a conversant to infer both of them if she hears an utterance such as (9). However, when one examines utterance (10), one observes immediately that there is a contradiction between the presuppositions carried by the individual components. Being a bachelor presupposes that *Chris is a male*, while being a spinster presupposes that *Chris is a female*. Normally, we would expect a conversant to notice this contradiction and to drop each of these elementary presuppositions when she interprets (10).

We now study how stratified logic and the model-ordering relation capture one's intuitions.

3.2.1 *Or* — non-cancellation

An appropriate formalization for utterance (9) and the necessary semantic and pragmatic knowledge is given in (12).

(12)

$$\begin{cases} uttered(\neg bachelor(Chris) \vee \\ \qquad regret(Chris, come(Mary, party))) \\ (\neg bachelor^u(Chris) \vee \\ \quad regret^u(Chris, come(Mary, party))) \rightarrow \\ \quad \neg(\neg bachelor^d(Chris) \wedge \\ \quad regret^d(Chris, come(Mary, party))) \\ \neg male(Mary) \\ (\forall x)(bachelor^u(x) \rightarrow \\ \quad male^u(x) \wedge adult^u(x) \wedge \neg married^u(x)) \\ (\forall^{Ut} x)(\neg bachelor^u(x) \rightarrow married^i(x)) \\ (\forall^{Ut} x)(\neg bachelor^u(x) \rightarrow adult^d(x)) \\ (\forall^{Ut} x)(\neg bachelor^u(x) \rightarrow male^d(x)) \\ (\forall^{Ut} x, y, z)(\neg regret^u(x, come(y, z)) \rightarrow \\ \qquad come^d(y, z)) \\ (\forall^{Ut} x, y, z)(regret^u(x, come(y, z)) \rightarrow \\ \qquad come^i(y, z)) \end{cases}$$

Besides the translation of the utterance, the initial theory contains a formalization of the defeasible implicature that natural disjunction is used as an exclusive *or*, the knowledge that *Mary* is not a name for males, the lexical semantics for the word *bachelor*, and the lexical pragmatics for *bachelor* and *regret*. The stratified semantic tableau generates 12 model schemata. Only four of them are kept as optimistic models for the utterance. The models yield *Mary came to the party; Chris is a male;* and *Chris is an adult* as pragmatic inferences of utterance (9).

3.2.2 *Or* — cancellation

Consider now utterance (10). The stratified semantic tableau that corresponds to its logical theory yields 16 models, but only *Chris is an adult* satisfies definition 2.2 and is projected as presupposition for the utterance.

3.3 Pragmatic inferences in sequences of utterances

We have already mentioned that speech repairs constitute a good benchmark for studying the generation and cancellation of pragmatic inferences along sequences of utterances (McRoy and Hirst, 1993). Suppose, for example, that Jane has two friends — John Smith and John Pevler — and that her roommate Mary has met only John Smith, a married fellow. Assume now that Jane has a conversation with Mary in which Jane mentions only the name John because she is not aware that Mary does not know about the other John, who is a five-year-old boy. In this context, it is natural for Mary to become confused and to come to wrong conclusions. For example, Mary may reply that *John is not a bachelor*. Although this is true for both Johns, it is more appropriate for the married fellow than for the five-year-old boy. Mary knows that John Smith is a married male, so the utterance makes sense for her. At this point Jane realizes that Mary misunderstands her: all the time Jane was talking about John Pevler, the five-year-old boy. The utterances in (13) constitute a possible answer that Jane may give to Mary in order to clarify the problem.

(13) a. No, John is not a bachelor.
b. I regret that you have misunderstood me.
c. He is only five years old.

The first utterance in the sequence presupposes (14).

(14) $\triangleright$ John is a male adult.

Utterance (13)b warns Mary that is very likely she misunderstood a previous utterance (15). The warning is conveyed by implicature.

(15) $\triangleright$ The hearer misunderstood the speaker.

At this point, the hearer, Mary, starts to believe that one of her previous utterances has been elaborated on a false assumption, but she does not know which one. The third utterance (13)c comes to clarify the issue. It explicitly expresses that John is not an adult. Therefore, it cancels the early presupposition (14):

(16) $\not\triangleright$ John is an adult.

Note that there is a gap of one statement between the generation and the cancellation of this presupposition. The behavior described is mirrored both by our theory and our program.

3.4 Conversational implicatures in indirect replies

The same methodology can be applied to modeling conversational implicatures in indirect replies (Green, 1992). Green's algorithm makes use of discourse expectations, discourse plans, and discourse relations. The following dialog is considered (Green, 1992, p. 68):

(17) Q: Did you go shopping?
A: a. My car's not running.
b. The timing belt broke.
c. (So) I had to take the bus.

Answer (17) conveys a "yes", but a reply consisting only of (17)a would implicate a "no". As Green notices, in previous models of implicatures (Gazdar, 1979; Hirschberg, 1985), processing (17)a will block the implicature generated by (17)c. Green solves the problem by extending the boundaries of the analysis to discourse units. Our approach does not exhibit these constraints. As in the previous example, the one dealing with a sequence of utterances, we obtain a different interpretation after each step. When the question is asked, there is no conversational implicature. Answer (17)a makes the necessary conditions for implicating "no" true, and the implication is computed. Answer (17)b reinforces a previous condition. Answer (17)c makes the preconditions for implicating a "no" false, and the preconditions for implicating a "yes" true. Therefore, the implicature at the end of the dialogue is that the conversant who answered went shopping.

4 Conclusions

Unlike most research in pragmatics that focuses on certain types of presuppositions or implicatures, we provide a global framework in which one can express all these types of pragmatic inferences. Each pragmatic inference is associated with a set of necessary conditions that may trigger that inference. When such a set of conditions is met, that inference is drawn, but it is assigned a defeasible status. An extended definition of satisfaction and a notion of "optimism" with respect to different interpretations yield the preferred interpretations for an utterance or sequences of utterances. These interpretations contain the pragmatic inferences that have not been cancelled by context or conversant's knowledge, plans, or intentions. The formalism yields an algorithm that has been implemented in Common Lisp with Screamer. This algorithm computes uniformly pragmatic inferences that are associated with simple and complex utterances and sequences of utterances, and allows cancellations of pragmatic inferences to occur at any time in the discourse.

Acknowledgements

This research was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada.

References

G. Frege. 1892. Über sinn und bedeutung. *Zeitschrift für Philos. und Philos. Kritik*, 100:373–394. reprinted as: On Sense and Nominatum, In Feigl H. and Sellars W., editors, *Readings in Philosophical Analysis*, pages 85–102, Appleton-Century-Croft, New York, 1947.

G.J.M. Gazdar. 1979. *Pragmatics: Implicature, Presupposition, and Logical Form.* Academic Press.

N. Green and S. Carberry. 1994. A hybrid reasoning model for indirect answers. In *Proceedings 32nd Annual Meeting of the Association for Computational Linguistics*, pages 58–65.

N. Green. 1990. Normal state implicature. In *Proceedings 28th Annual Meeting of the Association for Computational Linguistics*, pages 89–96.

N. Green. 1992. Conversational implicatures in indirect replies. In *Proceedings 30th Annual Meeting of the Association for Computational Linguistics*, pages 64–71.

J.B. Hirschberg. 1985. A theory of scalar implicature. Technical Report MS-CIS-85-56, Department of Computer and Information Science, University of Pennsylvania. Also published by Garland Publishing Inc., 1991.

G. Hirst, S. McRoy, P. Heeman, P. Edmonds, and D. Horton. 1994. Repairing conversational misunderstandings and non-understandings. *Speech Communication*, 15:213–229.

G. Hirst. 1991. Existence assumptions in knowledge representation. *Artificial Intelligence*, 49:199–242.

L. Karttunen and S. Peters. 1979. Conventional implicature. In Oh C.K. and Dinneen D.A, editors, *Syntax and Semantics, Presupposition*, volume 11, pages 1–56. Academic Press.

L. Karttunen. 1974. Presupposition and linguistic context. *Theoretical Linguistics*, 1:3–44.

P. Kay. 1992. The inheritance of presuppositions. *Linguistics & Philosophy*, 15:333–379.

D. Marcu. 1994. A formalism and an algorithm for computing pragmatic inferences and detecting infelicities. Master's thesis, Dept. of Computer Science, University of Toronto, September. Also published as Technical Report CSRI-309, Computer Systems Research Institute, University of Toronto.

D. Marcu and G. Hirst. 1994. An implemented formalism for computing linguistic presuppositions and existential commitments. In H. Bunt, R. Muskens, and G. Rentier, editors, *International Workshop on Computational Semantics*, pages 141–150, December.

S. McRoy and G. Hirst. 1993. Abductive explanation of dialogue misunderstandings. In *Proceedings, 6th Conference of the European Chapter of the Association for Computational Linguistics*, pages 277–286, April.

R.E. Mercer. 1987. *A Default Logic Approach to the Derivation of Natural Language Presuppositions.* Ph.D. thesis, Department of Computer Science, University of British Columbia.

W.V.O. Quine. 1949. Designation and existence. In Feigl H. and Sellars W., editors, *Readings in Philosophical Analysis*, pages 44–51. Appleton-Century-Croft, New York.

R. Reiter. 1980. A logic for default reasoning. *Artificial Intelligence*, 13:81–132.

B. Russell. 1905. On denoting. *Mind n.s.*, 14:479–493. reprinted in: Feigl H. and Sellars W. editors, *Readings in Philosophical Analysis*, pages 103–115. Appleton-Century-Croft, New York, 1949.

R.A. van der Sandt. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9:333–377.

J.M. Siskind and D.A. McAllester. 1993. Screamer: A portable efficient implementation of nondeterministic Common Lisp. Technical Report IRCS-93-03, University of Pennsylvania, Institute for Research in Cognitive Science, July 1.

R.M. Weischedel. 1979. A new semantic computation while parsing: Presupposition and entailment. In Oh C.K. and Dinneen D.A, editors, *Syntax and Semantics, Presupposition*, volume 11, pages 155–182. Academic Press.

H. Zeevat. 1992. Presupposition and accommodation in update semantics. *Journal of Semantics*, 9:379–412.