A formal and computational characterization of pragmatic infelicities

Daniel Marcu and Graeme Hirst[1]

Abstract. We study the logical properties that characterize pragmatic inferences and we show that classical understanding of notions such as entailment and defeasibility is not enough if one wants to explain infelicities that occur when a pragmatic inference is cancelled. We show that infelicities can be detected if a special kind of inference is considered, namely *infelicitously defeasible inference*. We also show how one can use *stratified logic*, a linguistically motivated formalism that accommodates indefeasible, infelicitously defeasible, and felicitously defeasible inferences, to reason about pragmatic inferences and detect infelicities associated with utterances. The formalism yields an algorithm for detecting infelicities, which has been implemented in Lisp.

1 Introduction

Consider the following utterances:

(1) $*$ John bought a new car but I don't believe that he did.

(2) $*$ Jane has three exams this term — in fact, twenty-four.

(3) $*$ Ed went to the airport and flew to California, but he flew to California first.

(4) $*$ John forgot to lock the door, but he did not intend to lock the door.

A linguist will say that utterances (1)–(4) are *infelicitous* and will notice that in all of them a pragmatic inference is infelicitously defeated: utterance (1) tries to cancel a conversational implicature that is triggered by the maxim of quality [6]; utterance (2) tries to cancel a conventional scalar implicature [7, 4]; utterance (3) tries to cancel a conversational implicature that is triggered by the maxim of manner [6]; and (4) tries to cancel a presupposition carried by an implicative verb [8].

The simple fact that cancelling a pragmatic inference can yield an infelicity or anomaly is puzzling, since it is widely acknowledged that most pragmatic inferences are felicitously defeasible. A possible explanation is that Gricean conversational principles do not carry the restriction that most syntactic grammar rules do: they are not prescriptive rules that conversants must obey, but rather merely default rules, or conversational hints. The looseness of these principles made Grice and other researchers in pragmatics consider *all* conversational implicatures to be cancellable. The same agreement has been reached for presuppositions triggered in negative environments: *all* of them are considered defeasible. In contrast, there is no agreement with respect to presuppositions triggered in positive environments: some researchers [9] consider them to be defeasible, others [13] treat them as entailments, and still others [4] consider that presuppositions carried by factives, clefts, and definite descriptions are entailments, but they do not commit themselves to a definitive position for other types of presuppositions.

Most pragmatic inferences *are* generally felicitously defeasible, i.e., it is *not* anomalous to defeat them. For example, utterance (5) cancels an implicature generated by the maxim of quantity; sentence (6) cancels a scalar implicature; and (7) cancels a presupposition generated in a negative environment.

(5) Jane has three exams this term — in fact, four.

(6) John says that some of the boys went to the theatre. John, Mike, and Jeff were there. Fred was there too. In fact *all* of them went to the theatre.

(7) John does not regret that Mary went to the party, because she did not go.

In this paper we argue that the classical understanding of notions such as entailment and defeasibility is not enough if one wants to characterize infelicities that occur when the speaker attempts to cancel certain implicatures and presuppositions. We show how *stratified logic* [14], a linguistically motivated formalism that distinguishes between different levels of defeasibility and different "levels" of satisfiability, can be used for expressing notions such as pragmatic inference and infelicity. We also propose an extension of the implementation described by Marcu and Hirst [15], which takes as input an utterance or a set of utterances expressed in terms of stratified formulae, evaluates them against a knowledge base that contains both semantic and pragmatic knowledge, and detects which utterances are felicitous and which are not.

2 A logical analysis of semantic and pragmatic inferences

We differentiate between two classes of inferences: *semantic inferences* are those that pertain to world knowledge, while

[1] Department of Computer Science, University of Toronto, Toronto, Ontario, Canada, M5S 3G4.

ECAI 96. *12th European Conference on Artificial Intelligence*
Edited by W. Wahlster
Published in 1996 by John Wiley & Sons, Ltd.

pragmatic inferences are those that pertain to knowledge about language use. Some semantic inferences, such as *All men are mortal* and *The sum of the angles in any triangle is 180°*, are indefeasible, and first-order logic provides an appropriate formal tool for modeling them. Other semantic inferences, such as *Birds typically fly* and *Politicians are liars*, are defeasible. Many different approaches have been designed to deal with this type of reasoning: those of Reiter [17], Ginsberg [5], and Kifer and Lozinskii [10] are only a few.

In contrast with semantic inferences, which are triggered by commonsense knowledge, pragmatic inferences are derived from general rules that govern the use of language. They can be lexical in nature, as in the case of a factive [11], syntactic, as in the case of a cleft sentence [1], semantic, as in the case of some conventional implicatures [7, 4], or purely pragmatic, as in the case of the conversational implicatures [6]. None of these inferences are derivable from purely semantic information. Rather, an utterance is necessary in order to bring them to life, and that is the reason we call them pragmatic.

When we study pragmatic inferences, we are faced at first sight with two possibilities: we can formalize pragmatic inferences either as entailments or as defeasible inferences in a classical logical system that handles defaults. We discuss now each of these possibilities in turn and show that they are inadequate for detecting infelicities.

Formalizing inferences that yield infelicities as logical entailments A possible attempt to explain the infelicity of utterances (1)–(4) is by means of inconsistency. For example, let a = *Ed went to the airport*, b = *Ed went to California*, and $and(a,b) \rightarrow before(a,b)$ be the pragmatic inference that is associated with uttering $and(a,b)$.[2] Utterance (3) asserts $and(a,b) \wedge before(b,a)$, and since *before* is antisymmetric – $before(b,a) \rightarrow \neg before(a,b)$ – the theory $\{and(a,b) \wedge before(b,a), and(a,b) \rightarrow before(a,b), before(b,a) \rightarrow \neg before(a,b)\}$ is inconsistent. Therefore, one can say that this is why utterance (3) is infelicitous. The only problem is that if $and(a,b) \rightarrow before(a,b)$ holds, according to the contrapositive law, $\neg before(a,b) \rightarrow \neg and(a,b)$ should hold as well. But it is not difficult to find readings where one can felicitously utter a conjunct such as $and(a,b)$, although a and b are not chronologically ordered (see example (8)). Because the contrapositive does not hold, it is not appropriate to formalize a conversational implicature as a logical implication. Moreover, since any conclusion follows from an inconsistent theory, it will be desirable to provide means for detecting infelicities that don't enforce inconsistencies. From a computational perspective, the detection of an infelicity in an approach based on inconsistency amounts to a disruption of the ability of the system to draw sound inferences.

Formalizing inferences that yield infelicities as defeasible rules Assume that we formalize pragmatic inferences as defeasible rules in a classical system that handles defaults. If we do so, infelicities will no longer be detectible because the information that is asserted or that is derivable through "modus ponens" will always cancel the information that is derived on pragmatic accounts. In other words, cancelling the pragmatic inferences associated with utterances (1)–(4) will be just like cancelling the fact that Tweety flies when we find out that Tweety is an ostrich, or just like the cancellations shown in examples (5)–(7), for example. Infelicities cannot be explained by irresolvable conflicts among defaults either, because as examples (1)–(4) show, infelicities occur because of the clash between inferences derived on pragmatic grounds (that in this paragraph are assumed to be represented as default rules) and indefeasible inferences, i.e., inferences that are asserted directly, or that are derivable through the use of classical entailment.

One can also attempt to detect infelicities by viewing utterances as sequences of communicative acts. In this case, cancelled defaults or inconsistencies will be seen to arise from conflicts among communicative acts that cannot be achieved at the same time. But even if we take this approach, we will still be unable to distinguish between communicative acts that can and cannot be felicitously cancelled (if we interpret pragmatic inferences as defeasible) and unable to distinguish between infelicitous utterances that are infelicitous due to some pragmatic factors (utterances (1)– (4)) and utterances that are inconsistent due to semantic factors, such as utterance "John bought a new car and John did not buy a new car".

In summary, it seems that what we need is a formalism that is able to account both for the pragmatic inferences that pertain to utterances (examples (5)–(7)) and the infelicities that occur when certain pragmatic inferences are cancelled (examples (1)–(4)). We believe that most formalisms that have been proposed in the knowledge representation literature have been designed to address only the first of these requirements. Both justification-based [17, 2, 3] and multivalued-based [5, 10] theories define satisfaction so that defeasible information is *always* cancelled by entailments or so that weaker defeasible information is *always* cancelled by stronger defeasible information. Although such a definition could account for the inferences that are most likely to be drawn from utterances of the kind shown in examples (5)–(7) (see Mercer [16] for an approach designed along these lines that handles presuppositions), it could not explain why defeating some pragmatic inferences is felicitous, while defeating others is not.

In this paper, we argue that a possible way to account for both these facets of cancelability is to use the multiple definition of satisfaction that was provided by Marcu and Hirst [15]. We show now how such a definition could be exploited in order to account for infelicities, without affecting the capability of the formalism to account for a large class of pragmatic inferences that ranges from conventional, conversational, scalar, clausal, and normal state implicatures to presuppositions.

3 Detecting infelicities

3.1 Informal intuitions

Linguistic knowledge and context play a major role in the detection of infelicities. For example, inferences associated with the maxim of manner, like those of sentence (3), can be felicitously defeated in some cases: the pattern of sentence (3) is similar to that in sentence (8), but for the latter, the temporal inference is felicitously defeasible.

(8) Last summer, John visited his parents and worked for Microsoft. In fact, he visited his parents in August and

[2] See [12] for a detailed discussion of the role of discourse structure on temporal relations.

worked in the earlier part of the summer.

The solution that we propose for detecting infelicities relies on formalizing *only* the necessary conditions that make a pragmatic inference infelicitous, and a refinement of traditional defaults that accounts for the different strength or commitment that seems to differentiate certain types of pragmatic inferences. When an implicature or a presupposition cannot be defeated felicitously, we say that that implicature or presupposition is *infelicitously defeasible*. Otherwise, we say that that pragmatic inference is *felicitously defeasible*. Our goal is to partition pragmatic inferences into classes of infelicitously defeasible and felicitously defeasible inferences. In the first class we will find implicatures derived using the maxim of quality and presuppositions that are triggered in positive environments. In the second class we will find implicatures derived using the maxims of quantity and relevance; the clausal, particularized, and floating implicatures; and presuppositions that are triggered in negative environments. Implicatures derived using the maxim of manner (order) are infelicitously defeasible when an enablement relation holds between the actions that are described in a sequence similar to the one given in (3). In summary, the knowledge of an agent can be divided as follows: semantic knowledge, which can be indefeasible or defeasible; and pragmatic knowledge, which can be felicitously defeasible or infelicitously defeasible.

3.2 Representing pragmatic inferences in stratified logic

In stratified logic, atomic formulas are labelled according to three levels of strengths in felicitously defeasible (e.g. $p^d(t_1,\ldots,t_n)$), infelicitously defeasible (e.g. $p^i(t_1,\ldots,t_n)$) or indefeasible (e.g. $p^u(t_1,\ldots,t_n)$). Compound formulas are obtained from atomic formulas by the application of classical logical connectors. The language of stratified logic also contains a meta-logical construct, $uttered(x)$, which takes as argument the logical translation of an utterance; and a quantifier, $\forall^{Ut}$, whose semantics is defined such that it instantiates only objects that belong to sentences that have been uttered.

Pragmatic inferences are formalized as felicitously or infelicitously defeasible according to the pragmatic class that they belong to. For example, an appropriate formalization (9) for the presupposition pertaining to an implicative verb like *forget* assigns different strengths to presuppositions that are associated with positive or negative environments.

(9) $$\left\{\begin{array}{l}(\forall^{Ut}x,y)(forgot(x,y)\rightarrow intended^i(x,y))\\(\forall^{Ut}x,y)(\neg forgot(x,y)\rightarrow intended^d(x,y))\end{array}\right.$$

In cases that are not clear-cut, such as those pertaining to inferences that are derived using the maxim of order, linguistic knowledge or contextual factors are formalized as necessary preconditions that need to hold if an infelicity is to be detected. A simplified representation for infelicities associated with the maxim of order will specify that if two events x_e, y_e constitute the main events $(me(x,x_e)\wedge me(y,y_e))$ in an utterance $and(x,y)$ that conjoins them; and if the occurrence of x_e enables the occurrence of y_e, $enablement(x_e,y_e)$; then one can derive that y_e follows x_e and can associate to this inference an infelicitously defeasible status:

(10) $$\begin{array}{l}(\forall^{Ut}x,y)(and(x,y)\wedge me(x,x_e)\wedge\\ \quad me(y,y_e)\wedge enablement(x_e,y_e)\rightarrow\\ \quad\quad temporal\_sequence^i(x_e,y_e))\end{array}$$

The trichotomy among felicitously defeasible, infelicitously defeasible, and undefeasible strength is crucial for both the semantics of the logic and the definition of satisfaction. Stratified logic defines three ways of satisfying stratified formulas: *u-satisfiability*, $\models^u$; *i-satisfiability*, $\models^i$; and *d-satisfiability*, $\models^d$. On one hand, the satisfiability relation that is associated with the undefeasible layer, $\models^u$, provides a high degree of liberty in satisfying sets of formulas that contain positive and negative information of different strengths; on the other hand, the satisfiability relations that are associated with the i and d layers, $\models^i$ and $\models^d$, are much stricter. They make a set of formulas *i-inconsistent* or *d-inconsistent* more easily, because the definition of satisfaction restricts the number of choices that may be used to construct an interpretation. For example, according to stratified logic, the theory $\{\neg p^u(t_1,\ldots,t_n), p^i(t_1,\ldots,t_n)\}$ is *u-satisfiable* but is neither *i-satisfiable* nor *d-satisfiable*.

An extension of the semantic tableaux method is used for computing interpretations for a given theory. Once a stratified tableau is built for a given theory, one can use the atomic formulas found on each branch that is not *u-closed* to construct model schemata for that theory. 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 (for details, see [14]).

4 Detecting infelicities in stratified logic

Consider the following utterance, its logical translation, and the requisite pragmatic knowledge:

(11) John did not forget to lock the door.
$$\left\{\begin{array}{l}uttered(\neg forgot^u(john,l\_d)))\\(\forall^{Ut}x,y)(forgot^u(x,y)\rightarrow intended^i(x,y))\\(\forall^{Ut}x,y)(\neg forgot^u(x,y)\rightarrow intended^d(x,y))\end{array}\right.$$

The sematic tableau that is built starting from the set of formulas given in (11) yields two model schemata (see figure 1); in both of them, it is defeasibly inferred that *John intended to lock the door*. 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 there are more ways to defeat this information in the future.

Consider now the following utterance:

(12) John did not forget to lock the door. He did not even intend to do it.

An appropriate formalization in stratified logic yields the two model schemata in figure 2. For each of the two schemata, a structure $\mathcal{SL}$ can be found that *u-satisfies* the initial theory. However there is a fundamental difference between the two schemata: m_0 is *i-satisfiable* while m_1 is not. For example,

(13) $\emptyset^u \cup \{\neg forgot^i(john,l\_d), \neg intended^i(john,l\_d)\} \cup \emptyset^d \models^i m_0$,

Schema #	Indefeasible	Infelicitously defeasible	Felicitously defeasible
m_0	$\neg forgot^u(john, l\_d)$		$intended^d(john, l\_d)$
m_1	$\neg forgot^u(john, l\_d)$	$intended^i(john, l\_d)$	$intended^d(john, l\_d)$

Figure 1. The model schemata for the utterance *John did not forget to lock the door.*

Schema #	Indefeasible	Infelicitously defeasible	Felicitously defeasible
m_0	$\neg forgot^u(john, l\_d)$ $\neg intended^u(john, l\_d)$		$intended^d(john, l\_d)$
m_1	$\neg forgot^u(john, l\_d)$ $\neg intended^u(john, l\_d)$	$intended^i(john, l\_d)$	$intended^d(john, l\_d)$

Figure 2. The model schemata for the utterance *John did not forget to lock the door. He did not even intend to do it.*

but there is no structure $\mathcal{SL}$ such that $\mathcal{SL} \models^i m_1$.

But rational agents tend to notice if something goes wrong. A model schema such as m_1 is a good example of this: any agent who notices that an infelicitously defeasible inference is cancelled, as happens in m_1, will treat that model as infelicitous. We will make this idea precise in definition 4.1 below. This intuition gives us a good reason to discard the infelicitous model schemata, if possible. In this case, we are left with model m_0 where some felicitously defeasible information is cancelled. This corresponds entirely to our expectations: the initial utterance is felicitously described by only one model in which a presupposition has been cancelled. Formally, m_0 is *u-satisfiable* but not *d-satisfiable.*

Contrast the above results with those that characterize the following utterance:

(14) $*$ John forgot to lock the door, but he did not intend to lock it.

An appropriate formalization in stratified logic yields only one model schema, which is *u-satisfiable*, but not *i-satisfiable* (see figure 3).

$$(15) \quad \begin{cases} uttered(forgot^u(john, l\_d) \wedge \\ \qquad \neg intended^u(john, l\_d)) \\ (\forall^{Ut} x, y)(forgot^u(x, y) \rightarrow intended^i(x, y)) \\ (\forall^{Ut} x, y)(\neg forgot^u(x, y) \rightarrow intended^d(x, y)) \end{cases}$$

The model we have obtained in stratified logic (see figure 3) is still *u-consistent.* For example,

(16) $\{forgot^u(john, l\_d), \neg intended^u(john, l\_d)\} \cup \emptyset^i \cup \emptyset^d \models^u (15)$

If one analyzes the above examples, one will notice that pragmatic inferences can be associated with defeasible information that is not cancelled in the preferred models that characterize an utterance or sequence of utterances. Infelicitous utterances are those that are not *i-satisfiable.* We now formalize our intuitions:

Definition 4.1 *Let Φ be a theory described in terms of stratified first-order logic that appropriately formalizes the semantics of lexical items and the pragmatics of lexical and syntactic constructs. Let $uttered(u)$ be the logical translation of a given utterance or set of utterances. We say that utterance u is* <u>infelicitous</u> *if and only if there is no stratified structure $\mathcal{SL}$ that i-satisfies $\Phi \cup uttered(u)$.*

Since our approach is intended to provide both the ability to detect infelicities and compute pragmatic inferences, we modify definition 2.2 of Marcu and Hirst [15] so that pragmatic inferences are determined only with respect to the models that are optimistic and felicitous. The modification yields the folllowing definition:

Definition 4.2 *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 was derived using pragmatic inferences in at least one felicitous 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 felicitous optimistic model schema of the theory. Symmetrically, one can define a negative pragmatic inference $(\neg p)$. In both cases, $\Phi \cup uttered(u)$ is u-consistent.*

The new algorithm that is defined on the basis of definitions 4.1 and 4.2 takes as input an utterance or a set of utterances expressed in terms of stratified formulas and evaluates them against a knowledge base that formalizes the semantic and pragmatic knowledge of a conversant. The algorithm computes the model schemata associated with a given utterance; if there is no model to *i-satisfy* the analyzed theory, the utterance is said to be infelicitous. Otherwise, the algorithm

Schema #	*Indefeasible*	*Infelicitously defeasible*	*Felicitously defeasible*
m_0	$forgot^u(john, l\_d)$ $\neg intended^u(john, l\_d)$	$intended^i(john, l\_d)$	$intended^d(john, l\_d)$

Figure 3. The model schema for the utterance *John forgot to lock the door, but he did not intend to lock it.*

computes the set of preferred schemata and the pragmatic inferences according to definition 4.2 (for details, see [15]). The new algorithm has been fully implemented in Common Lisp.

The definitions that account for infelicities and pragmatic inferences are general, i.e., they apply to all types of pragmatic inferences, and they yield the expected results for simple and complex utterances and sequences of utterances that could be expressed in first-order-like languages. Consider the following infelicitous sequence of utterances:

(17) a. I left for school and I forgot to lock the door.
b. In fact, I thought that Mary was still home.
c. I did not intend to lock the door.

Both the formalism and its implementation interpret this sequence of utterances as follows: When the speaker, let's say John, utters (17)a, the preferred interpretation will consist of some indefeasible information — the logical translation of the utterance; and the infelicitously defeasible information derived on pragmatic grounds — *John intended to lock the door*. Utterance (17)b only adds some indefeasible information to this interpretation. It is utterance (17)c that comes to cancel an infelicitously defeasible inference. There will be no interpretation to i-satisfy the sequence of utterances and the infelicity will be detected even though it depends on an utterance that occured earlier in the analyzed text. However, we can continue the interpretation of subsequent utterances because the theory is still u-consistent.

5 Conclusion

We argue that pragmatic infelicities can be detected if one considers a finer-grained taxonomy with respect to pragmatic inferences and a finer-grained definition of satisfiability. The finer-grained taxonomy of pragmatic inferences enables one to distinguish between defeasible inferences that are felicitous to cancel and defeasible inferences that are infelicitous to cancel. The finer-grained definition of satisfiability enables one to detect the infelicities that occur when some infelicitously defeasible inferences are cancelled and to compute the expected pragmatic inferences that pertain to an utterance or a sequence of utterances. We extend on previous work reported by Marcu and Hirst [15] and implement an algorithm that detects the infelicities that are associated with an utterance or a sequence of utterances.

Acknowledgements

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

REFERENCES

[1] J.D. Atlas and S. Levinson, 'It-clefts, informativeness and logical form: radical pragmatics (revised standard version)', in *Radical Pragmatics*, ed., Cole P., 1–61, Academic Press, (1981).

[2] G. Brewka, 'Reasoning about priorities in default logic', in *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 940–945, (1994).

[3] J.P. Delgrande, 'A preference-based approach to default reasoning: Preliminary report', in *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 902–908, (1994).

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

[5] M.L. Ginsberg, 'Multivalued logics: A uniform approach to reasoning in artificial intelligence', *Computational Intelligence*, 4, 265–316, (1988).

[6] H.P. Grice, 'Logic and conversation', in *Syntax and Semantics, Speech Acts*, eds., Cole P. and Morgan J.L., volume 3, 41–58, Academic Press, (1975).

[7] L.R. Horn, *On the Semantic Properties of Logical Operators in English*, Ph.D. dissertation, University of California, Los Angeles, 1972.

[8] L. Karttunen, 'Implicative verbs', *Language*, 47, 340–358, (1971).

[9] L. Karttunen and S. Peters, 'Conventional implicature', in *Syntax and Semantics, Presupposition*, eds., Oh C.K. and Dinneen D.A, volume 11, 1–56, Academic Press, (1979).

[10] M. Kifer and E.L. Lozinskii, 'A logic for reasoning with inconsistency', *Journal of Automated Reasoning*, 9 (2), 179–215, (November 1992).

[11] P. Kiparsky and C. Kiparsky, 'Fact', in *Semantics: an Interdisciplinary Reader in Philosophy, Linguistics and Psychology*, eds., Steinberg D. and Jakobovits L., 345–369, Cambridge University Press, (1971).

[12] A. Lascarides and N. Asher, 'Temporal interpretation, discourse relations, and commonsense entailment', *Linguistics and Philosophy*, 16, 437–493, (1993).

[13] S.C. Levinson, *Pragmatics*, Cambridge University Press, 1983.

[14] D. Marcu, *A Formalism and an Algorithm for Computing Pragmatic Inferences and Detecting Infelicities*, Master's thesis, Dept. of Computer Science, University of Toronto, September 1994. Also published as Technical Report CSRI-309, Computer Systems Research Institute, University of Toronto.

[15] D. Marcu and G. Hirst, 'A uniform treatment of pragmatic inferences in simple and complex utterances and sequences of utterances', in *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, 144–150, Cambridge, Massachussetts, (June 26–30 1995).

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

[17] R. Reiter, 'A logic for default reasoning', *Artificial Intelligence*, 13, 81–132, (1980).